

ANEXO III – Proposta de Preços

RAZÃO SOCIAL: CLEAR TECNOLOGIA DA INFORMAÇÃO S.A		
CNPJ: 30.088.923/0001-08	TELEFONE: (92) 3042-0123 / (11) 94937-3093	
E-MAIL: contratos@clearit.com.br		
ENDEREÇO: R Salvador, 440, Cond. ED. Soberane, Res. e Corp. e Mall Torre Corp. Unidade 1.810, Bairro Adrianopolis - Manaus - AM - CEP: 69.057-040		
BANCO: Bradesco	AGÊNCIA: 00301-1	CONTA CORRENTE: 0032817-0

GRUPO OU LOTE

ITEM	DESCRIÇÃO	UNIDADE	QUANTIDADE	VALOR UNITÁRIO	VALOR TOTAL (R\$)
01	Módulo de Computação para Nuvem Privada – Tipo 01	und.	02	R\$ 672.019,00	R\$ 1.344.038,00
02	Módulo de Computação para Nuvem Privada – Tipo 02	und.	06	R\$ 716.188,00	R\$ 4.297.128,00
03	Módulo de Armazenamento para Nuvem Privada	und.	04	R\$ 438.145,00	R\$ 1.752.580,00
04	Módulo de Computação para Inteligência Artificial	und.	02	R\$ 884.386,00	R\$ 1.768.772,00
05	Licença de Plataforma Computacional Ultimate (NCI)	und.	352	R\$ 15.775,00	R\$ 5.552.800,00
06	Licença de Plataforma Orquestração Avançada Central (NCM)	und.	288	R\$ 1.878,00	R\$ 540.864,00
07	Licença de Plataforma de Dados Heterogêneo (Mixed Mode/TiB)	und.	100	R\$ 9.884,00	R\$ 988.400,00

● MANAUS

Rua Salvador, 440 - sala 1810 -
Ed. Soberane - Torre Corporate
Adrianópolis - CEP 69057-040 -
Manaus - AM
+55 92 3042-0123

● BRASÍLIA

Centro Empresarial Parque Brasília
SIG QUADRA 1985,
Salas 256 e 257, Asa Sul
CEP 70610-410
+55 (61) 3020-2383

● SÃO PAULO

Rua Mergenthaler, 625 - Sala 61 e 71
Vila Leopoldina -
CEP 05311-030 -
São Paulo - SP
+55 11 4673-4999

08	Licença de Plataforma de Dados Dedicada (<i>Dedicated Mode/TiB</i>)	und.	264	R\$ 3.713,00	R\$ 980.232,00
09	Renovação de suporte de manutenção e licença para módulos computacionais de nuvem privada / NX-8035	und.	06	R\$ 1.088.463,00	R\$ 6.530.778,00
10	Renovação de suporte de manutenção e licença para módulos computacionais de nuvem privada / NX-8155	und.	06	R\$ 536.516,00	R\$ 3.219.096,00
11	Renovação de suporte de manutenção para módulos armazenamento de nuvem privada / NX-8155	und.	08	R\$ 58.940,00	R\$ 471.520,00
12	Renovação de Suporte Comutadores de Datacenter	und.	05	R\$ 89.343,00	R\$ 446.715,00
13	Serviço Técnico Especializado	UST	200	R\$ 520,00	R\$ 104.000,00
VALOR TOTAL (R\$) : R\$ 27.996.923,00 (vinte e sete milhões, novecentos e noventa e seis mil e novecentos e vinte e três reais)					

Validade: 60 (sessenta) dias.

● **MANAUS**

Rua Salvador, 440 - sala 1810 -
Ed. Soberane - Torre Corporate
Adrianópolis - CEP 69057-040 -
Manaus - AM
+55 92 3042-0123

● **BRASÍLIA**

Centro Empresarial Parque Brasília
SIG QUADRA 1985,
Salas 256 e 257, Asa Sul
CEP 70610-410
+55 (61) 3020-2383

● **SÃO PAULO**

Rua Mergenthaler, 625 - Sala 61 e 71
Vila Leopoldina -
CEP 05311-030 -
São Paulo - SP
+55 11 4673-4999



Declaramos que estão inclusos nos preços supramencionados todos os custos diretos e indiretos, inclusive de embalagens, transportes ou fretes, e ainda os resultantes da incidência de quaisquer tributos, contribuições ou obrigações decorrentes da legislação trabalhista, fiscal e previdenciária a que estiver sujeito.

Declaro que possuo capacidade operacional e técnica para atendimento a todos os requisitos deste Edital e seus anexos.

Manaus – AM, 06 de novembro de 2025.

RICARDO
CESAR
DIAS:0904560
7670

Assinado de forma
digital por RICARDO
CESAR
DIAS:09045607670
Dados: 2025.11.06
11:34:27 -03'00'

CLEAR TECNOLOGIA DA INFORMAÇÃO S.A.

CNPJ/MF sob o nº 30.088.923/0001-08

Ricardo Cesar Dias

CEO

● **MANAUS**

Rua Salvador, 440 - sala 1810 -
Ed. Soberane - Torre Corporate
Adrianópolis - CEP 69057-040 -
Manaus - AM
+55 92 3042-0123

● **BRAÍLIA**

Centro Empresarial Parque Brasília
SIG QUADRA 1985,
Salas 256 e 257, Asa Sul
CEP 70610-410
+55 (61) 3020-2383

● **SÃO PAULO**

Rua Mergenthaler, 625 - Sala 61 e 71
Vila Leopoldina -
CEP 05311-030 -
São Paulo - SP
+55 11 4673-4999



PROPOSTA TÉCNICA

DA GARANTIA E DO SUPORTE TÉCNICO

Conforme prescrito no **Edital do PREGÃO ELETRÔNICO Nº 90047.2025/LEI Nº 14.133/2021** – Tribunal De Justiça do Estado do Amazonas – TJAM nossa oferta contempla a garantia de 36 meses para os produtos apresentados nesta proposta comercial.

- E-mail para abertura de chamados (ClearIT): pos-venda@clearit.com.br
- Número para abertura de chamados (ClearIT): 0800 990 9000
- Website para abertura de chamados (ClearIT): <http://suporte.clearit.net.br/support/home>

SERVIÇOS CORRELATOS

Declaramos que a oferta dos nossos serviços está em conformidade com os requisitos, premissas e níveis de qualidade descritos para todo escopo de serviços presentes no Termo de Referência do Edital Nº 90047/2025.

ACESSÓRIOS PARA INSTALAÇÃO

Declaramos que todos os acessórios e demais componentes necessários a integração, a adequada instalação física e a implantação lógica e todos os itens garantindo o perfeito funcionamento da nossa solução ofertada com o parque tecnológico da Tribunal De Justiça do Estado do Amazonas – TJAM, estão inclusas em nossa proposta.

CABOS E TRANSCEIVERS

Declaramos que todos os cabos de alimentação elétrica, de gerência, de interconexão e de dados, bem como os transceivers, fibras e cabos de conexão em suas respectivas metragens, conforme solicitados no Termo de Referência do Edital, estão inclusos em nossa oferta.

● MANAUS

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● SÃO PAULO

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999



LOCAL E PRAZO DE ENTREGA OU APLICAÇÃO

Declaramos que estamos cientes e em conformidade com os prazos e locais de entrega preconizados no Termo do Edital Pregão Eletrônico N° 90047/2025.

Declaramos que atendemos todos os itens do edital e termo de referência e seus anexos, incluindo instalação, configuração, suporte e garantia ao longo de período de contrato.

DESCRIPTIVO TÉCNICO

ITEM 01 – Módulo de Computação para Nuvem Privada – Tipo 01

Marca: Nutanix

Modelo: NX-8155-G9

Lista de Part Numbers:

PART-NUMBER	DESCRIÇÃO	QUANTIDADE ENTREGUE
PIO-NX8155G9	Sistema	02
P4X-EMR6544Y-SRN6R-MCC	CPU	04
MEM-DR564L-SL02-NI22	Memória	32
HDD-KIT-12TB-EA-NI22	HDD	16
HDS-KIT-NVM7P68T-B-NI22	SSD NVMe	08
AOC-ATG-b2TM-NI22	Controladora de Rede	02
AOC-631102-ADAT-NI22	Controladora de Rede	02

● MANAUS

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● SÃO PAULO

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999



ITEM 02 – Módulo de Computação para Nuvem Privada – Tipo 02

Marca: Nutanix

Modelo: NX-8155-G9

Lista de Part Numbers:

PART-NUMBER	DESCRIÇÃO	QUANTIDADE ENTREGUE
PIO-NX8155G9	Sistema	06
P4X-EMR6544Y-SRN6R-MCC	CPU	12
MEM-DR564L-SL02-NI22	Memória	96
HDS-KIT-NVM7P68T-B-NI22	SSD NVMe	72
AOC-ATG-b2TM-NI22	Controladora de Rede	06
AOC-631102-ADAT-NI22	Controladora de Rede	06

ITEM 03 – Módulo de Armazenamento para Nuvem Privada

Marca: Nutanix

Modelo: NX-8155-G9

Lista de Part Numbers:

PART-NUMBER	DESCRIÇÃO	QUANTIDADE ENTREGUE
PIO-NX8155G9	Sistema	04
P4X-EMR4510-SRN60-LCC	CPU	08
MEM-DR532L-SL02-NI22	Memória	48
HDD-KIT-12TB-EA-NI22	HDD	40
HDS-KIT-NVM3P84T-B-NI22	SSD NVMe	08
AOC-KIT-S3816L-2U-NI22	Controladora de Rede	04

● MANAUS

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● SÃO PAULO

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999

AOC-631102-ADAT-NI22	Controladora de Rede	04

ITEM 04 – Módulo de Computação para Inteligência Artificial

Marca: Nutanix

Modelo: NX-3155-G9

Lista de Part Numbers:

PART-NUMBER	DESCRIÇÃO	QUANTIDADE ENTREGUE
PIO-NX3155G9	Sistema	02
P4X-EMR6542Y-SRN6P-MCC	CPU	04
MEM-DR564L-SL02-NI22	Memória	32
HDS-KIT-NVM7P68T-B-NI22	SSD NVMe	12
AOC-ATG-b2TM-NI22	Controladora de Rede	02
AOC-631102-ADAT-NI22	Controladora de Rede	02
GPU	GPU-KIT-L40S-48GBA-NI22	04

● MANAUS

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● SÃO PAULO

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999



ITEM 05 – Licença de Plataforma Computacional Ultimate (NCI)

Marca: Nutanix

Modelo: SW-NCI-ULT-PR

Lista de Part Numbers:

PART-NUMBER	DESCRIÇÃO	QUANTIDADE ENTREGUE
SW-NCI-ULT-PR	Subscription, Nutanix Cloud Infrastructure (NCI) Ultimate Software License & Production Software Support Service for 1 CPU Core	352
SW-NCM-STR-PR	Subscription, Nutanix Cloud Manager (NCM) Starter Software License & Production Software Support Service for 1 CPU Core	352
S-HW-PRD	24/7 Production Level HW Support for Nutanix HCI appliance	10

ITEM 06 – Licença de Plataforma Orquestração Avançada Central (NCM)

Marca: Nutanix

Modelo: SW-NCM-STR-PR

Lista de Part Numbers:

PART-NUMBER	DESCRIÇÃO	QUANTIDADE ENTREGUE
SW-NCM-STR-PR	Subscription, Nutanix Cloud Manager (NCM) Starter Software License & Production Software Support Service for 1 CPU Core	288

● MANAUS

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● SÃO PAULO

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999



ITEM 07 – Licença de Plataforma de Dados Heterogêneo (Mixed Mode/TiB)

Marca: Nutanix

Modelo: SW-NUS-PRO-PR

Lista de Part Numbers:

PART-NUMBER	DESCRIÇÃO	QUANTIDADE ENTREGUE
SW-NUS-PRO-PR	Subscription, Nutanix Unified Storage (NUS) Pro Software License & Production Software Support Service for 1 TiB of data stored	100

ITEM 08 – Licença de Plataforma de Dados Dedicada (Dedicated Mode/TiB)

Marca: Nutanix

Modelo: SW-NUS-STR-PR

Lista de Part Numbers:

PART-NUMBER	DESCRIÇÃO	QUANTIDADE ENTREGUE
SW-NUS-STR-PR	Subscription, Nutanix Unified Storage (NUS) Starter Software License & Production Software Support Service for 1 TiB of data stored	264
S-HW-PRD	24/7 Production Level HW Support for Nutanix HCI appliance	04

● MANAUS

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● SÃO PAULO

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999



ITEM 09 – Renovação de suporte de manutenção e licença para módulos computacionais de nuvem privada / NX-8035

Marca: Nutanix

Modelo: RS-HW-PRD-MY

Lista de Part Numbers:

PART-NUMBER	DESCRIÇÃO	QUANTIDADE ENTREGUE
RS-HW-PRD-MY	Multi Year Production Support renewal on hardware S/N: 22SH3G340400,22SH3G340401,22SH3G340403,22SH3G340414,22SH3G340415,22SH3G340417 - Start Date 2025-09-11 - End Date 2028-09-10	06
SW-NCI-ULT-PR	Subscription, Nutanix Cloud Infrastructure (NCI) Ultimate Software License & Production Software Support Service	06
SW-NCM-STR-PR	Subscription, Nutanix Cloud Manager (NCM) Starter Software License & Production Software Support Service	06

ITEM 10 – Renovação de suporte de manutenção e licença para módulos computacionais de nuvem privada / NX-8155

Marca: Nutanix

Modelo: RS-HW-PRD-MY

Lista de Part Numbers:

PART-NUMBER	DESCRIÇÃO	QUANTIDADE ENTREGUE
RS-HW-PRD-MY	Multi Year Production Support renewal on hardware S/N: 22SH5H350050,22SH5H350065,22SH5H350066,22SH5H350082,22SH5H350044,22SH5H350049 - Start Date 2025-09-11 - End Date	06

● **MANAUS**

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● **SÃO PAULO**

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999



	2028-09-10	
SW-NCI-ULT-PR	Subscription, Nutanix Cloud Infrastructure (NCI) Ultimate Software License & Production Software Support Service	06
SW-NCM-STR-PR	Subscription, Nutanix Cloud Manager (NCM) Starter Software License & Production Software Support Service	06

ITEM 11 – Renovação de suporte de manutenção para módulos armazenamento de nuvem privada / NX8155

Marca: Nutanix

Modelo: RS-HW-PRD-MY

Lista de Part Numbers:

PART-NUMBER	DESCRIÇÃO	QUANTIDADE ENTREGUE
RS-HW-PRD-MY	Multi Year Production Support renewal on hardware S/N: 23SH5H050734,23SH5H050735,23SH5H050692,23SH5H050721,23SH5H050726,23SH5H050727,23SH5H050730,23SH5H050728 - Start Date 2026-05-03 - End Date 2029-05-02	08
RSWA-NUS-SEC-PR	Subscription Renewal, Nutanix Unified Storage (NUS) Security add-on Software License & Production Software Support Service for 1 TiB of data stored	264
RSWA-NUS-ADR-PR	Subscription Renewal, Nutanix Unified Storage (NUS) Advanced Replication add-on Software License & Production Software Support Service for 1 TiB of data stored	264

● **MANAUS**

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● **SÃO PAULO**

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999



ITEM 12 – Renovação de Suporte Computadores de Datacenter

Marca: Arista

Modelo: DCS-7050SX3-48YC8-R

Lista de Part Numbers:

PART-NUMBER	DESCRIÇÃO	QUANTIDADE ENTREGUE
VC-7050SX3-48YC8- 1MNB	1Month A-Care Software & NBD Hardware Replacement/Same Day Ship for 7050SX3-48YC8 05/12/28 JMX2302A2LM	40
SVC-7050SX3- 48YC8-1MNB	1 Month A-Care Software & NBD Hardware Replacement/Same Day Ship for 7050SX3-48YC8 05/12/28 SGD2176125	40
SVC-7050SX3- 48YC8-1MNB	1 Month A-Care Software & NBD Hardware Replacement/Same Day Ship for 7050SX3-48YC8 05/12/28 SGD21330513	40
SVC-7050SX3-48Y	1 Month A-Care Software & NBD Hardware Replacement/Same Day Ship for 7050SX3-48YC8 05/12/28 SGD2133055	40
SVC-REINS-FIX	A-Care Service Reinstatement fee for Arista 1RUSwitches	1

● MANAUS

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● SÃO PAULO

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999



ITEM 13 – Serviço Técnico Especializado

Marca: ClearIT

Modelo: Serviço Técnico Especializado ClearIT

Lista de Part Numbers:

PART-NUMBER	DESCRIÇÃO	QUANTIDADE ENTREGUE
EA	Serviço Técnico Especializado	200

COMPROVAÇÃO TÉCNICA

ITENS	DOCUMENTAÇÃO	COMPROVAÇÃO
1.3.11. MÓDULO DE COMPUTAÇÃO PARA NUVEM PRIVADA - TIPO 1 (ITEM 1)		
1.3.11.1. Serão aceitas apenas módulos computacionais de nuvem hiperconvergente, ou seja, com recursos de computação, armazenamento e rede totalmente integrados fim a fim;	https://www.nutanix.com/products/hardware-are-platforms/specsheet?platformProvider=Nutanix	Hardware Platform Spec Sheets
1.3.11.2. Serão aceitos somente produtos homologados para uso com a solução de SDS (software defined storage) ofertada;	Ciente e de acordo Conforme Proposta Técnica e Comercial	N/A
1.3.11.3. Será aceita a entrega de equipamentos ou recursos extras para suprir com os recursos mínimos da solução, caso seja necessário e desde que todos os requisitos sejam atendidos;	Ciente e de acordo Conforme Proposta Técnica e Comercial	N/A
1.3.11.4. Cada chassi deverá possuir ventiladores e fontes redundantes para a operação adequada dos equipamentos;	https://portal.nutanix.com/page/documents/details?targetId=System-Specs-NX8155G9:System%20Specifications	N/A
1.3.11.5. Cada fonte de alimentação deve vir acompanhada do seu respectivo cabo elétricos e com conectores padrão C13/C-14;	Ciente e de acordo Conforme Proposta Técnica e Comercial	N/A

● MANAUS

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● SÃO PAULO

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999

1.3.11.6. Pertencer à família de produto Nutanix "NX" de forma a garantir compatibilidade e expansão das zonas de disponibilidade já existente, conforme o endereço de Internet https://www.nutanix.com/products/hardwareplatforms/specsheet?platformProvider=Nutanix ;	https://www.nutanix.com/products/hardware-platforms/specsheet?platformProvider=Nutanix	N/A
1.3.11.7. Ser entregue com no mínimo 02 (dois) processadores com arquitetura x86 para alto desempenho, projetado para utilização em servidores, cada um com as seguintes características técnicas mínimas:	Ciente e de acordo Conforme Proposta Técnica e Comercial	N/A
1.3.11.7.1. Ter 16 (dezesesseis) núcleos físicos;	https://www.intel.com.br/content/www/br/pt/products/sku/237569/intel-xeon-gold-6544y-processor-45m-cache-3-60-ghz/specifications.html	Vide especificação
1.3.11.7.2. Possuir frequência de operação nominal de, no mínimo, 3,5 GHz;	https://www.intel.com.br/content/www/br/pt/products/sku/237569/intel-xeon-gold-6544y-processor-45m-cache-3-60-ghz/specifications.html	Vide especificação
1.3.11.7.3. Em caso de atualização de processadores, para uma versão mais atual da fabricante, será admitida sua substituição por modelo equivalente, independentemente das demais métricas, exceto pela quantidade de núcleos existentes e suportados na pastilha;	Ciente e de acordo Conforme Proposta Técnica e Comercial	N/A
1.3.11.7.4. Utilizar como modelo de referência: Intel Xeon-Gold 6544Y.	Ciente e de acordo Conforme Proposta Técnica e Comercial	N/A
1.3.11.8. Possuir, no mínimo, 1024 GB (mil e vinte e quatro gigabytes) de memória RAM instalado, distribuídas em DIMMs DDR5 com 5600MHz RDM, ou superior, com quantidades idênticas entre os processadores;	Ciente e de acordo Conforme Proposta Técnica e Comercial	N/A
1.3.11.9. Conter no mínimo, 02 (duas) interfaces 25/10 GbE SFP28, presentes em uma placa de rede;	Ciente e de acordo Conforme Proposta Técnica e Comercial	N/A
1.3.11.10. Prover interface de gerência fora de banda e vir acompanhada de cabos UTP Cat-6 ou superior de, no mínimo, 05 (cinco) metros, com	Ciente e de acordo Conforme Proposta Técnica e Comercial	N/A

● MANAUS

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● SÃO PAULO

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999

conectores RJ-45, para interconexão a malha de comunicação existente;		
1.3.11.11. Possuir período de vigência do suporte técnico será de 36 (trinta e seis) meses, assegurando a continuidade operacional e a manutenção do desempenho das soluções contratadas;	Ciente e de acordo Conforme Proposta Técnica e Comercial	N/A
1.3.11.12. Fornecer na modalidade PRODUCTION (24x7) da Nutanix, garantindo atendimento técnico ininterrupto, disponível 24 (vinte e quatro) horas por dia, 7 (sete) dias por semana;	Ciente e de acordo Conforme Proposta Técnica e Comercial	N/A
1.3.11.13. Vir como no mínimo as seguintes configurações de armazenamento local:	Ciente e de acordo Conforme Proposta Técnica e Comercial	N/A
1.3.11.13.1. Ter 4 (quatro) discos NVME de, no mínimo, 7,68 TB de volumetria bruta individual;	Ciente e de acordo Conforme Proposta Técnica e Comercial	N/A
1.3.11.13.2. Ter 8 (oito) discos HDD de, no mínimo, 12 TB de volumetria bruta individual;	Ciente e de acordo Conforme Proposta Técnica e Comercial	N/A
1.3.12. MÓDULO DE COMPUTAÇÃO PARA NUVEM PRIVADA – TIPO 02 (ITEM 2)		
1.3.12.1. Serão aceitas apenas módulos computacionais de nuvem hiperconvergente, ou seja, com recursos de computação, armazenamento e rede totalmente integrados fim a fim;	https://www.nutanix.com/products/hardware-platforms/specsheet?platformProvider=Nutanix	N/A
1.3.12.2. Serão aceitos somente produtos homologados para uso com a solução de SDS (software defined storage) ofertada;	https://www.nutanix.com/products/hardware-platforms/specsheet?platformProvider=Nutanix	N/A
1.3.12.3. Será aceita a entrega de equipamentos ou recursos extras para suprir com os recursos mínimos da solução, caso seja necessário e desde que todos os requisitos sejam atendidos;	Ciente e de acordo Conforme Proposta Técnica e Comercial	N/A
1.3.12.4. Cada chassi deverá possuir ventiladores e fontes redundantes para a operação adequada dos equipamentos;	System-Specs-NX8155G9.pdf	Vide especificação
1.3.12.5. Cada fonte de alimentação deve vir acompanhada do seu respectivo cabo elétricos e com conectores padrão C13/C-14;	Ciente e de acordo Conforme Proposta Técnica e Comercial	N/A

● MANAUS

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● SÃO PAULO

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999

1.3.12.6. Pertencer à família de produto Nutanix "NX" de forma a garantir compatibilidade e expansão das zonas de disponibilidade já existente, conforme o endereço de Internet https://www.nutanix.com/products/hardware-platforms/specsheet?platformProvider=Nutanix ;	System-Specs-NX8155G9.pdf	
1.3.12.7. Ser entregue com no mínimo 02 (dois) processadores com arquitetura x86 para alto desempenho, projetado para utilização em servidores, cada um com as seguintes características técnicas mínimas:	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.12.7.1. Ter 16 (dezesesseis) núcleos físicos;	https://www.intel.com.br/content/www/br/pt/products/sku/237569/intel-xeon-gold-6544y-processor-45m-cache-3-60-ghz/specifications.html	Vide especificação
1.3.12.7.2. Possuir frequência de operação nominal de, no mínimo, 3,5 GHz;	https://www.intel.com.br/content/www/br/pt/products/sku/237569/intel-xeon-gold-6544y-processor-45m-cache-3-60-ghz/specifications.html	Vide especificação
1.3.12.7.3. Em caso de atualização de processadores, para uma versão mais atual da fabricante, será admitida sua substituição por modelo equivalente, independentemente das demais métricas, exceto pela quantidade de núcleos existentes e suportados na pastilha;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.12.7.4. Utilizar como modelo de referência: Intel Xeon-Gold 6544Y.	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.12.8. Possuir, no mínimo, 1024 GB (mil e vinte e quatro gigabytes) de memória RAM instalado, distribuídas em DIMMs DDR5 com 5600MHz RDM, ou superior, com quantidades idênticas entre os processadores;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.12.9. Conter no mínimo, 02 (duas) interfaces 25/10 GbE SFP28, presentes em uma placa de rede;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.12.10. Prover interface de gerência fora de banda e vir acompanhada de cabos UTP Cat-6 ou superior de, no mínimo, 05 (cinco) metros, com	Ciente e de acordo Conforme Proposta Técnica e Comercial	

● MANAUS

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● SÃO PAULO

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999

conectores RJ-45, para interconexão a malha de comunicação existente;		
1.3.12.11. Possuir período de vigência do suporte técnico será de 36 (trinta e seis) meses, assegurando a continuidade operacional e a manutenção do desempenho das soluções contratadas;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.12.12. Fornecer na modalidade PRODUCTION (24x7) da Nutanix, garantindo atendimento técnico ininterrupto, disponível 24 (vinte e quatro) horas por dia, 7 (sete) dias por semana;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.12.13. Vir como no mínimo as seguintes configurações de armazenamento local:		
1.3.12.13.1. Ter 12 (doze) discos NVME de, no mínimo, 7,68 TB de volumetria bruta individual ou outra combinação de discos que cumpra os requisitos mínimos descritos anteriormente.	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.13 MÓDULO DE ARMAZENAMENTO PARA NUVEM PRIVADA (ITEM 3)		
1.3.13.1. Serão aceitas apenas módulos de armazenamento de nuvem hiperconvergente, ou seja, com recursos de armazenamento e rede totalmente integrados fim a fim;	https://www.nutanix.com/products/hardware-platforms/specsheet?platformProvider=Nutanix	
1.3.13.2. Serão aceitos somente produtos homologados para uso com a solução de SDS (software defined storage) ofertada;	https://www.nutanix.com/products/hardware-platforms/specsheet?platformProvider=Nutanix	
1.3.13.3. Será aceita a entrega de equipamentos ou recursos extras para suprir com os recursos mínimos da solução, caso seja necessário e desde que todos os requisitos sejam atendidos;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.13.4. Cada chassi deverá possuir ventiladores e fontes redundantes para a operação adequada dos equipamentos;	Platform ANY - NX-3155-G9 System Specifications	
1.3.13.5. Cada fonte de alimentação deve vir acompanhada do seu respectivo cabo elétricos e com conectores padrão C13/C-14;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.13.6. Pertencer à família de produto Nutanix "NX" de forma a garantir compatibilidade e expansão das zonas de disponibilidade já existente, conforme o endereço de Internet https://www.nutanix.com/products/hardware-platforms/specsheet?platformProvider=Nutanix ;	https://www.nutanix.com/products/hardware-platforms/specsheet?platformProvider=Nutanix	

● MANAUS

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● SÃO PAULO

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999

1.3.13.7. Ser entregue com no mínimo 02 (dois) processadores com arquitetura x86 para desempenho mediano, projetado para utilização em servidores, cada um com as seguintes características técnicas mínimas:		
1.3.13.7.1. Ter 12 (doze) núcleos físicos;	https://www.intel.com.br/content/www/br/pt/products/sku/236637/intel-xeon-silver-4510-processor-30m-cache-2-40-ghz/specifications.html	Vide especificação
1.3.13.7.2. Possuir frequência de operação nominal de, no mínimo, 2,0 GHz;	https://www.intel.com.br/content/www/br/pt/products/sku/236637/intel-xeon-silver-4510-processor-30m-cache-2-40-ghz/specifications.html	Vide especificação
1.3.13.7.3. Em caso de atualização de processadores, para uma versão mais atual da fabricante, será admitida sua substituição por modelo equivalente, independentemente das demais métricas, exceto pela quantidade de núcleos existentes e suportados na pastilha;		
1.3.13.7.4. Utilizar como modelo de referência: Intel Xeon-Silver 4510.	Ciente e de acordo Conforme Proposta Técnica e Comercial	N/A
1.3.13.8. Possuir, no mínimo, 384 GB (trezentos e oitenta e quatro) de memória RAM instalado, distribuídas em DIMMs DDR5 com 5600MHz RDM, ou superior, com quantidades idênticas entre os processadores;	Ciente e de acordo Conforme Proposta Técnica e Comercial	N/A
1.3.13.9. Conter no mínimo, 02 (duas) interfaces 25/10 GbE SFP28, presentes em uma placa de rede;	Ciente e de acordo Conforme Proposta Técnica e Comercial	N/A
1.3.13.10. Prover interface de gerência fora de banda e vir acompanhada de cabos UTP Cat-6 ou superior de, no mínimo, 05 (cinco) metros, com conectores RJ-45, para interconexão a malha de comunicação existente;	Ciente e de acordo Conforme Proposta Técnica e Comercial	N/A
1.3.13.11. Possuir período de vigência do suporte técnico será de 36 (trinta e seis) meses, assegurando a continuidade operacional e a manutenção do desempenho das soluções contratadas;	Ciente e de acordo Conforme Proposta Técnica e Comercial	N/A

● MANAUS

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● SÃO PAULO

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999

1.3.13.12. Fornecer na modalidade PRODUCTION (24x7) da Nutanix, garantindo atendimento técnico ininterrupto, disponível 24 (vinte e quatro) horas por dia, 7 (sete) dias por semana;	Ciente e de acordo Conforme Proposta Técnica e Comercial	N/A
1.3.13.13. Vir como no mínimo as seguintes configurações de armazenamento local:		N/A
1.3.13.13.1. Ter 2 (dois) discos NVME de, no mínimo, 3,84 TB de volumetria bruta individual;	Ciente e de acordo Conforme Proposta Técnica e Comercial	N/A
1.3.13.13.2. Ter 10 (dez) discos HDD de, no mínimo, 12 TB de volumetria bruta individual.	Ciente e de acordo Conforme Proposta Técnica e Comercial	N/A
1.3.14 MÓDULO DE COMPUTAÇÃO PARA INTELIGÊNCIA ARTIFICIAL (ITEM 4)		
1.3.14.1. Serão aceitas apenas módulos computacionais de nuvem hiperconvergente, ou seja, com recursos de computação, armazenamento e rede totalmente integrados fim a fim;	https://www.nutanix.com/products/hardware-platforms/specsheet?platformProvider=Nutanix	Vide especificação
1.3.14.2. Serão aceitos somente produtos homologados para uso com a solução de SDS (software defined storage) ofertada;	https://www.nutanix.com/products/hardware-platforms/specsheet?platformProvider=Nutanix	N/A
1.3.14.3. Será aceita a entrega de equipamentos ou recursos extras para suprir com os recursos mínimos da solução, caso seja necessário e desde que todos os requisitos sejam atendidos;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.14.4. Cada chassi deverá possuir ventiladores e fontes redundantes para a operação adequada dos equipamentos;	System-Specs-NX8155G9.pdf	
1.3.14.5. Cada fonte de alimentação deve vir acompanhada do seu respectivo cabo elétricos e com conectores padrão C13/C-14;	Ciente e de acordo Conforme Proposta Técnica e Comercial	N/A
1.3.14.6. Pertencer à família de produto Nutanix "NX" de forma a garantir compatibilidade e expansão das zonas de disponibilidade já existente, conforme o endereço de Internet https://www.nutanix.com/products/hardware-platforms/specsheet?platformProvider=Nutanix ;	https://www.nutanix.com/products/hardware-platforms/specsheet?platformProvider=Nutanix	
1.3.14.7. Ser entregue com no mínimo 02 (dois) processadores com arquitetura x86 para alto desempenho, projetado para utilização em servidores, cada um com as seguintes características técnicas mínimas:	https://www.nutanix.com/products/hardware-platforms/specsheet?platformProvider=Nutanix	

● MANAUS

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● SÃO PAULO

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999

1.3.14.7.1. Ter 24 núcleos físicos;	https://www.intel.com.br/content/www/br/pt/products/sku/237559/intel-xeon-gold-6542y-processor-60m-cache-2-90-ghz/specifications.html	Vide especificação
1.3.14.7.2. Possuir frequência de operação nominal de, no mínimo, 2,5 GHz;	https://www.intel.com.br/content/www/br/pt/products/sku/237559/intel-xeon-gold-6542y-processor-60m-cache-2-90-ghz/specifications.html	Vide especificação
1.3.14.7.3. Em caso de atualização de processadores, para uma versão mais atual da fabricante, será admitida sua substituição por modelo equivalente, independentemente das demais métricas, exceto pela quantidade de núcleos existentes e suportados na pastilha;		
1.3.14.7.4. Utilizar como modelo de referência: Intel Xeon-Gold Intel® Xeon® Gold 6542Y	Ciente e de acordo Conforme Proposta Técnica e Comercial	N/A
1.3.14.8. Possuir, no mínimo, 1024 GB de memória RAM instalado, distribuídas em DIMMs DDR5 com 5600MHz RDM, ou superior, com quantidades idênticas entre os processadores;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.14.9. Conter no mínimo, 02 (duas) interfaces 25/10 GbE SFP28, presentes em uma placa de rede;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.14.10. Prover interface de gerência fora de banda e vir acompanhada de cabos UTP Cat-6 ou superior de, no mínimo, 05 (cinco) metros, com conectores RJ-45, para interconexão a malha de comunicação existente;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.14.11. Possuir período de vigência do suporte técnico será de 36 (trinta e seis) meses, assegurando a continuidade operacional e a manutenção do desempenho das soluções contratadas;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.14.12. Fornecer na modalidade PRODUCTION (24x7) da Nutanix, garantindo atendimento técnico ininterrupto, disponível 24 (vinte e quatro) horas por dia, 7 (sete) dias por semana;	Ciente e de acordo Conforme Proposta Técnica e Comercial	

● MANAUS

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● SÃO PAULO

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999

1.3.14.13. Vir como no mínimo as seguintes configurações de armazenamento local:		
1.3.14.13.1. Ter 6 (seis) discos NVME de, no mínimo, 7,68 TB de volumetria bruta individual.	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.14.14. Possuir no mínimo 02 (duas) GPUs:		
1.3.14.14.1. NVIDIA L40S com 48GB DDR6 com ECC;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.14.14.2. Largura de Banda da Memória de 864GB/s;	https://resources.nvidia.com/en-us-l40s/l40s-datasheet-28413	Technical Specifications
1.3.14.14.3. Núcleos CUDA 18.176;	https://resources.nvidia.com/en-us-l40s/l40s-datasheet-28413	Technical Specifications
1.3.14.14.4. Núcleo FP16 362;	https://resources.nvidia.com/en-us-l40s/l40s-datasheet-28413	Technical Specifications
1.3.14.14.5. TFLOPS do Núcleo Tensor BFLOAT16 de 362;	https://resources.nvidia.com/en-us-l40s/l40s-datasheet-28413	Technical Specifications
1.3.14.14.6. TFLOPS do Núcleo Tensor BFLOAT32 de 183;	https://resources.nvidia.com/en-us-l40s/l40s-datasheet-28413	Technical Specifications
1.3.14.14.7. Interface PCIe Gen4 x16: 64GB/s bidirecional;	https://resources.nvidia.com/en-us-l40s/l40s-datasheet-28413	Technical Specifications
1.3.14.14.8. Serão aceitos outros modelos equivalentes ou superiores, em face a escassez de peça ou atualização tecnológica.	ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.15. LICENCIAMENTO DE PLATAFORMA COMPUTACIONAL ULTIMATE - NCI (ITEM 5)		
1.3.15.1 Ofertar o licenciamento compatível com o ambiente de produção do Tribunal;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.15.2 Utilizar como referência, Nutanix Cloud Infrastructure Ultimate (SW-NCI-ULT-PR) ou licenciamento mais atual e compatível comercializável;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.15.3. Contemplar funcionalidades como:		

● MANAUS

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● SÃO PAULO

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999

1.3.15.3.1. Execução integral do hypervisor Nutanix AHV;	https://portal.nutanix.com/page/documents/solutions/details?targetId=TN-2038-AHV:TN-2038-AHV	
1.3.15.3.2. Gerenciamento centralizado através de console atual;	https://www.nutanixbible.com/3a-book-of-prism-architecture.html	Prism Central Prism is broken down into two main components: Prism Central (PC) Multi-cluster manager responsible for managing multiple Nutanix Clusters to provide a single, centralized management interface. Prism Central is an optional software appliance (VM) which can be deployed in addition to the AOS

● **MANAUS**

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● **SÃO PAULO**

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999



		Cluster (can run on it). 1-to-many cluster manager Prism Element (PE) Localized cluster manager responsibl e for local cluster managem ent and operations . Every Nutanix Cluster has Prism Element built-in. 1-to-1 cluster manager

● **MANAUS**

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● **SÃO PAULO**

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999

1.3.15.3.3. Criação de Snapshots e Clones das máquinas virtuais;

<https://www.nutanixbible.com/4c-book-of-aos-storage.html#snapshots-and-clones>

AOS provides native support for offloaded snapshots and clones which can be leveraged via VAAI, ODX, ncli, REST, Prism, etc. Both the snapshots and clones leverage the redirect-on-write algorithm which is the most effective and efficient. As explained in the Data Structure section above, a virtual machine consists of files (vmdk/vhdx) which are vDisks on the Nutanix

● **MANAUS**

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● **SÃO PAULO**

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999

platform.

A vDisk is composed of extents which are logically contiguous chunks of data, which are stored within extent groups which are physically contiguous data stored as files on the storage devices. When a snapshot or clone is taken, the base vDisk is marked immutable and another vDisk is created as read/write. At this point, both vDisks have the same block map, which is a

● **MANAUS**

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● **SÃO PAULO**

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999



metadata mapping of the vDisk to its corresponding extents. Contrary to traditional approaches which require traversal of the snapshot chain (which can add read latency), each vDisk has its own block map. This eliminates any of the overhead normally seen by large snapshot chain depths and allows you to take continuous snapshots without any performance impact.

● **MANAUS**

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● **SÃO PAULO**

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999

1.3.15.3.4. Tierização das áreas de armazenamento (equipamentos híbridos);	https://www.nutanixbible.com/4c-book-of-aos-storage.html#stargate-io-logic	For hybrid scenarios with flash and HDD, the flash would be Tier 0 with HDD being Tier 1.
1.3.15.3.5. Fator de replicação 2 e 3;	The Nutanix Cloud Bible	The Nutanix platform currently uses a resiliency factor, also known as a replication factor (RF), and checksum to ensure data redundancy and availability in the case of a node or disk failure or corruption . As explained above, the OpLog acts as a staging area to absorb incoming writes onto a low-

● **MANAUS**

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● **SÃO PAULO**

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999



latency SSD tier. Upon being written to the local OpLog, the data is synchronously replicated to another one or two Nutanix CVM's OpLog (dependent on RF) before being acknowledged (Ack) as a successful write to the host. This ensures that the data exists in at least two or three independent locations and is fault tolerant. NOTE: For RF3, a minimum of 5 nodes is required since

● **MANAUS**

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● **SÃO PAULO**

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999

		metadata will be RF5.
1.3.15.3.6. Deduplicação, compressão e erasure coding;	https://www.nutanixbible.com/4c-book-of-aos-storage.html	The following optimizations are leveraged: Erasure Coding (EC-X) Compression Deduplication

● **MANAUS**

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● **SÃO PAULO**

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999

1.3.15.3.7. QoS;

https://portal.nutanix.com/page/documents/details?targetId=Prism-Central-Guide-vpc_2024_2:mul-storage-qos-pc-c.html

Storage QoS provides administrators granular control to manage the performance of virtual machines and ensure that the system delivers consistent performance for all workloads. You can use a controllable knob to limit the IOPS that the storage layer would serve for individual virtual machines. IOPS is the number of requests the storage layer can serve in a second. You can

● **MANAUS**

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● **SÃO PAULO**

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999

		set throttle limits on a VM to prevent noisy VMs from over-utilizing the system resources.
1.3.15.3.8. Replicação assíncrona, de 1 em 1 minuto e de 1 em 1 hora para Nutanix AHV;	https://portal.nutanix.com/page/documents/details?targetId=Disaster-Recovery-DRaaS-Guide-vpc_2024_2:ecd-ecdr-procedure-nearsync-protectionpolicy-pc-c.html	Nearsync Replication (1 Minute to 15 Minutes RPO) Asynchronous Replication (1 Hour or More RPO)

● **MANAUS**

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● **SÃO PAULO**

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999

1.3.15.3.9. Orquestração das técnicas de Recuperação de Desastres;

https://portal.nutanix.com/page/documents/details?targetId=Disaster-Recovery-DRaaS-Guide-vpc_2024_2:Disaster-Recovery-DRaaS-Guide-vpc_2024_2

You can protect your entities (guest VMs, volume groups, and consistency groups) and orchestrate their DR to other Nutanix clusters when events that cause service disruption occur at the primary AZ. For the protection of your entities, protection policies with asynchronous, nearsync, or synchronous replication schedules generate and replicate recovery points to

● **MANAUS**

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● **SÃO PAULO**

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999

		other on-prem AZs. Recovery plans orchestrate DR from the replicated recovery points to other Nutanix clusters at the same or different on-prem AZs.
1.3.15.3.10. Criptografia de dados;	https://www.nutanixbible.com/4b-book-of-aos-security.html	Data Encryption Nutanix provides data encryption via three main options: Native software-based encryption (FIPS-140-

● **MANAUS**

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● **SÃO PAULO**

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999

		2 Level-1) *released in 5.5 Using self- encrypting drives (SED) (FIPS-140- 2 Level-2) Software + hardware encryption
1.3.15.3.11. Suporte a GPUs;	https://www.nutanixbible.com/5a-book-of-ahv-architecture.html	Dynamic GPU managem ent - Supports NVIDIA GPUs to provide specific vGPU profiles based on the VMs running in the system Creating remediati on plans based on the cost of movement to address each of the above case
1.3.15.3.12. Suporte a implementação de clusters de plataformas de containers, similares a Kubernetes;	https://www.nutanix.com/products/cloud-platform/software-options#nkp	Nutanix Kubernete s Platform

● **MANAUS**

Rua Salvador, 440 - Sala 1810
 Ed. Soberane - Torre Corporate
 Adrianópolis - Manaus - AM - 69057-040
 +55 92 3042-0123

● **SÃO PAULO**

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
 Ed. Monções da Gastão
 Vila Leopoldina - São Paulo - SP - 05314-000
 +55 11 4673-4999

1.3.15.3.13. Permitir microssegmentação do ambiente virtualizado;	https://www.nutanixbible.com/12a-book-of-network-services-flow-network-security.html	Supported Configurations The solution applies to the configurations below: Core Use Case(s): Microsegmentation Management interfaces(s): Prism Central (PC) Supported Environment(s): On-Premises: AHV Cloud: Nutanix Cloud Clusters on AWS

● **MANAUS**

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● **SÃO PAULO**

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999

1.3.15.3.14. Gerenciamento centralizado e gerenciamento do ciclo de vida dos componentes dos appliances (softwares e firmwares).	https://www.nutanix.com/content/dam/nutanix/en/resources/datasheets/ds-lcm.pdf	KEY BENEFITS • Easily accessible through Nutanix Prism. • One-click upgrade simplicity. • Upgrade software in a single process. • Automatically manage software and hardware dependencies. • Unifies firmware upgrade operations across multiple hardware vendors. • Included scheduler checks for updates automatically
1.3.15.4. Gerenciamento avançado de Nuvem Privada;		
1.3.15.5. Utilizar como referência, Nutanix Cloud Manager Starter (SW-NCM-STR-PR):	https://www.nutanix.com/products/cloud-platform/software-options#ncm	
1.3.15.6. Contemplar as funcionalidades como:	https://portal.nutanix.com/page/documents/details?targetId=Intelligent-Operations-	

● **MANAUS**

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● **SÃO PAULO**

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999

	Guide-vpc_7_3:mul-intelligent-operations-overview-c.html	
1.3.15.6.1. Emissão de relatório avançados do ambiente;	https://portal.nutanix.com/page/documents/details?targetId=Intelligent-Operations-Guide-vpc_7_3:mul-intelligent-operations-overview-c.html	
1.3.15.6.2. Personalização da console principal;	https://portal.nutanix.com/page/documents/details?targetId=Intelligent-Operations-Guide-vpc_7_3:mul-intelligent-operations-overview-c.html	
1.3.15.6.3. Planejamento e previsão de capacidade	https://portal.nutanix.com/page/documents/details?targetId=Intelligent-Operations-Guide-vpc_7_3:mul-intelligent-operations-overview-c.html	
1.3.15.6.4. Detecção de ineficiência de cargas de trabalho e dimensionamento ideal;	https://portal.nutanix.com/page/documents/details?targetId=Intelligent-Operations-Guide-vpc_7_3:mul-intelligent-operations-overview-c.html	
1.3.15.6.5. Planejamento de capacidade ou cargas de trabalho ao longo do tempo;	https://portal.nutanix.com/page/documents/details?targetId=Intelligent-Operations-Guide-vpc_7_3:mul-intelligent-operations-overview-c.html	
1.3.15.6.6. Permitir automação de operações de gerência do ambiente;	https://portal.nutanix.com/page/documents/details?targetId=Intelligent-Operations-Guide-vpc_7_3:mul-intelligent-operations-overview-c.html	
1.3.15.7. Licenciar cada módulo por unidade de processamento (Cores) de CPU, de acordo com a atual política do fabricante;	https://portal.nutanix.com/page/documents/details?targetId=Intelligent-Operations-Guide-vpc_7_3:mul-intelligent-operations-overview-c.html	
1.3.15.8. Possuir período de vigência do suporte técnico será de 36 (trinta e seis) meses, assegurando a continuidade operacional e a manutenção do desempenho das soluções contratadas;	https://portal.nutanix.com/page/documents/details?targetId=Intelligent-Operations-Guide-vpc_7_3:mul-intelligent-operations-overview-c.html	
1.3.15.9. Fornecer na modalidade PRODUCTION (24x7) da Nutanix, garantindo atendimento técnico ininterrupto, disponível 24 (vinte e quatro) horas por dia, 7 (sete) dias por semana;	https://portal.nutanix.com/page/documents/details?targetId=Intelligent-Operations-Guide-vpc_7_3:mul-intelligent-operations-overview-c.html	
1.3.15.10. Ser compatível com os servidores dos itens 1.3.11 (ITEM 1) e 1.3.12 (ITEM 2).	Ciente e de Acordo	
1.3.16. LICENCIAMENTO DE PLATAFORMA ORQUESTRAÇÃO AVANÇADA CENTRAL - NCM (ITEM 6)		

● MANAUS

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● SÃO PAULO

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999

1.3.16.1. Gerenciamento de orquestração avançada de Nuvem Privada;		
1.3.16.2. Utilizar como referência, Nutanix Cloud Manager Starter (SW-NCM-STR-PR):	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.16.3. Contemplar as funcionalidades como:	Nutanix Licensing and Cloud Platform Software Options Nutanix	
1.3.16.3.1. Emissão de relatório avançados do ambiente;	https://portal.nutanix.com/page/documents/details?targetId=Intelligent-Operations-Guide-vpc_7_3:mul-intelligent-operations-overview-c.html	Intelligent Operations comprises the following features: Analysis: Analyze cluster activity. For more information, see Performance Monitoring. Discovery: Discover applications running in a cluster. For more information, see Application Discovery. Monitoring Configurations: Monitor select applications.

● **MANAUS**

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● **SÃO PAULO**

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999



ns in a cluster. For more information, see Application Monitoring

Planning: Plan for resource capacity and usage. For more information, see Resource Planning.

Playbooks : Create playbooks for task automation. For more information, see Task Automation - Playbooks

Reports: Generate resource and activity reports for Nutanix and non-Nutanix environments. For

● MANAUS

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● SÃO PAULO

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999



more
informatio
n, see
Reports
Managem
ent.

● **MANAUS**

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● **SÃO PAULO**

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999

1.3.16.3.2. Personalização da console principal;

https://portal.nutanix.com/page/documents/details?targetId=Intelligent-Operations-Guide-vpc_7_3:Intelligent-Operations-Guide-vpc_7_3

Intelligent Operations comprises the following features:

Analysis: Analyze cluster activity. For more information, see Performance Monitoring.

Discovery: Discover applications running in a cluster. For more information, see Application

Discovery. Monitoring Configurations:

Monitor select applications in a cluster.

For more information, see Application Monitoring

● **MANAUS**

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● **SÃO PAULO**

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999



Planning: Plan for resource capacity and usage. For more information, see Resource Planning.

Playbooks : Create playbooks for task automation. For more information, see Task Automation - Playbooks

Reports: Generate resource and activity reports for Nutanix and non-Nutanix environments. For more information, see Reports Management.

● **MANAUS**

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● **SÃO PAULO**

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999

1.3.16.3.3. Planejamento e previsão de capacidade

https://portal.nutanix.com/page/documents/details?targetId=Intelligent-Operations-Guide-vpc_7_3:Intelligent-Operations-Guide-vpc_7_4

Intelligent Operations comprises the following features:

Analysis: Analyze cluster activity. For more information, see Performance Monitoring.

Discovery: Discover applications running in a cluster. For more information, see Application Discovery.

Monitoring Configurations: Monitor select applications in a cluster. For more information, see Application Monitoring.

● MANAUS

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● SÃO PAULO

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999



Planning:
Plan for resource capacity and usage. For more information, see Resource Planning.

Playbooks
: Create playbooks for task automation. For more information, see Task Automation - Playbooks

Reports:
Generate resource and activity reports for Nutanix and non-Nutanix environments. For more information, see Reports Management.

● **MANAUS**

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● **SÃO PAULO**

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999

1.3.16.3.4. Detecção de ineficiência de cargas de trabalho e dimensionamento ideal;

https://portal.nutanix.com/page/documents/details?targetId=Intelligent-Operations-Guide-vpc_7_3:Intelligent-Operations-Guide-vpc_7_5

Intelligent Operations comprises the following features:

Analysis: Analyze cluster activity. For more information, see Performance Monitoring.

Discovery: Discover applications running in a cluster. For more information, see Application Discovery.

Monitoring Configurations: Monitor select applications in a cluster. For more information, see Application Monitoring.

● MANAUS

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● SÃO PAULO

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999



Planning:
Plan for resource capacity and usage. For more information, see Resource Planning.

Playbooks
: Create playbooks for task automation. For more information, see Task Automation - Playbooks

Reports:
Generate resource and activity reports for Nutanix and non-Nutanix environments. For more information, see Reports Management.

● **MANAUS**

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● **SÃO PAULO**

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999

1.3.16.3.5. Planejamento de capacidade ou cargas de trabalho ao longo do tempo;

https://portal.nutanix.com/page/documents/details?targetId=Intelligent-Operations-Guide-vpc_7_3:Intelligent-Operations-Guide-vpc_7_6

Intelligent Operations comprises the following features:

Analysis: Analyze cluster activity. For more information, see Performance Monitoring.

Discovery: Discover applications running in a cluster. For more information, see Application

Discovery. Monitoring Configurations:

Monitor select applications in a cluster.

For more information, see Application

Monitoring

● MANAUS

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● SÃO PAULO

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999



Planning:
Plan for resource capacity and usage. For more information, see Resource Planning.

Playbooks
: Create playbooks for task automation. For more information, see Task Automation - Playbooks

Reports:
Generate resource and activity reports for Nutanix and non-Nutanix environments. For more information, see Reports Management.

● **MANAUS**

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● **SÃO PAULO**

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999

1.3.16.3.6. Permitir automação de operações de gerência do ambiente;

https://portal.nutanix.com/page/documents/details?targetId=Intelligent-Operations-Guide-vpc_7_3:Intelligent-Operations-Guide-vpc_7_7

Intelligent Operations comprises the following features:

Analysis: Analyze cluster activity. For more information, see Performance Monitoring.

Discovery: Discover applications running in a cluster. For more information, see Application

Discovery. Monitoring Configurations:

Monitor select applications in a cluster.

For more information, see Application Monitoring

● MANAUS

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● SÃO PAULO

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999



Planning:
Plan for resource capacity and usage. For more information, see Resource Planning.

Playbooks
: Create playbooks for task automation. For more information, see Task Automation - Playbooks

Reports:
Generate resource and activity reports for Nutanix and non-Nutanix environments. For more information, see Reports Management.

● **MANAUS**

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● **SÃO PAULO**

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999

1.3.16.4. Licenciar cada módulo por unidade de processamento (Cores) de CPU, de acordo com a atual político do fabricante;

https://portal.nutanix.com/page/documents/details?targetId=Intelligent-Operations-Guide-vpc_7_3:Intelligent-Operations-Guide-vpc_7_8

Intelligent Operations comprises the following features:

Analysis: Analyze cluster activity. For more information, see Performance Monitoring.

Discovery: Discover applications running in a cluster. For more information, see Application Discovery.

Monitoring Configurations: Monitor select applications in a cluster. For more information, see Application Monitoring.

● MANAUS

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● SÃO PAULO

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999



Planning:
Plan for resource capacity and usage. For more information, see Resource Planning.

Playbooks
: Create playbooks for task automation. For more information, see Task Automation - Playbooks

Reports:
Generate resource and activity reports for Nutanix and non-Nutanix environments. For more information, see Reports Management.

● **MANAUS**

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● **SÃO PAULO**

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999

1.3.16.5. Permitir o uso em cenários de renovação ou criação de novas instâncias;

https://portal.nutanix.com/page/documents/details?targetId=Intelligent-Operations-Guide-vpc_7_3:Intelligent-Operations-Guide-vpc_7_9

Intelligent Operations comprises the following features:

Analysis: Analyze cluster activity. For more information, see Performance Monitoring.

Discovery: Discover applications running in a cluster. For more information, see Application

Discovery. Monitoring Configurations:

Monitor select applications in a cluster.

For more information, see Application Monitoring

● MANAUS

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● SÃO PAULO

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999



Planning:
Plan for resource capacity and usage. For more information, see Resource Planning.

Playbooks
: Create playbooks for task automation. For more information, see Task Automation - Playbooks

Reports:
Generate resource and activity reports for Nutanix and non-Nutanix environments. For more information, see Reports Management.

● **MANAUS**

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● **SÃO PAULO**

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999

1.3.16.6. Possuir período de vigência do suporte técnico será de 36 (trinta e seis) meses, assegurando a continuidade operacional e a manutenção do desempenho das soluções contratadas;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.16.7. Fornecer na modalidade PRODUCTION (24x7) da Nutanix, garantindo atendimento técnico ininterrupto, disponível 24 (vinte e quatro) horas por dia, 7 (sete) dias por semana;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.16.8. Ser do tipo liquidação logo no início e comprovar a cobertura total do período descrito;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.17. LICENCIAMENTO DE PLATAFORMA DE DADOS HETEREGÊNEO - MIXED MODE/TiB (ITEM 7)		
1.3.17.1. Ser fornecido o licenciamento necessário para a implementação da plataforma/software de armazenamento de dados em clusters Nutanix, atendendo à especificação SW-NUS-PRO-PR (Nutanix Unified Storage – PRO), conforme https://www.nutanix.com/products/cloud-platform/software-options#nus ;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.17.2. Possuir licenciamento na modalidade Por Tebibyte Líquido (em base 2), correspondendo aproximadamente 1 (um) terabyte aplicado a um cluster dedicado ou já implementado;	ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.17.3. Contemplar as seguintes funcionalidades no licenciamento proposto:	https://www.nutanix.com/products/cloud-platform/software-options#nus	Vide: Nutanix Unified Storage Editions:
1.3.17.3.1. Erasure Coding	https://www.nutanix.com/products/cloud-platform/software-options#nus	Vide: Nutanix Unified Storage Editions:
1.3.17.3.2. Compressão e deduplicação Inline;	https://www.nutanix.com/products/cloud-platform/software-options#nus	Vide: Nutanix Unified Storage Editions:
1.3.17.3.3. Controle e gerenciamento eficiente de cotas para usuários e grupos, tanto para arquivos quanto para objetos.	https://www.nutanix.com/products/cloud-platform/software-options#nus	Vide: Nutanix Unified Storage Editions:

● MANAUS

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● SÃO PAULO

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999

1.3.17.3.4. Conformidade com normas de Write Once Read Many (WORM), garantindo a imutabilidade de arquivos e objetos para atender requisitos regulatórios e de segurança;	https://www.nutanix.com/products/cloud-platform/software-options#nus	Vide: Nutanix Unified Storage Editions:
1.3.17.3.5. Controle de versões para objetos armazenados, permitindo rastreamento de alterações e recuperação de versões anteriores;	https://www.nutanix.com/products/cloud-platform/software-options#nus	Vide: Nutanix Unified Storage Editions:
1.3.17.3.6 .Integração com soluções de antivírus para proteção avançada contra ameaças cibernéticas, assegurando a integridade e segurança dos dados;	https://www.nutanix.com/products/cloud-platform/software-options#nus	Vide: Nutanix Unified Storage Editions:
1.3.17.3.7. Migração integrada de arquivos a partir de sistemas de terceiros, como NFS e SMB, simplificando a transição e integração de dados;	https://www.nutanix.com/products/cloud-platform/software-options#nus	Vide: Nutanix Unified Storage Editions:
1.3.17.3.8. Controle de acesso baseado em funções, permitindo a administração granular de permissões de acordo com as necessidades organizacionais;	https://www.nutanix.com/products/cloud-platform/software-options#nus	Vide: Nutanix Unified Storage Editions:
1.3.17.3.9. Monitoramento contínuo e captura de todas as atividades realizadas em arquivos nos servidores de arquivos registrados, garantindo visibilidade e rastreamento detalhado;	https://www.nutanix.com/products/cloud-platform/software-options#nus	Vide: Nutanix Unified Storage Editions:
1.3.17.3.10. Capacidade de realizar o tiering (hierarquização) de arquivos e objetos para armazenamento em nuvem ou em repositórios de objetos, otimizando o uso do espaço e o desempenho.	https://www.nutanix.com/products/cloud-platform/software-options#nus	Vide: Nutanix Unified Storage Editions:
1.3.17.4. Incluir atualizações contínuas e correções para o software licenciado, além de acesso aos canais de suporte da fabricante por meio de telefone e plataforma web;	Ciente e de acordo Conforme Proposta Técnica e Comercial	Vide: Nutanix Unified Storage Editions:
1.3.17.5. Possuir período de vigência do suporte técnico será de 36 (trinta e seis) meses, assegurando a continuidade operacional e a manutenção do desempenho das soluções contratadas;	Ciente e de acordo Conforme Proposta Técnica e Comercial	Vide: Nutanix Unified Storage Editions:

● MANAUS

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● SÃO PAULO

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999

1.3.17.6. Compatibilizar com os servidores dos itens 1.3.11 (ITEM 1), 1.3.12 (ITEM 2).	Ciente e de acordo Conforme Proposta Técnica e Comercial	Vide: Nutanix Unified Storage Editions:
1.3.18. LICENCIAMENTO DE PLATAFORMA DE DADOS DEDICADA - DEDICATED MODE/TiB (ITEM 8)		
1.3.18.1. Ser fornecido o licenciamento necessário para a implementação da plataforma/software de armazenamento de dados em clusters Nutanix, atendendo à especificação SW-NUS-STR-PR (Nutanix Unified Storage – STARTER), conforme https://www.nutanix.com/products/cloud-platform/software-options#nus ;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.18.2. Possuir licenciamento na modalidade Por Tebibyte Líquido (em base 2), correspondendo aproximadamente 1 (um) terabyte aplicado a um cluster dedicado ou já implementado;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.18.3. Contemplar as seguintes funcionalidades no licenciamento proposto:	https://www.nutanix.com/products/cloud-platform/software-options#nus	Vide: Nutanix Unified Storage Editions:
1.3.18.3.1. Erasure Coding	https://www.nutanix.com/products/cloud-platform/software-options#nus	Vide: Nutanix Unified Storage Editions:
1.3.18.3.2. Compressão e deduplicação Inline;	https://www.nutanix.com/products/cloud-platform/software-options#nus	Vide: Nutanix Unified Storage Editions:
1.3.18.3.3. Controle e gerenciamento eficiente de cotas para usuários e grupos, tanto para arquivos quanto para objetos.	https://www.nutanix.com/products/cloud-platform/software-options#nus	Vide: Nutanix Unified Storage Editions:
1.3.18.3.4. Conformidade com normas de Write Once Read Many (WORM), garantindo a imutabilidade de arquivos e objetos para atender requisitos regulatórios e de segurança;	https://www.nutanix.com/products/cloud-platform/software-options#nus	Vide: Nutanix Unified Storage Editions:

● **MANAUS**

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● **SÃO PAULO**

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999

1.3.18.3.5. Controle de versões para objetos armazenados, permitindo rastreamento de alterações e recuperação de versões anteriores;	https://www.nutanix.com/products/cloud-platform/software-options#nus	Vide: Nutanix Unified Storage Editions:
1.3.18.3.6. Integração com soluções de antivírus para proteção avançada contra ameaças cibernéticas, assegurando a integridade e segurança dos dados;	https://www.nutanix.com/products/cloud-platform/software-options#nus	Vide: Nutanix Unified Storage Editions:
1.3.18.3.7. Replicação de através de fluxo de objeto;	https://www.nutanix.com/products/cloud-platform/software-options#nus	Vide: Nutanix Unified Storage Editions:
1.3.18.3.8. Encriptação de objetos.	https://www.nutanix.com/products/cloud-platform/software-options#nus	Vide: Nutanix Unified Storage Editions:
1.3.18.4. Incluir atualizações contínuas e correções para o software licenciado, além de acesso aos canais de suporte da fabricante por meio de telefone e plataforma web;	https://www.nutanix.com/products/cloud-platform/software-options#nus	Vide: Nutanix Unified Storage Editions:
1.3.18.5. Possuir período de vigência do suporte técnico será de 36 (trinta e seis) meses, assegurando a continuidade operacional e a manutenção do desempenho das soluções contratadas;	https://www.nutanix.com/products/cloud-platform/software-options#nus	Vide: Nutanix Unified Storage Editions:
1.3.18.6. Compatibilizar com os servidores do item 1.3.13 (ITEM 3).	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.19. RENOVAÇÃO DE SUPORTE DE MANUTENÇÃO E LICENÇA PARA MÓDULOS COMPUTACIONAIS DE NUVEM PRIVADA / NX-8035 (ITEM 9)		
1.3.19.1. Prover a cobertura de garantia do fabricante dos seguintes módulos computacionais por 36 meses, compatível com o nível do ambiente atual de produção:	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.19.1.1. Nutanix NX 8035 G8 - 22SH3G340400;	Ciente e de acordo Conforme Proposta Técnica e Comercial	

● MANAUS

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● SÃO PAULO

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999

1.3.19.1.2. Nutanix NX 8035 G8 - 22SH3G340401;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.19.1.3. Nutanix NX 8035 G8 - 22SH3G340403;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.19.1.4. Nutanix NX 8035 G8 - 22SH3G340414;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.19.1.5. Nutanix NX 8035 G8 - 22SH3G340415;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.19.1.6. Nutanix NX 8035 G8 - 22SH3G340417;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.19.2. Compatibilizar com as características dos módulos (appliances):	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.19.3. Modelo do bloco: NX-8035-G8;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.19.4. Quantidade de sub-módulo por entidade: 2 (dois);	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.19.5. Possibilitar a instalação de novos patches, a implementação de upgrades na solução e o acesso a suporte técnico especializado diretamente junto ao fabricante Nutanix, abrangendo eventuais falhas e novas configurações no ambiente;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.19.6. Manter os níveis de atendimento e qualidade dos serviços atualmente oferecidos para os equipamentos Nutanix existentes no ambiente;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.19.7. A renovação deverá abranger peças e sistemas existentes em cada servidor, inclusive a substituição do próprio equipamento, se for o caso;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.19.8. Disponibilizar técnico especializado na solução ofertada no regime de 24 horas por 7 dias, durante o período de vigência do contrato, para garantir o perfeito funcionamento da solução;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.19.9. Renovar as licenças:	Ciente e de acordo Conforme Proposta Técnica e Comercial	

● MANAUS

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● SÃO PAULO

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999

1.3.19.9.1. Nutanix Cloud Infrastructure Ultimate (RSW-NCI-ULT-PR)	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.19.9.2. Nutanix Cloud Manager Starter (SW-NCM-STR-PR)	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.20. RENOVAÇÃO DE SUPORTE DE MANUTENÇÃO E LICENÇA PARA MÓDULOS COMPUTACIONAIS DE NUVEM PRIVADA / NX-8155 (ITEM 10)		
1.3.20.1. Prover a cobertura de garantia do fabricante dos seguintes módulos computacionais por 36 meses, compatível com o nível do ambiente atual de produção:	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.20.1.1. Nutanix NX 8155 G8 - 22SH5H350044;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.20.1.2. Nutanix NX 8155 G8 - 22SH5H35004;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.20.1.3. Nutanix NX 8155 G8 - 22SH5H350050;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.20.1.4. Nutanix NX 8155 G8 - 22SH5H350065;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.20.1.5. Nutanix NX 8155 G8 - 22SH5H350066;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.20.1.6. Nutanix NX 8155 G8 - 22SH5H350082;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.20.2. Compatibilizar com as características dos módulos (appliances):	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.20.3. Modelo do bloco: NX-8155-G8;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.20.4. Quantidade de sub-módulo por entidade: 1 (um);	Ciente e de acordo Conforme Proposta Técnica e Comercial	

● **MANAUS**

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● **SÃO PAULO**

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999

1.3.20.5. Possibilitar a instalação de novos patches, a implementação de upgrades na solução e o acesso a suporte técnico especializado diretamente junto ao fabricante Nutanix, abrangendo eventuais falhas e novas configurações no ambiente;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.20.6. Fornecido por Hardware, no mínimo, os mesmos níveis de atendimento e qualidade dos serviços atualmente oferecidos para os equipamentos Nutanix existentes no ambiente.	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.20.7. A renovação deverá abranger peças e sistemas existentes em cada servidor, inclusive a substituição do próprio equipamento, se for o caso;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.20.8. Disponibilizar técnico especializado na solução ofertada no regime de 24 horas por 7 dias, durante o período de vigência do contrato, para garantir o perfeito funcionamento da solução;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.20.9. Renovar as licenças:	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.20.9.1. Nutanix Cloud Infrastructure Ultimate (RSW-NCI-ULT-PR)	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.20.9.2. Nutanix Cloud Manager Starter (SW-NCM-STR-PR)	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.21. RENOVAÇÃO DE SUPORTE DE MANUTENÇÃO PARA MÓDULOS ARMAZENAMENTO DE NUVEM PRIVADA / NX-8155 (ITEM 11)		
1.3.21.1. Prover a cobertura de garantia do fabricante dos seguintes equipamentos por 36 meses:	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.21.1.1. Nutanix NX 8155 G8 - 23SH5H050692;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.21.1.2. Nutanix NX 8155 G8 - 23SH5H050721;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.21.1.3. Nutanix NX 8155 G8 - 23SH5H050726;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.21.1.4. Nutanix NX 8155 G8 - 23SH5H050727;	Ciente e de acordo Conforme Proposta Técnica e Comercial	

● **MANAUS**

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● **SÃO PAULO**

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999

1.3.21.1.5. Nutanix NX 8155 G8 - 23SH5H050728;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.21.1.6. Nutanix NX 8155 G8 - 23SH5H050730;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.21.1.7. Nutanix NX 8155 G8 - 23SH5H050734;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.21.1.8. Nutanix NX 8155 G8 - 23SH5H050735;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.21.2. Compatibilizar com as características dos módulos (appliances):	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.21.3. Modelo do bloco: NX-8155-G8;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.21.4. Quantidade de sub-módulo por entidade: 1 (um);	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.21.5. Deve ser possível instalar novos patches, implementar atualizações na solução e contar com suporte técnico especializado diretamente junto ao fabricante Nutanix, cobrindo falhas e novas configurações no ambiente	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.21.6. Deve manter, no mínimo, os mesmos níveis de atendimento e qualidade atualmente oferecidos para os equipamentos Nutanix já em uso.	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.21.7. A renovação deve incluir todas as peças e sistemas existentes em cada servidor, inclusive a substituição do equipamento, se necessário.	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.21.8. É necessário garantir a disponibilidade de técnico especializado na solução ofertada em regime 24x7, durante toda a vigência do contrato, assegurando o pleno funcionamento da solução.	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.22. RENOVAÇÃO DE SUPORTE COMUTADORES DE DATA CENTER (ITEM 12)		

● MANAUS

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● SÃO PAULO

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999

1.3.22.1. Ser fornecida a renovação do suporte técnico e garantia para o switch modelo DCS-7050SX3-48YC8-R, identificado pelo número de série JMX2302A2LM, SGD21076125, SGD21330513, SGD21330556 pelo período de 36 meses. Essa renovação deverá incluir suporte técnico do fabricante, garantindo acesso a atualizações e correções de software, bem como assistência técnica conforme as políticas vigentes do fabricante;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.22.2. Contemplar o mesmo tipo de suporte contratado atualmente. Além disso, deverá incluir atualizações contínuas do sistema operacional e firmware durante 6.20.2 todo o período de vigência do contrato de suporte;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.22.3. Assegurar nível de atendimento prioritário para incidentes críticos, com solução adequada aos problemas reportados, em conformidade com as melhores práticas da fabricante. A renovação deverá ser registrada em nome do Tribunal, garantindo que todos os direitos e benefícios estejam vinculados exclusivamente à instituição;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.22.4. As funcionalidades e tecnologias atualmente disponíveis no equipamento deverão ser mantidas, sem qualquer limitação ou redução de capacidades operacionais;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.22.5. O suporte deverá assegurar também a compatibilidade total com o modelo em questão, garantindo eficiência e continuidade das operações no ambiente de produção do Tribunal;	Ciente e de acordo Conforme Proposta Técnica e Comercial	
1.3.22.6. O período de vigência do contrato de renovação deverá ser 36 meses, com início a partir do término do suporte e garantia atualmente vigente.	Ciente e de acordo Conforme Proposta Técnica e Comercial	

Documentos adicionais adicionados aos itens de comprovação técnica.

● **MANAUS**

Rua Salvador, 440 - Sala 1810
Ed. Soberane - Torre Corporate
Adrianópolis - Manaus - AM - 69057-040
+55 92 3042-0123

● **SÃO PAULO**

Av. Dr. Gastão Vidigal, 1305 - Salas 61 e 71
Ed. Monções da Gastão
Vila Leopoldina - São Paulo - SP - 05314-000
+55 11 4673-4999

ITENS	DOCUMENTAÇÃO
1.3.11. MÓDULO DE COMPUTAÇÃO PARA NUVEM PRIVADA - TIPO 1 (ITEM 1)	
1.3.11.1. Serão aceitas apenas módulos computacionais de nuvem hiperconvergente, ou seja, com recursos de computação, armazenamento e rede totalmente integrados fim a fim;	https://www.nutanix.com/prod
1.3.11.2. Serão aceitos somente produtos homologados para uso com a solução de SDS (software defined storage) ofertada;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.11.3. Será aceita a entrega de equipamentos ou recursos extras para suprir com os recursos mínimos da solução, caso seja necessário e desde que todos os requisitos sejam atendidos;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.11.4. Cada chassi deverá possuir ventiladores e fontes redundantes para a operação adequada dos equipamentos;	https://portal.nutanix.com/pag
1.3.11.5. Cada fonte de alimentação deve vir acompanhada do seu respectivo cabo elétricos e com conectores padrão C13/C-14;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.11.6. Pertencer à família de produto Nutanix "NX" de forma a garantir compatibilidade e expansão das zonas de disponibilidade já existente, conforme o endereço de Internet https://www.nutanix.com/products/hardwareplatforms/specsheet?	https://www.nutanix.com/prod
1.3.11.7. Ser entregue com no mínimo 02 (dois) processadores com arquitetura x86 para alto desempenho, projetado para utilização em servidores, cada um com as seguintes características técnicas mínimas:	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.11.7.1. Ter 16 (dezesseis) núcleos físicos;	https://www.intel.com.br/content/www/br/pt/products/sku/237569/intel-xeon-gold-6544y-processor-45m-cache-3-60-ghz/specifications.html
1.3.11.7.2. Possuir frequência de operação nominal de, no mínimo, 3,5 GHz;	https://www.intel.com.br/content/www/br/pt/products/sku/237569/intel-xeon-gold-6544y-processor-45m-cache-3-60-ghz/specifications.html
1.3.11.7.3. Em caso de atualização de processadores, para uma versão mais atual da fabricante, será admitida sua substituição por modelo equivalente, independentemente das demais métricas, exceto pela quantidade de núcleos existentes e suportados na pastilha;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.11.7.4. Utilizar como modelo de referência: Intel Xeon-Gold 6544Y.	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.11.8. Possuir, no mínimo, 1024 GB (mil e vinte e quatro gigabytes) de memória RAM instalado, distribuídas em DIMMs DDR5 com 5600MHz RDM, ou superior, com quantidades idênticas entre os processadores;	Ciente e de acordo Conforme Proposta Técnica e Comercial

1.3.11.9. Conter no mínimo, 02 (duas) interfaces 25/10 GbE SFP28, presentes em uma placa de rede;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.11.10. Prover interface de gerência fora de banda e vir acompanhada de cabos UTP Cat-6 ou superior de, no mínimo, 05 (cinco) metros, com conectores RJ-45, para interconexão a malha de comunicação existente;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.11.11. Possuir período de vigência do suporte técnico será de 36 (trinta e seis) meses, assegurando a continuidade operacional e a manutenção do desempenho das soluções contratadas;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.11.12. Fornecer na modalidade PRODUCTION (24x7) da Nutanix, garantindo atendimento técnico ininterrupto, disponível 24 (vinte e quatro) horas por dia, 7 (sete) dias por semana;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.11.13. Vir como no mínimo as seguintes configurações de armazenamento local:	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.11.13.1. Ter 4 (quatro) discos NVME de, no mínimo, 7,68 TB de volumetria bruta individual;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.11.13.2. Ter 8 (oito) discos HDD de, no mínimo, 12 TB de volumetria bruta individual;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.12. MÓDULO DE COMPUTAÇÃO PARA NUVEM PRIVADA – TIPO 02 (ITEM 2)	
1.3.12.1. Serão aceitas apenas módulos computacionais de nuvem hiperconvergente, ou seja, com recursos de computação, armazenamento e rede totalmente integrados fim a fim;	https://www.nutanix.com/prod
1.3.12.2. Serão aceitos somente produtos homologados para uso com a solução de SDS (software defined storage) ofertada;	https://www.nutanix.com/prod
1.3.12.3. Será aceita a entrega de equipamentos ou recursos extras para suprir com os recursos mínimos da solução, caso seja necessário e desde que todos os requisitos sejam atendidos;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.12.4. Cada chassi deverá possuir ventiladores e fontes redundantes para a operação adequada dos equipamentos;	System-Specs-NX8155G9.pdf
1.3.12.5. Cada fonte de alimentação deve vir acompanhada do seu respectivo cabo elétricos e com conectores padrão C13/C-14;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.12.6. Pertencer à família de produto Nutanix "NX" de forma a garantir compatibilidade e expansão das zonas de disponibilidade já existente, conforme o endereço de Internet https://www.nutanix.com/products/hardware-platforms/specsheet?platformProvider=Nutanix ;	System-Specs-NX8155G9.pdf
1.3.12.7. Ser entregue com no mínimo 02 (dois) processadores com arquitetura x86 para alto desempenho, projetado para utilização em servidores, cada um com as seguintes características técnicas mínimas:	Ciente e de acordo Conforme Proposta Técnica e Comercial

1.3.12.7.1. Ter 16 (dezesesseis) núcleos físicos;	https://www.intel.com.br/content/www/br/pt/products/sku/237569/intel-xeon-gold-6544y-processor-45m-cache-3-60-ghz/specifications.html
1.3.12.7.2. Possuir frequência de operação nominal de, no mínimo, 3,5 GHz;	https://www.intel.com.br/content/www/br/pt/products/sku/237569/intel-xeon-gold-6544y-processor-45m-cache-3-60-ghz/specifications.html
1.3.12.7.3. Em caso de atualização de processadores, para uma versão mais atual da fabricante, será admitida sua substituição por modelo equivalente, independentemente das demais métricas, exceto pela quantidade de núcleos existentes e suportados na pastilha;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.12.7.4. Utilizar como modelo de referência: Intel Xeon-Gold 6544Y.	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.12.8. Possuir, no mínimo, 1024 GB (mil e vinte e quatro gigabytes) de memória RAM instalado, distribuídas em DIMMs DDR5 com 5600MHz RDM, ou superior, com quantidades idênticas entre os processadores;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.12.9. Conter no mínimo, 02 (duas) interfaces 25/10 GbE SFP28, presentes em uma placa de rede;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.12.10. Prover interface de gerência fora de banda e vir acompanhada de cabos UTP Cat-6 ou superior de, no mínimo, 05 (cinco) metros, com conectores RJ-45, para interconexão a malha de comunicação existente;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.12.11. Possuir período de vigência do suporte técnico será de 36 (trinta e seis) meses, assegurando a continuidade operacional e a manutenção do desempenho das soluções contratadas;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.12.12. Fornecer na modalidade PRODUCTION (24x7) da Nutanix, garantindo atendimento técnico ininterrupto, disponível 24 (vinte e quatro) horas por dia, 7 (sete) dias por semana;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.12.13. Vir como no mínimo as seguintes configurações de armazenamento local:	
1.3.12.13.1. Ter 12 (doze) discos NVME de, no mínimo, 7,68 TB de volumetria bruta individual ou outra combinação de discos que cumpra os requisitos mínimos descritos anteriormente.	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.13 MÓDULO DE ARMAZENAMENTO PARA NUVEM PRIVADA (ITEM 3)	
1.3.13.1. Serão aceitas apenas módulos de armazenamento de nuvem hiperconvergente, ou seja, com recursos de armazenamento e rede totalmente integrados fim a fim;	https://www.nutanix.com/prod
1.3.13.2. Serão aceitos somente produtos homologados para uso com a solução de SDS (software defined storage) ofertada;	https://www.nutanix.com/prod
1.3.13.3. Será aceita a entrega de equipamentos ou recursos extras para suprir com os recursos mínimos da solução, caso seja necessário e desde que todos os requisitos sejam atendidos;	Ciente e de acordo Conforme Proposta Técnica e Comercial

1.3.13.4. Cada chassi deverá possuir ventiladores e fontes redundantes para a operação adequada dos equipamentos;	Platform ANY - NX-3155-G9 S
1.3.13.5. Cada fonte de alimentação deve vir acompanhada do seu respectivo cabo elétricos e com conectores padrão C13/C-14;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.13.6. Pertencer à família de produto Nutanix "NX" de forma a garantir compatibilidade e expansão das zonas de disponibilidade já existente, conforme o endereço de Internet https://www.nutanix.com/products/hardware-platforms/specsheet?	https://www.nutanix.com/products/hardware-platforms/specsheet?
1.3.13.7. Ser entregue com no mínimo 02 (dois) processadores com arquitetura x86 para desempenho mediano, projetado para utilização em servidores, cada um com as seguintes características técnicas mínimas:	
1.3.13.7.1. Ter 12 (doze) núcleos físicos;	https://www.intel.com.br/content/www/pt-br/products/processors/xeon/xeon-silver-4510.html
1.3.13.7.2. Possuir frequência de operação nominal de, no mínimo, 2,0 GHz;	https://www.intel.com.br/content/www/pt-br/products/processors/xeon/xeon-silver-4510.html
1.3.13.7.3. Em caso de atualização de processadores, para uma versão mais atual da fabricante, será admitida sua substituição por modelo equivalente, independentemente das demais métricas, exceto pela quantidade de núcleos existentes e suportados na pastilha;	
1.3.13.7.4. Utilizar como modelo de referência: Intel Xeon-Silver 4510.	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.13.8. Possuir, no mínimo, 384 GB (trezentos e oitenta e quatro) de memória RAM instalado, distribuídas em DIMMs DDR5 com 5600MHz RDM, ou superior, com quantidades idênticas entre os processadores;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.13.9. Conter no mínimo, 02 (duas) interfaces 25/10 GbE SFP28, presentes em uma placa de rede;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.13.10. Prover interface de gerência fora de banda e vir acompanhada de cabos UTP Cat-6 ou superior de, no mínimo, 05 (cinco) metros, com conectores RJ-45, para interconexão a malha de comunicação existente;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.13.11. Possuir período de vigência do suporte técnico será de 36 (trinta e seis) meses, assegurando a continuidade operacional e a manutenção do desempenho das soluções contratadas;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.13.12. Fornecer na modalidade PRODUCTION (24x7) da Nutanix, garantindo atendimento técnico ininterrupto, disponível 24 (vinte e quatro) horas por dia, 7 (sete) dias por semana;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.13.13. Vir como no mínimo as seguintes configurações de armazenamento local:	
1.3.13.13.1. Ter 2 (dois) discos NVME de, no mínimo, 3,84 TB de volumetria bruta individual;	Ciente e de acordo Conforme Proposta Técnica e Comercial

1.3.13.13.2. Ter 10 (dez) discos HDD de, no mínimo, 12 TB de volumetria bruta individual.	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.14 MÓDULO DE COMPUTAÇÃO PARA INTELIGÊNCIA ARTIFICIAL (ITEM 4)	
1.3.14.1. Serão aceitas apenas módulos computacionais de nuvem hiperconvergente, ou seja, com recursos de computação, armazenamento e rede totalmente integrados fim a fim;	https://www.nutanix.com/prod
1.3.14.2. Serão aceitos somente produtos homologados para uso com a solução de SDS (software defined storage) ofertada;	https://www.nutanix.com/prod
1.3.14.3. Será aceita a entrega de equipamentos ou recursos extras para suprir com os recursos mínimos da solução, caso seja necessário e desde que todos os requisitos sejam atendidos;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.14.4. Cada chassi deverá possuir ventiladores e fontes redundantes para a operação adequada dos equipamentos;	System-Specs- NX8155G9.pdf
1.3.14.5. Cada fonte de alimentação deve vir acompanhada do seu respectivo cabo elétricos e com conectores padrão C13/C-14;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.14.6. Pertencer à família de produto Nutanix "NX" de forma a garantir compatibilidade e expansão das zonas de disponibilidade já existente, conforme o endereço de Internet https://www.nutanix.com/products/hardware-platforms/specsheet?platformProvider=Nutanix ;	https://www.nutanix.com/prod
1.3.14.7. Ser entregue com no mínimo 02 (dois) processadores com arquitetura x86 para alto desempenho, projetado para utilização em servidores, cada um com as seguintes características técnicas mínimas:	https://www.nutanix.com/prod
1.3.14.7.1. Ter 24 núcleos físicos;	https://www.intel.com.br/cont
1.3.14.7.2. Possuir frequência de operação nominal de, no mínimo, 2,5 GHz;	https://www.intel.com.br/cont
1.3.14.7.3. Em caso de atualização de processadores, para uma versão mais atual da fabricante, será admitida sua substituição por modelo equivalente, independentemente das demais métricas, exceto pela quantidade de núcleos existentes e suportados na pastilha;	
1.3.14.7.4. Utilizar como modelo de referência: Intel Xeon-Gold Intel® Xeon® Gold 6542Y	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.14.8. Possuir, no mínimo, 1024 GB de memória RAM instalado, distribuídas em DIMMs DDR5 com 5600MHz RDM, ou superior, com quantidades idênticas entre os processadores;	Ciente e de acordo Conforme Proposta Técnica e Comercial

1.3.14.9. Conter no mínimo, 02 (duas) interfaces 25/10 GbE SFP28, presentes em uma placa de rede;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.14.10. Prover interface de gerência fora de banda e vir acompanhada de cabos UTP Cat-6 ou superior de, no mínimo, 05 (cinco) metros, com conectores RJ-45, para interconexão a malha de comunicação existente;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.14.11. Possuir período de vigência do suporte técnico será de 36 (trinta e seis) meses, assegurando a continuidade operacional e a manutenção do desempenho das soluções contratadas;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.14.12. Fornecer na modalidade PRODUCTION (24x7) da Nutanix, garantindo atendimento técnico ininterrupto, disponível 24 (vinte e quatro) horas por dia, 7 (sete) dias por semana;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.14.13. Vir como no mínimo as seguintes configurações de armazenamento local:	
1.3.14.13.1. Ter 6 (seis) discos NVME de, no mínimo, 7,68 TB de volumetria bruta individual.	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.14.14. Possuir no mínimo 02 (duas) GPUs:	
1.3.14.14.1. NVIDIA L40S com 48GB DDR6 com ECC;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.14.14.2. Largura de Banda da Memória de 864GB/s;	https://resources.nvidia.com/en-us-l40s/l40s-datasheet-28413
1.3.14.14.3. Núcleos CUDA 18.176;	https://resources.nvidia.com/en-us-l40s/l40s-datasheet-28414
1.3.14.14.4. Núcleo FP16 362;	https://resources.nvidia.com/en-us-l40s/l40s-datasheet-28415
1.3.14.14.5. TFLOPS do Núcleo Tensor BFLOAT16 de 362;	https://resources.nvidia.com/en-us-l40s/l40s-datasheet-28416
1.3.14.14.6. TFLOPS do Núcleo Tensor BFLOAT32 de 183;	https://resources.nvidia.com/en-us-l40s/l40s-datasheet-28417
1.3.14.14.7. Interface PCIe Gen4 x16: 64GB/s bidirecional;	https://resources.nvidia.com/en-us-l40s/l40s-datasheet-28418
1.3.14.14.8. Serão aceitos outros modelos equivalentes ou superiores, em face a escassez de peça ou atualização tecnológica.	ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.15. LICENCIAMENTO DE PLATAFORMA COMPUTACIONAL ULTIMATE - NCI (ITEM 5)	
1.3.15.1 Ofertar o licenciamento compatível com o ambiente de produção do Tribunal;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.15.2 Utilizar como referência, Nutanix Cloud Infrastructure Ultimate (SW-NCI-ULT-PR) ou licenciamento mais atual e compatível comercializável;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.15.3. Contemplar funcionalidades como:	

1.3.15.3.1. Execução integral do hypervisor Nutanix AHV;	https://portal.nutanix.com/pag
1.3.15.3.2. Gerenciamento centralizado através de console atual;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.15.3.3. Criação de Snapshots e Clones das máquinas virtuais;	https://www.nutanixbible.com/
1.3.15.3.4. Tierização das áreas de armazenamento (equipamentos híbridos);	https://www.nutanixbible.com/

1.3.15.3.5. Fator de replicação 2 e 3;	The Nutanix Cloud Bible
1.3.15.3.6. Deduplicação, compressão e erasure coding;	https://www.nutanixbible.com/

1.3.15.3.7. QoS;	https://portal.nutanix.com/pag
1.3.15.3.8. Replicação assíncrona, de 1 em 1 minuto e de 1 em 1 hora para Nutanix AHV;	https://portal.nutanix.com/pag

1.3.15.3.9. Orquestração das técnicas de Recuperação de Desastres;	https://portal.nutanix.com/pag
1.3.15.3.10. Criptografia de dados;	https://www.nutanixbible.com/

1.3.15.3.11. Suporte a GPUs;	https://www.nutanixbible.com/
1.3.15.3.12. Suporte a implementação de clusters de plataformas de containers, similares a Kubernetes;	https://www.nutanix.com/prod
1.3.15.3.13. Permitir microsegmentação do ambiente virtualizado;	https://www.nutanixbible.com/

1.3.15.3.14. Gerenciamento centralizado e gerenciamento do ciclo de vida dos componentes dos appliances (softwares e firmwares).

<https://www.nutanix.com/cont>

1.3.15.4. Gerenciamento avançado de Nuvem Privada;

https://portal.nutanix.com/page/documents/details?targetId=Intelligent-Operations-Guide-vpc_7_3:mul-intelligent-operations-overview-c.html

1.3.15.5. Utilizar como referência, Nutanix Cloud Manager Starter (SW-NCM-STR-PR):	https://portal.nutanix.com/pag

1.3.15.6. Contemplar as funcionalidades como:

<https://portal.nutanix.com/pag>

1.3.15.6.1. Emissão de relatório avançados do ambiente;	https://portal.nutanix.com/pag

1.3.15.6.2. Personalização da console principal;	https://portal.nutanix.com/pag

1.3.15.6.3. Planejamento e previsão de capacidade

<https://portal.nutanix.com/pag>

1.3.15.6.4. Detecção de ineficiência de cargas de trabalho e dimensionamento ideal;	https://portal.nutanix.com/page

1.3.15.6.5. Planejamento de capacidade ou cargas de trabalho ao longo do tempo;	https://portal.nutanix.com/pag
1.3.15.6.6. Permitir automação de operações de gerência do ambiente;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.15.7. Licenciamento de cada módulo por unidade de processamento (Cores) de CPU, de acordo com a atual política do fabricante;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.15.8. Possuir período de vigência do suporte técnico será de 36 (trinta e seis) meses, assegurando a continuidade operacional e a manutenção do desempenho das soluções contratadas;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.15.9. Fornecer na modalidade PRODUCTION (24x7) da Nutanix, garantindo atendimento técnico ininterrupto, disponível 24 (vinte e quatro) horas por dia, 7 (sete) dias por semana;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.15.10. Ser compatível com os servidores dos itens 1.3.11 (ITEM 1) e 1.3.12 (ITEM 2).	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.16. LICENCIAMENTO DE PLATAFORMA ORQUESTRAÇÃO AVANÇADA CENTRAL - NCM (ITEM 6)	
1.3.16.1. Gerenciamento de orquestração avançada de Nuvem Privada;	
1.3.16.2. Utilizar como referência, Nutanix Cloud Manager Starter (SW-NCM-STR-PR):	Ciente e de acordo Conforme Proposta Técnica e Comercial

1.3.16.3. Contemplar as funcionalidades como:

[Nutanix Licensing and Cloud I](#)

1.3.16.3.1. Emissão de relatório avançados do ambiente;

[https://portal.nutanix.com/
page/documents/details?
targetId=Intelligent-
Operations-Guide-
vpc_7_3:mul-intelligent-
operations-overview-c.html](https://portal.nutanix.com/page/documents/details?targetId=Intelligent-Operations-Guide-vpc_7_3:mul-intelligent-operations-overview-c.html)

1.3.16.3.2. Personalização da console principal;	https://portal.nutanix.com/pag

1.3.16.3.3. Planejamento e previsão de capacidade

<https://portal.nutanix.com/pag>

1.3.16.3.4. Detecção de ineficiência de cargas de trabalho e dimensionamento ideal;	<u>https://portal.nutanix.com/page</u>

1.3.16.3.5. Planejamento de capacidade ou cargas de trabalho ao longo do tempo;

<https://portal.nutanix.com/pag>

1.3.16.3.6. Permitir automação de operações de gerência do ambiente;

<https://portal.nutanix.com/pag>

1.3.16.4. Licenciar cada módulo por unidade de processamento (Cores) de CPU, de acordo com a atual política do fabricante;

<https://portal.nutanix.com/page>

1.3.16.5. Permitir o uso em cenários de renovação ou criação de novas instâncias;	https://portal.nutanix.com/pag
1.3.16.6. Possuir período de vigência do suporte técnico será de 36 (trinta e seis) meses, assegurando a continuidade operacional e a manutenção do desempenho das soluções contratadas;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.16.7. Fornecer na modalidade PRODUCTION (24x7) da Nutanix, garantindo atendimento técnico ininterrupto, disponível 24 (vinte e quatro) horas por dia, 7 (sete) dias por semana;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.16.8. Ser do tipo liquidação logo no início e comprovar a cobertura total do período descrito;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.17. LICENCIAMENTO DE PLATAFORMA DE DADOS HETEREGÊNEO - MIXED MODE/TiB (ITEM 7)	
1.3.17.1. Ser fornecido o licenciamento necessário para a implementação da plataforma/software de armazenamento de dados em clusters Nutanix, atendendo à especificação SW-NUS-PRO-PR (Nutanix Unified Storage – PRO), conforme https://www.nutanix.com/products/cloud-platform/software-options#nus ;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.17.2. Possuir licenciamento na modalidade Por Tebibyte Líquido (em base 2), correspondendo aproximadamente 1 (um) terabyte aplicado a um cluster dedicado ou já implementado;	ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.17.3. Contemplar as seguintes funcionalidades no licenciamento proposto:	https://www.nutanix.com/prod

1.3.17.3.1. Erasure Coding	https://www.nutanix.com/prod
1.3.17.3.2. Compressão e deduplicação Inline;	https://www.nutanix.com/prod
1.3.17.3.3. Controle e gerenciamento eficiente de cotas para usuários e grupos, tanto para arquivos quanto para objetos.	https://www.nutanix.com/prod
1.3.17.3.4. Conformidade com normas de Write Once Read Many (WORM), garantindo a imutabilidade de arquivos e objetos para atender requisitos regulatórios e de segurança;	https://www.nutanix.com/prod
1.3.17.3.5. Controle de versões para objetos armazenados, permitindo rastreamento de alterações e recuperação de versões anteriores;	https://www.nutanix.com/prod
1.3.17.3.6 .Integração com soluções de antivírus para proteção avançada contra ameaças cibernéticas, assegurando a integridade e segurança dos dados;	https://www.nutanix.com/prod
1.3.17.3.7. Migração integrada de arquivos a partir de sistemas de terceiros, como NFS e SMB, simplificando a transição e integração de dados;	https://www.nutanix.com/prod
1.3.17.3.8. Controle de acesso baseado em funções, permitindo a administração granular de permissões de acordo com as necessidades organizacionais;	https://www.nutanix.com/prod
1.3.17.3.9. Monitoramento contínuo e captura de todas as atividades realizadas em arquivos nos servidores de arquivos registrados, garantindo visibilidade e rastreamento detalhado;	https://www.nutanix.com/prod
1.3.17.3.10. Capacidade de realizar o tiering (hierarquização) de arquivos e objetos para armazenamento em nuvem ou em repositórios de objetos, otimizando o uso do espaço e o desempenho.	https://www.nutanix.com/prod
1.3.17.4. Incluir atualizações contínuas e correções para o software licenciado, além de acesso aos canais de suporte da fabricante por meio de telefone e plataforma web;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.17.5. Possuir período de vigência do suporte técnico será de 36 (trinta e seis) meses, assegurando a continuidade operacional e a manutenção do desempenho das soluções contratadas;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.17.6. Compatibilizar com os servidores dos itens 1.3.11 (ITEM 1), 1.3.12 (ITEM 2).	Ciente e de acordo Conforme Proposta Técnica e Comercial

1.3.18. LICENCIAMENTO DE PLATAFORMA DE DADOS DEDICADA - DEDICATED MODE/TiB (ITEM 8)	
1.3.18.1. Ser fornecido o licenciamento necessário para a implementação da plataforma/software de armazenamento de dados em clusters Nutanix, atendendo à especificação SW-NUS-STR-PR (Nutanix Unified Storage – STARTER), conforme https://www.nutanix.com/products/cloud-platform/software-options#nus ;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.18.2. Possuir licenciamento na modalidade Por Tebibyte Líquido (em base 2), correspondendo aproximadamente 1 (um) terabyte aplicado a um cluster dedicado ou já implementado;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.18.3. Contemplar as seguintes funcionalidades no licenciamento proposto:	https://www.nutanix.com/prod
1.3.18.3.1. Erasure Coding	https://www.nutanix.com/prod
1.3.18.3.2. Compressão e deduplicação Inline;	https://www.nutanix.com/prod
1.3.18.3.3. Controle e gerenciamento eficiente de cotas para usuários e grupos, tanto para arquivos quanto para objetos.	https://www.nutanix.com/prod
1.3.18.3.4. Conformidade com normas de Write Once Read Many (WORM), garantindo a imutabilidade de arquivos e objetos para atender requisitos regulatórios e de segurança;	https://www.nutanix.com/prod
1.3.18.3.5. Controle de versões para objetos armazenados, permitindo rastreamento de alterações e recuperação de versões anteriores;	https://www.nutanix.com/prod
1.3.18.3.6. Integração com soluções de antivírus para proteção avançada contra ameaças cibernéticas, assegurando a integridade e segurança dos dados;	https://www.nutanix.com/prod
1.3.18.3.7. Replicação de através de fluxo de objeto;	https://www.nutanix.com/prod
1.3.18.3.8. Encriptação de objetos.	https://www.nutanix.com/prod
1.3.18.4. Incluir atualizações contínuas e correções para o software licenciado, além de acesso aos canais de suporte da fabricante por meio de telefone e plataforma web;	https://www.nutanix.com/prod

1.3.18.5. Possuir período de vigência do suporte técnico será de 36 (trinta e seis) meses, assegurando a continuidade operacional e a manutenção do desempenho das soluções contratadas;	https://www.nutanix.com/prod
1.3.18.6. Compatibilizar com os servidores do item 1.3.13 (ITEM 3).	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.19. RENOVAÇÃO DE SUPORTE DE MANUTENÇÃO E LICENÇA PARA MÓDULOS COMPUTACIONAIS DE NUVEM PRIVADA / NX-8035 (ITEM 9)	
1.3.19.1. Prover a cobertura de garantia do fabricante dos seguintes módulos computacionais por 36 meses, compatível com o nível do ambiente atual de produção:	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.19.1.1. Nutanix NX 8035 G8 - 22SH3G340400;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.19.1.2. Nutanix NX 8035 G8 - 22SH3G340401;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.19.1.3. Nutanix NX 8035 G8 - 22SH3G340403;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.19.1.4. Nutanix NX 8035 G8 - 22SH3G340414;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.19.1.5. Nutanix NX 8035 G8 - 22SH3G340415;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.19.1.6. Nutanix NX 8035 G8 - 22SH3G340417;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.19.2. Compatibilizar com as características dos módulos (appliances):	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.19.3. Modelo do bloco: NX-8035-G8;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.19.4. Quantidade de sub-módulo por entidade: 2 (dois);	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.19.5. Possibilitar a instalação de novos patches, a implementação de upgrades na solução e o acesso a suporte técnico especializado diretamente junto ao fabricante Nutanix, abrangendo eventuais falhas e novas configurações no ambiente;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.19.6. Manter os níveis de atendimento e qualidade dos serviços atualmente oferecidos para os equipamentos Nutanix existentes no ambiente;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.19.7. A renovação deverá abranger peças e sistemas existentes em cada servidor, inclusive a substituição do próprio equipamento, se for o caso;	Ciente e de acordo Conforme Proposta Técnica e Comercial

1.3.19.8. Disponibilizar técnico especializado na solução ofertada no regime de 24 horas por 7 dias, durante o período de vigência do contrato, para garantir o perfeito funcionamento da solução;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.19.9. Renovar as licenças:	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.19.9.1. Nutanix Cloud Infrastructure Ultimate (RSW-NCI-ULT-PR)	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.19.9.2. Nutanix Cloud Manager Starter (SW-NCM-STR-PR)	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.20. RENOVAÇÃO DE SUPORTE DE MANUTENÇÃO E LICENÇA PARA MÓDULOS COMPUTACIONAIS DE NUVEM PRIVADA / NX-8155 (ITEM 10)	
1.3.20.1. Prover a cobertura de garantia do fabricante dos seguintes módulos computacionais por 36 meses, compatível com o nível do ambiente atual de produção:	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.20.1.1. Nutanix NX 8155 G8 - 22SH5H350044;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.20.1.2. Nutanix NX 8155 G8 - 22SH5H35004;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.20.1.3. Nutanix NX 8155 G8 - 22SH5H350050;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.20.1.4. Nutanix NX 8155 G8 - 22SH5H350065;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.20.1.5. Nutanix NX 8155 G8 - 22SH5H350066;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.20.1.6. Nutanix NX 8155 G8 - 22SH5H350082;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.20.2. Compatibilizar com as características dos módulos (appliances):	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.20.3. Modelo do bloco: NX-8155-G8;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.20.4. Quantidade de sub-módulo por entidade: 1 (um);	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.20.5. Possibilitar a instalação de novos patches, a implementação de upgrades na solução e o acesso a suporte técnico especializado diretamente junto ao fabricante Nutanix, abrangendo eventuais falhas e novas configurações no ambiente;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.20.6. Fornecido por Hardware, no mínimo, os mesmos níveis de atendimento e qualidade dos serviços atualmente oferecidos para os equipamentos Nutanix existentes no ambiente.	Ciente e de acordo Conforme Proposta Técnica e Comercial

1.3.20.7. A renovação deverá abranger peças e sistemas existentes em cada servidor, inclusive a substituição do próprio equipamento, se for o caso;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.20.8. Disponibilizar técnico especializado na solução ofertada no regime de 24 horas por 7 dias, durante o período de vigência do contrato, para garantir o perfeito funcionamento da solução;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.20.9. Renovar as licenças:	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.20.9.1. Nutanix Cloud Infrastructure Ultimate (RSW-NCI-ULT-PR)	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.20.9.2. Nutanix Cloud Manager Starter (SW-NCM-STR-PR)	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.21. RENOVAÇÃO DE SUPORTE DE MANUTENÇÃO PARA MÓDULOS ARMAZENAMENTO DE NUVEM PRIVADA / NX-8155 (ITEM 11)	
1.3.21.1. Prover a cobertura de garantia do fabricante dos seguintes equipamentos por 36 meses:	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.21.1.1. Nutanix NX 8155 G8 - 23SH5H050692;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.21.1.2. Nutanix NX 8155 G8 - 23SH5H050721;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.21.1.3. Nutanix NX 8155 G8 - 23SH5H050726;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.21.1.4. Nutanix NX 8155 G8 - 23SH5H050727;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.21.1.5. Nutanix NX 8155 G8 - 23SH5H050728;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.21.1.6. Nutanix NX 8155 G8 - 23SH5H050730;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.21.1.7. Nutanix NX 8155 G8 - 23SH5H050734;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.21.1.8. Nutanix NX 8155 G8 - 23SH5H050735;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.21.2. Compatibilizar com as características dos módulos (appliances):	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.21.3. Modelo do bloco: NX-8155-G8;	Ciente e de acordo Conforme Proposta Técnica e Comercial

1.3.21.4. Quantidade de sub-módulo por entidade: 1 (um);	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.21.5. Deve ser possível instalar novos patches, implementar atualizações na solução e contar com suporte técnico especializado diretamente junto ao fabricante Nutanix, cobrindo falhas e novas configurações no ambiente	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.21.6. Deve manter, no mínimo, os mesmos níveis de atendimento e qualidade atualmente oferecidos para os equipamentos Nutanix já em uso.	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.21.7. A renovação deve incluir todas as peças e sistemas existentes em cada servidor, inclusive a substituição do equipamento, se necessário.	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.21.8. É necessário garantir a disponibilidade de técnico especializado na solução ofertada em regime 24x7, durante toda a vigência do contrato, assegurando o pleno funcionamento da solução.	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.22. RENOVAÇÃO DE SUPORTE COMUTADORES DE DATACENTER (ITEM 12)	
1.3.22.1. Ser fornecida a renovação do suporte técnico e garantia para o switch modelo DCS-7050SX3-48YC8-R, identificado pelo número de série JMX2302A2LM, SGD21076125, SGD21330513, SGD21330556 pelo período de 36 meses. Essa renovação deverá incluir suporte técnico do fabricante, garantindo acesso a atualizações e correções de software, bem como assistência técnica conforme as políticas vigentes do fabricante;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.22.2. Contemplar o mesmo tipo de suporte contratado atualmente. Além disso, deverá incluir atualizações contínuas do sistema operacional e firmware durante 6.20.2 todo o período de vigência do contrato de suporte;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.22.3. Assegurar nível de atendimento prioritário para incidentes críticos, com solução adequada aos problemas reportados, em conformidade com as melhores práticas da fabricante. A renovação deverá ser registrada em nome do Tribunal, garantindo que todos os direitos e benefícios estejam vinculados exclusivamente à instituição;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.22.4. As funcionalidades e tecnologias atualmente disponíveis no equipamento deverão ser mantidas, sem qualquer limitação ou redução de capacidades operacionais;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.22.5. O suporte deverá assegurar também a compatibilidade total com o modelo em questão, garantindo eficiência e continuidade das operações no ambiente de produção do Tribunal;	Ciente e de acordo Conforme Proposta Técnica e Comercial
1.3.22.6. O período de vigência do contrato de renovação deverá ser 36 meses, com início a partir do término do suporte e garantia atualmente vigente.	Ciente e de acordo Conforme Proposta Técnica e Comercial

COMPROVAÇÃO
Hardware Platform Spec Sheets
N/A
N/A
N/A
N/A
N/A
N/A
N/A
Vide especificação
Vide especificação
N/A
N/A
N/A

N/A
N/A
N/A
N/A
N/A
N/A
N/A
N/A
N/A
N/A
N/A
Vide especificação
N/A

Vide especificação

Vide especificação

[illegible]

[illegible]

N/A
Vide especificação
N/A
N/A
Vide especificação
Vide especificação
N/A

Technical Specifications
Technical Specifications
Technical Specifications
Technical Specifications
Technical Specifications
Technical Specifications

As explained in the Data Structure section above, a virtual machine consists of files (vmdk/vhdx) which are vDisks on the Nutanix platform. A vDisk is composed of extents which are logically contiguous chunks of data, which are stored within extent groups which are physically contiguous data stored as files on the storage devices. When a snapshot or clone is taken, the base vDisk is marked immutable and another vDisk is created as read/write. At this point, both vDisks have the same block map, which is a metadata mapping of the vDisk to its corresponding extents. Contrary to traditional approaches which require traversal of the
For hybrid scenarios with flash and HDD, the flash would be Tier 0 with HDD being Tier 1.

currently uses a resiliency factor, also known as a replication factor (RF), and checksum to ensure data redundancy and availability in the case of a node or disk failure or corruption. As explained above, the OpLog acts as a staging area to absorb incoming writes onto a low-latency SSD tier. Upon being written to the local OpLog, the data is synchronously replicated to another one or two Nutanix CVM's OpLog (dependent on RF) before being acknowledged (Ack) as a successful write to the host. This ensures that the data exists in at least two or three independent locations and is fault tolerant. NOTE: For RF3, a minimum of 5 nodes is required since metadata will be

The following optimizations are leveraged:

Erasure Coding (EC-X)
Compression
Deduplication

Storage QoS provides administrators granular control to manage the performance of virtual machines and ensure that the system delivers consistent performance for all workloads. You can use a controllable knob to limit the IOPS that the storage layer would serve for individual virtual machines. IOPS is the number of requests the storage layer can serve in a second. You can set throttle limits on a VM to prevent noisy VMs from over-utilizing the system resources.

Nearsync Replication
(1 Minute to 15
Minutes RPO)

Asynchronous
Replication (1 Hour or
More RPO)

You can protect your entities (guest VMs, volume groups, and consistency groups) and orchestrate their DR to other Nutanix clusters when events that cause service disruption occur at the primary AZ. For the protection of your entities, protection policies with asynchronous, nearsync, or synchronous replication schedules generate and replicate recovery points to other on-prem AZs. Recovery plans orchestrate DR from the replicated recovery points to other Nutanix clusters at the same or different on-prem AZs.

Data Encryption
Nutanix provides data encryption via three main options:

Native software-based encryption (FIPS-140-2 Level-1) *released in 5.5

Using self-encrypting drives (SED) (FIPS-140-2 Level-2)

Software + hardware encryption

Dynamic GPU management - Supports NVIDIA GPUs to provide specific vGPU profiles based on the VMs running in the system Creating remediation plans based on the cost of movement to address each of the above case
Nutanix Kubernetes Platform
Supported Configurations The solution applies to the configurations below: Core Use Case(s): Microsegmentation Management interfaces(s): Prism Central (PC) Supported Environment(s): On-Premises: AHV Cloud: Nutanix Cloud Clusters on AWS

KEY BENEFITS

- Easily accessible through Nutanix Prism.
- One-click upgrade simplicity.
- Upgrade software in a single process.
 - Automatically manage software and hardware dependencies.
 - Unifies firmware upgrade operations across multiple hardware vendors.
- Included scheduler checks for updates automatically

Analysis: Analyze cluster activity. For more information, see Performance Monitoring.

Discovery: Discover applications running in a cluster. For more information, see Application Discovery. Monitoring

Configurations: Monitor select applications in a cluster. For more information, see Application Monitoring.

Planning: Plan for resource capacity and usage. For more information, see Resource Planning.

Playbooks: Create playbooks for task automation. For more information, see Task Automation - Playbooks.

Reports: Generate resource and activity reports for Nutanix and non-Nutanix

Analysis: Analyze cluster activity. For more information, see Performance Monitoring.

Discovery: Discover applications running in a cluster. For more information, see Application Discovery. Monitoring

Configurations:

Monitor select applications in a cluster. For more information, see Application Monitoring.

Planning: Plan for resource capacity and usage. For more information, see Resource Planning.

Playbooks: Create playbooks for task automation. For more information, see Task Automation - Playbooks.

Reports: Generate resource and activity reports for Nutanix and non-Nutanix

Analysis: Analyze cluster activity. For more information, see Performance Monitoring.

Discovery: Discover applications running in a cluster. For more information, see Application Discovery. Monitoring

Configurations:

Monitor select applications in a cluster. For more information, see Application Monitoring.

Planning: Plan for resource capacity and usage. For more information, see Resource Planning.

Playbooks: Create playbooks for task automation. For more information, see Task Automation - Playbooks.

Reports: Generate resource and activity reports for Nutanix and non-Nutanix

Analysis: Analyze cluster activity. For more information, see Performance Monitoring.

Discovery: Discover applications running in a cluster. For more information, see Application Discovery. Monitoring

Configurations:

Monitor select applications in a cluster. For more information, see Application Monitoring.

Planning: Plan for resource capacity and usage. For more information, see Resource Planning.

Playbooks: Create playbooks for task automation. For more information, see Task Automation - Playbooks.

Reports: Generate resource and activity reports for Nutanix and non-Nutanix

Analysis: Analyze cluster activity. For more information, see Performance Monitoring.

Discovery: Discover applications running in a cluster. For more information, see Application Discovery. Monitoring

Configurations:

Monitor select applications in a cluster. For more information, see Application Monitoring.

Planning: Plan for resource capacity and usage. For more information, see Resource Planning.

Playbooks: Create playbooks for task automation. For more information, see Task Automation - Playbooks.

Reports: Generate resource and activity reports for Nutanix and non-Nutanix

Analysis: Analyze cluster activity. For more information, see Performance Monitoring.

Discovery: Discover applications running in a cluster. For more information, see Application Discovery. Monitoring

Configurations:

Monitor select applications in a cluster. For more information, see Application Monitoring.

Planning: Plan for resource capacity and usage. For more information, see Resource Planning.

Playbooks: Create playbooks for task automation. For more information, see Task Automation - Playbooks.

Reports: Generate resource and activity reports for Nutanix and non-Nutanix

Analysis: Analyze cluster activity. For more information, see Performance Monitoring.

Discovery: Discover applications running in a cluster. For more information, see Application Discovery. Monitoring

Configurations:

Monitor select applications in a cluster. For more information, see Application Monitoring.

Planning: Plan for resource capacity and usage. For more information, see Resource Planning.

Playbooks: Create playbooks for task automation. For more information, see Task Automation - Playbooks.

Reports: Generate resource and activity reports for Nutanix and non-Nutanix

Analysis: Analyze cluster activity. For more information, see [Performance Monitoring](#).

Discovery: Discover applications running in a cluster. For more information, see [Application Discovery](#).

Monitoring
Configurations:
Monitor select
applications in a
cluster. For more
information, see
Application
Monitoring.

Planning: Plan for resource capacity and usage. For more information, see Resource Planning.

Playbooks: Create playbooks for task automation. For more information, see [Task Automation - Playbooks](#).

Reports: Generate resource and activity reports for Nutanix and non-Nutanix

and non-Nutanix

Analysis: Analyze cluster activity. For more information, see Performance Monitoring. Discovery: Discover applications running in a cluster. For more information, see Application Discovery. Monitoring Configurations: Monitor select applications in a cluster. For more information, see Application Monitoring. Planning: Plan for resource capacity and usage. For more information, see Resource Planning. Playbooks: Create playbooks for task automation. For more information, see Task Automation - Playbooks. Reports: Generate resource and activity reports for Nutanix and non-Nutanix

Analysis: Analyze cluster activity. For more information, see Performance Monitoring.

Discovery: Discover applications running in a cluster. For more information, see Application Discovery. Monitoring

Configurations:

Monitor select applications in a cluster. For more information, see Application Monitoring.

Planning: Plan for resource capacity and usage. For more information, see Resource Planning.

Playbooks: Create playbooks for task automation. For more information, see Task Automation - Playbooks.

Reports: Generate resource and activity reports for Nutanix and non-Nutanix

Analysis: Analyze cluster activity. For more information, see Performance Monitoring.

Discovery: Discover applications running in a cluster. For more information, see Application Discovery. Monitoring

Configurations:

Monitor select applications in a cluster. For more information, see Application Monitoring.

Planning: Plan for resource capacity and usage. For more information, see Resource Planning.

Playbooks: Create playbooks for task automation. For more information, see Task Automation - Playbooks.

Reports: Generate resource and activity reports for Nutanix and non-Nutanix

Analysis: Analyze cluster activity. For more information, see Performance Monitoring.

Discovery: Discover applications running in a cluster. For more information, see Application Discovery. Monitoring

Configurations:

Monitor select applications in a cluster. For more information, see Application Monitoring.

Planning: Plan for resource capacity and usage. For more information, see Resource Planning.

Playbooks: Create playbooks for task automation. For more information, see Task Automation - Playbooks.

Reports: Generate resource and activity reports for Nutanix and non-Nutanix

Analysis: Analyze cluster activity. For more information, see Performance Monitoring.

Discovery: Discover applications running in a cluster. For more information, see Application Discovery. Monitoring

Configurations:

Monitor select applications in a cluster. For more information, see Application Monitoring.

Planning: Plan for resource capacity and usage. For more information, see Resource Planning.

Playbooks: Create playbooks for task automation. For more information, see Task Automation - Playbooks.

Reports: Generate resource and activity reports for Nutanix and non-Nutanix

Analysis: Analyze cluster activity. For more information, see Performance Monitoring.

Discovery: Discover applications running in a cluster. For more information, see Application Discovery. Monitoring

Configurations:

Monitor select applications in a cluster. For more information, see Application Monitoring.

Planning: Plan for resource capacity and usage. For more information, see Resource Planning.

Playbooks: Create playbooks for task automation. For more information, see Task Automation - Playbooks.

Reports: Generate resource and activity reports for Nutanix and non-Nutanix

Analysis: Analyze cluster activity. For more information, see Performance Monitoring.

Discovery: Discover applications running in a cluster. For more information, see Application Discovery. Monitoring

Configurations:

Monitor select applications in a cluster. For more information, see Application Monitoring.

Planning: Plan for resource capacity and usage. For more information, see Resource Planning.

Playbooks: Create playbooks for task automation. For more information, see Task Automation - Playbooks.

Reports: Generate resource and activity reports for Nutanix and non-Nutanix

Analysis: Analyze cluster activity. For more information, see Performance Monitoring.

Discovery: Discover applications running in a cluster. For more information, see Application Discovery.

Monitoring

Configurations: Monitor select applications in a cluster. For more information, see Application Monitoring.

Planning: Plan for resource capacity and usage. For more information, see Resource Planning.

Playbooks: Create playbooks for task automation. For more information, see Task Automation - Playbooks.

Reports: Generate resource and activity reports for Nutanix and non-Nutanix

Vide: Nutanix Unified Storage Editions:

Vide: Nutanix Unified
Storage Editions:

Vide: Nutanix Unified
Storage Editions:

Vide: Nutanix Unified
Storage Editions:

Vide: Nutanix Unified
Storage Editions:

Vide: Nutanix Unified
Storage Editions:

Vide: Nutanix Unified
Storage Editions:

Vide: Nutanix Unified
Storage Editions:

Vide: Nutanix Unified
Storage Editions:

Vide: Nutanix Unified
Storage Editions:

Vide: Nutanix Unified
Storage Editions:

Vide: Nutanix Unified
Storage Editions:

Vide: Nutanix Unified
Storage Editions:

Vide: Nutanix Unified
Storage Editions:

[illegible]

Vide: Nutanix Unified Storage Editions:

[illegible]

[illegible]

[illegible]

[illegible]

Prism Central Infrastructure Guide

Prism pc.7.3
November 6, 2025

Contents

About this Publication.....	8
Prism Central Documentation Portfolio.....	10
Related Documentation.....	11
Getting Started.....	12
Prism Central Deployment.....	12
Prism Central in a Non-Nutanix Environment.....	13
Requirements for Prism Central Deployment.....	13
Limitations of Prism Central Deployment:.....	14
Deploy Multiple Prism Central Instances on a Prism Element Cluster.....	15
Prism Central Installation.....	15
Prism Central Upgrade.....	26
Service Manager for Installing Services and Applications.....	30
Prism Central Overview.....	31
Logging Into Prism Central.....	32
Prism Central Landing Page.....	35
Application Switcher Function.....	40
Prism Central Settings (Infrastructure).....	44
Understanding Displayed Statistics.....	46
Infrastructure Overview.....	48
Admin Center Overview.....	48
Prism Central GUI Organization.....	49
Prism Licensing.....	54
X-Small Prism Central.....	55
Registering a Cluster with Prism Central.....	56
Unregistering a Cluster from Prism Central.....	58
Application-specific Navigation Bar.....	61
Identifying the Prism Central Leader VM.....	64
Searching for Information.....	65
Main Dashboard - Infrastructure.....	76
Dashboard Management.....	78
Resetting Dashboard.....	80
Setting Data Density.....	81
Creating a New (Custom) Dashboard.....	81
Modifying a Dashboard.....	83
Sharing a Dashboard.....	84
Generating a Dashboard Summary Report.....	88
Widgets Management.....	91
Adding Dashboard Widgets.....	94
Compute Entities.....	99

VM Management.....	99
VMs Summary View.....	100
Viewing VM Details.....	113
Creating a VM through Prism Central (AHV).....	149
Managing a VM through Prism Central (AHV).....	161
Adding Multiple vGPUs to the Same VM.....	172
VM Migration Specifications.....	173
Creating a VM through Prism Central (ESXi).....	186
Managing a VM through Prism Central (ESXi).....	188
Nutanix Guest Tools.....	189
Storage Quality of Service (QoS).....	210
Memory Overcommit Management.....	212
Policies for VM Management.....	212
VM Template Management.....	212
VM Template Summary View.....	212
VM Template Details View.....	214
Limitations of VM Template Feature.....	217
Creating a VM Template.....	217
Deploying VM from a Template.....	220
Managing a VM Template.....	223
Kubernetes Clusters Management.....	226
Kubernetes Clusters Summary View.....	227
Kubernetes Cluster Details View.....	228
OVA Management.....	235
OVAs Summary View.....	236
Exporting a VM as an OVA.....	238
Uploading an OVA.....	239
Concatenating Upload using APIs.....	246
Deploying an OVA as VM.....	249
Downloading an OVA.....	256
Renaming an OVA.....	256
Deleting an OVA.....	257
Image Management.....	258
Images Summary View.....	259
Image Details View.....	262
Requirements.....	263
Limitations.....	263
Adding an Image.....	264
Image Upload Verification.....	279
Modifying an Image.....	280
Importing Images to Prism Central.....	281
Uploading Images to Objects.....	282
Policies for Image Management.....	283
Catalog Management.....	283
Catalog Items Summary View.....	283
Adding a Catalog Item.....	285
Deleting a Catalog Item.....	287
Storage Entities.....	289
Storage Container Management.....	289
Storage Containers Summary View.....	290
Storage Container Details View.....	296
Storage Components.....	301
Storage Efficiency.....	303
Capacity Reservation Best Practices.....	310

Limitations for Storage Containers.....	311
Creating a Storage Container.....	311
Modifying a Storage Container.....	316
Deleting a Storage Container.....	318
Volume Group Management.....	318
Volume Groups Summary View.....	319
Volume Group Details View.....	324
Creating a Volume Group.....	332
Modifying a Volume Group.....	336
Configuring Mutual CHAP Authentication.....	339
Attaching Volume Groups to Guest VMs.....	340
Limitations of Accessing a Volume Group using NVMeoF/TCP.....	340
Cluster RBAC for Volume Group.....	341
External vCenter Server Integration.....	341
External vCenter Datastores Summary View.....	342
External vCenter Datastore Details View.....	347
Registering a Cluster with a vCenter Server using Prism Central.....	351
Unregistering a Cluster from the External vCenter Server using Prism Central.....	352
Migrating a Nutanix Cluster between Two vCenter Servers using Prism Central.....	352
Updating the vCenter Service Account Credentials in Prism Central.....	353
Network and Security Entities.....	355
Virtual Switches.....	355
Viewing Virtual Switches Summary Page.....	357
Viewing Virtual Switch Details.....	360
Creating a Virtual Switch.....	362
Updating a Virtual Switch.....	363
Deleting a Virtual Switch.....	364
Assign a Subnet to a Virtual Switch.....	365
Subnets.....	365
Subnets Summary View.....	365
Subnet Details View.....	368
Network Configuration.....	369
Virtual Private Clouds.....	369
Viewing the Virtual Private Cloud Summary.....	370
Viewing the Virtual Private Cloud Details.....	371
Network Services.....	376
Network Load Balancer.....	376
Traffic Mirroring.....	390
Floating IPs.....	403
Floating IPs Summary View.....	403
Connectivity.....	404
Gateways Summary View.....	405
Gateway Details View.....	406
VPN Connections Summary View.....	408
VPN Connection Details View.....	409
Subnet Extensions Summary View.....	412
Subnet Extension Details View.....	414
BGP Sessions Summary View.....	416
BGP Session Details View.....	418
Security Policies.....	420
Security Dashboard.....	420
Data Protection and Recovery Entities.....	421

Protection Summary.....	421
Protection Policies.....	421
Recovery Plans.....	421
VM Recovery Points.....	421
VG Recovery Points.....	421
Consistency Groups.....	421
Hardware Entities.....	422
Cluster Management.....	422
Clusters Summary View.....	424
Cluster Details View.....	431
Creating a Cluster.....	435
Destroying a Cluster.....	437
Managing Cluster Fault Tolerance.....	439
Settings Profile for Clusters.....	439
Host Management.....	444
Hosts Summary View.....	445
Host Details View.....	449
Renaming an AHV Host.....	454
Expanding a Cluster through Prism Central.....	454
Removing a Node through Prism Central.....	461
Putting a Host into Maintenance Mode Using Prism Central.....	462
Exiting a Host from Maintenance Mode Using Prism Central.....	463
Disks Summary View.....	464
Disk Details View.....	468
GPUs Summary View.....	470
GPU Details View.....	471
NICs and NIC Profile Management.....	473
Viewing the NICs and NIC Profiles.....	473
Viewing the NIC Port Details.....	478
Viewing the NIC Profile Details.....	481
Creating a NIC Profile.....	485
Updating a NIC Profile.....	486
Deleting a NIC Profile.....	489
Attaching an SR-IOV NIC to a VM.....	490
Attaching a Network Offload NIC Profile to a VM.....	490
Activity Entities – Alert and Event Monitoring.....	492
Activity Entities – Tasks and Audits.....	493
Audits Summary View.....	493
Audit Details View.....	495
Tasks View.....	496
Operations Entities.....	499
Administration Entities.....	500
Availability Zones.....	500
Policies in Infrastructure.....	501

VM Policy Management.....	501
VM-Host Affinity Policies Defined in Prism Central.....	501
VM-VM Anti-Affinity Policies Defined in Prism Central.....	512
NGT Policies.....	522
VM Template Policy Management.....	526
Template Placement Policies.....	526
Image Policy Management.....	527
Image Placement Policies.....	527
Bandwidth Throttling Policies.....	535
Security Policy Management.....	540
Security Policies Summary View.....	540
Security Policy Details View.....	544
Storage Policy Management.....	545
Default Storage Policy.....	548
Storage Policies Summary View.....	550
Storage Policy Details View.....	553
Storage Policy Compliance.....	557
Storage Policy Based Replication Factor.....	558
Storage Policy-Based Encryption.....	558
Creating a Storage Policy.....	560
Managing Storage Policies.....	564
Deleting a Storage Policy.....	566
 Prism Self Service Administration.....	 569
Prism Self Service Setup.....	569
Creating a VM (Self Service).....	571
Creating a VM from Catalog Items (Self Service).....	571
Managing a VM (Self Service).....	578
Assigning a VM to a Project Member.....	578
 Additional Operations.....	 581
Managing Prism Central.....	581
Microservices Infrastructure.....	584
Prism Central Backup, Restore, and Migration.....	591
Intelligent Operations.....	612
Large Files Upload Using Objects Lite.....	612
Objects Lite Data Storage.....	612
Objects Lite Limitations.....	613
Objects Lite Access.....	613
High Availability in Prism Central.....	616
Finding the Prism Central Version.....	616
Finding the AHV Version on Prism Central.....	617
Finding the AOS Version Using Prism Central.....	618
Expanding (Scale Out) Prism Central.....	619
Shutting Down or Starting Up Prism Central VM.....	625
Shutting down or starting up all the PC VMs in a Prism Central Configuration.....	625
Shutting down or starting up a single PC VM in a scale-out PC configuration.....	626
Keyboard Shortcuts in Prism Central.....	626
IP Address Reconfiguration.....	627
Preparing to Reconfigure the IP Address and Gateway of PC VMs.....	627
Reconfiguring the IP Address and Gateway of Prism Central VMs.....	628
Logging Out of Prism Central.....	632
Power Usage.....	633
Power Usage Requirements.....	633

Power Usage Considerations.....	633
Activating Power Monitor Workflow.....	634
Deploying Power Monitor Application.....	634
Configuring Out-of-Band Management Credentials.....	635
Viewing Power Usage Metrics at a Cluster Level.....	636
Viewing Power Usage Metrics at a Host Level.....	636
Updating Out-of-Band Management Credentials.....	637
Removing Out-of-Band Management Credentials.....	637
Customer Support Services.....	639
Configuring Remote Connection Using CLI.....	639
Creating a Support Case.....	640
Viewing Case Status.....	643
Accessing the Nutanix Support Portal (Prism Central).....	644
Accessing the REST API Explorer.....	647
Troubleshooting.....	647
Collecting Logs.....	647
Viewing Log Collection History.....	649
Help Resources.....	651
Accessing Online Help (Prism Central).....	651
Accessing the Nutanix Next Community.....	652
Glossary.....	653
Copyright.....	654

ABOUT THIS PUBLICATION

This document provides the information about how to set up the site infrastructure using Prism Central.

- For information about all the documents applicable for Prism Central, see [Prism Central Documentation Portfolio](#) on page 10.
- For information about how to access this document using online help option, see [Accessing Online Help \(Prism Central\)](#) on page 651.
- To access other Nutanix documents, see [Nutanix support portal](#).

Scope of the Document

The scope of this document is limited to describe the following infrastructure-specific attributes and operations that you can perform using Prism Central:

- Monitor and Analyze system performance using exhaustive Dashboard. For more information, see [Main Dashboard - Infrastructure](#) on page 76.
- Compute and Storage Settings. For more information, see [Compute Entities](#) on page 99.
- Network and Security Settings. For more information, see [Network and Security Entities](#) on page 355.
- Data Protection and Recovery Settings. For more information, see [Data Protection and Recovery Entities](#) on page 421.
- Checking the Hardware Entities and managing Hosts. For more information, see [Hardware Entities](#) on page 422.
- Checking Alerts and Events for fault management. For more information, see [Activity Entities – Alert and Event Monitoring](#) on page 492.
- Checking ongoing and completed tasks and performing system Audits. For more information, see [Activity Entities – Tasks and Audits](#) on page 493.
- Performance Monitoring, Application Discovery, and Reports Management. For details, see [Operations Entities](#) on page 499.
- Define policies for configuration ease. For more information, see [Policies in Infrastructure](#) on page 501.
- Additional operations that include how to find the Prism Central version, AHV version, AOS version, shutdown Prism Central, perform disaster recovery, reconfigure Prism Central IP address and gateway, pair Availability Zones (AZ), logout from Prism Central, and use keyboard shortcuts. For more information, see [Additional Operations](#) on page 581.
- Prism Central self service setup. For more information, see [Prism Self Service Administration](#) on page 569.
- Customer support services from Nutanix for Prism Central. For more information, see [Customer Support Services](#) on page 639.

For a quick overview about all the infrastructure-related operations that you can perform through Prism Central, see [Infrastructure Overview](#) on page 48.

The admin-related operations are covered in the *Prism Central Admin Center Guide*. For information about how the Prism Central documentation is organized, see [Prism Central Documentation Portfolio](#) on page 10.

For a quick overview about all the admin-related operations that you can perform through Prism Central, see [Admin Center Overview](#) on page 48.

Intended Audience

The intended audience for this document are the site personnels, customer's operation teams and any tier or individual entity defined by the customer who is equipped to perform the site configurations and modifications. In addition to this, any individual with an intention to learn and develop understanding about Prism Central usage, can refer this document.

PRISM CENTRAL DOCUMENTATION PORTFOLIO

The following table describes how Prism Central documentation portfolio is organized:

Table 1: Prism Central Documentation Portfolio

Deliverable Name	Description
Prism Central Admin Center Guide	This document describes all the administrative-level tasks that you can perform from Prism Central using Admin Center application. For more information about Admin Center application, see Admin Center Overview on page 48. For more information about how to access the Admin Center application in Prism Central, see Application Switcher Function on page 40.
Prism Central Alerts and Events Reference Guide	This document describes how to manage alerts and events management from Prism Central. The Alerts and Events generated in Prism Central enables you to perform the fault management related operations.
Prism Central Infrastructure Guide	This document describes all the field-specific settings required to host a cluster in Prism Central. It covers all the operations that you can perform from Prism Central using Infrastructure application. For more information about Infrastructure application, see Infrastructure Overview on page 48. For more information about how to access the Infrastructure application in Prism Central, see Application Switcher Function on page 40.
Intelligent Operations Guide	This document describes the Prism Central intelligent operations functionality, and how to enable, disable, and perform configurations related to Intelligent Operations part of Cloud Manager module.

RELATED DOCUMENTATION

The [Nutanix Support Portal](#) provides software download pages, documentation, compatibility, and other information.

Note: For information about the ports that are used by Prism Central, Microservices Infrastructure, and other Nutanix Software, see [Ports and Protocols](#).

On the [Ports and Protocols](#) page, select the appropriate software or application such as **Prism Central** or **Microservices Infrastructure** in the **Software Type** dropdown menu to view the ports and protocols used by that Nutanix software or application.

Documentation	Description
v4 API Documentation	v4 API Reference Documentation.
v4 API Release Notes	New API features, known issues and resolved issues related to v4 APIs.
Release Notes Flow Virtual Networking	Flow Virtual Networking Release Notes.
Port And Protocols	Port Reference: See this page for details of ports that must be open in the firewalls to enable Flow Virtual Networking to function.
Nutanix Security Guide	Prism Element and Prism Central security, cluster hardening, and authentication.
Flow Network Security Next Gen	Flow Network Security Next-Gen is the next-generation Nutanix microsegmentation solution with an enhanced policy model, advance policy operation, and enterprise readiness features.
AOS guides and release notes	Covers AOS Administration, Hyper-V Administration for Acropolis, Command Reference, Powershell Cmdlets Reference, AOS Family Release Notes, and AOS release-specific Release Notes.
Life Cycle Manager Guide	How to upgrade firmware and software components.
AHV guides and release notes	Administration and release information about AHV.
Prism Central and Prism Element Web Console guides and release notes	Administration and release information about Prism Central and Prism Element. Select Release in the Release Version selector at the top of the page.

GETTING STARTED

This section provides an overview of Prism Central that involves the following information:

- Prism Central landing page, Infrastructure dashboard and entities, statistics, application switcher, and Prism Central GUI Organization.
- Installation and upgrade specifics.
- Licensing information.
- Cluster registration and de-registration with Prism Central.
- Application-specific navigation bar and searchability attributes.
- Prism Central GUI access.

Prism Central Deployment

Ensure that you are familiar with the Prism Central requirements, limitations, and the procedures to install and upgrade Prism Central.

A Prism Central instance consists of either a single VM or as a three-VM scale-out architecture that functions as a multicluster monitoring and control interface. For information about the maximum tested and supported values for entities in Prism Central, see [Nutanix Configuration Maximums](#).

You can install Prism Central in either a Nutanix or a non-Nutanix environment.

The following table describes the supported methods for Prism Central installation and upgrade:

Table 2: Prism Central Deployment - Supported Methods

Environment Type	Hypervisor	Prism Central Upgrade using Upgrade Prism Central option	Prism Central Install		
			1-Click Internet with Prism Element	1-Click No Internet with Prism Element	Manual installation using binaries
Nutanix (Nutanix AOS cluster)	AHV	Not Supported	Supported	Supported	Not supported
	ESXi				
Non-Nutanix	ESXi	Not Supported	Not supported	Not supported	Not supported for pc.2023.4. For more information see the note following this table

Note: For information about installing or upgrading to Prism Central pc.2024.1 and later releases on a three-tier cluster, see [Prism Central Installation in a Non-Nutanix ESXi Environment](#) on page 20 and [Prism Central Upgrade in a Non-Nutanix ESXi Environment without AOS](#) on page 29.

Important considerations for Prism Central deployment.

- Prism Central version 7.3 onwards, deployment of a dual stack enabled Prism Central is supported. However, IPv6 address is currently supported only for deployment, Prism Element registration, NTP server configuration, DNS server configuration, and Prism Central pairing (if both Prism Central instances are dual stack enabled). You must not disable IPv6 address in a dual stack configuration.
- Do not add any additional vDisks or NICs to a Prism Central VM other than what is specified in the installation or upgrade instructions.
For information about the dedicated Prism Central upgrade vdisk that the Prism Central upgrade process deploys and uses for upgrade operations, see [Prism Central Upgrade vDisk](#) on page 26.
- You can upgrade Prism Central only in Life Cycle Manager (LCM). For more information, see [Upgrading Prism Central in Life Cycle Manager](#) on page 27.
- To scale out an existing single VM Prism Central to multiple VMs, see [Expanding \(Scale Out\) Prism Central](#) on page 619.
- To upgrade one or more managed clusters through Prism Central, see [Prism Central-Managed Clusters Upgrade](#) on page 27.
- If you use a proxy server, see [Prism Central Configuration When a Cluster Uses Proxy Servers](#) information in *Prism Central Admin Center Guide*.

Prism Central in a Non-Nutanix Environment

Prism Central in a non-Nutanix environment such as ESX.

The following limitations apply when you install Prism Central in a non-Nutanix ESXi environment:

- Only the small (S), single-VM Prism Central instance is supported in a non-Nutanix ESXi environment.
Increasing the size of the small Prism Central instance to a larger size is not supported.
Scaling out the single-VM Prism Central to a three-VM instance is not supported.
Nutanix recommends that you consider deploying Prism Central on a Nutanix cluster if your environment requires a large or an Extra large (XL) Prism Central.
- Prism Central in a non-Nutanix environment does not support the Direct Upload upgrade method. Use Life Cycle Manager in Prism Central Admin Center to upgrade the Prism Central instance in a non-Nutanix environment.
- Prism Central on a non-nutanix (ESXi) environment supports only basic Nutanix Cloud Infrastructure functionality for VM management, Cluster management and Storage management. Prism Central in a non-Nutanix environment does not support the following Prism Central features and functionalities:
 - Prism Central Backup and Restore
 - Nutanix Disaster Recovery
 - Nutanix products and portfolio services including Nutanix Cloud Manager, Flow including Flow Network Security Next-Gen and Flow Virtual Networking, and services enabled in Nutanix Marketplace such as Nutanix Objects, Files, File Analytics, and Mine v4.

Prism Central deployed on Hyper-V is no longer supported. Prism Central VM on Hyper-V hypervisor was last supported on pc.2022.6.0.12. Nutanix recommends moving to Prism Central in a Nutanix environment.

Requirements for Prism Central Deployment

Ensure that the requirements are met when you deploy (install or upgrade) Prism Central deployment at your site.

- When using the 1-click method in the Prism Element cluster for installing Prism Central, select from a set of preconfigured Prism Central VM sizes. The *Prism Central Scalability* section in the [Prism Central Release](#)

[Notes](#) provides Prism Central VM resource and capacity specifications for the preconfigured sizes. In addition, certain optional features that you can enable such as application discovery and Self-Service (formerly Calm) require additional Prism Central VM resources (more memory and vCPUs). After installing Prism Central, you might need to increase the VM resources before enabling one or more of these optional features. The requirements for various feature combinations are listed in [KB 8932](#).

- For Microservices Infrastructure deployment, ensure that all the requirements provided in [Microservices Infrastructure Prerequisites and Considerations](#) on page 585 are fulfilled.

Note: If you are upgrading the Prism Central instance from pc.2022.6 to pc.2023.3 or later, Microservices Infrastructure is enabled by default even if it was not enabled before upgrade. With pc.7.3 and later versions, Microservices Infrastructure is mandatory and enabled by default during installation or upgrade.

For information on installing Prism Central through Prism Element, see [Installing Prism Central Through Prism Element web console](#) on page 17

- When upgrading Prism Central and AOS, upgrade Prism Central first, then upgrade AOS on the clusters managed by Prism Central. See [Acropolis Upgrade Paths](#) on the Nutanix Support portal to check the upgrade path from your current version to your target version for AOS and Prism Central.

- If the hosting Prism Element is dual stack enabled, Prism Central must also be dual stack enabled.
- Do not use the AOS binary and metadata JSON files to upgrade your existing Prism Central deployment.

Prism Central requires the use of specific Prism Central binary and metadata .json for upgrades and deployments. Using AOS files to upgrade Prism Central from any version is not supported and can result in Prism Central unavailability.

Do not use the Prism Central AHV or ESXi components on the Nutanix support portal to upgrade an existing Prism Central instance.

- Prism Central and each managed cluster are taking advantage of NCC features, ensure that:
 - Each node in your cluster is running the same NCC version.
 - Prism Central and each cluster managed by Prism Central are all running the same NCC version.

Limitations of Prism Central Deployment:

These limitations apply when you deploy (install or upgrade) Prism Central at your site.

- Prism Central deployment is not supported on a mixed cluster consisting of different hypervisors.
- Number of clusters limit based on the Prism Central release. For more information, see [Nutanix Configuration Maximums](#).
- A dual stack enabled Prism Central is supported, but you must not disable IPv6 on a dual stack enabled Prism Central.
- If Prism Central is not registered with the AOS cluster hosting it, refer to [KB-10596](#) for additional requirements before upgrading.
- Prism Central 2022.9 and later versions no longer support CLI-based PCBR. A preupgrade check blocks upgrade to Prism Central 2022.9 or later versions if a CLI-based PCBR setup is detected. Nutanix recommends removing the CLI-based PCDR configuration before upgrade and enabling Prism Central Backup and Restore (PCBR) post-upgrade. For more information, see [Prism Central Backup, Restore, and Migration](#) on page 591, [KB-13875](#) and [KB 9599](#).
- Nutanix does not support Prism Central and its managed clusters in NAT-deployed environments.

- If a single-node Prism Central VM is upgraded to Prism Central 2022.9 or later versions, but the underlying Prism Element is not running AOS 6.6, a precheck prevents the scaling out of Prism Central VM from 1-node to 3-node Prism Central.
- Prism Central versions where one-click upgrade is not enabled must use the legacy one-click upgrade method available on Prism Element as the LCM bundle is not available on the portal for download. For more information, see [Use Upgrade Software in the Prism Element Web Console \(Legacy 1-Click Upgrade\)](#) in *Prism Element Web Console Guide*.

Deploy Multiple Prism Central Instances on a Prism Element Cluster

This use case describes the deployment of multiple Prism Central instances on a single Prism Element cluster.

You can deploy multiple Prism Central instances (single-VM or three-VM) on a single Prism Element cluster. The following requirements apply to deployment of multiple Prism Central instances on one Prism Element cluster:

- Ensure that none of the Prism Central instances are registered to the Prism Element cluster until all the Prism Central instances are deployed. You can register only one Prism Central instance to the Prism Element cluster.
- Ensure that you register the Prism Element cluster to a Prism Central only after all the Prism Central instances are deployed, if you are deploying multiple instances of Prism Central on the same Prism Element cluster.

The Prism Central instance that is registered with the Prism Element cluster manages the Prism Element cluster.

Note: When you unregister a Prism Element cluster from a Prism Central instance, the Prism Element cluster is blocked from being registered again with the same or another Prism Central instance. To unblock the Prism Element cluster, contact Nutanix Support.

For information on deploying multiple Prism Central instances on a Prism Element cluster that is already registered to a Prism Central instance, see [Deploying Multiple Prism Central Instances When One Instance is already Registered](#).

Deploying Multiple Prism Central Instances When One Instance is already Registered

If the Prism Element cluster is already registered to a Prism Central instance, follow these steps to add more Prism Central instances.

About this task

To deploy more Prism Central instances on a Prism Element cluster already registered to a Prism Central instance, follow these steps:

Procedure

1. Unregister the Prism Element cluster from the existing Prism Central instance.
2. Deploy all the Prism Central instances that you need.
3. Register the Prism Element cluster with the previously unregistered Prism Central instance.

Note: When you unregister a Prism Element cluster from a Prism Central instance, the Prism Element cluster is blocked from being registered again to the same or another Prism Central instance. To unblock the Prism Element cluster, contact Nutanix Support.

Prism Central Installation

You can install a Prism Central instance in a Nutanix or non-Nutanix environment.

In a Nutanix environment, you can install a Prism Central instance using the one-click method through the Prism Element web console. In a non-Nutanix environment, the process to install a Prism Central instance involves

downloading the Prism Central VM installation files, installing the Prism Central VM, and configuring the Prism Central VM.

Prism Central Installation in Nutanix Environment

This section provides information about how to install Prism Central in Nutanix environment.

You can install a Prism Central VM using the one-click method for both connected sites (with internet connectivity) and dark sites (without internet connectivity) only through the Prism Element web console. For more information, see [Installing Prism Central Through Prism Element web console](#) on page 17. This method employs the Prism Element web console from a cluster of your choice and creates the Prism Central VM in that cluster.

For information on how to install Prism Central in the non-Nutanix ESXi environment, see [Prism Central Installation in a Non-Nutanix ESXi Environment](#) on page 20.

Prerequisites for Prism Central Installation

Ensure that you meet the following prerequisites before you install Prism Central.

- Check the [Requirements for Prism Central Deployment](#) and [Limitations of Prism Central Deployment](#).
- Check the port requirements between Prism Central and Prism Element. For more information, see [Ports and Protocols](#).
- Check the following requirements for the connected site (with internet connectivity) and dark site (without internet connectivity) environment.

Table 3: Prism Central Installation Requirements for Connected Site and Dark Site

Connected site	Dark Site
• The specified gateway must be reachable. • Ensure the port TCP port 2100 is open from the Prism Element cluster to the Prism Central VM IP address. For the complete list of required ports, see Ports and Protocols. • Ensure network connectivity between the VM VLAN and portgroup of the Prism Element cluster Controller VM and the Prism Central VM VLAN and portgroup. • No duplicate IP addresses are used. • The storage container used for deployment is mounted on all hypervisor hosts. • When installing on an ESXi cluster: • The vCenter and the ESXi cluster must be configured properly. For more information, see vSphere Administration Guide for Acropolis. • The vCenter instance must be registered in Prism. • DRS must be enabled in vCenter. • The vCenter instance is up and reachable during the deployment.	Download Prism Central binary .TAR and metadata .JSON files from the Nutanix Support portal from a connected machine. From the Downloads page for Prism Central, download the Prism Central 1-click deploy from Prism Element binary .TAR and metadata .JSON files to your local media You can also copy these files to a USB stick, CD, or other media. For more information, see the note following this table.

Note: Do not use the Prism Central OVA, ZIP, AHV image, or AOS binary .TAR .GZ and upgrade metadata .JSON files from the Nutanix support portal to create this new Prism Central instance. Use the .TAR format binary and metadata .JSON files.

The Prism Central OVA is used only to install or upgrade Prism Central in non-Nutanix ESXi environment. For more information, see [Prism Central Installation in a Non-Nutanix ESXi Environment](#) and [Prism Central Upgrade in a Non-Nutanix ESXi Environment without AOS](#).

Installing Prism Central Through Prism Element web console

You can install Prism Central in a connected site or dark site Nutanix environment.

Before you begin

Ensure that you meet the prerequisites provided in [Prerequisites for Prism Central Installation](#) on page 16 before you install Prism Central.

About this task

Perform this procedure for both connected site (with internet connectivity) and dark site (without internet connectivity).

Procedure

1. Log in to the Prism Element web console as the admin user for your cluster.
2. Run NCC as described in [Run NCC Checks](#).
3. Do one of the following:
 - » Click **Deploy New Prism Central** from the **Prism Central** widget on the **Home** dashboard
 - » Click the gear icon in the main menu. From the **Settings** menu, select **Prism Central Registration**, and then click **Deploy New Prism Central**.
4. In the **Prism Central (PC) Deployment** dialog box, in the **Version** section, enter the following information:
 - a. In **PC Details** > **Create PC Name** provide a name for the Prism Central instance.
 - b. [Applicable to Dark site only] In **Available Versions**, click the **Upload Installation Binary** link, select the Prism Central metadata file (.json) and Prism Central installation binary (.tar) files, and then click **Upload**.

If there is an image already uploaded, the system displays the available versions.
 - c. [Applicable to Connected site only] Select the required Prism Central version you must install.

Select **Only show compatible versions** checkbox to view the list of Prism Central versions compatible with the AOS cluster.

Note: If the Prism Central version you install does not appear in the list, typically because the cluster does not have Internet access (such as at a dark site), you can click the **Upload Installation Binary** link to upload an image from your local media.

5. Click **Next**.
6. In the **Size and Scale** tab, select the size and scale of the Prism Central deployment.

For more information on supported scalability and resource requirements of each size of Prism Central deployment, see the *Prism Central Scalability* topic in the release notes for [Prism Central](#) version to be installed.

 - » To deploy the smallest instance of Prism Central to explore the basic Prism Central capabilities, select **Extra-Small (XS)**. For more information about Extra-Small Prism Central instance, see [X-Small Prism Central](#) on page 55.
 - » To deploy a small instance of Prism Central, select **Small (S)**.
 - » To deploy a large instance of Prism Central, select **Large (L)**.
 - » To deploy an extra-large instance of Prism Central, select **Extra-Large (XL)**.

Select **Enable High Availability (HA)** to scale out your Prism Central deployment with three times resources and make it more resilient. The **Enable High Availability (HA)** checkbox appears only after you select the size of Prism Central deployment.
7. Click **Next**.

8. In the **Configuration > Networking Details** section, follow these steps to provide the information in the indicated fields:
 - a. **Network:** Select an existing network for this Prism Central instance from the list.
If the target network is not listed, click the **Create Network** link to create a new network. For more information, see [Network Management](#).
 - b. **Subnet Mask:** Enter the subnet mask value.
 - c. **Gateway IP:** Enter the IP address of the gateway.
 - d. **DNS Address(es):** Enter the IP address for one or more DNS servers in FQDN, IPv4 or IPv6 format. Enter multiple addresses in a comma separated list.
 - e. **NTP Address(es):** Enter the IP address for one or more NTP servers in FQDN, IPv4 or IPv6 format. Enter multiple addresses in a comma separated list.
 - f. **Container:** Select a container for the Prism Central VM from the drop-down list.
9. In the **Configuration > Networking Details** section, enter the following information:
 - **VM Name:** Enter a name for the Prism Central VM.
 - **IP Address:** Enter an IP address for Prism Central VM.

If you have selected **Enable High Availability (HA)** in the **Size and Scale** section, the IP addresses for the three Prism Central VMs are populated automatically when you enter virtual IP. You can keep those addresses or change them as desired.
10. In the **Internal Network Configuration** section, follow these steps to enter the necessary information in the indicated fields:
 - a. **Use default settings (recommended) :** Select the checkbox for Prism Central to use the default values for **Subnet Mask, Gateway IP Address** and **IP Address Range**. Clear the checkbox if Prism Central must use the specific (non-default or custom) values for **Subnet Mask, Gateway IP Address** and **IP Address Range**.
The following three fields are available only if you choose not to use the default settings.
 - b. **Subnet Mask:** Enter the subnet mask.
 - c. **Gateway IP Address:** Enter the IP address for the gateway.
 - d. **Gateway IP Address:** Enter a range of IP addresses that the network can use.
Enter a range of at least 10 available (unreserved) addresses. For a managed network, the range of addresses for microservices infrastructure must be outside the range of reserved IP addresses (for example, DHCP IP Pool) in the selected network.
 - e. After you check that the values entered for all the fields are correct, click **Deploy** to begin the deployment process.
On the **Home** page, the Prism Central widget displays **Deploying** until the installation is completed, then it displays **OK**. Click **OK** to launch the Prism Central web console in your browser.
11. Monitor the deployment progress from the **Tasks** page and view information about the deployed VMs through the **VM** dashboard. For more information, see [Tasks View](#).

What to do next

you can expand your cluster, remove nodes, and enable services such as Self-service, Flow Network Security, and Intelligent Operations without registering the hosting cluster with Prism Central.

Register this cluster with the Prism Central instance to manage it using Prism Central. For more information about how to connect to an existing Prism Central instance, see [Registering or Unregistering a Cluster with Prism Central](#).

Prism Central Installation in a Non-Nutanix ESXi Environment

You can install Prism Central in a non-Nutanix environment.

For information about conditions applicable to Prism Central in a non-Nutanix ESXi environment, see [Prism Central in a Non-Nutanix Environment](#) on page 13.

Use the procedures detailed in this section to deploy Prism Central only within an on-premises non-Nutanix environment.

Important: Ensure that you use the necessary ESXi credentials with the following conditions to install or upgrade a Prism Central instance in a non-Nutanix ESXi environment:

- The ESXi credentials must provide **Administrator** privileges.
- The user name in the ESXi credentials must be specified in the **user@domain** format if the credentials are for a local or LDAP user.
- The user name in the ESXi credentials is case-sensitive and must be specified in the lower case. For example if the user name is **administrator**, specify it as **administrator@vsphere.local**. The Microservices Infrastructure enablement fails if you specify it as, for example, **Administrator@vsphere.local**.
- The password in the user credentials must be a never-expiring password necessary to keep the pods running during the installation or upgrade operation. Ensure that you provide the password without double quotes (" ").

The vSphere CSI driver requires these credentials to perform the following VM updates and data store operations:

- VM update operations.
- Perform read and write operations on a data store.
- Query storage policies.

For information on permissions, key roles, and privileges for the vSphere User, see the *Preparing for Installation of vSphere Container Storage Plug-in* document in the VMware documentation.

ESXi Credentials for Non-Nutanix Environment

Ensure that you use the necessary ESXi credentials with the conditions to install or upgrade a Prism Central instance in a non-Nutanix ESXi environment.

- The ESXi credentials must provide **Administrator** privileges.
- The user name in the ESXi credentials must be specified in the **user@domain** format if the credentials are for a local or LDAP user.
- The user name in the ESXi credentials is case-sensitive and must be specified in the lower case. For example if the user name is **administrator**, specify it as **administrator@vsphere.local**. The Microservices Infrastructure enablement fails if you specify it as, for example, **Administrator@vsphere.local**.
- The password in the user credentials must be a never-expiring password necessary to keep the pods running during the installation or upgrade operation. Ensure that you provide the password without double quotes (" ").

The vSphere CSI driver requires these credentials to perform the following VM updates and data store operations:

- VM update operations.
- Perform read and write operations on a data store.
- Query storage policies

For information on permissions, key roles, and privileges for the vSphere User, see the *Preparing for Installation of vSphere Container Storage Plug-in* document in the VMware documentation.

Downloading Prism Central VM Installation Files (Manual Method)

This section describes how to download Prism Central VM installation files for manual installation (when the Prism Central 1-click deployment method is not available).

Procedure

1. Log on to the Nutanix Support portal and go to the [Prism Central Downloads](#) page.
The page lists the latest available version. You can download previous release bits from the **Other Versions** tab.
2. Locate the **Prism Central Install on ESX (Version: *pc_version*)** file, and click **Download** to get the Prism Central VM files, then save them to a folder, directory, or other location.
The system downloads a *pc_version.ova* file, where *pc_version* is the Prism Central version.
For example, the downloaded file for Prism Central pc.202x.x is named pc.202x.x.ova.

Installing the Prism Central VM (ESXi)

About this task

A special VM runs Prism Central from which you can monitor and manage multiple clusters. You can install the Prism Central VM on an ESXi hypervisor.

Before you begin

- Nutanix does not support deploying a large Prism Central in a non-Nutanix environment. Increasing the size of a Prism Central in a non-Nutanix ESXi environment to large Prism Central is also not supported. Nutanix recommends that you consider deploying Prism Central on a Nutanix cluster if your environment requires a large Prism Central.
- Expanding (scaling out) Prism central feature is not supported for a three-tier non-Nutanix Prism Central deployment.
- Nutanix recommends a VM size of 6 vCPUs, 28 GB of memory, and 500 GiB of storage for a small Prism Central VM deployment.
- Ensure to enable memory hot-plug for ESXi. For 1-click deployment methods, memory hot-plus is automatically enabled.
- Ensure that you have disabled Intelligent Operations if the **Preupgrade Steps** fail due to a memory issue. For information on disabling Intelligent Operations, see [Intelligent Operations Guide](#).

To install the Prism Central VM on an ESXi hypervisor, follow these steps:

Procedure

1. Download and extract the Prism Central files for ESXi as described in [Downloading Prism Central VM Installation Files \(Manual Method\)](#) on page 21.

2. Install the downloaded OVA file as follows:

- a. Connect to vCenter or an ESXi host using the vSphere client.
- b. Select the OVA file and deploy it.

Select an OVA file called `pc_version.ova`, where `pc_version` is the version of Prism Central. Deploy the OVA file from the download location on the workstation. See the vSphere documentation for instructions on how to do deploy an OVA file.

Note:

- Nutanix recommends a VM size of 6 vCPUs, 28 GB of memory, and 500 GiB of storage for a small Prism Central VM deployment.
Ensure that you have disabled AIOps if the **Preupgrade Steps** fail due to a memory issue. For information on disabling AIOps, see [Intelligent Operations Guide](#).
- In the OVA deployment wizard using vSphere web client, Nutanix recommends that you manually select **Thin Provision** as the virtual disk format while configuring the storage. Failing to manually select **Thin Provision** creates the default thick provisioned VMDK in vSAN.

3. Power on the Prism Central VM through the vSphere Client. Click **ACTIONS > Power > Power On** to start the VM.

Wait for the Prism Central VM to reboot and complete the `nutanix-rc-local` service start.

Caution: Do not configure the next steps before the Prism Central VM boots up completely and the `nutanix-rc-local` service completes successfully. If you do and the Prism Central VM does not boot up, the Prism Central installation fails, and you cannot complete the installation.

Run the following command to verify that the **Process** status of the `nutanix-rc-local` service as `<PID> ExecStart=/etc/nutanix/rc.local (code=exited, status=0/SUCCESS)`.

```
nutanix@pcvm$ sudo systemctl status nutanix-rc-local.service
```

4. Upgrade the Prism Central VM compatibility (VMX version).

- a. Right-click on the Prism Central VM, and navigate to **Compatibility > Upgrade VM Compatibility**.
- b. Set the **Compatible with** field to **ESXi 7.0 U2 and later**.
- c. Power on the Prism Central VM.

5. Log into the Prism Central VM through the vSphere console (user name nutanix).

6. Assign a static IP address to the Prism Central VM as follows:

- a. Open the `ifcfg-eth0` file for editing.

The following command opens the file using the vi editor:

```
nutanix@pcvm$ sudo vi /etc/sysconfig/network-scripts/ifcfg-eth0
```

- b. Add or update the `NETMASK`, `IPADDR`, `BOOTPROTO`, and `GATEWAY` entries as needed.

Note: Use the `route -n` command to retrieve the `NETMASK` and `GATEWAY` information. Ensure that the selected IP address is unreachable to the console IP address.

```
NETMASK="xxx.xxx.xxx.xxx"
IPADDR="xxx.xxx.xxx.xxx"
BOOTPROTO="none"
```

```
GATEWAY="xxx.xxx.xxx.xxx"
```

- Enter the desired netmask value in the **NETMASK** field. (Replace `xxx.xxx.xxx.xxx` with the appropriate value.)
- Enter the appropriate static IP address (usually assigned by your IT department) for the Prism Central VM in the **IPADDR** field.
- Enter **none** as the value in the **BOOTPROTO** field. (You might need to change the value from `dhcp` to `none` if you employ DHCP. Only a static address is allowed; DHCP is not supported.)
- Enter the IP address for your gateway in the **GATEWAY** field.

Warning: Carefully check the file to ensure there are no syntax errors, whitespace at the end of lines, or blank lines in the file.

- c. Save the changes.
- d. Remove any existing Nutanix Controller VM entries, that is ones which include `NTNX-<number>-CVM`, from the `/etc/hosts` file. (Ensure that you do not remove any other entries from the file.)

To edit the file using `vi`, enter

```
nutanix@pcvm$ sudo vi /etc/hosts
```

When the Prism Central VM starts the first time and DHCP is enabled in the network, an entry similar to the following is added to the `/etc/hosts` file. This entry (if present) needs to be removed from the `/etc/hosts` file before restarting the Prism Central VM in the next step, which generates a new entry if DHCP is enabled.

```
127.0.0.1 NTNX-10-3-190-99-A-CVM
```

- e. Restart the Prism Central VM.

```
nutanix@pcvm$ sudo reboot
```

Note: The Prism Central VM is powered in approximately a minute.

7. Set up the installation environment.

For more information, see [Setting Up the Installation Environment](#) on page 24.

After completing the steps in [Setting Up the Installation Environment](#) on page 24, continue to perform the following steps.

8. Log on to vCenter Server using an administrator account and do the following:

- Select the Prism Central VM.
- Right click **Edit Notes**.
- In the dialog box enter `NutanixPrismCentral` without any new line characters and then click **OK**.

Creating a Prism Central instance through a 1-click method automatically enables certain best practice restrictions on who can manage a Prism Central VM. This step enables the same set of best practice restrictions for a manually installed Prism Central VM. (Non-administrators may not perform any operations while administrators may perform a limited set of permissions including launch console, power off or on, pause or resume, migrate, and configure VM host affinity.)

What to do next

This completes Prism Central installation. The next step is to register clusters with Prism Central. Port 9440 needs to be open in both directions between the Prism Central VM and any registered clusters. For more information on ports, see [Ports and Protocols](#) for Prism Central and [Ports and Protocols](#) for

Microservices Infrastructure. For more information on how to connect to an existing Prism Central instance, see [Registering a Cluster with Prism Central](#) on page 56.

Setting Up the Installation Environment

You must set up the environment for installing Prism Central pc.2024.1 or later releases, when you are installing Prism Central for the first time.

Before you begin

Complete the steps mentioned before [step 7](#) in the [Installing the Prism Central VM \(ESXi\)](#) on page 21 section.

About this task

The installation environment setup involves the following phases:

- Installing the Prism Central pc.202x.x release using the OVA file. For more information, see [Installing the Prism Central VM \(ESXi\)](#) on page 21.
- Downloading the **vSphere CSI Bundle** from the Nutanix **Downloads** portal for [Microservices Platform](#).
- Setting up the installation tool.
- Setting up target Prism Central cluster configuration.
- Setting up the site proxy environment and defining storage policies.

Procedure

1. Log on to the [Nutanix Support](#) portal and browse to the [Microservices Platform \(MSP\)](#) page.
2. Download the CSI bundle and extract its contents into the /home/nutanix/data directory on the Prism Central VM.
The **vSphere CSI Bundle** contains essential tools such as `setup_three_tier`.

3. Extract the vSphere CSI bundle tar.xz file by running the following command:

```
nutanix@pcvm$ cd /home/nutanix/data
nutanix@pcvm:/data$ tar -Jxvf vsphere_csi_bundle.tar.xz
```

4. SSH to the Prism Central VM as a `nutanix` user.

5. Run the `setup_three_tier` tool.

The `setup_three_tier` tool creates the Prism Central environment configuration file.

For example:

```
nutanix@pcvm$ /home/nutanix/data/vsphere_csi_bundle/tools/setup_three_tier -e
```

6. Log on to the Prism Central VM again through the vSphere console and then run the command to create a cluster, specifying additional parameters such as DNS servers and NTP servers.

```
nutanix@pcvm$ cluster --cluster_function_list "multicluster" -s
static_ip_address --dns_servers "dns_server_ip_1,dns_server_ip_2" --ntp_servers
"0.centos.pool.ntp.org,1.centos.pool.ntp.org" create
```

The `static_ip_address` is the Prism Central VM IP address.

The `dns_server_ip_1` and `dns_server_ip_2` are the IP addresses of the DNS server.

When you enable Microservices Infrastructure, the services deployed on Microservices Infrastructure require dynamic storage allocation. To allocate the necessary storage when the services require it (on-demand storage),

you must log on to the Prism Central VM through the vSphere console. The Microservices Infrastructure CSI driver for vSphere performs the following essential storage operations when you log on to the Prism Central VM through the vSphere console only:

- Communicates with the vCenter to provision new storage
- Performs VM update API calls to attach it to the Prism Central VM
- Makes the attached storage available to services as persistent storage.

7. (Optional) Configure proxy settings.

If a proxy setup is in place, ensure to configure Microservices Infrastructure (CMSP) URLs in the allowedlist in your proxy environment to allow access to docker images and other configuration files. Microservices Infrastructure requires access to the following URLs:

- docker.io
- production.cloudflare.docker.com
- nutanix.github.io
- us-west-2.amazonaws.com
- quay.io

8. In the vSphere client, configure the vCenter settings:

- a. Create a tag and a category. Map the category to the tag.
 - b. Assign the tag to the datastore hosting the Prism Central VM.
 - c. Create a new VM storage policy with global permission, and assign it to the newly created tag.
- For more information, see *vSphere Tags and Attributes* information in *VMWare vSphere documentation*.

Note: Ensure that you select **Enable tag based placement rules** checkbox while defining the storage policy structure.

What to do next

Complete the steps after the [set-up the environment step](#) in [Installing the Prism Central VM \(ESXi\)](#) on page 21.

Configuring Prism Central VM in the Non-Nutanix ESXi Environment

This task describes how to configure Prism Central VM in the non-Nutanix environment running the ESXi hypervisor.

Before you begin

Ensure that you set up the install environment before you run the install tool (`setup_three_tier` tool). For more information, see [Setting Up the Installation Environment](#) on page 24.

Ensure that you use the necessary ESXi credentials. For more information, [Prism Central Installation in a Non-Nutanix ESXi Environment](#) on page 20.

Procedure

1. SSH to the Prism Central VM and run the `setup_three_tier` tool.

The `setup_three_tier` tool accepts the user configuration and sets up the environment for PC and CMSP enablement.

For example:

```
nutanix@pcvm$ /home/nutanix/data/vsphere_csi_bundle/tools/setup_three_tier
```

The system prompts you to enter the vCenter IP, vCenter port, user credentials, datacenter name, and storage policy name.

2. Enter the required details.
The system verifies the entered data, and displays the result.
3. Check that the Prism Central ENV (`pc_env_config.json`) configuration file is created at `/etc/nutanix`.
4. Run the `Enable Microservices Infrastructure` script.

Note: Before you run the `Enable Microservices Infrastructure` script, ensure that you have configured the DNS and NTP servers for the cluster.

SSH to the Prism Central VM and run the script as follows:

```
nutanix@pcvm$ python /home/nutanix/data/vsphere_csi_bundle/tools/enable_cmsp_3_tier.py
```

All the files required for deploying and configuring the Prism Central VM are packaged in the Prism Central deployment package.

Enabling Microservices Infrastructure could take up to 30 minutes for the supported single VM Prism Central deployment. To confirm if Microservices Infrastructure is enabled, run the `ecli task.list` command.

Prism Central reports any failure of the enablement after about two hours and five minutes.

Prism Central Upgrade

You can upgrade a Prism Central deployment in a Nutanix or non-Nutanix environment.

In a Nutanix environment, you can upgrade a Prism Central deployment only through the Life Cycle Manager (LCM) tool. In a non-Nutanix environment, the process to upgrade a Prism Central deployment differs based on the base Prism Central version from which you are upgrading.

Prism Central Upgrade in Nutanix Environment

To upgrade a Prism Central deployment in Nutanix environment, you must use the Life Cycle Manager (LCM) tool only.

Using Prism Central web console, you can upgrade AOS on some or all of the clusters registered to and managed by Prism Central. The upgrade procedure, known as 1-click centralized upgrade, enables you to upgrade each managed cluster to a specific version compatible with Prism Central.

For information on how to upgrade Prism Central in non-Nutanix ESXi environment, see [Prism Central Upgrade in a Non-Nutanix ESXi Environment without AOS](#) on page 29.

Prism Central Upgrade vDisk

Prism Central deploys a separate vDisk for upgrade on the Prism Element cluster hosting Prism Central.

Each Prism Central VM includes a `/home` disk partition that is used for storage during upgrade process. The Prism Central upgrade process can fail if the available disk space in this partition is near capacity.

For Nutanix clusters, the upgrade process uses a separate, dedicated vDisk of 30 GB capacity instead of the `/home` partition. Prism Central uses this dedicated vDisk to store the downloaded Prism Central software bundle, decompressing the bundle, and hosting the installer.

When you upgrade Prism Central version to the current version, the Prism Central upgrade process uses the vDisk image in the bundle to create the Prism Central upgrade vDisk. The upgrade process uses this vDisk for the upgrade operations that include any subsequent upgrade.

Requirements

- Ensure that the Prism Central instance is registered to a host in the cluster or microservices infrastructure is enabled on the Prism Central instance.
- Ensure that the Prism Element cluster hosting Prism Central is registered to this Prism Central instance.
- Ensure that sufficient disk resources to provide the 30 GB vDisk is available on the Prism Element cluster hosting Prism Central.

Upgrading Prism Central in Life Cycle Manager

You can upgrade a Prism Central deployment in Nutanix environment through the Life Cycle Manager (LCM) tool only.

About this task

To upgrade Prism Central in LCM, follow these steps:

Procedure

1. Log in to Prism Central.
2. Select the **Admin Center** application from [Application Switcher Function](#), and navigate to **LCM** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#).
3. Select the **Prism Central Cluster** tab and proceed with the next steps.
4. Choose either of the following upgrade methods based on the LCM connectivity status with the site:
 - » *Upgrade Prism Central from LCM with Connected site:* Perform the actions as described in [Performing Firmware and Software Updates in a Connected Site Setup](#) section of *Life Cycle Manager Guide*.
 - » *Upgrade Prism Central from LCM with Dark site using a web server:* Perform the actions as described in [Performing Firmware and Software Updates in a Dark Site Setup](#) section of *Life Cycle Manager Guide*.
 - » *Upgrade Prism Central from LCM with Dark site using Direct Upload:* Starting with LCM 3.0, Prism Central also supports the **Direct Upload** functionality. A minimum supported version of pc.2024.1 is required to use the Direct Upload functionality on Prism Central. Perform the actions as described in [Performing Firmware and Software Updates in a Dark Site Setup](#) section of *Life Cycle Manager Guide*.

Prism Central-Managed Clusters Upgrade

You can upgrade AOS on some or all of clusters registered to and managed by Prism Central through the Prism Central web console.

You can upgrade AOS on some or all of clusters registered to and managed by Prism Central. The upgrade procedure, known as 1-click centralized upgrade, is used to upgrade each managed cluster to a specific version compatible with Prism Central. [Acropolis Upgrade Paths](#) lists the upgrade path from your current version to your target version for AOS, Prism Central, and Nutanix Files.

Cluster Upgrade Requirements

- Upgrade clusters through the Prism Central web console.
- You can upgrade clusters if **Available Versions** is shown in the **Upgrade Software** dialog. It will display one or more AOS versions available to apply to clusters. If no versions are available (that is, none displayed), you might have to upgrade each cluster according to procedures described in this guide. Basically, if a version is available to Prism Central, it will be available to each cluster registered.
- Prism Central and each cluster (individually also known as Prism Element) must be connected to the Internet to access the Nutanix support portal.
- Prism Central and each cluster must be configured with a domain name server. For more information, see [Configuring Name Servers \(Prism Central\)](#) or [Configuring Name Servers](#) in the *Prism Element Web Console Guide*.
- Dark-site (internet-disconnected) upgrades cannot be performed using this feature. In this case, upgrade each cluster according to procedures described in this upgrade section.
- You cannot initiate an upgrade on a cluster that has an upgrade that is in process or is not completed.

Cluster Upgrade Features

- Sequential upgrade, in order, one cluster at a time. Select the order in which clusters are upgraded. Upgrade one cluster, then the next cluster is upgraded until all upgrades are complete. The clusters waiting be upgraded are placed in a Queued status. If one cluster fails to upgrade for some reason, the upgrade is cancelled for all remaining queued clusters.
- Parallel upgrade. Upgrade all clusters at the same time. If one cluster fails to upgrade for some reason, it does not affect the upgrade status of the other clusters. The upgrade status of each cluster is independent from every other cluster.
- Labeled clusters. You can apply a label to selected clusters and upgrade just the labeled clusters. The label also filters clusters by label.

Upgrading Multiple Managed Clusters

You can upgrade multiple Prism Central-managed Prism Element clusters.

About this task

To upgrade multiple Prism Central-managed clusters, follow these steps.

Procedure

1. Log in to the Prism Central web console.
2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Hardware > Clusters** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#).
3. Click the **List** tab and do one of the following:
 - » Select one or more clusters from the list.
 - » Click **Filters** and select a cluster label if you have already created a label to group specific clusters.

4. Click **Actions** > **Upgrade Software** to launch the **Upgrade Software** dialog box.
 - a. **Available Versions** lists AOS versions available to apply to clusters
 - b. **Upgrade Sequence** selects parallel upgrade (upgrade all selected clusters **At the same time**) or **One at a time (custom Order)** (upgrade one cluster at a time in order)
 - c. List of clusters to select for upgrade order if you selected **Custom Order**




Figure 1: Upgrade Software

5. Select an AOS version and **Upgrade Sequence** choice, then click **Upgrade**.
The upgrade begins. To view upgrade progress and status, click the blue task icon to open the **Tasks** page. The **Cluster** pane also displays **Upgrade Status**.
6. On the **LCM** page, click **Inventory** > **Perform Inventory** to enable LCM to check, update and display the inventory information.
For more information, see [Performing Inventory with the Life Cycle Manager](#) section in *Life Cycle Manager Guide*.

Prism Central Upgrade in a Non-Nutanix ESXi Environment without AOS

You can upgrade the Prism Central instance in a non-Nutanix environment consisting of ESXi hypervisors without any AOS deployment.

Upgrade Prism Central deployed in a non-Nutanix environment using the procedures provided in the [Upgrading Prism Central in Life Cycle Manager](#) on page 27 topic only when you upgrade the Prism Central VM from pc.2024.2 to later versions.

If you are upgrading Prism Central in a non-Nutanix environment from a version earlier than pc.2024.2, upgrade it in the following stages:

- Upgrade Prism Central to pc.2024.2. For information, see [Prism Central Deployment](#) in the *Prism Central Infrastructure Guide pc.2024.2*.
- Upgrade Prism Central from pc.2024.2 to the latest version using the procedures provided in the [Upgrading Prism Central in Life Cycle Manager](#) on page 27 topic.

Service Manager for Installing Services and Applications

The Service Manager feature is introduced to enable the installation of Prism Central services and application in a dark site. Service Manager enables the installation only when the Prism Central MSP Apps Bundle (pc-app-onprem bundle) packaged with the necessary charts and images is available in the LCM web server for the installation of the following services and applications:

- Security Dashboard
- Licensing application in Prism Central Admin Center
- Batch Service

Note: Ensure that the LCM web server is reachable by all the Nutanix clusters in the dark site. For information on preparing the LCM web server, see the [Local Web Server Setup](#) topic in the *Life Cycle Manager Guide*.

Service Manager downloads the necessary charts and images for the installation from the Prism Central MSP Apps Bundle that is placed in the LCM web server. For information on placing the Firmware and Software LCM Bundles on the local web server, see [Placing the Firmware and Software LCM Bundles on a Web Server for Prism Central](#) topic in the *Life Cycle Manager Guide*. The charts and images for the mentioned services are available in the LCM (Open Container Initiative-compliant) container registry, also called the LCM OCI registry.

Initiating Service Manager for a New Prism Central Installation

This section describes how to initiate Service Manager when you install Prism Central in a dark site for the first time.

Before you begin

Ensure that you have the following bundles in the dark site LCM web server.

- The Prism Central LCM Bundle
- The Prism Central MSP Apps Bundle that contains the necessary charts and images.

For information on preparing the LCM web server, see the [Life Cycle Manager Guide](#).

About this task

Perform the following steps to install Prism Central with the Prism Central services and applications

Procedure

1. Install Prism Central in the dark site using the Prism Central LCM bundle.
2. Configure the LCM web server in the Prism Central Admin Center as the source for the installation bundles.
On **Prism Central Admin Center > LCM > Settings**, select **Dark Site (Local Web Server)** in the **Source** field and provide the URL of the LCM web server in the **URL** field.

- SSH to the Prism Central VM as a `nutanix` user, and run the following commands to restart the `pc_platform_bootstrap` service on the dark site cluster:

```
nutanix@pcvm$ allssh genesis stop pc_platform_bootstrap
```

This command stops the `pc_platform_bootstrap` service.

```
nutanix@pcvm$ allssh cluster start
```

This command restarts the `pc_platform_bootstrap` service on all the Prism Central VMs.

Restarting the `pc_platform_bootstrap` starts the Service Manager which installs the services and applications.

Initiating Service Manager for a Prism Central Upgrade

This section describes how to initiate Service Manager when you upgrade existing Prism Central in a dark site.

Before you begin

- Ensure that you have the following bundles in the dark site LCM web server.
 - Prism Central LCM bundle
 - The Prism Central MSP Apps Bundle that contains the necessary charts and images.
- For information on preparing the LCM web server, see the [Life Cycle Manager Guide](#).
- Install Helm 3.0 for Service manager in the dark site.
 - Ensure that you have configured the LCM web server in the Prism Central Admin Center as source for the upgrade bundles.

About this task

Perform the following step to install Prism Central with the Prism Central services and applications

Procedure

- Upgrade Prism Central in the dark site using the Prism Central LCM bundle.

When you upgrade Prism Central to pc.2024.1 or later versions, Service Manager installs the services and applications using the Prism Central MSP Apps Bundle in the LCM web server.

Prism Central Overview

Prism Central provides a workspace to monitor and manage multiple clusters from a centralized environment. It runs as a separate instance that consists of either a single VM or as a three-VM scale-out architecture. For more information, see [Expanding \(Scale Out\) Prism Central](#) on page 619.

Prism Central provides the following functionalities:

- Platform-related functionalities that include:
 - Configuration and management of administrative activities such as Licensing, Life Cycle Manager (LCM), Identity and Access Management (IAM), Projects, Prism Central UI specific settings, and admin-specific settings, which are applicable across platforms and/or for various Nutanix Apps such as NCM Self-Service, Files, Move, Objects, Database Service, Foundation Central, or Kubernetes Management. IAM, LCM, or Projects are applicable for both Nutanix Apps and platforms.
 - Discover, deploy, and manage Nutanix Apps such as *NCM Self-Service*, *Files*, *Move*, *Objects*, *VolumesDatabase Service*, *Foundation Central*, or *Kubernetes Management* and other hybrid cloud applications (preferred partner apps) from Prism Central.

- Alerts and Notifications related configurations. For more information, see [Configuring Alerts in Prism Central](#) information in *Prism Central Admin Center Guide*.
- Security-related configurations.

For more information about all the admin-related activities that you can perform using Prism Central, see [Admin Center Overview](#) on page 48.

- Cloud Infrastructure-related functionalities that include:
 - Site-specific configurations and operations.
 - Fault and performance management operations. For more information, see [Prism Central Alerts and Events Reference Guide](#).

For more information, see [Infrastructure Overview](#) on page 48.

- Cloud Manager-related functionalities that include accessing the Cost Governance (formerly called as Beam), NCM Self-Service and Security Central applications, and configuring Prism Intelligent Operations. For more information, see the following documentation:
 - [Cost Governance](#)
 - [Self-Service](#)
 - [Security Central](#)
 - [Intelligent Operations](#)
- The Nutanix Intelligent Virtual Agent (NIVA) functionality that provides in-product contextual recommendations, tips, and help resources for blueprints, runbooks, playbooks, and other features of Self-Service and Intelligent Operations.

Important: Starting with Prism Central version pc.7.3.1, Nutanix rebranded Intelligent Assist to Nutanix Intelligent Virtual Agent (NIVA). While the user interface for Prism Central version pc.7.3 still shows the name Intelligent Assist, all current documentation and resources use the updated name, either Nutanix Intelligent Virtual Agent or NIVA.

For more information, see the [Nutanix Intelligent Virtual Agent Guide](#).

Note:

- Prism Central supports maximum 10 concurrent user sessions (UI or API or a combination of both).
- From Prism Central version pc.2024.3 onwards, the registration dependency between Prism Central and the hosting cluster has been removed for Prism Central scaling-out, infrastructure operations, and service enablement. Now, you can deploy Prism Central on a hosting cluster and then expand your cluster, remove nodes, and enable services such as Self-service, Flow Network Security, and Intelligent Operations without registering the hosting cluster with Prism Central. A few services such as Move, Nutanix Database Service (NDB), and Nutanix Kubernetes Engine (NKE) still need the cluster to be registered with Prism Central. You need to register your cluster with Prism Central if you want to manage your cluster through Prism Central.

Logging Into Prism Central

About this task

This section describes how to log into Prism Central. The Knowledge Base article [KB 1661](#) lists the default cluster credentials.

Procedure

To log in to Prism Central, perform the following steps:

1. Open a web browser, enter `https://prism_central_ip_addr:9440` in the address field, and press **Enter**.
Replace `prism_central_ip_addr` with the Prism Central VM IP address.

Note: Prism Central supports the latest version, and the two preceding major versions of Firefox, Chrome, Safari, and Microsoft Edge browsers. The browsers must support TLS 1.2.

The system redirects to the encrypted port (9440) and prompts you to accept the **Terms and Conditions**. The system might display an SSL certificate warning. Acknowledge the warning and proceed to the site. If user authentication is enabled and the browser does not have the correct certificate, a denied access message may appear. For the complete list of required ports, see [Port Reference](#).

2. Click **Accept terms and conditions**. The system displays the login screen.

3. Enter your Nutanix login credentials and press **Enter** or click **Log In**.

The login credentials might vary based on the authentication mechanism configured to access Prism Central. For more information about supported authentication mechanisms for Prism Central, see [Configuring Authentication](#) information in *Security Guide*.




Figure 2: Login Screen

Note:

- If LDAP authentication is used, enter the user name in *username@domain* format; the *domain\username* format is not supported. The user principle name [UPN] attribute is used to find the user account in Active Directory, so your user name must be in that format.

Note: If there is more than one LDAP/AD, the Prism Central login screen shows multiple boxes to select the configured domain names.

- The login page includes background animation that is enabled by default. You can modify it based on your requirement. For more information about how to configure Prism Central UI Settings, see *Prism Central Admin Center Guide*.
- If you are logging in as an administrator for the first time (user name admin and default password Nutanix/4u), the system prompts you to change the default password. Enter a new password in the **password** and **re-type password** fields and press **Enter**.

Note: The system also prompts you to accept the license agreement if the EULA agreement is changed or you are logging in for the first time. In this case, enter appropriate information in the

Name, Company, and Job Title fields, and accept the terms and conditions for the license agreement.

The password must meet the following complexity requirements:

- At least 8 characters long
- At least 1 lowercase letter
- At least 1 uppercase letter
- At least 1 number
- At least 1 special character (allowed special characters are: "\$%&'()*+,-./:;<=>@[^_`{|}~!)
- At least 4 characters different from the old password
- Should not be among the last 10 passwords

After you successfully change the password, the new password is synchronized across all Controller VMs and interfaces (Prism Element web console, nCLI, and SSH).

The default password expiration age for the admin user is 60 days. You can log in to PCVM with SSH and run the following commands to configure the minimum and maximum password expiration days based on your security requirement:

- `nutanix@pcvm$ sudo chage -M <MAX-DAYS> admin`
- `nutanix@pcvm$ sudo chage -m <MIN-DAYS> admin`
- After you upgrade from an AOS earlier version and then attempt to log in to Prism Central as the admin user, you are prompted to create a new admin user password.
- When you change the admin user password, ensure that you update any applications or scripts that uses the admin user credentials for authentication. Nutanix recommends that you create a service account with the admin role instead of using the admin user for authentication. For more information on authentication and roles, see [Security Management Using Prism Central \(PC\)](#) information in *Security Guide*.

Prism Central Landing Page

This section describes the functionalities available on the Prism Central landing page. The default application to be loaded in the Prism Central landing page is defined using the **Admin Center** application.

- For information about how to access the **Admin Center** application, see [Application Switcher Function](#) on page 40.
- For information about how to set the default landing page, see [Prism Central Admin Center Guide](#).

Note: The **Main Dashboard** is the first page that appears after logging into Prism Central if the default landing page set is **Infrastructure**. For more information about **Main Dashboard**, see [Main Dashboard - Infrastructure](#) on page 76.

An example Prism Central landing page when **Infrastructure** is set as default landing page:




Figure 3: Prism Central Landing Page

Icons in Prism Central UI

The following table describes the list of icons used in Prism Central UI for accessibility, navigation, configuration, and observation purposes:

Table 4: Icons - Prism Central UI

Icon	Name	Description
	Alerts icon	The Alerts icon appears on the right of the Prism Central landing page. Select the Alerts icon to display the alert messages that are raised by the system. For more information about Alerts, see Prism Central Alerts and Events Reference Guide .
	Tasks icon	The Tasks icon appears next to the alerts icon on the Prism Central landing page. Select the tasks icon to view the running or completed tasks. The system displays the tasks data of the last 48 hours. You can also select View All Tasks at the bottom of the list to display all the tasks. For more information, see Tasks View on page 496. You can also navigate to Activity > Tasks from the Navigation Bar to view the Tasks page. Note: If one or more of those tasks do not complete successfully, the tasks icon turns yellow (warning) or red (failure). The tasks icon turns blue color when all the tasks are running properly or completed successfully. Select the Tasks icon to see a list of the current tasks.
	Help icon	The help icon appears on the right side of the Prism Central landing page. For more information about the help resources available in Prism Central, see Help Resources on page 38.
	Navigation icon	The navigation icon is used to access the Navigation Bar for the selected Prism Central application. For more information, see Navigation Bar on page 38.

Icon	Name	Description
	Search icon	The Search icon appears on the right of the Prism Central landing page. Enter a string in this field to search for relevant content in Prism Central. For more information, see Searching for Information on page 65.
	Bookmark icon	The Bookmark icon is used for quicker access to a specific page or filtered entry. When you bookmark a page or a filtered entry, the system displays that page or filtered entry under Bookmarks in Navigation Bar .
	Settings icon	The Settings icon appears on the right side of the Prism Central landing page. Click the Settings icon to view the Prism Central Settings page for the Infrastructure application. For more information about list of settings available for the Infrastructure application, see Prism Central Settings (Infrastructure) on page 44. Note: The Settings icon also appears in some entities page such as Alerts page to enable alert emails and specify the email addresses to which alerts should be sent. For more information, see Prism Central Alerts and Events Reference Guide. You can also access the Prism Central Settings page from the Navigation Bar.
	NIVA icon	The NIVA icon appears in the bottom-right corner of the Prism Central user interface. The icon provides in-product contextual recommendations, tips, and help resources for blueprints, runbooks, playbooks, and other features of Self-Service and Intelligent Operations. These recommendations and tips are based on rules derived from telemetry data and various criteria, including license types, user roles, granted permissions, and environment details. You can use these recommendations and tips to take suggested actions and maximize the benefits of your subscribed features. For more information, see the Nutanix Intelligent Virtual Agent Guide. Important: Starting with Prism Central version pc.7.3.1, Nutanix rebranded Intelligent Assist to Nutanix Intelligent Virtual Agent (NIVA). While the user interface for Prism Central version pc.7.3 still shows the name Intelligent Assist, all current documentation and resources use the updated name, either Nutanix Intelligent Virtual Agent or NIVA.
	Add icon	The Add icon is used to add an entry in a list of allowed entries.
	Remove icon	The Remove icon is used to remove an entry from the list of allowed entries.
	Label icon	The Label icon is used to create a label for the selected entity items.
	Edit icon	The Edit icon is used to modify an existing entry.

Icon	Name	Description
	Delete icon	The Delete icon is used to delete an existing entry.
	Share icon	The Share icon is used to share the information with any Prism Central user.
	Power icon	The Power icon is used to select the following VM power options in the VM console: • Power off • Power cycle • Reset • Guest shutdown • Guest reboot. For more information about the VM Power options, see Managing a VM through Prism Central (AHV) on page 161.
	Launch icon	The Launch icon is used to launch the VM console.
	Reboot icon	The Reboot icon is used to send a Control-Alt-Delete command to the VM console.
	Screenshot-capture icon	The Screenshot-capture icon is used to take the screenshot of the VM console display.
	Back icon	The Back icon is used to return to the previous page.

Navigation Bar

The **Navigation Bar** enables you to view and select the required configuration options for the selected application.

For information about the **Navigation Bar** applicable for a specific Prism Central application, see [Application-specific Navigation Bar](#) on page 61.

For information about the applications available in Prism Central and how to access any application, see [Application Switcher Function](#) on page 40.

You can select **Lock Navigation Bar** to lock and **Unlock Navigation Bar** to unlock the **Navigation Bar**. It is a toggle. When the **Navigation Bar** is locked, it remains open all the time when you switch between multiple configurations on the Prism Central landing page.

Help Resources

The following table describes the help resources available when you select the help icon:

Table 5: Help Resources

Name	Description
Learn about search	Displays search guidelines. For more information, see Searching for Information on page 65.

Name	Description
Help with this page	Opens the online help that describes the current page. For more information, see Accessing Online Help (Prism Central) on page 651.
Online Documentation	Opens the online help at the introduction page. For more information, see Accessing Online Help (Prism Central) on page 651.
Support Portal	Opens a new browser tab (or window) at the Nutanix Support portal login page. For more information, see Accessing the Nutanix Support Portal (Prism Central) on page 644.
Nutanix Next Community	Opens a new browser tab (or window) at the Nutanix Next Community entry page. For more information, see Accessing the Nutanix Next Community on page 652. The portal is an online community site for customers and partners to exchange ideas, tips, and information about Nutanix technologies and related data center topics.
Create Support Case	Opens the Create a new support case page to view or create support cases with Nutanix customer support. For more information, see Creating a Support Case on page 640.

Application Switcher

The application switcher enables you to select the relevant application to perform any task or operation. For more information, see [Application Switcher Function](#) on page 40.

User Menu (<user_name>)

The <user_name> indicates the user identity such as admin, nutanix, or any configured user identity who is logged into the Prism Central. It appears on the far right side of the Prism Central landing page.

For information about how to define the Prism Central users, see [Controlling User Access \(RBAC\)](#) information in *Security Guide*. Select <user_name> to display a list of options to update your user account, log off from Prism Central, and other miscellaneous tasks.

The following table describes the options available in the <user_name> dropdown menu.

Table 6: User Menu Options

Name	Description
Change Password	Opens the Change Password window to update your password. For more information, see Updating My Account in <i>Security Guide</i> .
Update Profile	Opens the Update Profile window to update your user name and email address. For more information, see Updating My Account in <i>Security Guide</i> .
Download Cmdlets Installer	Downloads the PowerShell installer for the Nutanix cmdlets. For information about installing the cmdlets locally and for cmdlet descriptions, For more information, see Powershell Cmdlets Reference .
Download nCLI	Downloads the Nutanix command-line interface (nCLI) as a zip file to your local system. The download occurs immediately without additional prompts after you select this option. For more information about installing the nCLI locally and for nCLI command descriptions, see Command Reference .

Name	Description
REST API Explorer	Opens a new browser tab (or window) at the Nutanix REST API Explorer web page. For more information, see Accessing the REST API Explorer on page 647.
About Nutanix	Opens the About Nutanix window that displays Nutanix operating system (AOS) and other version information. For more information, see Finding the Prism Central Version on page 616.
Nothing To Do?	Opens a game that is strictly for entertainment. To quit the game, select "X" at the upper right of the screen.
Sign Out	Logs you out of Prism Central. For more information, see Logging Out of Prism Central on page 632.
Adjust Contrast (Chrome only)	Displays a contrast setting box at the bottom of the screen where you can set the Prism Central display to Normal (default) or High contrast.
Note: This option is available only when you access Prism Central using the Chrome browser.	

Application Switcher Function

The Application Switcher function facilitates you to seamlessly switch between following configuration module applications:

- **Platform Services** module applications that include:
 - **Admin Center** - Select this to perform admin-related tasks, activities, and configurations. For information about the administrative tasks that you can perform using **Admin Center**, see [Admin Center Overview](#) on page 48.
 - **Apps and Marketplace** - Select this to directly access the application deployment and application management functionalities in Admin Center. For more information about the application types and deployment, see the [Marketplace](#) information in *Prism Central Admin Center Guide*. For more information about application management, see the [My Apps](#) information in *Prism Central Admin Center Guide*.

Note: When you select **Apps and Marketplace** from the application switcher, the system redirects you to the **Admin Center** application with display of **Marketplace** menu.

- **Cloud Infrastructure** module applications that include:
 - **Infrastructure**. This document covers the complete information about **Infrastructure** application.
 - **Kubernetes Management** (*formerly called Karbon*)- Provides simplified provisioning and operations of Kubernetes clusters. Kubernetes is an open source container orchestration system for deploying and managing container-based applications. For more information, see [Nutanix Kubernetes Engine](#) documentation.
 - **Foundation Central** - Enables you to create clusters from factory-imaged nodes and re-image existing nodes that are already registered with Foundation Central, remotely from Prism Central. For more information, see [Foundation Central Documentation](#).

- **Cloud Manager** module applications that include:
 - **Intelligent Operations** (*formerly called Prism Pro/Ultimate*)- Optimizes capacity, proactively detects performance anomalies, and automates operations tasks with ease and confidence. For more information, see [Intelligent Operations Guide](#).
 - **NCM Self-Service** (*formerly called Calm*)- Allows you to seamlessly select, provision, and manage your business applications across your infrastructure for both the private and public clouds. Calm provides application automation, Life Cycle Management (LCM), monitoring, and remediation to manage the heterogeneous infrastructure, for example, VMs or bare-metal servers. For more information, see [Self-Service documentation](#).
 - **Cost Governance** (*formerly called Beam*)- Allows you to cost govern your resources using our pre-built TCO model to govern the spending of business units or cost centers. For more information, see [Cost Governance documentation](#).
 - **Security Central** (*formerly called Flow Security Central*)- A cloud based offering that provides complete visibility into the security posture and compliance of your multi-cloud environment (Nutanix, Azure and AWS). For more information, see [Security Central documents](#).
- **Unified Storage** module applications that include:
 - **Files** - A software-defined, scale-out file storage solution that lets you share files in a centralized and protected location to eliminate the requirement of a third-party file server. For more information, see [Files documentation](#).
 - **Objects** - A software-defined Object Store Service that is designed with an Amazon Web Services Simple Storage Service (AWS S3) compatible REST API interface capable of handling petabytes of unstructured and machine-generated data. Objects addresses storage-related use cases for backup, and long-term retention and data storage for your cloud-native applications by using standard S3 APIs. For more information, see [Objects documentation](#).
- **Database Services** module that includes:
 - **Move** - A cross-hypervisor mobility solution to move VMs with minimal downtime. For more information, see [Move User Guide](#).
 - **Database Service** - The Nutanix Database Service (NDB) automates and simplifies database administration, bringing one-click simplicity and invisible operations to database provisioning and life-cycle management. It also enables you to perform operations such as database registration, provisioning, cloning, patching, and restore. It allows administrators to define provisioning standards with end-state driven functionality that includes network segmentation and high availability (HA) database deployments. For more information, see [Nutanix Database Service User Guide](#).


Note:

- In a new deployment of Prism Central, the Application Switcher displays the **Admin Center**, **Infrastructure**, **Apps and Marketplace**, **Cost Governance**, and **Security Central** applications for users with the Nutanix admin role. You can view the Nutanix Apps in the Application Switcher after you enable them from the Marketplace.

For information about how to deploy apps from Prism Central, see [Nutanix Apps Deployment information in Prism Central Admin Center Guide](#).
- If you have upgraded to the current version of Prism Central, you can view all Nutanix Apps that you enabled in the previous version along with the **Admin Center**, **Infrastructure**, **Apps and Marketplace**, **Cost Governance**, and **Security Central** applications.

- For users with the non-admin role, the Application Switcher displays the applications based on the access policies. For example, a user with a system-defined consumer role can view only **Infrastructure, Apps and Marketplace**, and any other Nutanix Apps for which the access is provided to the user.
- You cannot add any custom applications to the Application Switcher.

Select the relevant application option to perform the configurations required at your site. The following is an example displaying the application options available through the Application Switcher function of Prism Central:



Important:

The Domain Manager container serves the Prism Central UI and API calls for the new Prism Central experience and is a repository for all deployed applications such as application switcher, Admin Center, Marketplace, deployed apps, and other Prism Central services.

If the Domain Manager container is not accessible, the application switcher function becomes non-functional. However, you can use the fallback URLs to access the deployed apps, **Infrastructure** application, **Marketplace** application, and other Prism Central services. For more information, see [KB-14217](#).

Prism Central Settings (Infrastructure)

The **Settings** page of the **Infrastructure** application enables you to configure the **Prism Central Settings** for infrastructure-specific services. You can access the **Settings** page using any of the following mechanisms:

- Click the [Settings icon](#) on the Prism Central Landing page. For information about Prism Central landing page, see [Prism Central Landing Page](#) on page 35.
- Navigate to **Prism Central Settings** from the **Navigation Bar** of **Infrastructure** application. For more information about the **Navigation Bar** of Prism Central applications, see [Application-specific Navigation Bar](#) on page 61).

The **Settings** page displays a list of tasks that you can perform for the **Infrastructure** application. Click the required task to open the window or page for that task.

The following is an example displaying the tasks available under **Settings** page:




Figure 5: Settings Page - Infrastructure

The following table provides the information about all the tasks available on **Settings** page:

Table 7: Settings Page Tasks - Infrastructure

Task Category	Task Name	Description
General	Entity Sync	Opens the Force Entity Sync page, which synchronizes entities with the specified availability zones. For more information, see Nutanix Disaster Recovery Guide .
	Licensing	Redirects to the Licensing page available under Admin Center application settings. For more information about how to access the Licensing page from Admin Center application, see Prism Central Admin Center Guide .
	Nutanix DRaaS	Opens the Nutanix DRaaS page to connect to DRaaS domain. For more information, see Nutanix Disaster Recovery Guide .
	Prism Central Management	Opens the Manage Prism Central page for viewing information about Prism Central and optionally expanding (scale out) Prism Central across multiple VMs. For more information, Managing Prism Central on page 581).
	Upgrade Prism Central	Opens the Upgrade Prism Central page to upgrade the Prism Central VM to a newer version. For more information, see Prism Central Deployment on page 12.
	Witness	Opens the Witness page. The witness service continuously monitors nodes and clusters to provide an automatic mechanism to handle any system failure. For more information about how to configure Witness service, see Configuring a Witness (Two-node Cluster) information in <i>Data Protection and Recovery with Prism Element</i> guide.
Setup	Enable Nutanix Disaster Recovery	Opens the Enable Nutanix Disaster Recovery page to enable the Disaster Recovery as a Service (DRaaS) capability. For more information, see Nutanix Disaster Recovery Guide .
	Project Manager Management	Opens the Project Manager Management page available under Identity and Access Management > Settings of Admin Center application. For information about how to access the Admin Center application in Prism Central, see Application Switcher Function on page 40. For information about the Admin Center application in Prism Central, see Admin Center Overview on page 48.
	OOB Management Credentials	Opens the Out of Band (OOB) Management Credentials page where you can manage the OOB credentials. For more information, see Configuring Out-of-Band Management Credentials on page 635.
	vCenter Registration	Opens the vCenter Registration window to register (or unregister) clusters with vCenter. For more information, see External vCenter Server Integration on page 341.

Task Category	Task Name	Description
Network	Advanced Networking	Opens the Advanced Networking page to enable the advanced networking for all On-Prem AHV Clusters in the Prism Central that runs AOS 6.1 or above acropolis version.
Flow	ID Based Security	Opens the ID Based Security page to add Active Directory domain services configurations and to import user groups for identity-based security policies. For more information, see Flow Microsegmentation Guide .
	Microsegmentation	Opens the Microsegmentation page to enable this feature, which is disabled by default. Before you can configure and use application security policies, isolation environment policies, and quarantine policies, you must enable the feature. For more information, see Flow Microsegmentation Guide .
User and Roles	Authentication	For information about User and Roles, see Security Guide .
	Local User Management	
	Role Mapping	
More Settings	For information about the settings available under More Settings , see Prism Central Admin Center Guide .	

Understanding Displayed Statistics

The Prism Element web console and Prism Central web console display various statistics that are derived from the following sources:

Note: Most displayed statistics appear in 30 second intervals. The values in the tables represent the most recent data point within the last 30 seconds. Prism Central collects the statistical data from each registered cluster, so the process of collecting that data could result in a longer lag time for some statistics displayed in Prism Central.

1. **Hypervisor:** Hypervisor provides usage statistics only. The support to provide usage statistics is available only in ESXi, and not in Hyper-V and AHV hypervisors. If the cluster consists of Hyper-V or AHV hypervisor, the controller provides the usage statistics.

Note: Ensure that you consider the usage statistics reported from ESXi hypervisor in both Prism Central and Prism Element, only when it matches with the usage statistics in vCenter.
2. **Controller (Stargate):** When hypervisor statistics are unavailable or inappropriate, the Controller VM (CVM) provides the statistics from Stargate. For more information about Stargate, see [Nutanix Bible](#). The Controller-reported statistics might differ from those reported by the hypervisor for the following reasons:
 - An NFS client might break up large I/O requests into smaller I/O units before issuing them to the NFS server, thus increasing the number of operations reported by the controller.
 - The hypervisor might read I/O operations from the cache in the hypervisor which are not counted by the controller.
3. **Disk (Stargate):** Stargate can provide statistics from both controller and disk perspective. The controller perspective includes reading both I/O operations from memory and disk I/O operations, but the disk perspective includes only disk I/O operations.

Note: The difference in statistics derived from the sources: Hypervisor, Controller, and Disk, only applies to storage-related statistics such as IOPS, latency, and bandwidth.

The following field naming conventions are used in Prism Central to identify the information source:

- A field name with **Controller** word indicates the statistic is derived from the controller (for example *Controller IOPS*).
- A field name with **Disk** word indicates the statistic is derived from the disk (for example *Disk IOPS*).
- A field name without **Controller** or **Disk** word indicates the statistic is derived from the hypervisor. For example **IOPS**.

For VM statistics in a mixed ESXi/AHV cluster, the statistics source depends on the type of hypervisor that hosts the VM. If the Hypervisor is:

- ESXi - The hypervisor is the source for statistics.
- AHV - The controller is the source for statistics.

Note:

- The overview, VM, and storage statistics are derived from either the hypervisor or controller.
- Hardware statistics are derived from disk.
- Metrics in the analysis page are derived from any of the sources: hypervisor, controller, or disk, based on the type of metric.

The following table provides the information about the source for various statistics based on hypervisor type:

Table 8: Source for Displayed Statistics

Hypervisor Type	Statistics	Source	Analysis
ESXi	Overview, VM, and Storage	Both Hypervisor and controller <i>Controller for some storage statistics only</i>	Metric dependent
	Hardware	Disk	
Hyper-V	Overview, VM, and Storage	Controller	
	Hardware	Disk	
AHV	Overview, VM, and Storage	Controller	
	Hardware	Disk	
Citrix Hypervisor	Overview, VM, and Storage	Controller	
	Hardware	Disk	
Mixed (ESXi + AHV)	Overview, VM, and Storage	Hypervisor	

Hypervisor Type	Statistics	Source	Analysis
	Hardware	Disk	

Infrastructure Overview

The **Infrastructure** application of [Application Switcher Function](#) on page 40 enables you to configure the field-specific settings required to host a cluster in Prism Central. It involves the following configurations:

- Customizable main dashboard that displays summary information across the registered clusters (see [Main Dashboard - Infrastructure](#) on page 76).
- Dashboard and configuration options to manage VMs, VM templates, OVA, image placement, catalog items, storage, volume groups, Nutanix Guest Tools (NGT), and vCenter server across the registered clusters ([Compute Entities](#) on page 99).
- Dashboard and configuration options to manage subnets, network connections, Virtual Private Clouds (VPC), floating IPs, Virtual Private Network (VPN) connections, and security policies across the registered clusters ([Network and Security Entities](#) on page 355).
- Dashboard and configuration options to manage protection policies, recovery plans, VM recovery points, VG recovery points, and consistency groups (see [Data Protection and Recovery Entities](#) on page 421).
- Hardware component dashboards with drill-down options to view detailed information about individual clusters, hosts, disks, and GPUs across the registered clusters (see [Hardware Entities](#) on page 422).
- Activity monitors for alerts and events (see [Activity Entities – Alert and Event Monitoring](#) on page 492).
- Activity monitors for audits and tasks (see [Activity Entities – Tasks and Audits](#) on page 493).
- Dashboards to manage categories and availability zones (see [Administration Entities](#) on page 500).
- Settings menu from which you can configure Prism Central functions (see [Prism Central Settings \(Infrastructure\)](#) on page 44).

The **Infrastructure** application also provides the following additional functionalities for cluster and field-specific health checks:

- Fault Management based on Alerts and Events Customization. For more information on Alerts and Events, see [Prism Central Alerts and Events Reference Guide](#).
- Detailed information about the cluster health.

Note: The **Services** menu available in the earlier Prism Central versions is now retired. You can use the [Application Switcher Function](#) on page 40 to access Nutanix applications and services.

Admin Center Overview

The **Admin Center** application enables you to perform the following administrative-level tasks in Prism Central:

- Enable Marketplace in Prism Central. For more information, see [Enabling Marketplace](#) information in *Prism Central Admin Center Guide*.
- Enable Prism Intelligent Operations. For more information, *Enabling Prism Intelligent Operations* in [Intelligent Operations Guide](#).
- Discover and deploy Nutanix apps such as NCM Self-Service (formerly known as Calm), *Files, Move, Objects, Database Service, Foundation Central*, and *Kubernetes Management*.
- Discover and deploy preferred partner apps from the Marketplace.

- Manage all deployed apps from a common workspace.
- Define multiple projects as per the field requirements.
- Manage Nutanix applications License.
- Define users, roles, and set up authentication mechanism for Identify and Access Management (IAM).

Note: Prism Central provides the Self Service feature that allows you to create projects where the consumers of the IT infrastructure can provision and manage VMs in a self-service manner, without engagement of IT in day-to-day operations. The consumers of the IT infrastructure within an enterprise include individual users, development teams, test teams, and DevOps teams. Only a Prism Central administrator can create a self-service administrator role.

For information about the roles and functions for Prism Central Self Service, see [Prism Self Service Setup](#) on page 569.

- Perform inventory and software updates as part of the Life Cycle Manager (LCM) functionality.
- Set welcome banner, language, default landing page, and other UI settings for Prism Central.
- Security management using Cluster Lockdown and SSL Certificate mechanisms.
- Define email aliases, SMTP server, and Syslog server to receive alerts and notification from Prism Central.
- Define network management guidelines for Prism Central using HTTP Proxy, Name server, NTP server, and SNMP settings.

For more information about **Admin Center**, see [Prism Central Admin Center Guide](#).

Prism Central GUI Organization

This section provides the information about the GUI-related functionalities of Prism Central, how the Prism Central GUI is organized, and how to perform GUI-specific operations in Prism Central.

Entity Layout

In Prism Central GUI, you can access the entities from the **Navigation Bar** using the [Navigation icon](#). For information about the **Navigation Bar**, see [Prism Central Landing Page](#) on page 35.

The entities are categorised into the following types:

- [Compute Entities](#) on page 99
- [Network and Security Entities](#) on page 355
- [Data Protection and Recovery Entities](#) on page 421
- [Hardware Entities](#) on page 422
- [Activity Entities – Alert and Event Monitoring](#) on page 492
- [Activity Entities – Tasks and Audits](#) on page 493
- [Operations Entities](#) on page 499
- [Administration Entities](#) on page 500

Note: The **Services** menu available in the earlier Prism Central versions is now retired. You can use the [Application Switcher Function](#) on page 40 to access Nutanix applications and services.

When a list appears for the selected entity, the display may include:

- Entity tabs, **Actions** options (if any), **Group by** options, and **View by** options for the list, and a **Modify Filters** option on the far right. When you click **Modify Filters** option, the **Filters** pane appears.
- A query field that identifies the filters (if any) being applied to the list. This field displays all filter options that are currently in use. It also allows for basic filtering on the entity name.
- A table (list) of entities. The table content is determined based on the **Group by**, **View by**, and **Modify Filters** options.

Note: By default, the system displays 20 rows of entity items per page. The system also provides you an option to configure the display with a different number of rows per page between 10 to 60. To specify a different number for rows per page, click **X Rows**, where x indicates the number such as 10, 20, 30, 40, 50, and 60.

- An export mechanism to download the table contents in CSV format with a maximum of 1000 rows. Click **Export** to download the table contents on your local machine.




Figure 6: Entity Layout (example)

Entity Tabs

Each entity provides relevant tabs such as **Summary**, **List**, **Policies**, **Alerts**, **Events**, **Audit**, **Metrics**, **Scenarios**, **Gateway**, **VPN connections**, **Subnet Extensions**, and **BGP Sessions**, to perform entity-related configurations and observe entity related-activities.

Note: For some entities, no tab appears and the relevant content appears directly in the main page of the selected entity.

Sort Function

The sort function allows you to sort the entries in alphabetic or ascending-descending order. This function appears for the Tiles and Circles views only; ordering in the Grid (tabular) view is accomplished when you click a column header.

Filters Pane

Click **Modify Filters** option to display the **Filters** pane.

The **Filters** pane include a set of fields that vary according to the selected entity. When you type the desired field values in the **Filters** pane, an entry appears in the search field for each value you enter.

Note:

- You can bookmark (or save) a filter using the **Bookmark icon** in the search field. The selected filter is saved under **Bookmarks**. To remove the bookmarked filter, navigate to **Bookmarks** from the **Navigation Bar**, and click the **Bookmark icon** for the saved filter.

- Numeric filters have **To/From** fields to specify a range. You can specify the numeric values along with units in these fields. For example, the filter adjusts the scale accordingly when you type in "10 K" or "100 M".

To apply a filter, enter a value or select the checkbox of the desired value (or multiple values) for the field you want to use as a filter. You can apply filters across multiple parameters. Some parameter filters require additional context such as a constraint string or a range.

The following is an example displaying the **Filters** pane:




Figure 7: Filters Pane

Group by

The **Group by** option enables you to organize the entity information based on attributes such as Cluster, Hypervisor, Power State, vCPU, Health, Type, Name, and so on. The **Group by** options vary based on the selected entity.

View by

The **View by** option enables you to specify the following view types for the displayed information:

- **General** - Displays a set of general information parameters.
- **Performance** - Displays a set of performance-specific parameters.
- **Anomalous Behavior** - Displays the Anomalous behavior. For example, Anomaly count of the VMs.
- **Efficiency** - Displays a set of efficiency-related parameters. **Efficiency** is determined through the VM behavioral learning engine. For more information, see [Behavioral Learning Tools](#) information in *Intelligent Operations Guide*.
- **GPU** - Displays a set of GPU-specific parameters.
- **Data Protection** - Displays a set of data protection-related parameters.
- **Storage Configuration** - Displays a set of storage configuration-related parameters.

Note: The **General** view type is available for all entities, but the other options are available only if it is applicable for an entity.

In addition to the pre-defined view types, you can also create one or more custom view types.

To create a custom view type:

1. Click the **View by** option and select **+ Add custom**.

2. In the **<entity> Columns** window that is displayed, perform the following actions:

1. Enter a name for the custom view type in the first (top) field.
2. Click the [Add icon](#) in the left column for each entity property to be included in the custom view type. The selected entity property appears in the right column.

Note:

- By default, the **Name** column is already selected.
- A maximum of 10 entity properties are allowed in a custom **View by** option including the **Name** column.
- To arrange the order of the selected columns, hover on the column name and click the up or down arrow as appropriate.
- To search and filter the required entity property from the list, enter the string in the Search field above the left column.
- To remove an entity property from the custom **View by** option, click the Remove icon associated with the entity property.

3. Click **Save** after the entity-property list in the right column is finalized.

The following following is an example displaying the **VM Columns** window:




Figure 8: VM Column Window (example)

The custom view type appears (by name) in the available view types under **View By** option.

Note: The custom view type appears only for the user who created it, and the administrative user and other users cannot view it.

Actions Dropdown Menu

The **Summary** page provides options to administer and categorize the entity items. The **Actions** dropdown menu is available for the selected entities in Prism Central. You can select an entity item, a set of entity items, or all entity items using the relevant checkbox available in the table list on **Summary** page. You can clear the relevant checkbox to clear a selected item from the table list.

You can also perform some administrative tasks for a cluster through Prism Element. For information about the tasks to be performed from Prism Element, see the [Prism Element Web Console Guide](#).

Note: The **Actions** dropdown list does not appear if there are no relevant actions available for the selected entity type.




Figure 9: Action Dropdown Menu

Label

Labels enable you to create custom groupings for the entities. To add a label, click the [Label icon](#), and enter the label name in the field. The label is applied to all the selected entity items in the table list.

Note:

- Only super admin and Prism admin roles have permissions to perform this operation.
- Labels are available for VMs and clusters only.

Export

The **Export** option enables you to download the table, that contains the entity items list to a file. Click the **Export** option available on an entity item.

Entity-Related Observation Activities

The entity-related observation activities involve the following fault and performance related statistics:

- Alarms - For information on all the alarms generated in Prism Central, see [Activity Entities – Alert and Event Monitoring](#) on page 492.
- Events - For information on all the events generated in Prism Central, see [Activity Entities – Alert and Event Monitoring](#) on page 492.
- Tasks - For information on how to observe a task and its statistics in Prism Central, see [Activity Entities – Tasks and Audits](#) on page 493.
- Audits - For information on an audit activity in Prism Central, see [Activity Entities – Tasks and Audits](#) on page 493.

Prism Licensing

Nutanix provides licenses you can apply to enable a variety of features.

The Prism Central web console and Nutanix Support Portal provide the most current information on your licenses. For more information on licenses, see the [License Manager Guide](#).

X-Small Prism Central

X-Small Prism Central is a single-node Prism Central VM that is intended for small-scale deployments with limited workload.

Note: You can deploy x-small Prism Central on clusters running AOS 6.8 or later versions only. X-Small Prism Central lets you manage clusters running AOS versions earlier than AOS 6.8.

For information on the type of Prism Central deployments, see the *Prism Central Scalability* topic in the release notes for the [Prism Central](#) version to be installed.

X-Small Prism Central is ideal for users who have five or fewer clusters on their site and use only the basic or essential capabilities of Prism Central. X-Small Prism Central provides the following capabilities:

- [Multi-cluster management \(up to 5 clusters\)](#)
- [VM management](#)
- [Host management](#)
- [Infrastructure management, monitoring, and health](#)
- [Enterprise authentication and role-based access control \(RBAC\)](#)
- [REST APIs](#)
- [Comprehensive search](#)
- [Pulse Insights](#)
- [Life Cycle Manager \(LCM\)](#)
- [Prism Central backup and restore](#)
- [Categories](#)
- [Projects](#)
- [Microservices infrastructure](#)
- [Identity and access management \(IAM\)](#)

Specifications

X-Small Prism Central has the following specifications:

Table 9: X-Small PC Specifications

Number of vCPU	4
Memory	18 GB
Total Storage Capacity	100 GiB

For information about the maximum tested and supported values for x-small Prism Central, see [Nutanix Configuration Maximums](#).

Limitations

- X-Small Prism Central does not support the scale-out Prism Central feature.

For information on Prism Central scalability, see the *Prism Central Scalability* topic in the release notes for the [Prism Central](#) version to be installed.

- Adding entities to a x-small Prism Central deployment beyond the recommended number of VMs, clusters, or nodes results in an unsupported configuration. Nutanix recommends that you scale up to small, large, or x-large Prism Central deployment to increase its capacity. For more information, contact Nutanix Support.
- X-Small Prism Central does not allow you to enable the following features:

- Flow Virtual Networking
- Flow Network Security
- Self-Service
- Intelligent operations
- Objects
- Nutanix Kubernetes Engine
- Nutanix Kubernetes Platform
- Nutanix Data Services for Kubernetes (NDK)
- Nutanix Enterprise AI
- GPT-in-a-Box
- Files
- Foundation
- Foundation Central
- Quotas
- Multi-site DR (available in PRO customers)
- Marketplace
- Reporting and Dashboards
- Security Dashboard

To enable these features, you must upgrade to a small, large or x-large Prism Central deployment. For more information, contact Nutanix Support.

- X-Small Prism Central does not support nearsync and synchronous replication.

Registering a Cluster with Prism Central

You must register a cluster with Prism Central only if you want to manage the cluster through Prism Central.

Before you begin

- If you have never logged into Prism Central as the admin user, you must log in and change the password before you attempt to register a cluster with Prism Central. For more information, see [Logging into Prism Central](#).

- Do not enable client authentication in combination with ECDSA certificates on a registered cluster because it causes interference when communicating with Prism Central.
- Port 9440 must be open in both directions between the Prism Central VM and all the Controller VMs (and the cluster virtual IP address if configured) in each registered cluster. For the complete list of required ports, see [Ports and Protocols](#).
- If you have a proxy server configured and you want the cluster communication between the cluster and Prism Central to go through the proxy, open the relevant ports on the proxy. If you do not want the communication to go through the proxy, add the Prism Central IP address to the proxy allowlist in the cluster settings. For more information about configuring proxy, see [Configuring HTTP Proxy](#) in *Prism Element Web Console Guide*. For the complete list of required ports, see [Ports and Protocols](#).
- You can register a cluster with only one Prism Central instance at a time. To register with a different Prism Central instance, first unregister the cluster.
- If the cluster is dual stack enabled, the Prism Central instance must also be dual stack enabled to register it.

About this task

To register a cluster with Prism Central, follow these steps:

Note: To perform this task, ensure that you log in to the Prism Element web console as an admin user.

Procedure

1. Log in to the Prism Element web console on the target cluster.
2. To run Nutanix Cluster Checks, go to the **Health** dashboard, and from the **Actions** dropdown menu, click **Run Checks**.
3. From the dropdown menu on the left of the main menu, select **Settings**.
The system displays the **Global Settings** page.
4. From the **Setup** section in the left navigation pane, click **Prism Central Registration**.
The system displays the **Prism Central** page.
5. Click **Register** or deploy new Prism Central.
The Prism Central dialog box opens displaying options to deploy a new Prism Central or connect to an existing Prism Central instance
6. Click **Connect**.
The system displays the Connect Info tab outlining the services available through Prism Element and Prism Central
7. After reviewing the message, click **Next**.
The system displays the **Configuration** tab.
8. Perform the following steps:
 - a. In the **Prism Central IP/FQDN** field, enter the IP address or FQDN of the Prism Central VM.
 - b. In the **Port** field, enter the port number to allow communication between Prism Element and Prism Central.
For the complete list of required ports, see [Ports and Protocols](#).
 - c. In the **Username** field, enter admin as the Prism Central user name.
 - d. In the **Password** field, enter the password for the Prism Central user.

9. Click **Connect** to save the details.

The cluster is now registered with the specified Prism Central VM. After successful registration, the specified cluster and Prism Central start communicating with each other.

Note:

- The user credentials provided when registering a cluster (Prism Element) with Prism Central are only used once. After registration, modifying the admin password would not impact any communication between Prism Central and the cluster.
- On small, large, and x-large Prism Central deployments, when you register a new cluster to Prism Central, Prism Central synchronises the past 90 days of data (including multiple metrics) from the cluster. On x-small Prism Central deployments, Prism Central synchronises the past 2 hours of data (including multiple metrics) from the cluster. To view the list of metrics that are synced during registration, see the file `/home/nutanix/config/arithmos/data_sender/arithmos_history.json` in the Controller VM. To view the list of metrics that are synced during a regular synchronisation between Prism Central and the cluster, see the file `/home/nutanix/config/arithmos/data_sender/arithmos.json` in the Controller VM.

What to do next

After the cluster is registered with Prism Central, you can configure various settings for network, security, and alerts. For more information, see [Admin Center Settings Options](#).

Unregistering a Cluster from Prism Central

You can unregister a cluster from Prism Central if you no longer need to manage the cluster through Prism Central.

About this task

Caution: Unregistering a cluster from Prism Central is not a supported workflow. The unregistered cluster might be disallowed for re-registration with a Prism Central instance.

You can use the destroy Cluster feature of Prism Central, which implicitly unregisters the cluster. For more information, see [Destroying a Cluster](#) on page 437.

If you still want to go ahead with unregistration of the cluster, consider the following points:

Before you begin

- Unregistering a cluster through the Prism Element web console is no longer available. This option is removed to reduce the risk of accidentally unregistering a cluster. Several features such as role-based access control, application management, micro-segmentation policies, and self-service capability require Prism Central to run your clusters. If a cluster is unregistered from Prism Central, it leads to features unavailability and configuration erasure. You can only use the following procedure from Controller VM (CVM) to unregister a cluster.
- Perform the entire registration process, followed by the cleanup process.
- Do not remove the IPs of the cluster and Prism Central from the whitelists on both sides until the unregistration process completes successfully.

If you have enabled additional applications or features in Prism Central, see the following table for recommendations before you unregister a cluster. For more information, see [KB 4944](#).

Table 10: Unregistering a cluster with additional applications or features enabled

Nutanix Disaster Recovery (Leap)	Before unregistering the cluster from Prism Central, you must remove any Nutanix Disaster Recovery (Leap) configuration involving virtual machines or volume groups for the cluster being unregistered. Otherwise you will not be able to manage the snapshot creation or replication policies configured on the cluster. For more information, see KB 12749. Note: Do not proceed with unregistration, if the stretch config (AHV Synchronous replication) is present.
Flow Networking	You must disable Flow Networking on the cluster before unregistering it from Prism Central, using the steps mentioned in Unregistering a PE from the PC in <i>Flow Virtual Networking Guide</i>. If Flow Networking is not disabled on the cluster prior to unregistering from Prism Central, attempts to enable Flow Networking in the same cluster does not work as expected. For more information, see KB 12449.
NuCalm/App Management	You must clean up Calm entities after unregistration. Contact Nutanix Support for assistance in cleaning up the CALM entities.
Prism Self Service configuration	Changes that have been made to the Prism Self Service configuration in Prism Central are lost after unregistration. Ensure that you follow the extra cleanup steps mentioned in KB 4944.
NKE Kubernetes	Do not unregister the cluster hosting an NKE Kubernetes cluster from Prism Central. Unregistration of cluster from Prism Central will prevent the management of the NKE clusters.
Service Microservices Platform (SMSP)	Do not unregister a cluster hosting an SMSP cluster, if you still need the SMSP cluster. If you no longer need the SMSP cluster, contact Nutanix Support for assistance with the unregistration process.
Nutanix Objects	Do not unregister the cluster that hosts a Nutanix Objects cluster from Prism Central. Unregistering the cluster from Prism Central prevents the management of the Nutanix Objects clusters.

To unregister a cluster from an existing Prism Central instance, perform the following steps:

Procedure

1. Log in to any CVM of the registered cluster through an SSH session.
2. Run the cluster status command and verify that all services are in a healthy state.
3. Run the following command to unregister the cluster from Prism Central.

```
nutanix@cvm$ ncli multicluster remove-from-multicluster external-ip-address-or-svm-ips=pc-name-or-ip username=pc-username password=pc-password
```

Replace *pc-name-or-ip* with the Prism Central name or IP address and *pc-username* and *pc-password* with the login credentials for your Prism Central administrator account. If the password contains any special

characters, ensure to enclose the password in single quotes. This step can take some time (though typically just a few seconds).

If the unregistration process is successful at Prism Element cluster, you will see a task UUID in the command output. After you see the task UUID in the command output, wait for about 2 minutes to let the unregistration process complete at Prism Central.

4. Run the following command on Prism Element to verify if the cluster unregistration process is successful on Prism Element:

```
nutanix@cvm$ ncli multicluster get-cluster-state
```

The command output should list 0 as registered cluster count: Registered Cluster Count: 0

5. Run the following command on Prism Central to verify if the cluster unregistration process is successful on Prism Central:

```
nutanix@pcvm$ ncli multicluster get-cluster-state
```

If the command output does not include the name and UUID of the cluster you want to unregister, the unregistration process is successful. If the cluster information is still visible in the command output, refer to [KB-4944](#) to troubleshoot the issue.

6. Run the following command to retrieve the UUID for the cluster:

```
nutanix@cvm$ ncli cluster info
```

The following output shows the **Cluster UUID** value in the cluster information output:

```
Cluster Id      : 586dd889-84dc-422f-b8e8-8ce3cd781dcf::4100682370753502671
Cluster Uuid    : 586dd889-84dc-422f-b8e8-8ce3cd781dcf
Cluster Name    : PC_10.51.147.185
Cluster Version : pc.2024.1
Cluster Full Version : e18.5-release-fraser-2024.1-stable-
ca1136cd25915d9b5635483235e1fd7af3ccddfb
External IP address : 10.51.147.185
Is LTS           : false
External Data Services... :
Support Verbosity Level : BASIC_COREDUMP
Lock Down Status : Disabled
Password Remote Login ... : Enabled
Timezone         : Atlantic/Reykjavik
NCC Version      : ncc-5.0.0
Degraded Node Monitoring : Enabled
```

7. Log in to the Prism Central VM through an SSH session (as the *nutanix* user) and perform the following steps:

- a. Run the unregistration clean-up script.

```
[pcvm]$ python /home/nutanix/bin/unregistration_cleanup.py uuid
```

Replace *uuid* with the value you obtained in step 6. This script removes all remaining registration information about that cluster and completes the unregistration process with the Prism Central VM.

Note: If you do not run the clean-up script, some artifacts continue to retain references to entities that are no longer managed by the cluster. Some artifacts that might have lost references due to the unregistration process might not be able to recover their references.

- b. Run the following command to retrieve the UUID for Prism Central:

```
[pcvm]$ ncli cluster info
```

Find the **Cluster UUID** value in the displayed information (see step 6), which in this case is the UUID for Prism Central.

- Go back to the CVM, and run the `unregistration_cleanup.py` script to complete the unregistration process on the cluster.

```
nutanix@cvm$ python /home/nutanix/bin/unregistration_cleanup.py uuid
```

In this case the `uuid` is the Prism Central UUID obtained in substep 7.b on page 60.

If you do not encounter any error after running the cleanup script, you can consider that the cleanup is successful.

Application-specific Navigation Bar

Infrastructure

Select the **Infrastructure** application from the [Application Switcher Function](#) on page 40, and click [Navigation icon](#) to display the **Navigation Bar** on the left side.

The following is an example showing the **Navigation Bar** for the **Infrastructure** application:




Figure 10: Navigation Bar - Infrastructure

Note:

- The **Navigation Bar** consists of **Dashboard**, **Bookmarks**, entity lists (**Compute**, **Storage**, **Network & Security**, **Data Protection**, **Hardware**, **Activity**, **Operations**, **Administration**, and **Prism Central Settings**).
- Click an entity list (for example, **Compute**) to display the entity items available in it.
- Click the entity item (for example **VMs**) to display the **Lists** page or the main page for that entity item.

- For quicker access to one or more entities in the entity list, you can bookmark the required entity to promote to the primary view of the **Navigation Bar**. Click the [Bookmark icon](#) displayed at the right of the entity name to bookmark the entity. The **Bookmarks** category appears below the **Dashboard**.

Note: To remove a bookmarked entry, navigate to **Bookmarks**, and click the [bookmark icon](#) displayed at the right of the selected bookmarked item name.

The following table describes the entity list available in the **Navigation Bar** of the **Infrastructure** application:

Table 11: Entities - Infrastructure

Entities	Entity Items	Description
Dashboard		Displays the Main Dashboard (see Main Dashboard - Infrastructure on page 76).
Compute & Storage	VMs	Displays the VMs configuration page (see VM Management on page 99).
	Templates	Displays the Templates configuration page (see VM Template Management on page 212).
	OVA's	Displays the OVA's configuration page (see OVA Management on page 235).
	Images	Displays the Images configuration page (see Image Management on page 258).
	Catalog Items	Displays the Catalog configuration page (see Catalog Management on page 283).
	Storage Containers	Displays the Storage Containers configuration page (see Limitations for Storage Containers on page 311).
	Volume Groups	Displays the Volume Groups configuration page (see Volume Group Management on page 318).
	Storage Policies	Displays the Storage Policies configuration page (see Storage Policy Management on page 545).
	vCenter Datastores	Displays the vCenter Datastores configuration page (see External vCenter Server Integration on page 341).
Network & Security	Subnets	Displays the subnets configuration page (see Subnets Summary View on page 365).
	Virtual Private Clouds	Displays the virtual private clouds configuration page (see Virtual Private Clouds Summary View on page 370).
	Floating IPs	Displays the floating IPs configuration page (see Network and Security Entities on page 355).
	Connectivity	Displays the Connectivity configuration page (see Connectivity on page 404).
	Security Policies	Displays the Security Policies configuration page (see Security Policies Summary View on page 540).

Entities	Entity Items	Description
Data Protection	Security Dashboard	Displays the Security Dashboard (see Security Dashboard on page 420).
	Protection Summary	Displays the Protection Summary dashboard for disaster recovery. (See Protection Summary on page 421).
	Protection Policies	Displays the Protection Policies configuration page (see Protection Policies on page 421).
	Recovery Plans	Displays the Recovery Plans configuration page (see Recovery Plans on page 421).
	VM Recovery Points	Displays the VM Recovery Points configuration page (see VM Recovery Points on page 421).
	VG Recovery Points	Displays the VG Recovery Points configuration page (see VG Recovery Points on page 421).
Hardware	Consistency Groups	Displays the Consistency Groups configuration page (see Consistency Groups on page 421).
	Clusters	Displays the Clusters management page (see Clusters Summary View on page 424).
	Hosts	Displays the Hosts management page (see Host Management on page 444).
	Disks	Displays the Disks management page (see Disks Summary View on page 464).
Activity	GPUs	Displays the GPUs management page (see GPUs Summary View on page 470).
	Alerts	Displays the Alerts dashboard (see Prism Central Alerts and Events Reference Guide).
	Events	Displays the Events dashboard (see Prism Central Alerts and Events Reference Guide).
	Audits	Displays the Audits dashboard (see Audits Summary View on page 493).
Operations	Tasks	Displays the Tasks dashboard (see Tasks View on page 496).
	Analysis	Displays the Analysis management page (see Analysis Dashboard in <i>Intelligent Operation Guide</i>).
	Discovery	Displays the App Discovery dashboard or the enable wizard if application discovery is not enabled (see Application Discovery in <i>Intelligent Operation Guide</i>).
	Monitoring Configurations	Displays the Monitoring Integrations dashboard or the enable wizard if monitoring integrations is not enabled (see Application Monitoring in <i>Intelligent Operation Guide</i>).
	Operations Policies	Displays the Operations Policies dashboard (see Operations Policy Management in <i>Intelligent Operation Guide</i>).

Entities	Entity Items	Description
	Planning	Displays the capacity Planning dashboard (see Resource Planning in <i>Intelligent Operation Guide</i>).
	Playbooks	Displays the list of playbooks (see Task Automation-Playbooks in <i>Intelligent Operation Guide</i>).
	Reports	Displays the Reports dashboard (see Reports Management in <i>Intelligent Operation Guide</i>).
	Settings & Configurations	Displays the Settings and Configurations dashboard (see Settings & Configurations in <i>Intelligent Operation Guide</i>).
Administration	LCM	Displays the Life Cycle Manager (LCM) dashboard (see Life Cycle Manager Guide).
	Projects	Displays the Projects dashboard (see Project Overview information in <i>Prism Central Admin Center Guide</i>).
	Roles	Displays the Roles dashboard (see Security Guide).
	Users	Displays the Users dashboard (see Security Guide).
	Availability Zones	Displays the Availability Zones dashboard (see Availability Zones on page 500).
Prism Central Settings		Displays the Settings menu for the Infrastructure application (see Prism Central Settings (Infrastructure) on page 44).

Admin Center

For information about the navigation bar of **Admin Center** application, see [Prism Central Admin Center Guide](#).

Nutanix_Apps and Other_apps

For more information about all the navigation bar of Nutanix_Apps and Other_apps that you can discover, deploy, and manage from Prism Central, refer to the application-specific documentation as described in [Application Switcher Function](#) on page 40.

Apps and Marketplace

No specific navigation bar is available for **Apps and Marketplace** option that appears in the [Application Switcher](#). When you select **Apps and Marketplace** from the [Application Switcher](#), the system redirects you to the **Admin Center** application with display of **Marketplace**.

Identifying the Prism Central Leader VM

This section provides the steps to identify the leader VM in a scale-out cluster of Prism Central VMs.

About this task

Use the following steps to identify the leader Prism Central VM in a scale-out cluster of Prism Central VMs.

Procedure

1. To identify the leader VM on the Prism Central UI, follow these steps:

- a. Access the **Settings** page using any of the following mechanisms:
 - » On the Prism Central landing page, click the [Settings icon](#) on the Prism Central Landing page.
For information on the Prism Central landing page, see [Prism Central Landing Page](#) on page 35.
 - » From the **Navigation Bar** of **Infrastructure** application, go to **Prism Central Settings** .
For more information on the **Navigation Bar** of Prism Central applications, see [Application-specific Navigation Bar](#) on page 61).
- b. From the menu (see [Prism Central Settings \(Infrastructure\)](#) on page 44), select the **Settings > Prism Central Management** .
This displays the **Manage Prism Central** page.
- c. Note the **Virtual IP** address configured for the Prism Central cluster displayed in the **Prism Central Summary** widget.
- d. Go to **Compute > VMs** from the **Navigation Bar**.
The Prism Central VMs are listed on the **List** tab of the **VMs** dashboard.
The Prism Central VM with the **Virtual IP** address listed in **Compute > VMs > List** tab is the Prism Central Leader VM.

2. To identify the leader VM on the CLI, follow these steps:

- a. SSH to a Prism Central VM using the virtual IP assigned to the Prism Central cluster.
If you have not assigned a virtual IP address to the Prism Central cluster, SSH to any Prism Central VM using the IP address of the Prism Central VM.
- b. Use one of the following options to identify the Prism Central VM leader:
 - »

```
nutanix@pcvm$ curl localhost:2019/prism/leader && echo
```


The following sample output shows the IP address of the leader.

```
{"leader":"X.X.X.X:9080", "is_local":true, "serve_api_locally":true}
```
 - »

```
nutanix@pcvm$ panacea_cli show_leaders | grep -i prism
```


The following sample output shows the IP address of the leader.

```
prism_monitor:9080 X.X.X.X 2 leader
```


The Prism Central VM with the IP address displayed in the CLI output is the Prism Central Leader VM.

Searching for Information

The **Infrastructure** application page includes a search field on the right (see [Prism Central Landing Page](#) on page 35) that allows you to find information about selected entities in various ways.

If you need help at any time to navigate Prism Central or apply a search filter, click the [Help icon](#) in the **Infrastructure** application page, and select **Learn about search**. This displays the **Search Guidelines** page that explains the search rules and options.

Search Basics

An *entity* is an object type such as a VM, cluster, node, security policy, project, report, event, alert, or audit. The search field is context-sensitive, which means it is populated automatically based on where you are in Prism Central. The first screen after logging on is the main dashboard, so *Dashboard* appears in the search field initially.




Figure 11: Search Field (Prism Central Landing page - Infrastructure)

When you navigate to different Prism Central entity pages, the search string changes automatically to match the current entity.

The following is an example showing the example of search string that appears on the VMs page:




Figure 12: Search Field (Entity page)

When you enter a string in the search field, a dropdown menu appears with relevant matches across Prism Central.

Note: The search strings are case insensitive.

The following is an example showing the sample dropdown menu that appears when you enter the string VM:




Figure 13: Search Results (for an entity type)

- *Category Value*: Select this option to display the **Category** page with any VM-related entries.
- *VM Type=User VM*: Select this option to display the filtered **Summary** page for user VMs.
- *[blank]*: Select this option to display the unfiltered **Summary** page for any user VM.
- *VM Type=User VM > List*: Select this option to display the **List** page of the user VMs.
- *VM Type=User VM > Alerts*: Select this option to display the **Alerts** page of the user VMs .
- *Search in Prism*: Select this option to search across Prism Central for any information about the target entity.

In the following example, two tables appear in the results, a list of the top VMs and a list of the top VM-related alerts. The top VMs list includes a link to the full list of VMs, and the top alerts list includes a link to the full list

of alerts (which in this case is the same). When you click the VM or alert name, the system redirects you to the details page for that VM or alert.




Figure 14: "Search in Prism" Results

The search is context-sensitive. You can do a quick search from your current location without entering an explicit string. For example, if you are on the user VMs summary page, clicking in the search field (or clicking the forward slash [/] character) displays the search results for user VMs. In the following example, clicking the List, Alerts, or Events entry is the same as clicking those tabs on the page. Clicking *Memory Swap* displays the Memory Swap metrics page, and *Search in Prism* displays a results page for user VMs.




Figure 15: Search Results (for the current location)

Search provides intelligent responses based on whatever you enter in the field. For example, entering *scale out* returns a link to the Manage Prism Central page from which you can scale out Prism Central. If you enter a more generic string such as *version*, the search engine returns results with links to pages across entities with version information.




Figure 16: Search Results (for generic string)

Click the [Star icon](#) in the search field to bookmark a filtered page. This adds a bookmark for that page in the **Navigation Bar**. For more information about **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61).




Figure 17: Bookmark Entry

On some pages, the applied filters are not reflected in the search field. For those pages, a local filters field appears. When you select one or more filters on such a page, the applied filters appear in the local filters field.

You can do the following from this field:

- Remove all filters by clicking **Clear** (on the right).
- Remove a filter by clicking **X** for that filter.
- Bookmark the filter list by clicking the [Bookmark icon](#). You can save a maximum of 20 filter lists per entity type.
- Use a saved filter list by selecting from the dropdown menu.




Figure 18: Local Filters Field

Query Syntax

The search field supports a range of simple and complex query options in the following syntax forms:

- [entity|metric]: Enter an entity or metric type such as *cluster*, *vm*, or *alerts* to return information about the instances of that entity or metric across the registered clusters.
- <entity> <metric>: To refine an entity search, add a metric to the query. For example, *vm iops* returns IOPS information for the VMs.
- <entity> <attribute> <metric>: To further refine the search, add an attribute for the metric. For example, *node failure alerts* returns any failure alerts about the nodes.
- [metric|attribute] operator <value>: Enter an expression for a metric or attribute to return instances that match the expression. For example, *"block model"=1050* returns information about the NX-1050 nodes. See the following "Filter Expressions" section for the supported expression operators.
- <complex expression>: Combine the syntax forms to create complex queries. For example, *clusters hypervisor = AHV "cpu usage" < 30* returns information about clusters running the AHV hypervisor with CPU usage below 30%. Complex expressions have an implied AND so only instances that satisfy all conditions are returned.
- <action expression>: In addition to search queries, you can initiate certain actions from the query bar. For example, *<vm_name> launch console* launches a console window for the specified VM (if you are allowed to do so).

The following table describes the syntax rules for search queries.

Table 12: Syntax Rules

Rule	Example
RULES THAT GO TO WELL DEFINED PAGES	
entity type	vms
entity type + metric perspective (io, cpu, memory)	vms io
entity type + alerts	vm alerts
entity type + alerts + alert filters	vm alerts severity=critical
entity type + events	vm events

Rule	Example
entity type + events + event filters	vm events classification=anomaly
entity type + filters (both metric and attribute)	vm "power state"=on
entity type + filters + metric perspective (io, cpu, memory)	vm "power state"=on io
entity type + filters + alerts	vm "power state"=on alerts
entity type + filters + alerts + (alert filters)	vm "power state"=on alerts severity=critical
entity type + filters + events	vm "power state"=on events
entity type + filters + events + event filters	vm "power state"=on events classification=anomaly
entity instance (name, ip address, disk serial, ...)	vm1, 10.1.3.4, BHTXSPWRM
entity instance + metric perspective (io, cpu, memory)	vm1 io
entity instance + alerts	vm1 alerts
entity instance + alerts + alert filters	vm1 alerts severity=critical
entity instance + events	vm1 events
entity instance + events + event filters	vm1 events classification=anomaly
entity instance + pages	vm1 nics, c1 capacity
parent instance + entity type	c1 vms
alert title search	disk bad alerts
page name search	analysis, tasks
RULES THAT ONLY GO TO SEARCH RESULT PAGES (exclusively)	
entity type + metric	vm iops
entity type + attribute	vm power state
entity instance + metric	vm1 iops
entity instance + attribute	vm1 power state
entity type + filters + metric	vm "power state"=on iops
entity type + filters + attribute	vm "power state"=on hypervisor
help text search	upgrade cluster

Keywords

The following table lists the keywords or phrases you can use when formulating a query.

- The *Object* column identifies the type of object.
 - *Entities*: Lists the entities you can specify. The list is limited. For example, you cannot search for information about remote sites or network switches.
 - *Fields*: Lists the parameters (fields) you can specify.
 - *Alerts*: Lists the alert conditions you can specify.
 - *Events*: Lists the event conditions you can specify.
 - *Actions*: Lists the actions you can specify.
- The *Entity* column specifies the entities for which you can use that keyword in a query. For example, Alert queries can apply to any entity, but Fields has multiple rows broken down by entity type (such as cluster, VM, and disk) because there are select keywords that apply to each entity type.
- The *Keywords* column is a comma-separated list of the keywords or phrases you can use in a query for the specified object/entity type.

Table 13: Search Keywords

Object	Entity	Keywords
Entities	(n/a)	vm, cluster, node, container, disk
Fields	(not specified)	cpu usage, memory usage, disk usage, free physical storage, storage logical usage, saving ratio, savings, iops, read iops, write iops, io bandwidth, read io bandwidth, write io bandwidth, io latency, read io latency, write io latency, memory capacity, hypervisor
	Cluster	cluster name, ip address, version, number of hosts, cpu count, memory capacity, runway, storage runway, cpu runway, memory runway
	VM	vm name, ip address, host ip, virtual cpus count, power state, reserved memory, os, virtual hardware version
	Disk	serial, tier, mode, iops
	Container	container name, rf, on disk deduplication, perf-tier deduplication, compression, erasure coding
	Node	host name, ip address, service vm, hypervisor name, cpu capacity, cpu model, cpu sockets count, cpu cores count, cpu thread count, serial number, block model

Object	Entity	Keywords
Alerts	(any)	alert, alert title - Severity levels: critical, warning, info - Categories (impact types): capacity, performance, configuration, availability, system indicator - Dispositions types: resolved, unresolved, acknowledged, unacknowledged
Events	(any)	event type, title
Actions	VM	clone, migrate, delete, power on, power off, suspend, create vm, launch console, create network config, resume, snapshot, update, configure vm host affinity
	Cluster	launch prism element, unregister

Filter Expressions

You can use any of the following operators in an expression.

Table 14: Query Operators

Description	Operator
contains	~
does not contain	!~
starts with ("starts with" x)	=x*
ends with ("ends with" x)	=*x
equal to	=
not equal to	!=
greater than	>
greater than or equal to	>=
less than	<
less than or equal to	<=
a metric range, for example CPU usage between 10 and 30	=[10 to 30]

The following implicit operator rules apply to expressions:

- In most cases, an AND is applied to all filters in an expression. For example, *hypervisor=AHV iops>100* means the hypervisor is AHV and IOPS are over 100.
- If the filters are on the same attribute, an OR is applied. For example, *hypervisor=AHV hypervisor=ESXi* means the hypervisor is either AHV or ESXi.

- If the filters are on the same metric, an AND is applied to create a range. For example, *iops>100 iops<500* means IOPS in the range 100-500.
- If multiple range filters are defined in a single query, an OR is applied across the ranges. For example, *iops>100 iops<500 cpu>20 cpu<40* means IOPS in the 100-500 range or CPU in the 20-40 range.

Search Units

The data for search queries comes from the following parameters.

Table 15: Data Parameters

Data Type	Unit	Parameter
bytes usage related properties	GiBiBytes	memory_capacity_bytes, storage.usage_bytes, storage.free_bytes, storage.logical_usage_bytes, data_reduction.saved_bytes, storage.usage_bytes, storage.user_container_own_usage_bytes, data_reduction.saved_bytes, memory_size_bytes
runway properties	days	capacity.runway, capacity.storage_runway, capacity.cpu_runway, capacity.memory_runway
disk capacity related properties	TebiBytes	storage.capacity_bytes, storage.user_capacity_bytes
CPU/memory usage properties	percent	hypervisor_cpu_usage_ppm, hypervisor_memory_usage

Example Queries

Here are examples of various query types.

- Entity queries:

```
<cluster_name>
<cluster_ip_address>
<disk_serial#>
VMs "Power State"=On List
VMs Hypervisor=AHV List
powered on vms "memory capacity" > 32
```

- Performance queries:

```
clusters running out of cpu
clusters hypervisor = AHV "cpu usage" < 30
vm iops
```

- Alert queries:

```
node failure alerts
<cluster_name> alerts
<cluster_name> critical availability alerts
<alert title>
Alerts "Create Time"="08-Nov-2018 9:46 AM to 08-Nov-2018 10:46 AM" Severity=Critical
```

- Action queries:

```
<vm_name> launch console
<cluster_name> launch prism element
```

```
create vm
```

- Exploration queries:

```
clusters hypervisor=AHV  
vm os=Linux  
<cluster_name> vms  
"block model"=1050  
"cpu model"=Intel  
containers Rf > 2  
clusters version=4.6.2  
hosts iops < 1000  
<cluster_name> powered off vms  
disks tier=ssd  
vms "cpu usage"  
vms "power state"
```

MAIN DASHBOARD - INFRASTRUCTURE

The **Main Dashboard** provides a dynamic summary view of all the registered clusters in Prism Central.

To view the **Main Dashboard**, select **Infrastructure** application from the [Application Switcher Function](#) on page 40, and click **Dashboard** in the **Navigation Bar**.

Important: Nutanix recommends using Google Chrome web browser to access Prism Central. The display of dashboard and widgets is optimized for Google Chrome browser.

Dashboard Layout

The **Dashboard** provides a collection of widgets that appear as tiles with targeted information about the registered clusters in each tile. The **Dashboard** includes the following options:

- **Main Dashboard** tab: Displays the Infrastructure dashboard. The additional tabs appear if you create any custom dashboards.
- **Manage Dashboard:** Enables you to create a custom dashboard, edit the custom dashboard, or delete a custom dashboard. For more information about how to create a custom dashboard and modify the custom dashboard, see [Creating a New \(Custom\) Dashboard](#) on page 81 and [Modifying a Dashboard](#) on page 83.
- **Reset Dashboard:** Enables you to reset the dashboard to the default set of widgets.
- **Generate Report:** Enables you to generate the complete **Main Dashboard** and **Custom Dashboard** widgets as a report of all the registered clusters. For information about how to generate the dashboard report, see [Generating a Dashboard Summary Report](#) on page 88.

Note: Generating a report requires a valid license - Nutanix Cloud Manager (NCM) Starter, Pro or Ultimate. For more information, see [Nutanix Cloud Manager Licensing](#). The **Generate Report** option does not appear if the appropriate license is not available.

- **Add Widget:** Enables you to add a widget to the displayed dashboard. For information about how to add a widget to a dashboard, see [Adding Dashboard Widgets](#) on page 94.

Note: Adding a widget to a dashboard requires a Prism Pro license. The **Add Widgets** does not appear if Prism Pro is disabled.

- **Data Density** - Enables you to select the density of data representation in the widgets. Select one of the following three options:
 - **Light:** The information is lightly packed in the widget with more spacing between the elements.
 - **Default:** The information is moderately packed with sufficient spacing between the elements. This scheme of data density is based on user research by Nutanix.
 - **Dense:** The information is densely packed with minimal spacing between the elements for essential legibility.




Figure 19: Data Density

Note: The **Manage Dashboards** and **Add Widget** are Prism Pro license features. If Prism Pro is disabled, these options disappear, and you cannot customize or add custom dashboards. Any custom dashboards created before Prism Pro was disabled exists, but clicking the **Reset All** (which replaces the **Reset Dashboard** when custom dashboard exists) deletes all custom dashboards and returns the **Main Dashboard** to the default state.

The following is an example showing the main dashboard:




Figure 20: Main Dashboard

Dashboard Management

The **Main Dashboard** displays an array of information tiles (widgets) by default. You can customize this view by adding or deleting widgets as desired. For information about how to modify a dashboard, see [Modifying a Dashboard](#) on page 83.

The widgets are lined up in rows of four widgets per row. The default widgets are described in this section as they appear from left to right, top to bottom on the **Main Dashboard** page.

The following is an example view of the **Main Dashboard** page with the default widgets:




Figure 21: Main Dashboard with Default Widgets

The following is an example displaying a widget with a period (time) selection:




Figure 22: Sample Widget

- **Alerts:** This widget displays colored bar graph representing the alerts raised in the selected period. You can select the period on the drop down list at the top right corner of the widget. Each category of alerts is represented as a colored bar. Click any bar on the graph to see a list of the alerts represented in the bar. Click **View All Alerts** link to see all the alerts on the **Alerts** page.

The Alerts are categorized as **Critical**, **Warning** and **Info**, and are represented as red, yellow, and blue bars respectively in the graphs.

- **Cluster Quick Access:** This widget displays a list of the clusters managed by the Prism Central. It allows you to open the Prism Element for a registered cluster in a new window when you click the name of the cluster. The cluster **Home** page appears in Prism Element. For more information about the **Home** page of Prism Central, see [Prism Central Web Console Guide](#).
- **Cluster Storage:** This widget displays storage and resiliency information for the highest usage clusters. Each line includes the cluster name, a usage column with a bar that visually indicates how much capacity is used currently, and a data resiliency column that displays the current status (critical, warning, OK, or unknown). The cluster list is ordered by data resiliency status with critical clusters at the top of the list. Click the cluster name to view the **Summary** page for that cluster. Clicking the arrow button displays a dialog box with usage details including data reduction ratio, usage percentage, used space, total space, fault domain (disk, node, or rack), and fault tolerance level (0, 1, or 2). The data reduction ratio indicates the data savings due to data reduction techniques such as de-duplication, compression, and erasure coding. (A 1:1 ratio indicates none of these data reduction techniques are in use currently.)

- **Cluster Latency:** This widget displays the total (read and write) IO latency average for the highest latency clusters. Click *<cluster name>* to display the **Summary** page for that cluster.
- **Cluster Memory Usage:** This widget displays the percentage of total memory in use currently for the highest usage clusters. Click *<cluster name>* to display the **Summary** page for that cluster.
- **Cluster CPU Usage:** This widget displays the percentage of total CPU in use currently for the highest usage clusters (or all clusters if there are fewer than five). Click *<cluster name>* to display the **Summary** page for that cluster.
- **Controller IOPS:** This widget displays the total (read and write) controller IOPS for the highest volume clusters. Click *<cluster name>* to display the **Summary** page for that cluster. The IOPS number comes from the controller when the hypervisor is AHV or Hyper-V and from the hypervisor when the hypervisor is ESXi.
- **Power Usage:** This widget displays the accumulated power consumption over the selected period of time at the cluster level.

For more information on how to enable power usage widget, see [Power Usage](#) on page 633.

- **Cluster Runway:** This widget alerts you to potential storage, CPU, or memory resource constraints across the clusters and provides an estimated *runway* (time remaining) before the resources are maxed out based on current usage. Click *<cluster name>* to display the **Capacity** page for that cluster.
- **VM Efficiency:** This widget displays the number of VMs that are considered inefficient and are broken down by category: over-provisioned, inactive, constrained, and bully. It provides a link to the VMs dashboard. For information about these VMs, see [VMs Summary View](#) on page 100.

For more information, see [Behavioral Learning Tools](#) information in *Intelligent Operations Guide*.

- **Plays:** This widget displays a list of plays running on the cluster. The list is categorised into completed, failed, and paused plays. The completed plays are displayed in the center widget. The number for each category is linked to the **Plays** page.
- **Tasks:** This widget displays a list of recent tasks with the current status of each task. Click the **View All Tasks(s)** link to view the **Tasks** page. For more information, see [Tasks View](#) on page 496.
- **Reports:** This widget displays a table that lists the number of total and scheduled reports with a link to the **Reports** page. For more information, see [Reports Management](#) in *Intelligent Operation Guide*.
- **Data Protection Status:** This widget displays the protection summary and recovery plans if created. If no recovery plan is created, it displays a **How to setup?** link to create a create recovery plan.
- **Security:** This widget displays the summary of **Security Dashboard**. Click the **View All Issues** link to view the **Security Dashboard** page. For more information, see [Security Dashboard](#) on page 420.

Resetting Dashboard

This section describes how to reset an existing dashboard in Prism Central.

About this task

Perform the following steps to reset a dashboard:

Procedure

1. Log in to Prism Central, and select the **Infrastructure** application from the [Application Switcher Function](#) on page 40.
The infrastructure dashboard is displayed.
2. Click **Reset Dashboard** to reset the main dashboard to the default set of widgets.
The system prompts you to confirm the reset dashboard action.

3. Click **OK** to confirm.
The main dashboard is reset to default.

Setting Data Density

This section describes how to set up the data density in Prism Central.

About this task

Perform the following steps to set up the data density:

Procedure

1. Log in to Prism Central, and select the **Infrastructure** application from the [Application Switcher Function](#) on page 40.
The infrastructure dashboard is displayed.
2. Click **Data Density** option to select the density of data representation in the widgets.
The system prompts you to select the data density option.




Figure 23: Data Density Options

Select one of the following three options:

Light: The information is lightly packed in the widget with more spacing between the elements.

Default: The information is moderately packed with sufficient spacing between the elements. This scheme of data density is based on user research by Nutanix.

Dense: The information is densely packed with minimal spacing between the elements for essential clarity.

Creating a New (Custom) Dashboard

About this task

The **Main Dashboard** provides a default view into the clusters, but you can add custom views by creating custom dashboards. To create a custom dashboard, do the following:

Note: Creating a dashboard requires a valid NCM license. For more information, see [License Manager Guide](#). The **Manage Dashboards** does not appear if the appropriate license is not available.

Procedure

1. Log in to Prism Central, and select the **Infrastructure** application from the [Application Switcher Function](#) on page 40.
The infrastructure dashboard is displayed.
2. Click **Manage Dashboards**, and perform the following steps:
 - a. Click **New Dashboard**.
 - b. Enter a name for the dashboard in the displayed field.
 - c. Click **Save**.




Figure 24: Manage Dashboards Windows (new dashboard)

3. Click **Close** to close the **Manage Dashboards** window.
A new dashboard tab appears next to the **Main Dashboard** (or other custom dashboard) tab.

Note: The new dashboard is empty (no widgets) initially.

4. Click the new dashboard tab, and then click the **Add Widgets**.

The system displays the **Add Widgets** window.

5. In the **Add Widgets** window, select a widget and add it to the dashboard. For more information on how to add a widget, see [Adding Dashboard Widgets](#) on page 94.

Repeat this step for all the desired widgets to be added to the new dashboard.

Modifying a Dashboard

About this task

The **Main Dashboard** provides a default view into the registered clusters, but you can customize that view at any time. To modify the Main dashboard or any other dashboard you create, do the following:

Note: Customizing a dashboard requires a Prism Pro license. The **Add Widgets** and **Manage Dashboards** do not appear if Prism Pro is disabled.

Procedure

1. Log in to Prism Central, and select the **Infrastructure** application from the [Application Switcher Function](#) on page 40.
The infrastructure dashboard is displayed.
2. Click **Main Dashboard** tab or previously added custom dashboard tab.

3. Perform either of the following operations:

- To add a widget to the displayed **Main Dashboard** or custom dashboard, click **Add Widgets**, select a widget from the **Add Widgets** window, and add it to the **Main dashboard** or custom dashboard.

For information about how to add a widget, see [Adding Dashboard Widgets](#) on page 94.

- To delete a widget, click *X* in the upper right of the displayed tile. The system prompts you to confirm the delete action.

Click **OK**. The widget disappears from the dashboard.

- To reset the **Main Dashboard** to the default set of widgets (after you have previously added or deleted widgets), click **Reset Dashboard**. The system prompts you to confirm the reset action.

Click **OK**. The **Main Dashboard** returns to its default view.

- To rename a custom dashboard, perform the following steps:

1. Click **Manage Dashboards**.
2. Click the [Edit icon](#) for that dashboard.
3. Enter a new name in the displayed field.
4. Click **Save**.

- To delete a custom dashboard, click **Manage Dashboards** and then click *X* for that dashboard. The system prompts you to confirm the delete action.

Click **OK**. The tab for that custom dashboard disappears from the **Dashboard** page.




Figure 25: Manage Dashboards Window

Sharing a Dashboard

About this task

Any user with edit permission can share a custom dashboard with other users and provide view or edit access to them.

Note: Enable CMSP to use this feature.

Procedure

1. Log in to Prism Central, and select the **Infrastructure** application from the [Application Switcher Function](#) on page 40.
The infrastructure dashboard is displayed.
2. Click **Manage Dashboard**.
The systems displays the **Manage Dashboard** window. This window lists all the existing custom dashboards.

3. Select the custom dashboard you want to share, and click the [Share icon](#).




Figure 26: Manage Dashboards

The **Sharing Dashboard** window appears.

Figure 27: Sharing Dashboard

4. In the **Add Users & Groups** tab, perform the following steps:

- a. In **Share with** field, add the users or groups to share the custom dashboard. You can type a first few characters of a user or group, and as you type, a dropdown menu appears based on the typed keywords. Select the user or group to be added from the displayed list.

Note: The system also allows you to add more than one user or group.

- b. In the **Permission Type** dropdown menu, select the permission as **View only** or **Edit**.

The system assigns the selected permission to the user or group entered in the **Share with** field.

Note:

- Only the user with **Edit** permission can view, edit and share the custom dashboard. The user with **View only** permission can only view the custom dashboard.
- The **Manage Users & Groups** tab allows you to change the permission type for users or groups with whom you have shared the dashboard earlier. You can also remove the users or groups as per the requirement.

5. Click **Save**.

A message appears that indicates the dashboard is shared successfully.

After you share the dashboard, wait for few minutes for changes to take effect.

Note: You cannot view any information about entities in the dashboard chart without entity permissions.

Generating a Dashboard Summary Report

In Prism Central, you can monitor infrastructure widgets by visualizing the data on the dashboards. However, the data is available only in the Prism Central interface. Generating a Report feature allows you to export the infrastructure data as a report for further analysis, archival, compliance, and other purposes.

Note: Generating a report requires a valid license - Nutanix Cloud Manager (NCM) Starter, Pro or Ultimate. For more information, see [Nutanix Cloud Manager Licensing](#). The **Generate Report** option does not appear if the appropriate license is not available.

You can generate a report in the following ways:

- Download the dashboard data as a report in PDF format or CSV format.
- Email the dashboard data as a report to the recipients in PDF format or CSV format, or both, as an attachment.
- Save the dashboard data as a report configuration to schedule, share, and export reports.

The following is an example showing the dropdown menu of the **Generate Report** option from the Main Dashboard:




Figure 28: Main Dashboard - Generate Report Dropdown Menu

The following table describes the dropdown menu list of the **Generate Report** option:

Table 16: Generate Report Dropdown List

Options	Description
Download Report (csv, pdf)	Click this action to download the report. For more information, see Downloading a Report on page 89.
Email Report (csv, pdf)	Click this action to email the report. For more information, see Emailing a Report on page 89.
Save as Report Configuration	Click this action to save the report as a configuration. For more information, see Saving as a Report Configuration on page 90.

Downloading a Report

This task describes how to download the main and custom dashboard widgets data as a report, either in PDF format or CSV format.

About this task

Downloading a report has the following workflow:

- The report instance is generated in your preferred format and notified on the main and custom dashboards.
- The generated report instance is available for you to download under **Generated Reports** tab of the Reports page.

Important:

- You can also resend the generated report instance as an email from the **Generated Reports** tab of the Reports page. For more information, see [Resending a Report Instance](#).
- You must download or resend the report within 24 hours from the time you generated the report instance. After 24 hours, Prism Central removes the report instance from the **Generated Reports** tab.

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from **Application Switcher** function. The infrastructure dashboard is displayed.
3. Click **Main Dashboard** tab or previously added custom dashboard tab.
4. Click **Generate Report**.
5. Select **Download Report (csv, pdf)** and enter the following information:
 - a. **Report Name:** Enter the report name.
 - b. **Report Format:** Select PDF format or CSV format.
6. Click **Download**.
Prism Central displays a confirmation message on the dashboard screen that the report was generated.

What to do next

Select the **Intelligent Operations** application from **Application Switcher** function, and navigate to **Reports** from the **Navigation Bar**. Select **Generated Reports** tab to view and download the generated report instance. For more information, see [Downloading a Report Instance](#).

Emailing a Report

This task describes how to email the main and custom dashboard widgets as a report.

About this task

You can email the main and custom dashboard reports as an attachment in PDF format, CSV format, or both. After you email the report, it is also available as a report instance under **Generated Reports** tab of the Reports page to resend or download the report.

Important: Ensure to resend or download the report within 24 hours from the time you generated the report instance. After 24 hours, Prism Central removes the report instance from the **Generated Reports** tab.

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from **Application Switcher** function.
The infrastructure dashboard is displayed.
3. Click **Main Dashboard** tab or previously added custom dashboard tab.
4. Click **Generate Report**.
5. Select **Email Report (csv, pdf)** and enter the following information:
 - a. **Recipients:** Enter the recipients (comma-separated) email address of the report.
 - b. **Subject:** Enter the subject of the email.
 - c. **Body:** Enter the desired text.
 - d. **Report Format:** Select the **PDF** checkbox or **CSV** checkbox or both to set a report format.
6. Click **Send**.

What to do next

Select the **Intelligent Operations** application from **Application Switcher** function, and navigate to **Reports** from the **Navigation Bar**. Select **Generated Reports** tab to view and manage the generated report instance. For more information, see [Managing a Report Instance from Generated Reports Tab](#).

Saving as a Report Configuration

This task describes how to save the main and custom dashboards data as a report configuration.

About this task

You can save the main and custom dashboards data as a report configuration. The configuration of the report that you save is available under **Configurations** tab of the Reports page to perform report management actions.

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from **Application Switcher** function.
The infrastructure dashboard is displayed.
3. Click **Main Dashboard** tab or previously added custom dashboard tab.
4. Click **Generate Report**.
5. Select **Save as Report Configuration** and enter the following information:
 - **Report Configuration Name:** Enter the report configuration name.
6. Click **Send**.

What to do next

Select the **Intelligent Operations** application from **Application Switcher** function, and navigate to **Reports** from the **Navigation Bar**. Select **Configurations** tab to view and manage the saved report configuration. For more information, see [Managing a Custom Report Configuration](#).

Widgets Management

Widgets provides a quick access to the fault, security, and performance-related data for a cluster such as Alerts, memory usage, CPU usage, CVM IOPS, VM efficiency, Reports, Plays, canceled and succeeded Tasks, Data Protection status, and Security dashboard. Each widget provides unique information, and can be added to the dashboard based on the observation requirements.

Resize the Widgets

You can click and drag down the resize handle at the bottom of the widget to vertically resize a widget.

Note: You cannot resize a widget horizontally.

The following is an example of the resize handle:




Figure 29: Vertically Resizing the Widget

When you click and drag up or down the resize handle to resize a widget, the widget placed below the resized widget shifts upward or downward.

You can resize (increase or decrease) a widget length in multiples of half of the default length of the widget, such as half, one-and-half, twice, and two-and-half times the default length. Even if you attempt to resize it in any other variation of the default length, the widget auto-sizes itself to the nearest half length.

The following is an example of widget resizing:




Figure 30: Widget Resize to One-and-half Times the Default Length




Figure 31: Widget Resize to Twice the Default Length

Drill-down or Expanding Data Elements in a Widget

You can click the data element of a widget to expand or drill-down the information presented in the widget.

Note:

- Expanding a data element in a widget is not possible in all widgets.
- In widgets that provide a list, if you select the list item or row, the system opens an entity page in Prism Central or another application such as Prism Element web console in a new browser tab. For example:
 - In the **Cluster Quick Access** widget, if you click *<cluster name>*, the system opens the Prism Element Web Console for that cluster in a new browser tab.
 - In the **Cluster Storage** widget, if you click *<cluster name>*, the system displays the cluster details view for that cluster.
- In widgets that provide numbers, if you click a number, the system opens the respective entity page in Prism Central. For example:

In the **Reports** widget, if you click **Total Reports** or the **Scheduled Reports** number, the system displays the **Reports** page.

- In widgets that display charts, if you click the chart, the system displays an expanded or drill-down view of the chart.

Note: If you click the entity name (such as `<cluster name>`), instead of displaying the drill-down view of the chart, the system displays the entity details page.

- The following is an example of expanding a chart in a widget.




Figure 32: Widget Data Drill-down

At the bottom of the drill-down view, you can select **Full View** to open the details page of the selected entity (cluster or VM). Click **Add To Analysis** to add the chart to the **Analysis** dashboard. For more information, see [Analysis Dashboard](#) in *Intelligent Operation Guide*.

Hover Details in Drill-down Views

After you select a chart in a widget, the drill-down view of the chart is displayed. You can hover on any point on the time chart to display the parameters at that point.

In the same example as presented in the [Drill-down or Expanding Data Elements in a Widget](#) on page 93, if you hover over the chart, the sample details are displayed.

Adding Dashboard Widgets

Prism Central provides a selection of widgets that can be added to the **Main Dashboard** or a custom dashboard. Each widget provides information about a specific resource or usage topic that is displayed in a tile on the dashboard.

To add a widget to a dashboard, perform the following steps:

Note: Adding a widget to a dashboard requires a Prism Pro license. The **Add Widgets** does not appear if Prism Pro is disabled.

1. Log in to Prism Central, and select **Infrastructure** application from [Application Switcher Function](#) on page 40. The infrastructure dashboard is displayed.
2. In the **Main Dashboard** or custom dashboard, click **Add Widgets**.
3. In the **Add Widgets** window, select the desired widget from the left column.

Note: To find a widget, enter the name in the search field. The system displays a preview of the selected widget.

4. Enter values for the configurable parameters in the right column.

Note: The parameters vary by widget; some widgets do not have any configurable parameters.

5. Click **Add to Dashboard** (or **Add & Return to Dashboard**) to add the widget to the dashboard.

The following table describes the widgets that you can add to the **Main dashboard**:

Table 17: Widget Parameters

Widget	Description	Parameter	Values
Custom Widgets			
Custom Alerts Widget	Displays a custom list of alerts.	Widget Name: Enter a name for the widget. A name (which you can keep or overwrite) is provided after selecting the other options.	(user entered name or settings-based name)
		Select a Cluster: Select the cluster(s) to monitor from the dropdown menu. The default is all clusters.	All Clusters, <cluster name>
		Choose Size: Select the size (width/height dimensions) of the widget. The default is 1 x 1.	1 x 1, 2 x 1, 2 x 2, 2 x 3, 2 x 4, 4 x 4
		Severity. The default is critical.	Critical, Warning, Info
		Entity: Select the entity of interest. The default is cluster.	Cluster, Storage, VM, Hardware, DR
Top Lists Widget	Displays a list of the top performers for a selected metric.	Widget Name: Enter a name for the widget that you can keep or overwrite.	(user entered name or settings-based name)
		Select a Cluster: Select the cluster(s) to monitor from the dropdown menu. The default is all clusters.	All Clusters, <cluster name>
		Choose Size: Select the size (width/height dimensions) of the widget. The default is 2 x 1.	2 x 1, 2 x 2, 2 x 3
		Entity: Select the entity of interest. The default is host.	host, VM
		Metric: Select the metric of interest. The default is IOPS.	IOPS, Memory Usage, CPU Usage, Bandwidth, Latency
Custom Chart Widget	Displays a performance graph for a selected entity and metric pair.	Widget Name: Enter a name for the widget. The default is <i>New Chart Widget</i> .	(name)
		Entity Type: The default is host.	Host, Disk, Storage Pool, Storage Container, Virtual Machine, Virtual Disk

Widget	Description	Parameter	Values
		Entity: There is no default value. The system displays the list of entity items based on the Entity Type field.	(entity name)
		Metric: There is no default value. For information about metrics available for the selected entity item, see the information available in Metrics Tab section of that individual entity . For example, for VM-related metrics information, see Metrics Tab section in VMs Summary View on page 100 and Viewing VM Details on page 113.	(metric name)
Cluster Info Widget	Displays cluster summary information about alerts, anomalies, runway, and inefficient VMs.	Widget Name: Enter a name for the widget. The default is <i>New Widget</i> .	(n/a)
		Select a Cluster: Select the target cluster from the dropdown menu.	(registered cluster names)
Virtual Infrastructure Widgets: These widget appears on the Main dashboard by default. For details, see Main Dashboard - Infrastructure on page 76.			
VM Efficiency	Displays the number of overprovisioned, inactive, constrained, and bully VMs in the cluster with links to the details.	(no customizable parameters)	(n/a)
Storage Over-provisioning	Displays the configured threshold for storage over-provisioning	Storage Over-provisioning Ratio	
Hardware Widgets: All these widgets (except Impacted Cluster and Performance) appear on the Main dashboard by default. For details, see Main Dashboard - Infrastructure on page 76.			
Cluster CPU Usage	Displays the percentage of total CPU in use currently for the highest usage clusters.	(no customizable parameters)	(n/a)
Cluster Latency	Displays the total (read and write) IO latency average for the highest latency clusters.	(no customizable parameters)	(n/a)
Cluster Memory Usage	Displays the percentage of total memory in use currently for the highest usage clusters.	(no customizable parameters)	(n/a)

Widget	Description	Parameter	Values
Cluster Quick Access	Displays a list of registered clusters. Health and alert icons appear for each cluster. Clicking the cluster line opens Prism (element) for that cluster in a separate tab or window.	(no customizable parameters)	(n/a)
Cluster Runway	Displays storage, CPU, and memory runway estimates (time remaining before the resource reaches capacity). For more information, see Capacity Runway Summary View in <i>Intelligent Operation Guide</i> .	(no customizable parameters)	(n/a)
Cluster Storage	Displays storage statistics for the highest usage clusters.	(no customizable parameters)	(n/a)
Controller IOPS	Displays the total (read and write) controller IOPS for the highest volume clusters.	(no customizable parameters)	(n/a)
Impacted Cluster	Displays information about any clusters that are impacted (performance, capacity, or other potential issues) and may need attention.	(no customizable parameters)	(n/a)
Performance	Displays latency, bandwidth, and IOPS statistics for the highest usage clusters.	(no customizable parameters)	(n/a)
Alerts	Displays colored bar graph that represents the alerts raised in the period selected. Click View All Alerts link to see all the alerts on the Alerts page. For more information about the alerts generated in Prism Central, see Prism Central Alerts and Events Reference Guide .	(no customizable parameters)	(n/a)
Activity Widgets: The Tasks widget appears on the Main Dashboard by default. For more information, see Main Dashboard - Infrastructure on page 76.			
Tasks	Displays a list of recent tasks with the current status of each task and a link to the Tasks page.	(no customizable parameters)	(n/a)

Widget	Description	Parameter	Values
Discovered Apps	Displays the number of applications discovered in a specified set of monitored clusters. Click the App instances number to display the application discovery page. Click the Identified or Unidentified number to display the application discovery page filtered for that condition. This widget displays the discovered apps information only if you enable App Discovery.	(no customizable parameters)	(n/a)
Operations Widgets: These widgets appear on the Main Dashboard by default. For more information, see Main Dashboard - Infrastructure on page 76.			
Plays	Displays a list of the Plays running on the cluster. The list is categorised into completed, failed and paused plays. The completed plays are displayed in the center widget. The number for each category is linked to the Plays page.	(no customizable parameters)	(n/a)
Reports	Displays a table that lists the number of total and scheduled reports with a link to the Reports page.	(no customizable parameters)	(n/a)
Nutanix Disaster Recovery Widgets: These widgets appear on the Main Dashboard by default. For more information, see Main Dashboard - Infrastructure on page 76.			
Data Discovery Status	Displays the Protection summary and recovery plans if created. If no recovery plan is created, it displays a How to setup? link to create a create recovery plan.	(no customizable parameters)	(n/a)
Security Widgets: These widgets appear on the Main Dashboard by default. For more information, see Main Dashboard - Infrastructure on page 76.			
Security	Displays the summary of Security Dashboard. Select the View All Issues link to view the Security Dashboard page. For more information, see Security Dashboard on page 420.	(no customizable parameters)	(n/a)

COMPUTE ENTITIES

You can access the following entity items from the **Compute** entity of the **Infrastructure** application:

- VMs (see [VM Management](#) on page 99)
- Templates (see [VM Template Management](#) on page 212)
- Kubernetes Clusters (see [Kubernetes Clusters Management](#) on page 226)
- OVAs (see [OVA Management](#) on page 235)
- Images (see [Image Management](#) on page 258)
- Catalog Items (see [Catalog Management](#) on page 283)
- vCenter Datastores (see [External vCenter Server Integration](#) on page 341)

For information about how to access the entities items available in **Compute** entities, see [Application-specific Navigation Bar](#) on page 61.

VM Management

You can perform the following actions to manage a VM from Prism Central:

- Create and manage VMs directly from Prism Central when the hypervisor is either ESXi or AHV. For more information, see the following sections:
 - [Creating a VM through Prism Central \(ESXi\)](#) on page 186.
 - [Managing a VM through Prism Central \(ESXi\)](#) on page 188.
 - [Creating a VM through Prism Central \(AHV\)](#) on page 149.
 - [Managing a VM through Prism Central \(AHV\)](#) on page 161.
- Create and manage a VM from Prism Central using Prism Self Service. For more information, see the following sections:
 - [Creating a VM \(Self Service\)](#) on page 571.
 - [Managing a VM \(Self Service\)](#) on page 578.
 - [Creating a VM from Catalog Items \(Self Service\)](#) on page 571

For more information about Prism Self service setup, see [Prism Self Service Setup](#) on page 569.

- Add multiple vGPUs to the same VM, and perform a live migration of vGPU-enabled VMs within or outside the cluster. For more information, see the following sections:
 - [Adding Multiple vGPUs to the Same VM](#) on page 172
 - [Live Migration of vGPU-enabled VMs](#) on page 179
- Export VM as an OVA. For more information, see [Exporting a VM as an OVA](#) on page 238.
- Manage NGT. For more information, see [Nutanix Guest Tools Overview](#) on page 189.
- Define the QoS for Storage. For more information, see [Storage Quality of Service \(QoS\)](#) on page 210.
- Perform Memory Overcommit for new and existing VMs. For more information, see [Memory Overcommit Management](#) on page 212.

- Define VM-Host Affinity, VM-VM Anti-affinity, and NGT Policies for VM Management. For more information, see [Policies for VM Management](#) on page 212.

VMs Summary View

The **Summary** tab on the **VMs** page provides a dashboard of all the VMs across registered clusters.

To access the summary view of all VMs, perform the following steps:

1. Log in to the Prism Central web console.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Compute > VMs** from the **Navigation Bar**. For information on the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

The system displays the **List** tab by default with all the VMs across registered clusters in **Nutanix** environment.

3. Click the **Summary** tab.

The system displays the **Summary** view of all the VMs across registered clusters:

Note: To access the summary view of all the non-nutanix VMs that are managed by an external vCenter, select **Non-Nutanix** environment from the dropdown menu in the **VMs** page. If you have opted-in for decoupled NCM, you can access the summary view of all the non-nutanix VMs in the **Infra Browser** in Intelligent Operations. For more information, see [Intelligent Operations Overview](#) in *Intelligent Operations Guide*.

The **Summary**, **Alerts**, and **Events** tabs display the same information as for Nutanix-managed VMs. The **Metrics** tab displays a subset of the full list (10 of the 12 metrics). The **List** tab displays fields for name, node name, hypervisor, memory, IP addresses, power state, and cluster name. While displaying the VM name, vCenter encodes any special characters included in the VM name to utf-8.

Using playbooks, you can manage the VMs in **Non-Nutanix** environment. For more information on playbooks, see [Task Automation - Playbooks](#) in *Intelligent Operations Guide*.




Figure 33: Summary View - All VMs

The **VMs** page includes six tabs on the left (**Summary**, **List**, **Policies**, **Alerts**, **Events**, and **Metrics**).

Note: This section describes the information and options that appear in the **Summary** page for all VMs. For instructions on how to view and organize that information in various ways, see [Prism Central GUI Organization](#) on page 49.

The **Summary** tab for all VMs displays the following three widgets:

- **Highlighted Entities:** Displays a list of the VMs with the highest usage of the *<parameter>* you select from the dropdown menu on the right of the widget. The *<parameter>* involves **CPU Usage**, **IO Latency**, **Memory Usage**, and **IOPS**. Click **View All XX VMs** link at the bottom to display the **List** tab.

- **Alerts:** Displays a list of VM-related alerts that are generated during the specified *<interval>* you select from the dropdown menu on the right of the widget. The *<interval>* involves **Last week**(default), **Last 24 hours**, and **Last 1 hour**. When an alert appears, you can click the graph to view a list of those alerts. Click any alert to display the details page for that alert.
- **Anomalies:** Displays a graph of memory, I/O, CPU, networking, or disk anomalies that occurred during the specified *<interval>* you select from the dropdown menu on the right of the widget. The *<interval>* involves **Last week** (default), **Last 24 hours**, and **Last 1 hour**. When an anomaly appears, you can click the graph to view a list of those anomalies. Click any anomaly to display the event page for that anomaly.

List Tab

The **List** tab displays the list of VMs across all registered clusters.




Figure 34: VMs List Tab

The following table describes the fields that appear in the VMs **List** tab.

Note: The fields vary based on the **View by** and **Group by** options. A dash (-) is displayed in a field when a value is not available or applicable. For more information on **View by** and **Group by** options, see [Prism Central GUI Organization](#) on page 49.

Table 18: VMs List Tab - Field Descriptions

Field	Description	Values
Select General from View by option.		
Name	Displays the VM name. You can click the VM name to view the detailed information for that VM. For information on VM details view, see Viewing VM Details on page 113.	(VM name)
vCPU	Displays the virtual CPU count of the VM.	(vCPU count)
Memory	Displays the total amount of memory available to this VM.	xxx [MB GiB]

Field	Description	Values
IP addresses	Displays one or more IPv4 and IPv6 IP addresses assigned to the VM. For conditions applicable to IPv6 IP addresses, see the note placed after this table.	(comma-separated list of IPv4 and IPv6 addresses)
Cluster	Displays the name of the cluster in which the VM resides.	(cluster name)
Hypervisor	Displays the hypervisor type on which the VM is running.	AHV, ESX, or Hyper-V
OS	Displays the guest operating system of the VM.	(OS)
NGT	Displays whether Nutanix Guest Tools is installed on the VM.	Installed, Not Installed
Project	Displays the name of the project to which this VM belongs.	(project name)
Owner	Displays the owner (user name) of this VM.	(user name)
Select Performance from View by option.		
Name	Displays the VM name.	(VM name)
Memory Overcommit	Displays the Memory Overcommit state of the VM.	Enabled / Disabled
Memory Usage	Displays the percentage of allocated memory capacity currently being used by this VM.	0 -100%
CPU Usage	Displays the percentage of allocated CPU currently being used by this VM	0 -100%
Read IOPS	Displays read I/O operations per second (IOPS) for this VM.	(number)
Write IOPS	Displays write I/O operations per second for this VM.	(number)
IO Bandwidth	Displays I/O bandwidth used per second for this VM.	xxx [MBps KBps]
IO Latency	Displays the average I/O latency for this VM.	xxx [ms]
Cluster	Displays the name of the cluster in which the VM resides.	(cluster name)
Select Anomalous Behavior from View by option.		
Name	Displays the VM name.	(VM name)
Anomaly Count	Displays the Anomaly count for Memory, I/O, CPU, Network, and Disks	(Integer number)
Anomaly Disabled	Displays whether Anomaly reporting is Disabled.	Yes/No
Select Efficiency from View by option.		
Name	Displays the VM name.	(VM name)

Field	Description	Values
Efficiency	Displays the assessed efficiency of the VM. If a VM is performing in an expected range, the efficiency is listed as <i>Good</i> . If not, the type of inefficiency is displayed as determined by the VM behavioral learning engine. For more information, see Behavioral Learning Tools in <i>Intelligent Operations Guide</i> .	Good, Constrained, Overprovisioned, Inactive, Bully; (NA if insufficient baseline data for categorization, which is typically 21 days)
Reason	Displays the reason why a VM is considered inefficient. A dash (-) appears for a <i>Good</i> VM.	(text message)
Project	Displays the name of the project to which this VM belongs.	(project name)
Owner	Displays the owner (user name) of this VM.	(user name)
Cluster	Displays the name of the cluster in which the VM resides.	(cluster name)
Select GPU from View by option.		
Name	Displays the VM name.	(VM name)
GPU Configuration	Displays the GPU board (in case of Passthrough mode), vGPU software, type and number of vGPU instances in brackets (in case you configured the vGPU mode).	(GPU board/software-type name) Example: Passthrough - Nvidia Tesla M60 Example: vGPU - Nvidia GRID M60-8Q (2)
GPU Type	Displays the GPU configuration made as Passthrough or vGPU or if you did not configure GPU then a dash (-) appears.	[Passthrough vGPU -]
GPU Usage	Displays the percentage of GPU capacity being used by the VM.	(percentage)
GPU Framebuffer Usage	Displays the percentage of GPU framebuffer (RAM) capacity being used by the VM.	(percentage)
Select Data Protection from View by option.		
Name	Displays the VM name. You can click the VM name to view the detailed information for that VM. For information on VM details view, see Viewing VM Details on page 113.	(VM name)
Categories	Displays the category mapped to the VM. For information on categories in Prism Central, see Category Management in the <i>Prism Central Admin Center Guide</i> .	String (Category name)
Consistency Group	Displays the consistency group of the VM. For information on Consistency Groups in Prism Central, see Consistency Groups information in <i>Nutanix Disaster Recovery Guide</i> .	(Yes/No)

Field	Description	Values
Protection Status	Displays the Protection Status of the VM.	(Protected/Unprotected)
Sync Status	Indicates the VM synchronization status between the Availability Zones in the Disaster Recovery setup.	(Sync Status)
Protection Type	Displays the type of protection defined for the VM.	(Protection Policy/ Legacy Protection Domain)
Protection Policy	Displays the protection policy that is applicable for the VM. For more information on protection policies, see Nutanix Disaster Recovery Guide .	(String) Protection Policy Name
Recovery Plans	Displays the recovery plans that is applicable for the VM. For more information on VM recovery plans, see Nutanix Disaster Recovery Guide .	(String) Recovery Plans Name
Select Storage Configuration from View by option.		
Name	Displays the VM name.	(VM name)
Replication Factor	Displays the Replication Factor setting for the VM in the storage policy applied to the VM.	Inherit from Container or 2 or 3
Encryption	Displays the Encryption setting for the VM in the storage policy applied to the VM.	Enabled or Inherit from Cluster
Compression	Displays the Compression setting for the VM in the storage policy applied to the VM.	Inline or Post Process (if enabled) , Disabled or Inherit from Cluster
Throttled Throughput (IOPS)	Displays the throttled throughput value in terms of IOPS set for the VM in the storage policy applied to the VM.	(Integer number)
Throttled Throughput (MB/s)	Displays the throttled throughput value in terms of MB/s set for the VM in the storage policy applied to the VM.	(Integer number)
Associated Policy	Displays the name of the storage policy applied to the VM.	(String) Storage Policy Name

Field	Description	Values
Compliance State	Displays the compliance state of the VM based on realization of the storage policy applied to the VM. Note: The compliance state is displayed as In Progress until the extent groups are synchronized by the Curator service at full scan intervals. Full scan intervals are intervals of 6 hours. Depending up on the time elapsed in the full scan interval at the time when the Replication Factor was set or updated, it may take up to 6 hours for the In Progress state to change to Compliant or Non Compliant.	In Progress, Compliant or Non Compliant
Cluster	Displays the name of the cluster that the VM resides in.	(String) Cluster Name
Select Add custom from View by option.		
Acropolis VM	Displays whether the VM is running on the AHV host.	(Yes/No)
Anomalies (Last 24 hours)	Displays the sum of all the anomaly counts for 20 metrics such as memory, I/O, CPU, network, and disks in the last 24 hours.	(Integer number)
AZ Name	Displays the availability zone specified to the VM.	AZ name
Cluster UUID	Displays the UUID of the cluster in which the VM resides.	(cluster UUID)
Constrained Level	Displays the VM constrained level along with the metric that constrains the VM.	“high”:”memory”, “moderate”:”cpu” or “high”:”cpu”, “moderate”:”memory”
Controller IOPS	Displays the total (read and write) I/O operations per second (IOPS) for the VM.	(Integer number)
Controller Read IO Bandwidth	Displays the read I/O bandwidth of the VM.	(Integer number (KBps))
Controller Read I/O Latency	Displays the read I/O latency of the VM.	(Integer number (ms))
Controller Write IO Bandwidth	Displays the write I/O bandwidth of the VM.	(Integer number (KBps))
Controller Write IO Latency	Displays the write I/O latency of the VM.	(Integer number (ms))
Created Timestamp	Displays the date and time of VM creation.	(date and time), Timestamp in MM DD, YYYY, hh: mm AM/PM format. For example, Oct 27, 2023, 06:51 AM
Description	Displays the description specified for the VM.	(description text)

Field	Description	Values
Disk Capacity	Displays the total disk capacity of the VM.	(Integer number (GiB))
Disk Usage	Displays the disk usage percentage for the VM.	(0 - 100%)
Disk Usage Bytes	Displays the amount of used disk space.	(Integer number (GiB))
Efficiency Detail	Displays the reason why a VM is considered inefficient.	(text message)
Health	Displays the health state of the VM.	(Critical/Warning/Good)
Host	Displays the host name. This field might be blank if the VM is powered off and a host is not assigned.	(hostname)
Host IP	Displays the host IP address.	(IP address)
Host UUID	Displays the UUID of the host.	(host UUID)
Hypervisor IOPS	Displays the total (read and write) I/O operations per second (IOPS) from the hypervisor for the VM.	(Integer number)
IS CVM	Displays whether the VM is a CVM or not.	(Yes/No)
Network Adapters	Displays the number of network adapters in the VM.	(Integer number)
Network Rx Bytes	Displays the number of bytes received from the network.	(Integer number)
Network Rx Packets Dropped	Displays number of packets received from the network and dropped.	(Integer number)
Network Tx Packets Dropped	Displays number of packets sent to the network and dropped.	(Integer number)
NGT communication status	Displays the NGT communication status between CVM and guest VMs. For more information, see Nutanix Guest Agent and Controller VM Communication on page 206.	(Active/Inactive)
NGT communication type	Displays the transport channel established between CVM and guest VMs. For more information, see Verifying Communication Type Using Prism Central on page 208.	(IP-based/IP-less)
NGT Installed Version	Displays if NGT is installed in the VM. For more information, see Setting up NGT on page 193.	(Installed/Not Installed/Latest)
Overprovisioned	Displays the type of over-provisioned VMs.	(Good/High/Moderate)
	Note: This and the following field filters based on a VM efficiency algorithm, which is part of the VM behavioral learning capabilities (see Behavioral Learning Tools).	
Overprovisioned level	Displays the VM overprovisioned level along with the metric that overprovisions the VM.	“high”:”memory”, “moderate”:”cpu” or “high”:”cpu”, “moderate”:”memory”

Field	Description	Values
Power State	Displays whether the VM is powered off or powered on.	(Off/On)
Services Enabled	Displays the name of the NGT services enabled in the VM.	(file_level_restore, vss_snapshot)
Shared Usage	Displays the shared data usage by the VM.	(Integer number (GiB))
Snapshot Usage	Displays the amount of storage used for snapshots.	(Integer number (GiB))
Total Working Set Size	Displays the total working set size of the VM. Working set is the amount of memory that a VM requires in a given time interval.	(Integer number (GiB/MiB))
Virtual NICS	Displays the name of the virtual NIC that allows the VM to connect to a network	(NIC name)
VM Type	Displays the type of VM.	(Guest, PCVM, MSP, AFS)
VPC Name	Displays the VPC Name specified for the VM.	(VPC name)
Working Set Size Read	Displays the read working set size of the VM.	(Integer number (GiB/MiB))
Working Set Size Write	Displays the write working set size of the VM.	(Integer number (GiB/MiB))

Note: Assignment of IPv6 IP addresses to VM is subject to the following conditions:

- IPv6 IP addressing is not supported on Basic VLAN Subnets. It is supported only on unmanaged (Network Controller-based) VLAN Subnets. For more information about the VLAN types, see [Network Types](#) in the *Flow Virtual Networking Guide*.
- The underlying physical infrastructure must have an IPv6 network configured.
- VMs created on hosts that are networked with an IPv6 network display the IPv6 IP address assigned to them.

You can perform the following actions for the VMs in the **List** tab:

- Access the detailed information about an individual VM. For more information, see [Viewing VM Details](#) on page 113.
- Filter the VMs list based on available parameter values using **Filters** pane. For more information, see [Filters Pane - VMs Page](#) on page 108.
- Export the table that contains the list of VMs and their information to a file in a CSV format. For more information, see [Export](#) on page 54.
- Group the VMs based on pre-defined criteria. For information on how to group the VMs, see [Group by](#) on page 51.
- View VMs based on pre-defined criteria or create a custom view. For information on available views and how to create a custom view, see [View by](#) on page 51.
- Assign a label to the filtered criteria. For information on how to define a label, see [Label](#) on page 54.

- Create a VM. For more information on how to create a VM, see [Creating a VM through Prism Central \(AHV\)](#) on page 149 or [Creating a VM through Prism Central \(ESXi\)](#) on page 186.
- Perform the VM-specific actions on a single or multiple VMs using the **Actions** dropdown menu. For instructions on how to perform these actions, see [Managing a VM through Prism Central \(AHV\)](#) on page 161 or [Managing a VM through Prism Central \(ESXi\)](#) on page 188.

Note: The available actions appear in bold; other actions are grayed out. For grayed out options, a tool tip explaining the reason is provided. The available actions depend on the current state of the selected VM(s).

Filters Pane - VMs Page

The following table describes the fields available in the **Filters** pane:

Table 19: Filter Pane Fields

Parameter	Description	Values
Labels	Filters based on label name. Select one or more labels from the Labels dropdown menu. If there are no labels currently, the system displays a message about how to create labels.	(label names)
Name	Filters based on the VM name. Select a condition from the dropdown menu and enter a string in the field. The system returns a list of VMs that satisfy the VM name condition/string. Note: In this field, the condition menu options are Contains , Doesn't contain , Starts with , Ends with , and Equal to .	(VM name string)
Description	Filters based on the description specified for the VM. Enter a string in the field to filter the result. The system returns a list of VMs that satisfy the description string. Note: In this field, the condition menu options are Contains , Doesn't contain , Starts with , Ends with , and Equal to .	(Description string)
IP Addresses	Filters based on the IPv4 and IPv6 IP addresses specified for the VM. Enter a valid IP address in the field to filter the result. The system returns a list of VMs that satisfy the IP address string. Note: In this field, the condition menu options are Contains , Doesn't contain , Starts with , Ends with , and Equal to .	(A Valid IP Address)

Parameter	Description	Values
VPC Name	Filters based on the VPC Name specified for the VM. Enter a string in the field to filter the result. The system returns a list of VMs that satisfy the VPC string. Note: In this field, the condition menu options are Contains , Doesn't contain , Starts with , Ends with , and Equal to .	(VPC name string)
Host	Filters based on the host name. Select a condition from the dropdown menu and enter a string in the field. The system returns a list of VMs that satisfy the host name condition/string.	(host name string)
Cluster	Filters based on the cluster name. Select a condition from the dropdown menu and enter a string in the field. The system returns a list of VMs that satisfy the cluster name condition/string.	(cluster name string)
Categories	Filters based on category names. Enter a category name in the field and then check the box. As you type a dropdown menu appear to help you select the correct category. A new field appears where you can add more categories to the filter. The number of VMs tagged to each selected category are displayed in the List tab.	(category name)
Quarantined	Filters based on the Quarantined mechanism. Select either Strict or Forensics or both the checkboxes. The system displays the number of VMs currently quarantined as Strict or Forensics.	Strict, Forensics
Hypervisor	Filters based on the hypervisor type. Select one or more of the checkboxes to filter on those hypervisors. The system displays the number of VMs currently on each hypervisor type.	AHV, ESX, Hyper-V
Health	Filters based on the VM health state (good, warning, or critical). Select one or more states to return a list of VMs in that state(s). The system displays the number of VMs currently in each state.	Critical, Warning, Good
Nutanix Guest Tools	Filters based on the NGT status. Select one or more NGT status to return a list of VMs with selected NGT status. For more information on Nutanix Guest Tools filters, see Setting up NGT on page 193.	Installed, SSR, VSS, VSS and SSR, Not installed
Power State	Filters based on the VM power state. Select one or more states to return a list of VMs in that state(s). The system displays the number of VMs currently in each state.	On, Off, Suspended, Paused, Unknown
Is CVM	Filters based on the whether the VM is a CVM or not. The system displays the number of CVMs, if filter selected is Yes .	No, Yes

Parameter	Description	Values
VM Type	Filters based on the type of VM. Select either Guest, PCVM, MSP, AFS or multiple checkboxes. The system displays the number of VMs of each type.	Guest, PCVM, MSP, AFS
Memory Overcommit	Filters based on the Memory Overcommit status. Select one or more Memory Overcommit status to return a list of VMs with selected status.	Enabled, Disabled, Not Available
Memory Usage	Filters based on the amount of memory capacity being used. Select the checkbox(es) for the desired percentage range or enter a percentage range in the from <low> to <high> % field. The system returns a list of VMs utilizing memory in that range.	([xx] to [yy]% range)
CPU Usage	Filters based on the processor usage. Select the checkbox(es) for the desired percentage range or enter a percentage range in the From <low> to <high> % field. The system returns a list of VMs utilizing memory in that range.	([xx] to [yy]% range)
Read IOPS	Filters based on the read IOPS. Select the checkbox(es) for the desired range or enter a range in the From <low> to <high> IOPS field. The system returns a list of VMs with read IOPS in that range.	([xx] to [yy] IOPS range)
Write IOPS	Filters based on the write IOPS. Select the checkbox(es) for the desired range or enter a range in the From <low> to <high> IOPS field. The system returns a list of VMs with write IOPS in that range.	([xx] to [yy] IOPS range)
I/O Bandwidth	Filters based on the I/O bandwidth used. Select the checkbox(es) for the desired range or enter a range in the From <low> to <high> Mbps field. The system returns a list of VMs with I/O bandwidth usage in that range.	([xx] to [yy] Mbps range)
I/O Latency	Filters based on the average I/O latency. Select the checkbox(es) for the desired range or enter a range in the From <low> to <high> ms field. The system returns a list of VMs with average I/O latency in that range.	([xx] to [yy] ms range)
Over Provisioned	Filters for over-provisioned VMs. Select the checkbox(es) for the desired type (high and moderate).	High, Moderate
Note: This and the following two fields filter based on a VM efficiency algorithm, which is part of the VM behavioral learning capabilities. For more information, see Behavioral Learning Tools in <i>Intelligent Operations Guide</i>.		

Parameter	Description	Values
Constrained	Filters for constrained VMs. Select the checkbox(es) for the desired type (high and moderate).	High, Moderate
Efficiency	Filters for certain VM profiles. Check the boxes for the desired profile types. There is one for efficient VMs (good) and four for inefficient VMs (bully, over-provisioned, constrained, and inactive).	Bully, Over Provisioned, Constrained, Inactive VM, Good
GPU Configuration	Filters for GPU configuration information such as model name. Enter the GPU configuration information in the field and then select the checkbox. As you type, a dropdown menu appears to help you select the correct configuration information.	(configuration info)
GPU Type	Filters based on GPU operational mode. Select the checkbox for one or more of the GPU types.	vGPU, Passthrough, Passthrough(Compute)
vCPU	Filters based on the number of vCPUs allocated to the VM. Select one or more numeric checkboxes to return a list of VMs with selected number of vCPUs.	vCPU Integer value
GPU Usage	Filters based on the amount of GPU capacity being used. Enter a percentage range in the from <low> to <high> % field. The system returns a list of GPUs in that range.	([xx] to [yy]% range)
GPU Framebuffer Usage	Filters based on the amount of GPU framebuffer (RAM) capacity being used. Enter a percentage range in the from <low> to <high> % field. The system returns a list of GPUs in that range.	([xx] to [yy]% range)
Storage Configuration	Filters based on the storage policy configuration parameters set for VM. This view displays the storage configurations parameters set for the VMs by storage policies applied to the VMs. For more information on storage policies, see Storage Policies Summary View on page 550.	Storage policy name
vGPU Guest driver Version	Filters on the guest driver version. Enter the guest driver version number in the field.	(guest driver version number)
AZ Name	Filters based on the availability zone specified of the VM. Enter a string in the field to filter the result. The system returns a list of VMs that satisfy the AZ string. Note: In this field, the condition menu options are Contains, Doesn't contain, Starts with, Ends with, and Equal to .	AZ String
NGT Communication Status	Filters based on the NGT communication status between CVM and guest VMs. Select the communication status to filter the result. The system returns a list of VMs that satisfy the selected status.	Active, Inactive

Parameter	Description	Values
NGT Communication Type	Filters based on the transport channel established between CVM and guest VMs. Select the required transport channel checkbox to filter the result. The system returns a list of VMs that satisfy the selected transport channel.	IP-Less, IP-Based

Note: Assignment of IPv6 IP addresses to VM is subject to the following conditions:

- IPv6 IP addressing is not supported on VLAN Basic Subnets. It is supported only on unmanaged (Network Controller-based) VLAN Subnets. For more information about the VLAN types, see [Network Types](#) in the *Flow Virtual Networking Guide*.
- The underlying physical infrastructure must have an IPv6 network configured.
- A VM displays the assigned IPv6 IP address when it is attached to an unmanaged VLAN Subnet that has IPv6 enabled by the underlying physical infrastructure.

Policies Tab

The **Policies** tab displays **VM-Host Affinity Policies**, **VM-VM Anti-affinity Policies**, and **NGT Policies**.

For more information on how to configure **VM-Host Affinity Policies**, **VM-VM Anti-affinity Policies**, and **NGT Policies**, see [VM-Host Affinity Policies Defined in Prism Central](#) on page 501, [VM-VM Anti-Affinity Policies Defined in Prism Central](#) on page 512, and [NGT Policies](#) on page 522.

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options that are available on the **Activity > Alerts** page, however it only displays the VM-related alerts across the registered clusters. For more information on alerts, see [Prism Central Alerts and Events Reference Guide](#).

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options that are available on the **Activity > Events** page, however it only displays the VM-related events across the registered clusters. For more information on events, see [Prism Central Alerts and Events Reference Guide](#).

Metrics Tab

The **Metrics** tab allows you to view performance metrics across the VMs. Click the **Metrics** tab to display dropdown menu of available metrics, and select the metric name to display the relevant performance information.

Note: The **Metrics** dropdown menu is hypervisor-specific, and might vary based on the hypervisors used in the cluster.

The following table describes the dropdown menu list of the **Metrics** tab:

Table 20: Metrics Tab Fields

Metric	Description
CPU Usage	Displays a CPU usage table that lists the current values and total VMs (number). The current values are split into percentile intervals (for example, less than 25%, 25-50, 50-75, more than 75%). You can click a percentile interval to view the Summary tab filtered with selected VMs. Note: The same format also applies to the other metrics in this table with either percentile or quantity intervals.
CPU Ready Time	Displays a CPU ready time percentage usage table.
Memory Usage	Displays a memory percentage usage table.
Memory Swap	Displays memory swap-out and swap-in rate tables.
IOPS	Displays total, read, and write IOPS tables.
IO Latency	Displays total, read, and write I/O latency rate tables.
I/O Bandwidth	Displays total, read, and write I/O bandwidth rate tables.
Storage Usage	Displays total, snapshot, and shared storage size tables.
Working Set Size	Displays total, write, and read working set size tables. Working set is the amount of memory that a VM requires in a given time interval.
Network Packets Dropped	Displays tables for the number of transmitted and received packets dropped.
Network Bytes	Displays tables for the amount of transmitted and received bytes (in GiB).
Disk Usage	Displays a disk usage table that lists the current values and total VMs (number). The current values are split into percentile intervals, for example, 13.68% - 4.1%, more than 4.93%. Click a percentile interval to display the Summary tab filtered to just those VMs. Note: The same format also applies to the other metrics in this table with either percentile or quantity intervals.

Viewing VM Details

View the details of a VM.

About this task

To view the details view of an individual VM, follow these steps:

Procedure

1. Log in to the Prism Central web console.
2. From [Application Switcher Function](#) on page 40, select the **Infrastructure** application.

3. From the **Navigation Bar**, select **Compute > VMs**.

For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61. The system displays the **List** tab by default with all the VMs across registered clusters in **Nutanix** environment. For more information on how to access the list of VMs in **Non-Nutanix** environment (non-nutanix VMs managed by external vCenter), see [VMs Summary View](#) on page 100.

4. Click the target <VM_Name>.

Replace <VM_Name> with the actual VM name at your site.

The **Summary** tab of the selected VM is displayed. The following is an example showing the **Summary** tab of an individual VM:




Figure 35: Summary Tab - Individual VM

5. Click any of the following tabs to perform operations on the VM:

- **Summary**

For more information, see [Viewing the list of VM instances](#) on page 120.

- **Console**

For more information, see [Console Tab](#) on page 120.

- **Recovery Points**

For more information, see [Recovery Points Tab](#) on page 122.

- **NICs**

For more information, see [NICs Tab](#) on page 140.

- **Alerts**

For more information, see [Alerts Tab](#) on page 125.

- **Events**

For more information, see [Events Tab](#) on page 129.

- **Metrics**

For more information, see [Metrics Tab](#) on page 129.

- **Disks**

For more information, see [Disks Tab](#) on page 135.

- **Categories**

For more information, see [Categories Tab](#) on page 138.

- **Hardware**

For more information, see [Hardware Tab](#) on page 138.

- **Apps & Relationships**

For more information, see [Apps & Relationships Tab](#) on page 145.

Summary Tab

The **Summary** tab of an individual VM provides the following widgets:

- **Details** : Displays summary information about the VM. For information about the fields available in **Details** widget, see [Details Widget- Parameter Details](#) on page 116.
- **Policies and Management**: Displays the categories and any associated affinity policies for the VM. For more information, see [Policies and Management Widget - Parameter Details](#) on page 117. You can click **View Mapping** to view the categories and any associated affinity policies in the **Categories** tab. For more information, see [Categories Tab](#) on page 138.
- **Performance**: Displays usage metrics for the selected VM. For information about the fields available in **Performance** widget, see [Performance Widget - Parameter Details](#) on page 117. Displays graphs showing CPU usage, memory usage, I/O bandwidth, I/O latency, and IOPS. To analyse the graphs, click **View All Metrics**. The graphs are displayed in the **Metrics** tab. For more information, see [Metrics Tab](#) on page 129.
- **Compute**: Displays the virtual CPU and memory allocated to the VM. For more information, see [Compute Widget- Parameter Details](#) on page 117.

- **Storage:** Displays the storage properties as defined in the storage policy associated with the VM. The storage properties involve the actual values for Replication Factor, Encryption, and Compression that is applied to the VM. For information, see [Storage Widget - Parameter Details](#) on page 118.
- **Infrastructure:** Displays the details of the registered Prism Central . For more information, see [Infrastructure Widget - Parameter Details](#) on page 119.
- **External Devices:** Displays the attached external devices such as GPUs and vGPUs. For more information, see [External Devices Widget- Parameter Details](#) on page 119.
- **Guest Tools:** Displays information about the Nutanix Guest Tools (NGT) if NGT is installed on the VM. For more information, see [Guest Tools Widget- Parameter Details](#) on page 119.
- **Alerts:** Displays a list of related alerts that occurred during the specified interval. Select either **Last week** (default), **Last 1 hour**, or **Last 24 hours** from the dropdown menu on the top right corner of the widget.
- **Anomalies:** Displays a graph of memory, I/O, CPU, network, or disk anomalies that occurred during the specified interval. Select either **Last week** (default), **Last 1 hour**, or **Last 24 hours** from the dropdown menu on the top right corner of the widget. When an anomaly appears, you can click the graph to display a list of those anomalies. If you click an individual anomaly, the system displays the event page for that anomaly.

Details Widget- Parameter Details

The following table describes the fields available in the **Details** widget:

Table 21: Details Widget - Parameter Details

Parameter	Description
Power State	Displays whether the VM is powered on or powered off
Description	Displays the description specified for the VM
Unique ID	Displays the UUID used to identify the VM throughout its lifecycle.
IP Addresses	Displays one or more IP addresses configured in the VM.
Boot Device	Displays the device which attempts to be booted first, overriding the boot configuration.
Boot Configuration	Displays whether the VM will be booted using BIOS booting or UEFI.
Guest OS	Displays the type of operating system installed in the guest, returned by Nutanix Guest Tools.
Virtual Trusted Platform Module (vTPM)	Displays whether a vTPM device is presented to the VM.
Windows Defender Credential Guard	Displays if Windows Defender Credential Guard is enabled.
Secure Boot	Displays whether the Guest OS is required to support secure boot.
Owner	Displays the VM owner name.
Project	Displays the name of the project to which the VM is mapped.
Agent VM	Displays if the VM is an agent VM.

Policies and Management Widget - Parameter Details

The following table describes the fields available in the **Policies and Management** widget:

Note: A dash (-) in a field indicates that there is not enough data to evaluate or a value is not assigned.

Table 22: Policies and Management Widget Fields

Parameter	Description
Categories	Displays the categories associated with the VM.
Policies	Displays the policies like VM-Host affinity and VM-VM anti affinities that are associated with the VM.

Performance Widget - Parameter Details

The following table describes the fields available in the **Performance** widget:

Table 23: Performance Widget- Parameter Details

Parameter	Description
CPU Usage	Displays the percentage of allocated CPU currently being used by this VM.
Storage IOPS	Displays the current I/O operations per second (IOPS).
Memory Usage	Displays the percentage of allocated memory capacity currently being used by this VM.
Storage IO Bandwidth	Displays I/O bandwidth used per second for this VM.
Storage Avg. IO Latency	Displays the average I/O latency for this VM.
Network Rx Bytes	Displays the number of bytes received by the VMs network interfaces.
Network Tx Bytes	Displays number of packets sent by the VMs network interfaces.

Note: VirtIO must be installed in a VM for AHV to display correct VM memory statistics. For more information about VirtIO drivers, see [Nutanix VirtIO for Windows](#) in *AHV Administration Guide*.

Compute Widget- Parameter Details

The following table describes the fields available in the **Compute** widget:

Table 24: Table

Parameter	Description
vCPU	Displays the number of virtual CPUs assigned to this VM.
Memory	Displays the amount of memory available to this VM.
Memory Overcommit	Displays whether Memory Overcommit is enabled or disabled. For more information on Memory Overcommit feature, see Memory Overcommit
Advanced processor compatibility	Displays whether Advanced processor compatibility is enabled or disabled.
Advanced processor CPU generation	Displays the CPU generation.

Storage Widget - Parameter Details

The following table describes the fields available in the **Storage** widget:

Note: A dash (-) in a field indicates that there is not enough data to evaluate or a value is not assigned. For more information about the Storage properties, see [Storage Policy Details View](#) on page 553.

Table 25: Storage Widget - Parameter Details

Parameter	Description
Storage Properties Defined via	Displays the name of the storage policy that the VM is associated with.
Containers	Displays the self storage containers associated with the VM and the link to view the details of the self storage container.
Total Virtual Disks	Displays the number of virtual disks associated with the VM.
Total Storage (Logical)	Displays the total disk capacity available to this VM.
Encryption	Displays the status of encryption that the storage policy applies to the VM. The values can be Enabled or Inherit from Cluster.
Replication Factor	Displays the Replication Factor setting for the VM in the storage policy applied to the VM. The values can be Inherit from Container, 2, or 3.
Compression	Displays the status and type of compression that the storage policy applies to the VM. The values can be Inline, Post Process, Off or Inherit from Container.
QOS Metric (Throttled IOPS)	Displays the type of QOS metric that the storage policy applies to the VM with the throttled throughput value in IOPS.

Parameter	Description
QOS Metric (Throttled Throughput)	Displays the type of QOS metric that the storage policy applies to the VM with the throttled throughput value in MBps.

Infrastructure Widget - Parameter Details

The following table describes the fields available in the **Infrastructure** widget:

Table 26: Infrastructure Widget - Parameter Details

Parameter	Description
Domain	Displays the registered instance of Prism Central.
Cluster	Displays the name of the cluster in which the VM resides.
Host	Displays the host name. This field may be blank if the VM is powered off and a host is not assigned.
Host IP	Displays the host IP address.

External Devices Widget- Parameter Details

The following table describes the fields available in the **External Devices** widget.

Table 27: External Devices Widget- Parameter Details

Parameter	Description
Type	Displays the type as GPU.
Mode	Displays the operational mode of the GPU or vGPU.
GPU Configuration	Displays the configuration of the GPU or vGPU .
GPU Slice	Displays the vGPU slice allotted for the VM. This field is only displayed for vGPU.

Guest Tools Widget- Parameter Details

The following table describes the fields available in the **Guest Tools** widget.

Table 28: Guest Tools Widget- Parameter Details

Parameter	Description
NGT Status	Displays whether the installed NGT version is the latest or an upgrade is available.
NGT Version	Displays the installed NGT version.

Parameter	Description
NGT Applications Enabled	Displays the NGT applications that are installed such as Self Service Restore (SSR) and Volume Snapshot Service (VSS).

Viewing the list of VM instances

View the list of VM instances.

About this task

To view the list of VM instances, follow these steps:

Note: You can perform administrative actions on VMs in Acropolis-managed clusters only.

Procedure

1. Log in to the Prism Central web console.
2. From [Application Switcher Function](#) on page 40, select the **Infrastructure** application.
3. From the **Navigation Bar**, select **Compute > VMs**.
For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61. The system displays the **List** tab by default with all the VMs across registered clusters in **Nutanix** environment. For more information on how to access the list of VMs in **Non-Nutanix** environment (non-nutanix VMs managed by external vCenter), see [VMs Summary View](#) on page 100.
4. Click the target <VM_Name>.
Replace <VM_Name> with the actual VM name at your site.
The **Summary** tab of the selected VM is displayed.
5. View the list of VM instances in the selected cluster.
 - » Click the **VM** dropdown menu.
 - » Search the VM instance name residing on the target cluster.
6. View the list of VM instances on the selected host.
 - » Click the **VM** dropdown menu
 - » Search the VM instance name belonging to the target host.

Console Tab

The **Console** tab accesses and manages VM consoles.

You can perform the following actions to manage and troubleshoot the VM:

- Powering on the VM console. For more information, see [Powering On the VM Console](#) on page 121.
- Capture a screenshot of the VM console. For more information, see [Capturing a Screenshot of a VM Console](#) on page 121.
- Power off or reboot the VM console. For more information, see [Powering Off or Rebooting the VM Console](#) on page 122.

Powering On the VM Console

Power on VM consoles.

About this task

To power on the VM console, follow these steps:

Procedure

1. Log in to the Prism Central web console.
2. From [Application Switcher Function](#) on page 40, select the **Infrastructure** application.
3. From the **Navigation Bar**, select **Compute > VMs**.
For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61. The system displays the **List** tab by default with all the VMs across registered clusters in **Nutanix** environment. For more information on how to access the list of VMs in **Non-Nutanix** environment (non-Nutanix VMs managed by external vCenter), see [VMs Summary View](#) on page 100.
4. Click the target <VM_Name>.
Replace <VM_Name> with the actual VM name at your site.
The **Summary** tab of the selected VM is displayed.
5. Click the **Console** tab.
6. Click **Power On**.
The VM console screen is displayed.
7. (Optional) To open the console in a new window, click the [Launch icon](#).

Capturing a Screenshot of a VM Console

Capture a screenshot of a VM console.

About this task

To capture a screenshot, follow these steps:

Procedure

1. Log in to the Prism Central web console.
2. From [Application Switcher Function](#) on page 40, select the **Infrastructure** application.
3. From the **Navigation Bar**, select **Compute > VMs**.
For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61. The system displays the **List** tab by default with all the VMs across registered clusters in **Nutanix** environment. For more information on how to access the list of VMs in **Non-Nutanix** environment (non-Nutanix VMs managed by external vCenter), see [VMs Summary View](#) on page 100.
4. Click the target <VM_Name>.
Replace <VM_Name> with the actual VM name at your site.
The **Summary** tab of the selected VM is displayed.
5. Click the **Console** tab.

6. Click **Power On**.

The VM console screen is displayed.

7. Click the [Screenshot-capture icon](#).

Powering Off or Rebooting the VM Console

Power off the VM console.

About this task

To power off the VM console, follow these steps:

Procedure

1. Log in to the Prism Central web console.

2. From [Application Switcher Function](#) on page 40, select the **Infrastructure** application.

3. From the **Navigation Bar**, select **Compute > VMs**.

For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61. The system displays the **List** tab by default with all the VMs across registered clusters in **Nutanix** environment. For more information on how to access the list of VMs in **Non-Nutanix** environment (non-Nutanix VMs managed by external vCenter), see [VMs Summary View](#) on page 100.

4. Click the target <VM_Name>.

Replace <VM_Name> with the actual VM name at your site.

The **Summary** tab of the selected VM is displayed.

5. Click the **Console** tab.

6. Click **Power On**.

The VM console screen is displayed.

7. (Optional) To send a **Control-Alt-Delete** command to the console, click the [Reboot icon](#).

8. Click **Power Off**.

9. Select any of the following:

- **Power Off**
- **Power Cycle**
- **Reset**
- **Guest Shutdown**
- **Guest Reboot**

10. Click **Submit**.

Recovery Points Tab

The **Recovery Points** tab displays a list of recovery points (backup snapshots) when recovery points are created for the selected VM.

For information about how to create a recovery point, see [Managing a VM through Prism Central \(AHV\)](#) on page 161.

The protection policy associated with the VM also creates the VM recovery points for Disaster Recovery. For more information about VM recovery points, see [Nutanix Disaster Recovery Guide](#).

You can perform the following actions:

- View recovery points. For more information, see [Viewing Recovery Points](#) on page 123.
- Clone a recovery point. For more information, see [Cloning Recovery Points](#) on page 123.
- Replicate a recovery point. For more information, see [Replicating Recovery Points](#) on page 124.
- Restore a recovery point. For more information, see [Restoring Recovery Points](#) on page 124.
- Delete a recovery point. For more information, see [Deleting Recovery Points](#) on page 125.

Viewing Recovery Points

View recovery points for a VM.

About this task

To view recovery points, follow these steps:

Procedure

1. Log in to the Prism Central web console.
2. From [Application Switcher Function](#) on page 40, select the **Infrastructure** application.
3. From the **Navigation Bar**, select **Compute > VMs**.
For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61. The system displays the **List** tab by default with all the VMs across registered clusters in **Nutanix** environment. For more information on how to access the list of VMs in **Non-Nutanix** environment (non-Nutanix VMs managed by external vCenter), see [VMs Summary View](#) on page 100.
4. Click the target <VM_Name>.
Replace <VM_Name> with the actual VM name at your site.
The **Summary** tab of the selected VM is displayed.
5. Click the **Recovery Points** tab.
The recovery points are displayed.

The list is blank if no recovery points available. The total number of recovery points and the latest and oldest recovery points are listed on the left. A list of all recovery points appears in a table on the right with the create time, location, expiry time, and recovery point type provided for each recovery point.

Cloning Recovery Points

Clone recovery points for a VM.

About this task

To clone recovery points, follow these steps:

Procedure

1. Log in to the Prism Central web console.
2. From [Application Switcher Function](#) on page 40, select the **Infrastructure** application.

3. From the **Navigation Bar**, select **Compute > VMs**.

For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61. The system displays the **List** tab by default with all the VMs across registered clusters in **Nutanix** environment. For more information on how to access the list of VMs in **Non-Nutanix** environment (non-Nutanix VMs managed by external vCenter), see [VMs Summary View](#) on page 100.

4. Click the target <VM_Name>.

Replace <VM_Name> with the actual VM name at your site.

The **Summary** tab of the selected VM is displayed.

5. Click the **Recovery Points** tab.

6. Select a recovery point.

7. Click **Actions > Clone**.

Replicating Recovery Points

Replicating recovery points for a VM to send a copy of the VM to another Nutanix cluster from the selected recovery point,

About this task

To replicate recovery points, follow these steps:

Procedure

1. Log in to the Prism Central web console.

2. From [Application Switcher Function](#) on page 40, select the **Infrastructure** application.

3. From the **Navigation Bar**, select **Compute > VMs**.

For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61. The system displays the **List** tab by default with all the VMs across registered clusters in **Nutanix** environment. For more information on how to access the list of VMs in **Non-Nutanix** environment (non-Nutanix VMs managed by external vCenter), see [VMs Summary View](#) on page 100.

4. Click the target <VM_Name>.

Replace <VM_Name> with the actual VM name at your site.

The **Summary** tab of the selected VM is displayed.

5. Click the **Recovery Points** tab.

6. Select a recovery point.

7. Click **Actions > Replicate**.

Restoring Recovery Points

Restore recovery points for a VM to bring the VM back to the state it was at selected recovery point.

About this task

To restore recovery points, follow these steps:

Procedure

1. Log in to the Prism Central web console.
2. From [Application Switcher Function](#) on page 40, select the **Infrastructure** application.
3. From the **Navigation Bar**, select **Compute > VMs**.
For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61. The system displays the **List** tab by default with all the VMs across registered clusters in **Nutanix** environment. For more information on how to access the list of VMs in **Non-Nutanix** environment (non-Nutanix VMs managed by external vCenter), see [VMs Summary View](#) on page 100.
4. Click the target <VM_Name>.
Replace <VM_Name> with the actual VM name at your site.
The **Summary** tab of the selected VM is displayed.
5. Click the **Recovery Points** tab.
6. Select a recovery point.
7. Click **Actions > Restore**.

Deleting Recovery Points

Deleting recovery points for a VM.

About this task

To delete recovery points, follow these steps:

Procedure

1. Log in to the Prism Central web console.
2. From [Application Switcher Function](#) on page 40, select the **Infrastructure** application.
3. From the **Navigation Bar**, select **Compute > VMs**.
For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61. The system displays the **List** tab by default with all the VMs across registered clusters in **Nutanix** environment. For more information on how to access the list of VMs in **Non-Nutanix** environment (non-nutanix VMs managed by external vCenter), see [VMs Summary View](#) on page 100.
4. Click the target <VM_Name>.
Replace <VM_Name> with the actual VM name at your site.
The **Summary** tab of the selected VM is displayed.
5. Click the **Recovery Points** tab.
6. Select the recovery point.
7. Click **Actions > Delete**.

Alerts Tab

Alerts tab is similar to the **Alerts** dashboard; however, it is filtered to display the alerts only for the selected VM or an individual VM

You can perform the following actions:

- View alerts. For more information, see [Viewing Alerts in the Alerts Tab](#) on page 126.
- Resolve an alert. For more information, see [Resolving an Alert in the Alerts Tab](#) on page 127.
- Acknowledging an alert. For more information, see [Acknowledging an Alert in the Alerts Tab](#) on page 126.
- Configuring Alert emails. For more information, see [Configuring Alert Emails](#) on page 128.
- Delete an alert. For more information, see [Deleting an Alert in the Alerts Tab](#) on page 128.

Viewing Alerts in the Alerts Tab

View alerts in the **Alerts** tab.

About this task

To view alerts, follow these steps:

Procedure

1. Log in to the Prism Central web console.
2. From [Application Switcher Function](#) on page 40, select the **Infrastructure** application.
3. From the **Navigation Bar**, select **Compute > VMs**.
For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61. The system displays the **List** tab by default with all the VMs across registered clusters in **Nutanix** environment. For more information on how to access the list of VMs in **Non-Nutanix** environment (non-Nutanix VMs managed by external vCenter), see [VMs Summary View](#) on page 100.
4. Click the target <VM_Name>.
Replace <VM_Name> with the actual VM name at your site.
The **Summary** tab of the selected VM is displayed.
5. Click the **Alerts** tab.
The **Alerts** tab displays the alerts for the selected VM. For more information about alerts, see [Prism Central Alerts and Events Reference Guide](#).
6. (Optional) To filter the alerts, click **Filters** and update the following fields:
 - **Severity:** Select a severity level.
 - **Resolved:** Specify whether the alert was resolved and how the alert was resolved.
 - **Acknowledged:** Specify whether the alert was acknowledged.
 - **Impact Type:** Select the impact of the issue on the VM.
 - **Created Time:** Specify the interval when the alert first occurred.
 - **Last Occurred:** Specify the interval when the alert last occurred.
 - **Title:** Enter the title of the alert.
 - **Source Entity:** Select the filter criteria and enter the name of the VM.

Acknowledging an Alert in the Alerts Tab

Acknowledge alerts in the **Alerts** tab.

About this task

To acknowledge an alert, follow these steps:

Procedure

1. Log in to the Prism Central web console.
2. From [Application Switcher Function](#) on page 40, select the **Infrastructure** application.
3. From the **Navigation Bar**, select **Compute > VMs**.
For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61. The system displays the **List** tab by default with all the VMs across registered clusters in **Nutanix** environment. For more information on how to access the list of VMs in **Non-Nutanix** environment (non-Nutanix VMs managed by external vCenter), see [VMs Summary View](#) on page 100.
4. Click the target <VM_Name>.
Replace <VM_Name> with the actual VM name at your site.
The **Summary** tab of the selected VM is displayed.
5. Click the **Alerts** tab.
The **Alerts** tab displays the alerts for the selected VM. For more information about alerts, see [Prism Central Alerts and Events Reference Guide](#).
6. Select an alert.
7. Click **Acknowledge**.

Resolving an Alert in the Alerts Tab

Resolve alerts in the **Alerts** tab.

About this task

To resolve an alert, follow these steps:

Procedure

1. Log in to the Prism Central web console.
2. From [Application Switcher Function](#) on page 40, select the **Infrastructure** application.
3. From the **Navigation Bar**, select **Compute > VMs**.
For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61. The system displays the **List** tab by default with all the VMs across registered clusters in **Nutanix** environment. For more information on how to access the list of VMs in **Non-Nutanix** environment (non-nutanix VMs managed by external vCenter), see [VMs Summary View](#) on page 100.
4. Click the target <VM_Name>.
Replace <VM_Name> with the actual VM name at your site.
The **Summary** tab of the selected VM is displayed.
5. Click the **Alerts** tab.
The **Alerts** tab displays the alerts for the selected VM. For more information about alerts, see [Prism Central Alerts and Events Reference Guide](#).
6. Select an alert.

7. Click **Resolve**.

Configuring Alert Emails

Configure the alert messages that are sent from Prism Central in the **Alerts** tab.

About this task

To configure the alert messages that are sent from Prism Central, follow these steps:

Procedure

1. Log in to the Prism Central web console.
2. From [Application Switcher Function](#) on page 40, select the **Infrastructure** application.
3. From the **Navigation Bar**, select **Compute > VMs**.
For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61. The system displays the **List** tab by default with all the VMs across registered clusters in **Nutanix** environment. For more information on how to access the list of VMs in **Non-Nutanix** environment (non-Nutanix VMs managed by external vCenter), see [VMs Summary View](#) on page 100.
4. Click the target <VM_Name>.
Replace <VM_Name> with the actual VM name at your site.
The **Summary** tab of the selected VM is displayed.
5. Click the **Alerts** tab.
The **Alerts** tab displays the alerts for the selected VM. For more information on alerts, see [Prism Central Alerts and Events Reference Guide](#).
6. Click **Email Configuration**.
You can configure alert settings, reporting rules, and message templates. For more information on configuring alerts, see [Configuring Alert Emails in Prism Central](#).

Deleting an Alert in the Alerts Tab

Delete alerts in the **Alerts** tab.

About this task

To delete an alert, follow these steps:

Procedure

1. Log in to the Prism Central web console.
2. From [Application Switcher Function](#) on page 40, select the **Infrastructure** application.
3. From the **Navigation Bar**, select **Compute > VMs**.
For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61. The system displays the **List** tab by default with all the VMs across registered clusters in **Nutanix** environment. For more information on how to access the list of VMs in **Non-Nutanix** environment (non-nutanix VMs managed by external vCenter), see [VMs Summary View](#) on page 100.

4. Click the target <VM_Name>.

Replace <VM_Name> with the actual VM name at your site.

The **Summary** tab of the selected VM is displayed.

5. Click the **Alerts** tab.

The **Alerts** tab displays the alerts for the selected VM. For more information about alerts, see [Prism Central Alerts and Events Reference Guide](#).

6. Select an alert.

7. Click **Delete**.

Events Tab

The **Events** tab displays a table of events.

This tab provides the same features and options as the Events dashboard, however it is filtered to display the events only for the selected VM (individual VM). For more information about events, see [Prism Central Alerts and Events Reference Guide](#).

Metrics Tab

The **Metrics** tab allows you to view usage metrics for the selected VM.

You can perform the following actions:

- View metrics. For more information, see [Viewing Metrics of a VM](#) on page 129.
- Create an alert policy for a metric. For more information, see [Setting Alerts for a Metric](#) on page 134.
- Add the selected metric to **Analysis** dashboard. For more information, see [Adding a Metric to the Analysis Dashboard](#) on page 134.

Viewing Metrics of a VM

View Metrics of a VM

About this task

To view the metrics for a VM, follow these steps:

Procedure

1. Log in to the Prism Central web console.
2. From [Application Switcher Function](#) on page 40, select the **Infrastructure** application.
3. From the **Navigation Bar**, select **Compute > VMs**.


For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61. The system displays the **List** tab with all the VMs across registered clusters in **Nutanix** environment. For more information on how to access the list of VMs in **Non-Nutanix** environment (non-Nutanix VMs managed by external vCenter), see [VMs Summary View](#) on page 100.

4. Click the target <VM_Name>.

Replace <VM_Name> with the actual VM name at your site.

The **Summary** tab of the selected VM is displayed.

5. Click the **Metrics** tab.



The graph is a rolling time interval performance or usage monitor. The baseline range appears as a blue band in the graph.

Note: The baseline range and identified anomalies are based on sophisticated machine-learning capabilities. For more information, see [Behavioral Learning Tools](#) in *Intelligent Operations Guide*. The machine-learning algorithm uses 21 days of data to monitor and predict performance. A graph or baseline band may not appear if less than 21 days of data is available.

For information about the widgets in the **Metrics** tab, see [Metrics Tab Widgets](#) on page 132.

6. Click the **Filters** option to select one or more appropriate metric checkboxes to display the selected metrics in the **Metrics** page. For more information, see [Metrics Tab](#) on page 129.
7. (Optional) Select the duration (time interval) from the dropdown menu on the right (last hour, last 24 hours, or last week).
8. (Optional) Place the cursor anywhere on the horizontal axis to display the value at that time.

Metrics Tab Widgets

The following table describes the metrics available in **Metrics** tab:

Note: Metrics are hypervisor-dependant, and might not be available on all hypervisors.

Table 29: Metrics Tab Widgets

Metric	Description
CPU Usage (%)	Displays the percentage of CPU capacity currently being used by the VM (0–100%).
CPU Ready Time	Displays the current, high, and low percentage of CPU wait time (0–100%).
Memory Usage (%)	Displays the percentage of memory capacity currently being used by the VM (0–100%).
Storage Controller IOPS	Input/Output operations per second from the Storage Controller. For more information, see https://portal.nutanix.com/page/documents/details?targetId=Intelligent-Operations-Guide-vpc_2024_3_1:wc-analysis-metrics-r.html .
Storage Controller Read IOPS	Input/Output read operations per second from the Storage Controller For more information, see https://portal.nutanix.com/page/documents/details?targetId=Intelligent-Operations-Guide-vpc_2024_3_1:wc-analysis-metrics-r.html .
Storage Controller Write IOPS	Input/Output write operations per second from the Storage Controller For more information, see https://portal.nutanix.com/page/documents/details?targetId=Intelligent-Operations-Guide-vpc_2024_3_1:wc-analysis-metrics-r.html .

Metric	Description
Storage Controller Latency	I/O latency in milliseconds from the Storage Controller. For more information, see https://portal.nutanix.com/page/documents/details?targetId=Intelligent-Operations-Guide-vpc_2024_3_1:wc-analysis-metrics-r.html .
Storage Controller Read Latency	Storage Controller read latency in milliseconds. For more information, see https://portal.nutanix.com/page/documents/details?targetId=Intelligent-Operations-Guide-vpc_2024_3_1:wc-analysis-metrics-r.html .
Storage Controller Write Latency	Storage Controller write latency in milliseconds. For more information, see https://portal.nutanix.com/page/documents/details?targetId=Intelligent-Operations-Guide-vpc_2024_3_1:wc-analysis-metrics-r.html .
Storage Controller I/O Bandwidth	Displays the graph for total I/O bandwidth used per second (MBps or KBps) by storage controller.
Storage Controller Read Bandwidth	Read data transferred in KB/second from the Storage Controller.
Storage Controller Write Bandwidth	Write data transferred in KB/second from the Storage Controller.
Disk Usage Bytes	Displays the amount of used disk space.
Snapshot Usage	Displays the amount of storage used for snapshots.
Shared Data	Displays the shared data usage by the VM.
I/O Working Set Size	Displays graph for total bandwidth used per second (MBps or KBps) for physical disk requests by the VM.
Write I/O Working Set Size	Displays graph for write (only) I/O bandwidth used per second (MBps or KBps) for physical disk requests by the VM.
Read I/O Working Set Size	Displays graph for read (only) I/O bandwidth used per second (MBps or KBps) for physical disk requests by the VM.
Disk Usage (%)	Displays the disk usage percentage for the VM.
IOPS related metrics	Displays separate graphs for total, write, and read I/O operations per second (IOPS) for the VM.
IO Latency related metrics	Displays separate graphs for total, write, and read average I/O latency (in milliseconds) for physical disk requests by the VM.
Usage related metrics	Displays separate graphs for disk, snapshot, and shared data usage (in GiBs) by the VM.
Working Set Size related metrics	Displays separate graphs for total, write, and read storage usage (in GiBs) for the VM working set size.
Network Receive Packets Dropped	Displays graph for the number of received packets dropped.

Metric	Description
Network Transmit Packets Dropped	Displays graph for the number of transmitted packets dropped.
Network Bytes related metrics	Displays separate graphs for the amount of transmitted and received bytes (in GiBs).

Setting Alerts for a Metric

Create an alert policy for a metric.

About this task

To create an alert policy for a metric, follow these steps:

Procedure

1. Log in to the Prism Central web console.
2. From [Application Switcher Function](#) on page 40, select the **Infrastructure** application.
3. From the **Navigation Bar**, select **Compute > VMs**.
For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61. The system displays the **List** tab with all the VMs across registered clusters in **Nutanix** environment. For more information on how to access the list of VMs in **Non-Nutanix** environment (non-nutanix VMs managed by external vCenter), see [VMs Summary View](#) on page 100.
4. Click the target <VM_Name>.
Replace <VM_Name> with the actual VM name at your site.
The **Summary** tab of the selected VM is displayed.
5. Click the **Metrics** tab.
6. Select **Actions > Alert Settings**.
For more information about how to create alert policies, see [Prism Central Alerts and Events Reference Guide](#).

Adding a Metric to the Analysis Dashboard

Add a metric to the **Analysis** dashboard.

About this task

To add a metric to the **Analysis** dashboard, follow these steps:

Procedure

1. Log in to the Prism Central web console.
2. From [Application Switcher Function](#) on page 40, select the **Infrastructure** application.
3. From the **Navigation Bar**, select **Compute > VMs**.
For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61. The system displays the **List** tab with all the VMs across registered clusters in **Nutanix** environment. For more information

on how to access the list of VMs in **Non-Nutanix** environment (non-Nutanix VMs managed by external vCenter), see [VMs Summary View](#) on page 100.

4. Click the target <VM_Name>.

Replace <VM_Name> with the actual VM name at your site.

The **Summary** tab of the selected VM is displayed.

5. Click the **Metrics** tab.

6. Select **Actions > Add to Analysis**.

The metric is displayed on the Analysis dashboard. For more information, see [Analysis Dashboard](#) in *Intelligent Operation Guide*.

Disks Tab

The **Disks** tab displays the virtual disks attached to a VM.

You can manage a disk in the following ways:

- View virtual disks. For more information, see [Viewing Virtual Disks](#) on page 135.
- Attach a virtual disk. For more information, see [Attaching Virtual Disks](#) on page 136.
- Update a virtual disk. For more information, see [Updating Virtual Disks](#) on page 136.
- Delete a virtual disk. For more information, see [Deleting Virtual Disks](#) on page 137.

Viewing Virtual Disks

View virtual disks in the **Disks** tab.

About this task

To view virtual disks, follow these steps:

Procedure

1. Log in to the Prism Central web console.

2. From [Application Switcher Function](#) on page 40, select the **Infrastructure** application.

3. From the **Navigation Bar**, select **Compute > VMs**.

For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61. The system displays the **List** tab with all the VMs across registered clusters in **Nutanix** environment. For more information on how to access the list of VMs in **Non-Nutanix** environment (non-Nutanix VMs managed by external vCenter), see [VMs Summary View](#) on page 100.

4. Click the target <VM_Name>.

Replace <VM_Name> with the actual VM name at your site.

The **Summary** tab of the selected VM is displayed.

5. Click the **Disks** tab.

The **Disks** tab displays the virtual disks attached to a VM. The **Disks** tab displays the following fields:

- **Disk Address:** Displays the disk address such as ide.0 or scsi.1.
- **Capacity:** Displays the disk capacity in MiB or GiB.
- **Type:** Displays the disk type.

6. (Optional) To filter the disks, click **Filters** and update the following fields: Disk Address, Capacity, or Type.

Attaching Virtual Disks

Attaching virtual disks in the **Disks** tab.

About this task

To attach virtual disks, follow these steps:

Procedure

1. Log in to the Prism Central web console.
2. From [Application Switcher Function](#) on page 40, select the **Infrastructure** application.
3. From the **Navigation Bar**, select **Compute > VMs**.
For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61. The system displays the **List** tab with all the VMs across registered clusters in **Nutanix** environment. For more information on how to access the list of VMs in **Non-Nutanix** environment (non-nutanix VMs managed by external vCenter), see [VMs Summary View](#) on page 100.
4. Click the target <VM_Name>.
Replace <VM_Name> with the actual VM name at your site.
The **Summary** tab of the selected VM is displayed.
5. Click the **Disks** tab.
6. Click **Attach Disk**.
7. In the **Type** drop-down, select **Disk** or **CD-ROM**.
8. In the **Operation** drop-down, select how to create the disk.
9. (Optional) Select **Allocate on Storage Container**.
 - a. Select a storage container.
 - b. Enter the size of the disk in GiB.
10. (Optional) Select **Clone from Image**.
 - a. Select an image.
 - b. Enter the size of the disk in GiB.
11. (Optional) Select **Clone from another VM Disk**.
 - a. Select the VM.
 - b. Select the disk.
12. Select the bus type.
13. Click **Save**.

Updating Virtual Disks

Update virtual disks in the **Disks** tab.

About this task

To update virtual disks, follow these steps:

Procedure

1. Log in to the Prism Central web console.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40.
3. Select **Compute > VMs** from the **Navigation Bar**.
For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61. The system displays the **List** tab with all the VMs across registered clusters in **Nutanix** environment. For more information on how to access the list of VMs in **Non-Nutanix** environment (non-nutanix VMs managed by external vCenter), see [VMs Summary View](#) on page 100.
4. Click the target <VM_Name>.
Replace <VM_Name> with the actual VM name at your site.
The **Summary** tab of the selected VM is displayed.
5. Click the **Disks** tab.
6. Select a disk.
7. Click **Update**.
8. From the image dropdown menu, select the required image.
9. Click **Save**.

Deleting Virtual Disks

Deleting virtual disks in the **Disks** tab.

About this task

To delete virtual disks, follow these steps:

Procedure

1. Log in to the Prism Central web console.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40.
3. Select **Compute > VMs** from the **Navigation Bar**.
For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61. The system displays the **List** tab with all the VMs across registered clusters in **Nutanix** environment. For more information on how to access the list of VMs in **Non-Nutanix** environment (non-nutanix VMs managed by external vCenter), see [VMs Summary View](#) on page 100.
4. Click the target <VM_Name>.
Replace <VM_Name> with the actual VM name at your site.
The **Summary** tab of the selected VM is displayed.
5. Click the **Disks** tab.
6. (Optional) To filter the disks, click **Filters** and update the following fields: Disk Address, Capacity, or Type.

7. Click **Delete**.
8. In the confirmation dialog box, click **Delete**.

Categories Tab

The **Categories** tab displays the categories and policies associated with the VM. Each VM can have a one-to-many relationship with categories and the categories can have a many-to-many relationship with any of the policies.

For more information on categories management and policies, see [Category Management](#) and [Image Placement Policies](#) on page 527.

You can assign categories to a VM. For more information, see [Managing Categories in the Categories Tab](#) on page 138.

Managing Categories in the Categories Tab

Assign categories to a VM.

About this task

To assign categories to a VM, follow these steps:

Procedure

1. Log in to the Prism Central web console.
2. From [Application Switcher Function](#) on page 40, select the **Infrastructure** application.
3. From the **Navigation Bar**, select **Compute > VMs**.
For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61. The system displays the **List** tab with all the VMs across registered clusters in **Nutanix** environment. For more information on how to access the list of VMs in **Non-Nutanix** environment (non-Nutanix VMs managed by external vCenter), see [VMs Summary View](#) on page 100.
4. Click the target <VM_Name>.
Replace <VM_Name> with the actual VM name at your site.
The **Summary** tab of the selected VM is displayed.
5. Click the **Categories** tab.
6. Click **Manage Categories**.
For more information, see [Assigning a Category to an Entity](#).

Hardware Tab

The **Hardware** tab displays the GPUs configured with the VM in your cluster. You can attach GPUs to a new or existing VM.

You can manage a GPU in the following ways:

- View a GPU. For more information, see [Viewing GPUs](#) on page 139.
- Attach a GPU. For more information, see [Attaching a GPU to a VM](#) on page 139.
- Delete a GPU. For more information, see [Deleting a GPU from a VM](#) on page 140.

Viewing GPUs

View the GPUs attached with the VM in your cluster.

About this task

To view GPUs, follow these steps:

Procedure

1. Log in to the Prism Central web console.
2. From [Application Switcher Function](#) on page 40, select the **Infrastructure** application.
3. From the **Navigation Bar**, select **Compute > VMs**.
For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61. The system displays the **List** tab by default with all the VMs across registered clusters in **Nutanix** environment. For more information on how to access the list of VMs in **Non-Nutanix** environment (non-Nutanix VMs managed by external vCenter), see [VMs Summary View](#) on page 100.
4. Click the target <VM_Name>.
Replace <VM_Name> with the actual VM name at your site.
The **Summary** tab of the selected VM is displayed.
5. (Optional) To view either virtual GPUs or passthrough GPUs, click **Filters** and select the required option.
6. (Optional) To view the **Summary** tab of an individual GPU, click the target *GPU_name*. Replace *GPU_name* with the actual GPU name at your site.

Attaching a GPU to a VM

Attach a GPU to a VM.

About this task

To attach a new GPU to an existing VM, follow these steps:

Procedure

1. Log in to the Prism Central web console.
2. From [Application Switcher Function](#) on page 40, select the **Infrastructure** application.
3. From the **Navigation Bar**, select **Compute > VMs**.
For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61. The system displays the **List** tab with all the VMs across registered clusters in **Nutanix** environment. For more information on how to access the list of VMs in **Non-Nutanix** environment (non-Nutanix VMs managed by external vCenter), see [VMs Summary View](#) on page 100.
4. Click the target <VM_Name>.
Replace <VM_Name> with the actual VM name at your site.
The **Summary** tab of the selected VM is displayed.
5. Select the **Hardware** tab.
6. Click **Attach GPU**.
7. Specify **vGPU** or **Passthrough**.

8. Select the GPU.
9. Click **Save**.
The GPUs are displayed in the **Hardware** tab.

Deleting a GPU from a VM

Delete a GPU from a VM.

About this task

To delete a GPU, follow these steps:

Procedure

1. Log in to the Prism Central web console.
2. From [Application Switcher Function](#) on page 40, select the **Infrastructure** application.
3. From the **Navigation Bar**, select **Compute > VMs**.
For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61. The system displays the **List** tab with all the VMs across registered clusters in **Nutanix** environment. For more information on how to access the list of VMs in **Non-Nutanix** environment (non-nutanix VMs managed by external vCenter), see [VMs Summary View](#) on page 100.
4. Click the target <VM_Name>.
Replace <VM_Name> with the actual VM name at your site.
The **Summary** tab of the selected VM is displayed.
5. Select the **Hardware** tab.
6. Select a GPU.
7. Click **Delete**.

NICs Tab

The **NICs** tab displays the virtual NICs (vNICs) added to the VM only after you add NICs to a VM.

You can manage a vNIC as follows:

- View vNICs. For more information, see [Viewing Virtual NICs Attached to a VM](#) on page 141.
- Add any of the following vNICs:
 - Normal NICs. For more information, see [Attaching to a Subnet](#) on page 141.
 - Traffic Mirror Destination NICs. For more information, see [Adding Traffic Mirror Destination vNICs to a VM](#) on page 142.
 - Network function vNICs. For more information, see [Creating Network Function vNICs](#) on page 143.
 - NIC with SR-IOV NIC profile. For more information, see [Adding SR-IOV NICs to a VM](#) on page 143.
- Update a vNIC. For more information, see [Updating Virtual NICs Attached to a VM](#) on page 144.
- Delete a vNIC. For more information, see [Deleting Virtual NICs Attached to a VM](#) on page 145.

Viewing Virtual NICs Attached to a VM

View virtual NICs (vNICs) attached to a VM in Prism Central.

About this task

To view vNICs attached to a VM, follow these steps:

Procedure

1. Log in to the Prism Central web console.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40.
3. Select **Compute > VMs** from the **Navigation Bar**.
For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61. The system displays the **List** tab by default with all the VMs across registered clusters in **Nutanix** environment. For more information on how to access the list of VMs in **Non-Nutanix** environment (non-Nutanix VMs managed by external vCenter), see [VMs Summary View](#) on page 100.
4. Click the target <VM_Name>.
Replace <VM_Name> with the actual VM name at your site.
The **Summary** tab of the selected VM is displayed.
5. Click the **NICs** tab.
The NICs added to the VM are displayed. For more information, see [NICs Fields](#) on page 141.

NICs Fields

The following table describes the fields displayed in the **NICs** tab:

Table 30: NIC Fields

Parameter	Description
VLAN ID	Displays the VLAN ID of the NIC.
MAC Address	Displays the virtual NIC MAC address.
Network Connection State	Displays whether the NIC is connected to the network currently.
VPC	Displays the Virtual Private Cloud (VPC) of the NIC. For more information about VPC, see Virtual Private Clouds on page 369.
Virtual Switch	Displays the virtual switch of the NIC. For more information about how to create a virtual switch, see Creating a Virtual Switch .
IP Address	Displays the virtual NIC IP address.

Attaching to a Subnet

Adding virtual NICs (vNICs) to a VM in Prism Central.

About this task

To add vNICs to a VM, follow these steps:

Procedure

1. Log in to the Prism Central web console.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40.
3. From the **Navigation Bar**, select **Compute > VMs**.
For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61. The system displays the **List** tab by default with all the VMs across registered clusters in **Nutanix** environment. For more information on how to access the list of VMs in **Non-Nutanix** environment (non-Nutanix VMs managed by external vCenter), see [VMs Summary View](#) on page 100.
4. Click the target <VM_Name>.
Replace <VM_Name> with the actual VM name at your site.
The **Summary** tab of the selected VM is displayed.
5. Click the **NICs** tab.
6. Click **Attach to Subnet**.
7. From the **Subnet** dropdown menu, select a subnet .
8. Select the **Attachment Type** as **Trunked** or **Access**.
9. Click **Create**.
The **NICs** tab displays the vNICs of type Normal that you added to the VM.

Adding Traffic Mirror Destination vNICs to a VM

Adding traffic mirror destination virtual NICs (vNICs) to a VM in Prism Central.

About this task

To add traffic mirror destination to a VM, follow these steps:

Procedure

1. Log in to the Prism Central web console.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40.
3. Select **Compute > VMs** from the **Navigation Bar**.
For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61. The system displays the **List** tab by default with all the VMs across registered clusters in **Nutanix** environment. For more information on how to access the list of VMs in **Non-Nutanix** environment (non-Nutanix VMs managed by external vCenter), see [VMs Summary View](#) on page 100.
4. Click the target <VM_Name>.
Replace <VM_Name> with the actual VM name at your site.
The **Summary** tab of the selected VM is displayed.
5. Click the **NICs** tab.
6. Click **More Network Options > Add Traffic Mirroring**.
The **NICs** tab displays the Traffic Mirror Destination NICs that are added to the VM.

Creating Network Function vNICs

Create network function vNICs for the service insertion feature.

Before you begin

- Service Insertion feature requires a minimum of two network function vNICs for each network function service. Therefore, you must create a minimum of two vNICs for every network function service.
- Ensure that your cluster is running the following software:
 - AOS 7.3 or later
 - Flow Network Security Next-Gen 5.2.0 or later

About this task

To create network function vNICs for the Service Insertion feature, follow these steps:

Procedure

1. Log in to the Prism Central web console.
2. From [Application Switcher Function](#) on page 40, select the **Infrastructure** application.
3. From the **Navigation Bar**, select **Compute > VMs**.
For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61. The system displays the **List** tab by default with all the VMs across registered clusters in **Nutanix** environment. For more information on how to access the list of VMs in **Non-Nutanix** environment (non-Nutanix VMs managed by external vCenter), see [VMs Summary View](#) on page 100.
4. Click the target <VM_Name>.
Replace <VM_Name> with the actual VM name at your site.
The **Summary** tab of the selected VM is displayed.
5. Click the **NICs** tab.
6. Select **More Network Options > Add Network Function**.
7. Create a minimum of two vNICs for every network function service. Repeat Step 6.
The **NICs** tab displays the network function vNIC.
8. From the **Navigation Bar**, select **Network & Security > Security Policies**.
The **Security Policies** page appears.
9. Click **Network Functions**.
The **Network Functions** page appears.
10. In the **Create Network Function** page, assign a vNIC as ingress and assign the other vNIC as egress.
For more information, see [Creating a Network Function Entity](#) in the *Flow Network Security Next-Gen Guide*.
For more information about see [Service Insertion](#) in the *Flow Network Security Next-Gen Guide*.
The **NICs** tab displays the vNICs for the Service Insertion feature that are added to the VM.

Adding SR-IOV NICs to a VM

Add SR-IOV virtual NICs (vNICs) to a VM in Prism Central.

About this task

To add SR-IOV vNICs to a VM, follow these steps:

Procedure

1. Log in to the Prism Central web console.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40.
3. Select **Compute > VMs** from the **Navigation Bar**.
For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61. The system displays the **List** tab by default with all the VMs across registered clusters in **Nutanix** environment. For more information on how to access the list of VMs in **Non-Nutanix** environment (non-Nutanix VMs managed by external vCenter), see [VMs Summary View](#) on page 100.
4. Click the target <VM_Name>.
Replace <VM_Name> with the actual VM name at your site.
The **Summary** tab of the selected VM is displayed.
5. Click the **NICs** tab.
6. Select **More Network Options > Add SR-IOV**.
7. Select a NIC profile.
8. Enter the VLAN ID.
9. Click **Save**.
For more information about SR-IOV vNICs, see [Single Root Input Output Virtualization](#)

Updating Virtual NICs Attached to a VM

Update virtual NICs (vNICs) attached to a VM in Prism Central.

About this task

To update vNICs attached to a VM, follow these steps:

Procedure

1. Log in to the Prism Central web console.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40.
3. Select **Compute > VMs** from the **Navigation Bar**.
For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61. The system displays the **List** tab by default with all the VMs across registered clusters in **Nutanix** environment. For more information on how to access the list of VMs in **Non-Nutanix** environment (non-Nutanix VMs managed by external vCenter), see [VMs Summary View](#) on page 100.
4. Click the target <VM_Name>.
Replace <VM_Name> with the actual VM name at your site.
The **Summary** tab of the selected VM is displayed.
5. Click the **NICs** tab.
6. Select the required vNIC.

7. Click **Actions** > **Update**.

Deleting Virtual NICs Attached to a VM

Delete virtual NICs (vNICs) attached to a VM in Prism Central.

About this task

To delete vNICs attached to a VM, follow these steps:

Procedure

1. Log in to the Prism Central web console.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40.
3. Select **Compute** > **VMs** from the **Navigation Bar**.
For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61. The system displays the **List** tab by default with all the VMs across registered clusters in **Nutanix** environment. For more information on how to access the list of VMs in **Non-Nutanix** environment (non-Nutanix VMs managed by external vCenter), see [VMs Summary View](#) on page 100.
4. Click the target <VM_Name>.
Replace <VM_Name> with the actual VM name at your site.
The **Summary** tab of the selected VM is displayed.
5. Click the **NICs** tab.
6. Select the required vNIC.
7. Click **Actions** > **Delete**.

Apps & Relationships Tab

The **Apps & Relationships** tab displays a dropdown menu with two options: **Discovered Apps** and **App Relationships**.

Note: Application Discovery function needs to be enabled to view the **Discovered Apps** and **App Relationships**. For more information about Application discovery function and how to enable it in Prism Central, see [Application Discovery](#) in *Intelligent Operations Guide*.

The **Discovered Apps** page displays (up to) the top 20 applications that are communicating with other entities (usually client VMs). The discovered applications are sorted based on the number of clients communicating with the application. This menu option provides similar information as the application discovery dashboard, except it is filtered to display the details only for this VM. For more information, see [Application Discovery](#) in *Intelligent Operation Guide*.

The **App Relationships** page displays three tabs: **Visualization**, **Incoming List**, and **Outgoing List**. All the three tabs display (up to) the top 20 results with respect to the visualization, incoming or outgoing client communication.

The **Visualization** tab displays a visual representation of the VM and the applications running on it, along with its incoming and outgoing communication with other entities. Data for last 24 hours is fetched to display the information on the visualization tab.

You can perform the following actions in **Apps & Relationships** tab:

- Hover on an entity to see the communication details with the applications.

- Hover on an application to see the incoming communication of the application with external entities and outgoing communications of the VM on which the application is running.

Note: In case of dual NIC VMs, click and expand the App listing to view the specific IP that is being used by the application.




Figure 37: Visualization Tab

Incoming List: Displays a list of entities having incoming communication with this VM. Each line represents an incoming communication and includes the following fields.

Table 31: Incoming List Fields

Parameter	Description	Values
Source Entity Name	Displays the name of the communicating entity. If the entity is not identified by Prism Central (through registered PE or through vCenter using monitoring configurations feature), the name is displayed as unknown.	(VM name), Unknown
Source IP Address	Displays the IP Address of the communicating entity.	(IP Address)
Destination App	Displays the name of the application that is being communicated to. If the application is not an identified one, the Destination App is displayed as Unknown.	(application name), Unknown
Destination IP Address	Displays the IP Address of the VM.	(IP address)
Destination Port	Displays the port used by the application for communication.	(port number)

Parameter	Description	Values
Destination App Identification Status	Indicates whether the application is an identified application type. An identified application is one that is a known type.	Identified, Unidentified

Click **Modify Filters** option to view the **Filters** pane. You can use the **Filters** pane to select a filter option.

You can apply filters with multiple fields. The filters fetch the latest discovery and communication data and accordingly display the result set, to a maximum of 20 results.

The following table describes the filter options for **Incoming List**.

Table 32: Filters Pane - Incoming List

Field	Description	Values
Source Entity Name	Filters based on the source entity name. Select a condition from the dropdown menu (Contains, Equal to, Not equal to, Doesn't contain, Starts with, Ends with) and enter a string in the field.	(source entity name string)
Entity Type	Filters based on the entity type. A VM belongs to a cluster managed by Prism Central or vCenter (only if using monitoring configurations feature). An Unknown entity is one that is not identified by Prism Central.	VM, Unknown
Source IP Address	Filters based on the source IP address.(same options as Source Entity Name).	(source IP address string)
Destination App	Filters based on the destination app (same options as Source Entity Name).	(Destination App string)
Destination IP Address	Filters based on the destination IP address (same options as Source Entity Name).	(destination IP address string)
TCP Port	Filters based on the TCP port number.	(TCP port number)
UDP Port	Filters based on the UDP port number.	(UDP port number)
Time Range	Filters based on the time range.	Last 30 Mins, Last 1 Hour, Last 12 Hours, and Last 24 Hours

Outgoing List: Displays a list of entities having outgoing communication from this VM. Each line represents an outgoing communication and includes the following fields.

Table 33: Field Description - Outgoing List

Field	Description	Values
Destination Entity Name	Displays the name of the entity this VM is communicating to. If the entity is not identified by Prism Central (through registered PE or through vCenter using monitoring configurations feature), the name is displayed as unknown.	(entity name), Unknown
Destination IP Address	Displays the IP Address of the entity this VM is communicating to.	(IP address)
Destination Port	Displays the port on which this VM is communicating with the destination entity.	(port number)
Source IP Address	Displays the IP Address of the VM.	(IP address)
Source App Identification Status	Indicates whether the communicating application is an identified application type. An identified application is one that is a known type.	Identified, Unidentified

You can filter this list by opening the Filter pane to select a filter option. The following table describes the filter options for **Outgoing List**. You can apply filters with multiple fields. The filters fetch the latest discovery and communication data and accordingly display the result set, to the maximum of 20 results.

The following table describes the filter options for **Outgoing List**:

Table 34: Filter Pane - Outgoing List

Fields	Description	Values
Destination Entity Name	Filters based on the destination entity name. Select a condition from the dropdown menu (Contains, Equal to, Not equal to, Doesn't contain, Starts with, Ends with) and enter a string in the field.	(destination entity name string)
Entity Type	Filters based on the entity type. A VM belongs to a cluster managed by Prism Central or vCenter (only if using monitoring configurations feature). An Unknown entity is one that is not identified by Prism Central.	VM, Unknown
Destination IP Address	Filters based on the destination IP address. (same options as Destination Entity Name).	(destination IP address string)
TCP Port	Filters based on the TCP port number.	(TCP port number)
UDP Port	Filters based on the UDP port number.	(UDP port number)
Source IP Address	Filters based on the source IP address (same options as Destination Entity Name).	(source IP address string)
Time Range	Filters based on the time range.	Last 30 Mins, Last 1 Hour, Last 12 Hours, and Last 24 Hours

Creating a VM through Prism Central (AHV)

You can create virtual machines (VMs) in Acropolis-managed clusters through Prism Central.

About this task

Note:

- For instructions about how to install Nutanix VirtIO on Windows VMs, see [Windows VM Provisioning](#) in the *AHV Administration Guide*.
- If you are logged in as a project manager or a project member, see [Creating a VM \(Self Service\)](#) on page 571.

To create a VM in an AHV cluster, perform the following steps:

Procedure

To create a VM in AHV cluster, perform the following steps:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and go to **Compute > VMs** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.
The system displays the **List** tab by default with all the VMs across registered clusters in **Nutanix** environment. For information about how to access the list of non-nutanix VMs managed by an external *vCenter*, see [VMs Summary View](#) on page 100.

3. Click **Create VM**, and enter the following information in the **Configuration** step:

- a. **Name:** Enter a name for the VM.
- b. **Description** (optional): Enter a description for the VM.
- c. **Project:** Select the project to which you want to associate the VM. This field displays the projects to which your role has access.
- d. **Cluster:** Select the target cluster from the dropdown menu on which you intend to create the VM.
- e. **Number of VMs:** Enter the number of VMs you intend to create. The created VM names are suffixed sequentially.
- f. **vCPU(s):** Enter the number of virtual CPUs to allocate to this VM.

Note: For the *IDE* Bus Type, the maximum limit for vCPU is 240.

- g. **Number of Cores per vCPU:** Enter the number of cores assigned to each virtual CPU.
- h. **Memory:** Enter the amount of memory (in GiBs) to allocate to this VM.
- i. **Enable Memory Overcommit:** Select this checkbox to enable memory overcommit for the VM.
Before you turn memory overcommit on or off, see [Memory Overcommit](#) in the *AHV Administration Guide*. For more information on configuring memory overcommit, see [Memory Overcommit Management](#) on page 212.

Note: Ensure that the status of the **Memory Overcommit** property under **Configuration** in the **Review** tab is set as **Enabled**.

- j. **Enable Advanced Processor Compatibility:** Select this checkbox to enable advanced processor compatibility for the VM.

After you select the **Enable Advanced Processor Compatibility** checkbox, the system provides you the option to select the required CPU generation.

Select the required CPU generation name in the **Select CPU generation** dropdown menu.

For more information, see [Advanced Processor Compatibility in AHV](#) section in the *AHV Administration Guide*.

Note: Advanced processor compatibility setting for VMs with automatic cluster selection configuration is not supported. For more information on automatic cluster selection, see [Automatic Cluster Selection for VM Placement](#) section in the *AHV Administration Guide*.

The following is an example showing the **Configuration** step:

Create VM

1 Configuration

2 Resources

3 Management

4 Review

General

Name

Description (Optional)

Project

Cluster

Number of VMs

VM Properties

CPUsvCPU

Cores Per CPU

Memory

Advanced Settings

☐ Enable memory overcommit

☐ Enable advanced processor compatibility

Cancel

Next

Figure 38: Create VM - Configuration

4. In the **Resources** step, perform the following actions to attach a Disk to the VM:

Disks: Click **Attach Disk**, and enter the following information:

- a. **Type:** Select the type of storage device, **Disk** or **CD-ROM**, from the dropdown menu.
- b. **Operation:** Specify the device contents from the dropdown menu.
 - Select **Empty CD-ROM** to create a blank CD-ROM device. A CD-ROM device is needed when you intend to provide a system image from CD-ROM.

Note: The **Empty CD-ROM** option is available only when *CD-ROM* is selected as the storage device in the **Type** field.

- Select **Allocate on Storage Container** to allocate space without specifying an image. Selecting this option means you are allocating space only. You have to provide a system image later from a CD-ROM or another source.

Note: The **Allocate on Storage Container** option is available only when *Disk* is selected as the storage device in the **Type** field.

- Select **Clone from Image** to copy an image that you have imported by using the image service feature onto the disk.

- c. If you select:

- **Allocate Storage Container** in the **Operation** field, the system prompts you to specify the **Storage Container**.

Note: If the VM is assigned to a category associated with a storage policy, then the storage container that you select might not govern the storage properties of the VM. The storage properties that the storage policy provides override the storage properties that the storage container provides.

- **Clone from Image** in the **Operation** field, the system prompts you to specify the **Image**.

- d. Enter either of the following information based on your selection in the **Operation** field:

- **Storage Container:** Select the appropriate storage container.
- **Image:** Select the image you created using the image service feature. For information about Image management, see [Image Management](#) on page 258.

Note: The image transfer can trigger image bandwidth throttling if a bandwidth throttling policy is associated with the image. For more information, see [Bandwidth Throttling Policies](#) on page 535.

- e. **Bus Type:** Select the bus type from the dropdown menu.

The options displayed in the dropdown menu vary based on the storage device **Type** selected in the **Type** field.

If the storage device **Type** is:

- **Disk:** The available choices are **SCSI**, **SATA**, **PCI**, or **SATA**.
- **CD-ROM:** The available choices are **IDE**, or **SATA**

- f. **Capacity:** Enter the disk size in GiB.

- g. When all the field entries are correct, click **Save** to attach the disk to the VM and return to the **Create VM** window.

Repeat this step to attach additional devices to the VM.

5. In the **Resources** tab, perform the following actions to create a network interface for the VM:

Networks: Click **Attach to Subnet**. The **Attach to Subnet** window appears.

Attach to Subnet

Subnet Attachment

Subnet

VLAN-50

VLAN ID	IPAM	Virtual Switch
50	Managed	br0

Network Connection State

Connected

Private IP Assignment

Network Address / Prefix	Free IPs (in Subnet)	Free IPs (in Pool)
10.0.0.0/24	252	21

Assignment Type: Assign with DHCP IP Address:

NIC Configuration

Attachment Type: Trunked

VLANs for Trunk: 100 200 300

If no VLANs are specified then all VLANs from 1-4094 will be trunked

Cancel Save

Figure 39: Attach to Subnet Window

- Subnet:** Select the target subnet from the dropdown menu.
The list includes all the defined networks. For more information on how to perform the network configuration, see [Network Configuration for VM Interfaces](#) in *Prism Element Web Console Guide*.
- VLAN ID:** This is a read-only field that displays the VLAN ID.
- IPAM:** This is a read-only field that informs you if the subnet is IPAM managed or not.
- Virtual Switch:** This is a read-only field that displays the name of the virtual switch associated with the subnet.
- Network Connection State:** Select the state for the network after VM creation:
 - Connected:* If the VM needs to be connected to the network to operate.
 - Disconnected:* If the VM needs to be in disconnected state after creation.
- Network Address/Prefix:** This is a read-only field that displays the network IP address and prefix.

- g. **Free IPs (in Subnet):** This is a read-only field that displays the free IP addresses in the selected subnet.
- h. **Free IPs (in Pool):** This is a read-only field that displays the free IP addresses in the network IP pool.
- i. **Assignment Type:** Select the IP assignment type from the dropdown menu. The options are *Assign with DHCP*, *Assign Static IP*, and *No Private IP*.

Note: This field is displayed only if the selected **Subnet** is managed by IPAM, in other words, a subnet such as a managed Basic VLAN Subnet, VLAN Subnet or a Virtual Private Cloud (VPC) subnet.

If the selected **Subnet** is an unmanaged **Subnet**, the **Assignment Type** field is not displayed.

Important: The **Floating IP Assignment** fields in the **Create VM > Attach to Subnet** dialog box are disabled, so floating IP addresses cannot be assigned to the VM vNIC during VM creation.

After the VM is created, update the VM vNIC to assign the floating IP address. For the steps to update the vNIC, see [Updating a VM through Prism Central \(AHV\)](#) on page 165.

- j. **IP Address:** Enter the static IP address for VLAN.
This field is enabled only if *Assign Static IP* is selected in the **Assignment Type** field.
- k. **Attachment Type:** Select the NIC configuration attachment type from the dropdown menu. The options are *Access* and *Trunked*.
If you select *Trunked* attachment type, **VLANS for Trunk** field appears where you can specify the VLANs ID for the trunk.
- l. **VLANS for Trunk:** Specify the VLAN ID for the trunk.
If you do not specify the VLAN ID for the trunked attachment type, then all VLANs from 1- 4094 are trunked.

Note: VLANS for Trunk is available only for basic VLANs.

- m. Click **Save** to create a network interface for the VM, and return to the **Create VM** window.
Repeat this step to create additional network interfaces for the VM.

6. (Optional) If you want to add Traffic Mirror Destination NICs, click **Add Mirror Destination NIC**.

Note: NIC MAC Address is assigned after the creation of the VM.

You can add additional Traffic Mirror Destination NICs by clicking **Add New NIC**.

7. In the **Resources** step, perform the following actions for boot configuration.

Boot Configuration: Select one of the following firmware to boot the VM.

- » **Legacy BIOS Mode:** Used to start the VM with legacy BIOS firmware.
- » **UEFI BIOS Mode:** Used to start the VM with UEFI firmware. UEFI firmware supports larger hard drives, faster boot time, and provides more security features. For more information on UEFI firmware, see [UEFI Support for VM](#) information in the *AHV Administration Guide*.

If you select **UEFI BIOS Mode** for **Boot Configuration**, the system enables you to define the **Shield VM Security Settings**:

- **Secure Boot:** Select this check box to enable UEFI secure boot policies for your guest VMs. For more information on Secure Boot functionality, see [Secure Boot for VMs](#) information in *AHV Administration Guide*.

Note: Secure Boot is not enabled if the **Bus Type** is *IDE*.

- **Attach vTPM:** Select this checkbox to enable vTPM. For information about vTPM functionality, see [Securing AHV VMs with Virtual Trusted Platform Module \(vTPM\)](#) information in *Security Guide*.

8. In the **Resources** step, perform the following actions to add a GPU:

GPUs: Click **Add GPU**. The **Add GPU** dialog box appears. Do the following in the indicated fields:

Note: This step is applicable

- a. only to AHV clusters with GPU-enabled mode: Select the option for the desired mode, either **vGPU** or **Passthrough**.
- b. If you select **vGPU**, perform the following configuration settings:
 - **NVIDIA Virtual GPU License:** Select a license type from the dropdown menu. This action sets (filters the list of) available profiles. Click the [Help icon](#) for information about the license types.

Note:

You can add multiple vGPUs to the VM only if you select the license for NVIDIA Virtual GPU software version 10.1 (440.53) or later.

Before you add multiple vGPUs to the VM, see [Multiple Virtual GPU Support](#) and [Restrictions for Multiple vGPU Support](#) in the *AHV Administration Guide*.

- **vGPU Profile:** Select the option for the desired profile. Click the help links (Virtual Slice column and end of line) for more information on the profiles.

Note:

Multiple vGPUs are supported on the same VM only if you select the highest vGPU profile type. After you add the first vGPU, to add multiple vGPUs, see [Adding Multiple vGPUs to the Same VM](#) on page 172




Figure 40: Attach GPUs Window: vGPU

- c. If you select **Passthrough**, click the option for the desired type.



Figure 41: Attach GPUs Window: Passthrough

- d. Click **Save**.

Create VM

1 Configuration 2 **Resources** 3 Management 4 Review

Disks [Attach Disk](#)

#	Type	Image	Size	Bus Type	Actions
1	Disk	-	1GiB	SCSI	

Networks [Attach to Subnet](#)

Subnet	VLAN ID / VPC	Private IP	Public IP	Actions
vlan.2144	2144	None	None	

Boot Configuration

☒ **Legacy BIOS Mode**

Set Boot Priority

Default Boot Order (CD-ROM, Disk, Network)

☐ **UEFI Mode**

UEFI BIOS Mode supports enhanced Shield VM security settings.

Shield VM Settings

GPUs

vGPU

GPUs are available as vGPU or Passthrough and are allotted at power on.

[Attach GPUs](#)

[Back](#) [Next](#)

Figure 42: Create VM Window (Resources)

9. In the **Management** step, perform the following actions to define categories and timezone:

- a. Turn on the **Enable 'Default-Storage' policy** switch to apply the default storage policy to the VM that you are creating.

The **Enable 'Default-Storage' policy** switch is turned off by default.

For more information on the default storage policy, see [Default Storage Policy](#) on page 548.

- b. **Categories:** Search for the category to be assigned to the VM. Select the checkboxes of the categories that you want to assign to the VM.

If you enabled the **Enable 'Default-Storage' policy** switch for the VM in the previous step, then the **Storage:\$Default** is displayed in the **Categories** dropdown menu. To assign more categories to the VM, select the categories from the dropdown menu.

Note:

Do not assign any other category that is already associated with another storage policy to the entity, such as a VM or VG, if you want to enable the **Enable 'Default-Storage' policy** switch for that entity. If you enable the **Enable 'Default-Storage' policy** switch for an entity that is already

associated with another category, the **Storage: \$Default** category is enabled, but the storage policy associated with the other category overrides the default storage policy.

- c. **Timezone:** Select the local timezone to use from the dropdown menu.
- d. **Use this VM as an agent VM:** Select this option to make this VM as an agent VM.

You can use this option for the VMs that must be powered on before the rest of the VMs (for example, to provide network functions before the rest of VMs are powered on the host) and must be powered off and migrated after rest of the VMs (for example, during maintenance mode operations).

Create VM

Configuration Resources **Management** Review

Enable 'Default-Storage' policy to manage the storage configurations across all VM disks. The policy applies via category 'Storage:\$Default'.

☐ Enable 'Default-Storage' policy [How does it work?](#)

☐ Applies to VMs on clusters with AOS 6.1 or above only.

Categories

Type to search...

Tag the VM with Category: Value to assign policies associated with value

Timezone

(UTC) UTC

Use UTC timezone for Linux VMs and local timezone for Windows VMs.

☐ Use this VM as an Agent VM

Guest Customization

Script Type Configuration Method

No Customization Custom Script

Back Cancel Next

Figure 43: Create VM Window (Management)

10. In the **Management** step, perform the following actions for VM customization:

Guest Customization: To customize the VM, select Cloud-init (for Linux VMs) or Sysprep (for Windows VMs).

Script Type: Select the VM script customization type from the dropdown menu. The options are **No Customization**, **Sysprep(Windows)**, or **Cloud-init (Linux)**.

Configuration Method: Select the VM script customization method from the dropdown menu. This field is activated when the **Script Type** is either **Sysprep(Windows)** or **Cloud-init (Linux)**. The options are **Custom Script** or **Guided Script**. A custom script allows you to use an externally created Windows Sysprep or Linux Cloud-init script to customize the operating system of the guest VM. A guided script provides additional configuration options that you can use to customize the operating system of the guest VM, such as username, password, locale, or hostname.

The system displays the options required to configure Cloud-init and Sysprep, such as options to specify a configuration script and options to upload a script.

11. To specify a user data file (Linux VMs) or answer file (Windows VMs) script for unattended provisioning, perform either of the following actions:

- » If the file is available on your local computer, click **Upload Script**, choose and upload the file.
- » Create or paste the contents of the file in the text box below **Upload Script**.

Note: The script type supports the following file formats.

- Sysprep: XML
- Cloud-init (Linux): YAML, JSON, or Shell.

Figure 44 shows the 'Create VM' window, Management tab. The window displays a progress bar with four steps: Configuration, Resources, Management (active), and Review. Below the progress bar, there is a section for enabling the 'Default-Storage' policy, a search bar for categories, a timezone selector set to '(UTC) UTC', and a checkbox for 'Use this VM as an Agent VM'. At the bottom, there is a 'Guest Customization' section with 'Script Type' set to 'No Customization' and 'Configuration Method' set to 'Custom Script'. Navigation buttons 'Back', 'Cancel', and 'Next' are at the bottom right.

Figure 44: Create VM Window (Management)

12. In the **Review** step, when all the field entries are correct, click **Create VM** to create the VM, and close the **Create VM** window.

The new VM appears in the VMs **Summary** page and **List** page.

Managing a VM through Prism Central (AHV)

About this task

After you create a VM, you can perform various operations to manage it based on your requirement.

Note:

- You can perform some of the available actions on a single VM at a time, while others can be performed on multiple VMs simultaneously. The available actions appear in bold; the unavailable actions are grayed out. The available actions depend on the current state of the VM and user permissions.

- You can perform only those operations for which you have permissions from the admin.
- Do not install any external software on the Prism Central VM.

To manage a VM, perform the following steps:

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Compute > VMs** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.
The system displays the **List** tab by default with all the VMs across registered clusters in **Nutanix** environment. For information about how to access the list of non-nutanix VMs managed by an external vCenter, see [VMs Summary View](#) on page 100.
3. Select the required action using any of the following methods:
 - Select the target VM checkbox in the **List** tab, click the horizontal ellipsis (three-dot icon), and choose the required action from the dropdown menu.
 - Right-click on the target VM in the **List** tab, and select the required action from the displayed dropdown menu.
 - Select the required action in the **Summary** page of an individual VM. For information about how to access the **Summary** page of an individual VM, see [Viewing VM Details](#) on page 113.

The following table provides the information about all the actions that you can perform to manage the VM:

Table 35: VM Actions

Action	Description	Applicable to Multiple VMs
Update	Update the VM configuration. For more information, (see Updating a VM through Prism Central (AHV) on page 165.	No
Delete	Delete the VM.	Yes
Clone	Clone a VM. For more information, see Cloning a VM through Prism Central (AHV) on page 169.	No
Create VM Template	Create a VM template. For more information, see Creating a VM Template on page 217.	No
Launch Console	Launch the VM console. For more information, see Launching a VM Console through Prism Central (AHV) on page 171.	No
Migrate Across Clusters	On-demand migrate the VM across AHV clusters without the need to protect them in synchronous replication schedules or set up Nutanix Disaster Recovery from the Prism Central web console. For more information about on-demand VM migration, see On-Demand Cross-Cluster Live Migration on page 180.	Yes
Power On	Power on the VM.	Yes
Power Off	Power off the VM.	Yes

Action	Description	Applicable to Multiple VMs
Power Cycle	power off the VM followed by a power on.	Yes
Reset	Perform an ACPI reset action through the BIOS of the VM.	Yes
Guest Shutdown	Perform a graceful shutdown of the operating system of the VM.	Yes
Guest Reboot	Perform a graceful restart of the operating system of the VM.	Yes
Pause/Suspend	Pause or suspend the VM. This option is available only when the VM is powered on.	Yes
Resume	Resume a paused or suspended VM.	No
Disable Efficiency Measurement	Disable machine-learning capability of efficiency measurement for the VM. You can also create a category of VMs and assign the category in Inefficiency Measurement Exclusion policy to control this function through operations policies. For more information, see Operations Policy Management in <i>Intelligent Operations Guide</i>. Note: By default, the efficiency measurement function is enabled for all VMs.	Yes
Enable Efficiency Measurement	Enable machine-learning capability of efficiency measurement. Note: If the VM is part of the VM category that is mapped to the Inefficiency Measurement Exclusion operations policy, the policy takes precedence and the system does not allow you to enable efficiency measurement. For information about how to assign or remove the VM category in Inefficiency Measurement Exclusion policy, see Managing Operations Policies in <i>Intelligent Operations Guide</i>. You cannot enable efficiency measurement on individual VMs if the efficiency measurement is disabled on the cluster level.	Yes
Disable Anomaly Detection	Disable machine-learning capability of Anomaly Detection for the VM. You can also create a category of VMs and assign the category in Anomaly Detection Exclusion policy to control this function through operations policies. For more information, see Operations Policy Management in <i>Intelligent Operations Guide</i>. Note: By default, the anomaly detection function is enabled for all VMs.	Yes
Enable Anomaly Detection	Enable machine-learning capability of anomaly detection. Note: If the VM is part of the VM category that is mapped to the Anomaly Detection Exclusion operations policy, the policy takes precedence and the system does not allow you to enable anomaly detection. For information about how to assign or remove the VM category in Anomaly Detection Exclusion policy, see Managing Operations Policies in <i>Intelligent Operations Guide</i>.	Yes

Action	Description	Applicable to Multiple VMs
Protect	Assign the VM to a protection policy. Selecting Protect opens a window to specify the protection policy to which this VM should be assigned. You can create a protection policy for a VM or set of VMs that belong to one or more categories by enabling Nutanix Disaster Recovery and configuring the Availability Zone.	Yes
Unprotect	Remove the VM from a protection policy.	Yes
Create Recovery Point	Create VM recovery point. The VM can be restored or replicated from a recovery point either locally or remotely in a state of a chosen recovery point. In the Create VM Recovery Point window, enter a name for the VM recovery point. You can choose to create an app consistent VM recovery point by enabling the App Consistent checkbox. For more information about app consistent recovery point, see Terminology section in the <i>Data Protection and Recovery with Prism Element</i> guide.	No
Migrate	Migrate the VM within the cluster to another host or outside the cluster. In the Migrate VM window, choose any of the following options: • Migrate Within Cluster to select the target host from the dropdown menu (or select the System will automatically select a host option to let the system choose the host). • Migrate Outside Cluster (Protected VMs) to live migrate the VM protected with synchronous replication outside the cluster. Note: This option is disabled if the VM is not protected with synchronous replication schedule or VM synchronization is not completed at the recovery AZ. Note: Nutanix recommends to live migrate VMs when they are under light load. If they are migrated while heavily utilized, migration may fail because of limited bandwidth. For more information about VM migration, see VM Migration Specifications on page 173.	No
Add to Recovery Plan	Add the VM to a recovery plan you created previously. For more information, see the Adding Guest VMs Individually to a Recovery Plan section in the <i>Nutanix Disaster Recovery Guide</i> .	Yes
Run Playbook	Run a playbook you created previously. For more information, see Running a Playbook (Manual Trigger) in <i>Intelligent Operations Guide</i> .	Yes
Manage Categories	Assign the VM a category value. For more information, see Assigning a Category to an Entity in the <i>Prism Central Admin Center Guide</i> .	Yes
Install NGT	Install Nutanix Guest Tools (NGT). For more information, see Setting up NGT on page 193.	Yes

Action	Description	Applicable to Multiple VMs
Manage NGT Applications	Enable or disable NGT. For more information, see Managing NGT Applications on page 196.	Yes
Upgrade NGT	Upgrade NGT. For more information, see Upgrading NGT on page 196.	Yes
Configure VM Host Affinity	This option is disabled and cannot be used to define VM-Host affinity in Prism Central. The VM-Host affinity policies can be created using VM categories and host categories. For more information about affinity policies Prism Central, see VM-Host Affinity Policies Defined in Prism Central on page 501. For information about how to create an affinity policy, see Creating a VM-Host Affinity Policy on page 509. It is recommended to create an affinity to multiple hosts (at least two) to protect against downtime due to a node failure.	No
Add to Catalog	Add the VM to the catalog. For more information, see Adding a Catalog Item on page 285.	No
Manage Ownership	Specify a project and user who owns the VM. In the Manage VM Ownership window, select the target project from the dropdown menu and enter a user name as owner of the VM.	No
Set QoS Attributes	Configure quality of service (QoS) settings. For more information, see Setting QoS for an Individual VM on page 211.	Yes
Export as OVA	Export the VM as OVA. For more information, see OVA Management on page 235 and Uploading an OVA on page 239.	No

4. Check the status of the selected action under tasks widget.

Updating a VM through Prism Central (AHV)

You can use Prism Central to update the VM configuration.

About this task

To update a VM, perform the following steps:

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Compute > VMs** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.
The system displays the **List** tab by default with all the VMs across registered clusters in the **Nutanix** environment. For information about how to access the list of non-Nutanix VMs managed by an external vCenter, see [VMs Summary View](#) on page 100.

3. Select the **Update** action using any of the following methods:

- Select the target VM checkbox in the **List** tab, and choose **Update** from the **Actions** dropdown menu.
- Right-click on the target VM in the **List** tab, and select **Update** from the displayed dropdown menu.
- Click **Update** in the **Summary** page of an individual VM. For information about how to access the **Summary** page of an individual VM, see [Viewing VM Details](#) on page 113.

The **Update VM** window appears, which includes the same fields as the **Create VM** window. For information about how to create a VM, see [Creating a VM through Prism Central \(AHV\)](#) on page 149.

4. Select an option from the Update list to perform actions for updating VM. You will be redirected to the corresponding window based on the option selected.

Note: If you select **General Configuration** the **Update VM Configuration** window appears. You can select any of the options from the left-pane to perform update actions on the VM.

Figure 45: Update VM Window (Configuration)

Modify the VM configuration based on your requirement as indicated in the following steps, and click **Save**.

5. (Optional) **Enable Memory Overcommit:** Select this checkbox to enable memory overcommit for the VM. Before you turn memory overcommit on or off, see [Memory Overcommit](#) in the *AHV Administration Guide*. For more information on configuring memory overcommit, see [Memory Overcommit Management](#) on page 212.

Note: Verify that the status of the **Memory Overcommit** property under **Advanced Configuration** in the **VM Properties** is set as **Enabled**.

6. (Optional) **Enable advanced processor compatibility:** Select this checkbox to enable advanced processor compatibility for the VM.

After you select the **Enable advanced processor compatibility** checkbox, the system provides you the option to select the required CPU generation.

Select the required CPU generation name in the **Select CPU generation** dropdown menu.

For more information, see [Advanced Processor Compatibility in AHV](#) section in the *AHV Administration Guide*.

Note: Advanced processor compatibility setting for VMs with automatic cluster selection configuration is not supported. For more information on automatic cluster selection, see [Automatic Cluster Selection for VM Placement](#) section in the *AHV Administration Guide*.

7. (Optional) Select the **Use this VM as an Agent VM** to mark this VM as an agent VM.
8. (Optional) **Boot Configuration:** You cannot modify the BIOS mode after the VM is created. However, if you have enabled the Secure Boot option under UEFI Mode while creating the VM, you can modify this setting during the VM update.
9. (Optional) Select the **Timezone** from the dropdown. The default timezone is UTC.
10. (Optional) Select **Storage** from the **Update** list or you can select from the **Update VM Configuration** window. You can enable the **Flash Mode** settings for the VM and VM disks to on or off.
 - a. You are redirected to the **Disks** page to update the existing VM settings. From the **Disks** page, you can add new disks to the VM using the **Attach Disk** option. You can modify the existing disk attached to the VM using the controls under the **Actions** column.
11. (Optional) Select **Network** from the **Update** list or you can select from the **Update VM Configuration** window.
 - a. You are redirected to the NICs page to update the existing VM settings. You can attach a new network to the VM using the **Attach to Subnet** option. You can also modify the existing subnet and Traffic Mirror destination NICs attached to the VM from the **More Network Options**. Before you modify the NIC network or create a new NIC for a VM, see [Creating a VM through Prism Central \(AHV\)](#) on page 149 and [Limitation for vNIC Hot-Unplugging](#) information in the *AHV Administration Guide*.
12. (Optional) Select **GPU** from the **Update** list or you can select from the **Update VM Configuration** window.
13. Select the **Use this VM as an Agent VM** to mark this VM as an agent VM.
14. **Boot Configuration:** You cannot modify the BIOS mode after the VM is created. However, if you have enabled the Secure Boot option under UEFI Mode while creating the VM, you can modify this setting during the VM update.
15. Select the **Timezone** from the dropdown. The default timezone is UTC.
16. Select **Storage** from the **Update** list or you can select from the **Update VM Configuration** window. You can enable the **Flash Mode** settings for the VM and VM disks on or off.

17. You are redirected to the **Disks** page to update the existing VM settings. From the Disks page, you can add new disks to the VM using the **Attach Disk** option. You can also modify the existing disk attached to the VM using the controls under the actions column.

See [Creating a VM through Prism Central \(AHV\)](#) on page 149 before you modify the disk or create a new disk for a VM.




Figure 46: Disks Update

Note: Starting with Prism Central 2023.4, the individual VM disks can be migrated using Prism Central UI. If more granular control is required, perform the steps through acli as described in [Migrating a vDisk to Another Container](#) section in *AHV Administration Guide*.

18. You can turn the flash mode settings for the VM and VM disks on or off. To update the flash mode for individual virtual disks, click the update disk icon in the **Disks** pane and deselect the **Enable Flash Mode** check box.




Figure 47: Update VM Window (Resources)

19. Select **Network** from the **Update** list or you can select from the **Update VM Configuration** window.

20. You are redirected to the NICs page to update the existing VM settings. You can attach a new network to the VM using the **Attach to Subnet** option. You can also modify the existing subnet and Traffic Mirror destination NICs attached to the VM from the **More Network Options**. Before you modify the NIC network or create a new NIC for a VM, see [Creating a VM through Prism Central \(AHV\)](#) on page 149 and [Limitation for vNIC Hot-Unplugging](#) information in the *AHV Administration Guide*.
21. Select **GPU** from the **Update** list or you can select from the **Update VM Configuration** window.
 - a. You are redirected to **Hardware** page to update the existing VM settings. You can add pass-through GPUs if a VM is already using GPU pass-through. You can also change the GPU configuration from pass-through to vGPU or vGPU to pass-through, change the vGPU profile, add more vGPUs, and change the specified vGPU license.

Note:

- You must shut down the VM before you perform GPU configuration-related operations.
- Before you add multiple vGPUs to the VM, see [Multiple Virtual GPU Support](#) and [Restrictions for Multiple vGPU Support](#) information in the *AHV Administration Guide*.
- Multiple vGPUs are supported on the same VM only if you select the highest vGPU profile type.
- After you add the first vGPU, to add multiple vGPUs, see [Adding Multiple vGPUs to the Same VM](#) on page 172.

22. (Optional) Select the VM and go to **Categories** or you can select from the **Update VM Configuration** window.
 - a. Select **Manage Categories** to associate categories with the existing VM. The default storage policy uses the **Storage:\$Default** category.
 - b. The **Enable 'Default-Storage' policy** switch on the **Update VM Configuration** is turned off by default. When you turn on the **Enable 'Default-Storage' policy** switch, the VM is associated with the **Storage:\$Default** category.

For more information on the default storage policy, see [Default Storage Policy](#) on page 548.

Cloning a VM through Prism Central (AHV)

You can use Prism Central to clone a VM.

About this task

To clone a VM, perform the following steps:

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Compute > VMs** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

The system displays the **List** tab by default with all the VMs across registered clusters in **Nutanix** environment. For information about how to access the list of non-nutanix VMs managed by an external vCenter, see [VMs Summary View](#) on page 100.

3. Perform the **Clone** action using any of the following methods:

- Select the target VM checkbox in the **List** tab, and choose **Clone** from the **Actions** dropdown menu.
- Right-click on the target VM in the **List** tab, and select **Clone** from the displayed dropdown menu.
- Click **Clone** in the **Summary** page of an individual VM. For information about how to access the **Summary** page of an individual VM, see [Viewing VM Details](#) on page 113.

The **Clone VM** window appears, which includes the same fields as the **Create VM** window. For information about how to create a VM, see [Creating a VM through Prism Central \(AHV\)](#) on page 149.

A cloned VM inherits most the configurations (except the name) of the source VM.

Note:

- You can clone up to 250 VMs at a time.
- You cannot override the secure boot setting while cloning a VM, unless the source VM already had secure boot setting enabled.

The following is an example showing the **Clone VM** window:

Clone VM

General Configuration

Number Of Clones: 2

Prefix Name: test-vm

Starting Index Number: 1

Compute Details

vCPU(s): 1

Number Of Cores Per vCPU: 1

Memory: 1 GB

Boot Configuration

The BIOS mode cannot be updated after the VM is created.

☒ Legacy BIOS

Set Boot Priority: Default: Boot Order (CD-ROM, Disk, Network)

☐ UEFI

Disks

Type	Address	Parameters
CD-ROM	ide:0	EMPTY=true; BUS=ide
DISK	scsi:0	SIZE=1GIB; CONTAINER=...

Volume Groups

You haven't added any volume groups yet.

Network Adapters (NIC)

VLAN ID	VIRTUAL SWITCH	PROMOTE SP	MAC
100	vs0		

☐ CloneVM Host Affinity

You haven't pinned the VM to any hosts yet.

☐ Custom Script

Cancel Save

Figure 48: Clone VM Window

4. Enter a name for the clone, and click **Save**.

You can optionally override some of the configurations before you click **Save**. For example, you can override the number of vCPUs, memory size, boot priority, NICs, or the guest customization.

The VM clone is created successfully.

Launching a VM Console through Prism Central (AHV)

You can use Prism Central to launch a VM console.

Procedure

1. Log in to Prism Central.

2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Compute > VMs** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

The system displays the **List** tab by default with all the VMs across registered clusters in **Nutanix** environment. For information about how to access the list of non-nutanix VMs managed by an external vCenter, see [VMs Summary View](#) on page 100.

3. Select the **Launch Console** action using any of the following methods:

- Select the target VM checkbox in the **List** tab, and choose **Launch Console** from the **Actions** dropdown menu.
- Right-click on the target VM in the **List** tab, and select **Launch Console** from the displayed dropdown menu.
- Click **Launch Console** in the **Summary** page of an individual VM. For information about how to access the **Summary** page of an individual VM, see [Viewing VM Details](#) on page 113.

This opens a Virtual Network Computing (VNC) client and displays the console in a new tab or window. This option is available only when the VM is powered on. The VM power options that you access from the **Power On Actions** (or **Power Off Actions**) action link below the VM table can also be accessed from the VNC console window. To access the VM power options, click [Power icon](#) at the top-right corner of the console window.




Figure 49: Console Window (VNC)

Note:

- A VNC client may not function properly on all browsers. Some keys are not recognized when the browser is Chrome. It is observed that the VNC client works better on Firefox browser.
- To enable or disable the console support for a VM with one vGPU configured, see [Enabling or Disabling Console Support for vGPU VMs](#).

Adding Multiple vGPUs to the Same VM

About this task

You can add multiple vGPUs of the same vGPU type when you create a new VM or update an existing VM.

- For information on how to create a VM, see [Creating a VM through Prism Central \(AHV\)](#).
- For information on how to update an existing VM, see [Updating a VM through Prism Central \(AHV\)](#).

For more information on multiple vGPU support, see [Multiple Virtual GPU Support](#).

Before you begin

Ensure that the following prerequisites are met before you add multiple vGPUs to the VM:

- Select the license for NVIDIA Virtual GPU (vGPU) software version 10.1 (440.53) or later.
- Observe the guidelines and restrictions specified in [Multiple Virtual GPU Support](#) and [Restrictions for Multiple vGPU Support](#).

Procedure

To add multiple vGPUs to the same VM, perform the following steps:

1. Click **Add GPU** in the **Resources** step of create VM workflow or update VM workflow. For more information, see [Creating a VM through Prism Central \(AHV\)](#) or [Updating a VM through Prism Central \(AHV\)](#)

2. In the **Add GPU** window, click **Add**.

The License field is grayed out and you cannot select a different license when you add a vGPU for the same VM.

The **VGPU Profile** is auto-selected. The system allows you to select additional vGPU of the same vGPU type as indicated by the message at the top of the **Add GPU** window.

The newly added vGPU appears in the **Create VM** or **Update VM** window.

3. Repeat the steps for each vGPU addition to the VM.

VM Migration Specifications

VM migration can be performed in the following ways:

- *Live Migration* – This method is used to transfer the VM from one physical host to another host without affecting the normal functions and operations of the VM, and with minimal and often no interruption to the users of the VM.

Live migration is supported for the following scenarios:

- Intra-cluster live migration of guest VMs from one host to another host. For more information, see [Migrating Within the Cluster](#) in the *Prism Central Infrastructure Guide*.
- Cross-cluster live migration (CCLM) of guest VMs protected with synchronous replication schedules. For more information, see [Cross-Cluster Live Migration](#) in the *Nutanix Disaster Recovery Guide*.
- On-demand cross-cluster live migration of guest VMs across AHV clusters registered to the same or different Prism Central AZs. For more information, see [On-Demand Cross-Cluster Live Migration](#) on page 180.
- *Cold Migration* - This method is used to transfer the VM from one physical host to another host with minimal VM downtime. You can use **Nutanix Move** for Inter-cluster VM migration <Third-Party Hypervisor> to the AHV cluster. For more information, see [Move User Guide](#).

Note: The <Third-party hypervisors> are Xen (Amazon EC2), Microsoft Hyper-V, and VMware ESXi.

Live Migration Cases

VM live migration is applicable for the following cases:

- Maintenance activity required on the cluster - This scenario is applicable for a planned event (for example, scheduled maintenance of guest VMs) at the primary AZ when you put node into maintenance mode. For information on how to configure a planned failover, see [Cross-Cluster Live Migration](#) section in *Nutanix Disaster Recovery Guide*.

- AHV or Firmware Upgrades - When you upgrade the AHV or Firmware, the VMs need to be live migrated to another AHV cluster.
- Load balancing or to isolate specific VMs on hosts.
- Changes in Host affinity policy.
- Acropolis Dynamic Scheduling (ADS)- If ADS detects a resource contention in the cluster, it creates a migration plan to eliminate the hotspots in the cluster by migrating VMs from one host to another. For more information about ADS, see [Acropolis Dynamic Scheduling in AHV](#) section in *AHV Administration Guide*.
- Defragmentation activity to create resources for VM power-on operation. This activity is performed when the cluster can accommodate resources for VM power-on operation after moving some VMs to another host.

Live Migration Restrictions

This section provides the list of factors that affect VM live migration.

The following conditions affect VM live migration;

- VM is powered off.
- VM is configured as an agent VM.

You can select or clear the **Use this VM as an agent VM** checkbox to enable or disable this function when you create or update a VM. For more information, see the [Creating a VM through Prism Central \(AHV\)](#).

- For guest VMs with Windows Credential Guard enabled, the VM live migration is only supported between the nodes with Skylake or newer Intel CPU generations. While expanding the cluster, if you add a node with CPU generation prior to Skylake, the system generates an alert *VMs became non-migratable due to node addition* after cluster expansion. Nutanix recommends that you remove the nodes with older CPU generation, power-off, then power-on affected guest VMs to re-enable VM live migration support. For more information, see [KB-15227](#).

For information about Windows Credential Guard, see the [Windows Defender Credential Guard Support](#) section in the *AHV Administration Guide*.

- For guest VMs with WSL2 enabled, VM live migration is only supported between the nodes with Skylake or newer Intel CPU generations. While expanding the cluster, if you add a node with CPU generation prior to Skylake, the system generates an alert *VMs became non-migratable due to node addition* after cluster expansion. Nutanix recommends that you remove the nodes with older CPU generation, power-off, then power-on affected guest VMs to re-enable VM live migration support. For more information, see [KB-15227](#).

For information about WSL2, see the [Windows Subsystem for Linux \(WSL2\) Support on AHV](#) section in the *AHV Administration Guide*.

- GPU passthrough is enabled.

For information on how to verify if GPU passthrough is enabled for the VM, see [Checking Live Migration Status of a VM](#) on page 175.

- CPU passthrough is enabled.
- VM-Host affinity is set from the Prism Element web console with one host.

For information on how to verify the VM-Host affinity policy from the Prism Element web console, see [Verifying Affinity Policy Association](#) on page 176.

- VM affinity policies are defined from Prism Central using VM categories and host categories.

For information on how to verify the affinity policies mapped to a VM from Prism Central, see [Checking Affinity Policies of a VM From Prism Central](#) on page 177.

- Defragmentation activity to create resources for VM-power-on is in progress.

Note:

- Starting with AOS 6.8 release, AHV supports the live migration of the guest VMs that have Windows Credential Guard enabled. When you upgrade to AOS 6.8 or later release and enable HA reservation for the cluster, the system reserves the resources to accommodate the HA event. However, if you have the existing VMs in the cluster that have Windows Credential Guard enabled, some guest VMs (both with or without Windows Credential Guard enabled) might fail to start due to the unavailability of sufficient resources in the cluster. Ensure that you have sufficient resources available in the cluster to accommodate the HA reservation.

For more information on Windows Credential Guard, see the [Windows Defender Credential Guard Support in AHV](#) section in the *AHV Administration Guide*.

For more information on how to enable HA reservation, see the [Enabling High Availability for the Cluster](#) section in the *Prism Element Web Console Guide*.

- Starting from AOS 7.0 release, with a minimum required AHV version of 10.0, AHV supports the live migration of the guest VMs that have WSL2 enabled. When you upgrade to AOS 7.0 or later release and enable HA reservation for the cluster, the system reserves the resources to accommodate the HA event. However, if you have the existing VMs in the cluster that have WSL2 enabled, some guest VMs (both with or without WSL2 enabled) might fail to start due to the unavailability of sufficient resources in the cluster. Ensure that you have sufficient resources available in the cluster to accommodate the HA reservation.

For more information on WSL2, see [Windows Subsystem for Linux \(WSL2\) Support on AHV](#) in the *AHV Administration Guide*.

For information on how to enable HA reservation, see [Enabling High Availability for the Cluster](#) section in the *Prism Element Web Console Guide*.

Checking Live Migration Status of a VM

This section describes how to check whether the live migration is allowed for the VM.

Procedure

To check if a VM can be live migrated, perform the following steps:

- Log on to the CVM as a nutanix user using SSH.
- Run the following command:

```
nutanix@cvm$ acli vm.get <VM_NAME> | grep 'allow_live_migrate'
```

Replace `<VM_NAME>` with the actual VM name at your site.

The following attribute value confirms the live migration is allowed:

```
allow_live_migrate: True
```

What to do next

If `allow_live_migrate: False` is returned for the above command, then check the status of the VM properties that restricts live migration.

The following table provides the information about the verification procedure to check the status of VM properties:

VM Properties	Verification Procedure
WSL2	Log on to the CVM as a Nutanix user using SSH, and run the following command: <pre>nutanix@cvm\$ acli vm.get <VM_NAME> grep 'hardware_virtualization'</pre> Replace <VM_NAME> with the actual VM name at your site. The following attribute value confirms WSL2 is enabled: <pre>hardware_virtualization: True</pre>
GPU	1. Log in to Prism Central. 2. Select the Infrastructure application from Application Switcher Function, and navigate to Compute > VMs from the Navigation Bar. For information on the Navigation Bar, see Application-specific Navigation Bar. The system displays the List tab by default with all the VMs across registered clusters in Nutanix environment. 3. Apply the GPU TYPE filter from the Modify Filter option. The system displays the VMs based on the selected GPU criteria.



Figure 50: GPU TYPE Filter

Verifying Affinity Policy Association

This section describes how to verify the VM-Host affinity from Prism Element web console and affinity policies mapped to the VM from Prism Central.

Procedure

To verify affinity policy association for a VM, perform the following checks:

- [Checking VM-Host Affinity From Prism Element](#) on page 176.
- [Checking Affinity Policies of a VM From Prism Central](#) on page 177.

Checking VM-Host Affinity From Prism Element

About this task

Using Prism Element, you can define the VM-host affinity policy that controls the placement of a VM. You can use this policy to specify that a selected VM can only run on the members of the affinity host list.

Procedure

To verify the VM-Host affinity, perform the following steps:

1. Log in to Prism element web console.
2. Navigate to **VM** > **Table**, and double-click the target VM name. The system displays the **Update VM** window.
3. Under **VM Host Affinity**, check the hosts specified for the VM.

Note:

VMs with Host affinity policies can only be migrated to the hosts specified in the affinity policy during an HA event. If only one host is specified, the VMs cannot be migrated.

Checking Affinity Policies of a VM From Prism Central

About this task

Using Prism Central, you can define the affinity policy based on VM Categories and Host categories.

Procedure

To verify the affinity policy associated with a VM, perform the following steps:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Compute** > **VMs** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#).

The system displays the **List** tab by default with all the VMs across registered clusters in **Nutanix** environment.

3. Click the target VM name.
The system displays the details view of the target VM.
4. Click the **Categories** tab in the target VM's details view.
The system displays the Categories and Affinity Policy associated with the VM.

Defining Behaviour for Non-migratable VMs (GPU/CPU/PCI/Host Affinity configured VMs)

About this task

When a host is put in maintenance mode, VMs are moved from that host to other hosts in the cluster. After exiting maintenance mode, those VMs are automatically returned to the original host, eliminating the need to manually move them. However, some VMs cannot be migrated to other hosts. For more information, see [Non-Migratable VMs](#).

When an attempt is made for a host to enter maintenance mode, the VMs with GPU passthrough, CPU passthrough, and PCI passthrough are not migrated to the other hosts in the cluster and therefore may block the attempt to enter maintenance mode.

Note: VMs with host affinity policies are also not migrated to other hosts in the cluster, if any of the following condition is met:

- The hosts that are configured as part of VM-Host affinity policy are not available.
- The hosts that are part of VM-Host affinity policy does not have the sufficient resources to run the VM.

To overcome this situation, Nutanix recommends you define either of the following behavior for these non-migratable VMs when host enters the maintenance mode:

- Automatically shutdown these non-migratable VMs.
- Set a block migration indication for these non-migratable VM.

Procedure

To define the above behaviors for non-migratable VMs, perform the following steps:

1. Log on to CVM with SSH.
2. Run the following command to put the host into maintenance mode with default action set for non-migratable VMs:

```
nutanix@cvm$ acli host.enter_maintenance_mode <Hypervisor-IP-address> [wait="{ true | false }" ]  
[non_migratable_vm_action="{ acpi_shutdown | block }"]
```

Replace:

- **<Hypervisor-IP-address>** - with the actual host name at your site.
- **wait**: Set the wait parameter to true to wait for the host evacuation attempt to finish.
- **non_migratable_vm_action**: By default, this parameter is set to *block*, that prevents the host from entering maintenance mode, if non-migratable VMs exist. For more information on non-migratable VMs, see [Non-Migratable VMs](#).

If you want to automatically shut down such VMs for the duration of the href="..ahv/vm-not-migrated-to-other-hosts-r.dita"/>If you want to automatically shut down such VMs for the duration of the maintenance mode, set this parameter to *acpi_shutdown*.

Migrating Within the Cluster

You can live-migrate a VM to another host within the cluster.

About this task

The **Migrate Within Cluster** option in the **Migrate VM** window allows you to migrate VMs to other hosts within the cluster.

Before you begin

Ensure that you observe the prerequisites as described in [Live Migration of vGPU-enabled VMs](#) on page 179 section, and the limitations described under [Limitations for Live Migration](#) information in *AHV Administration Guide*.

Procedure

To live migrate the VMs within the cluster, perform the following steps:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Compute > VMs** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.
The system displays the **List** tab by default with all the VMs across registered clusters in **Nutanix** environment. For information about how to access the list of non-nutanix VMs managed by an external vCenter, see [VMs Summary View](#) on page 100.

3. Select the target VM checkbox in the **List** tab, and choose **Migrate** from the **Actions** dropdown menu.
The system displays the **Migrate VM** window with **Migrate Within Cluster** option selected as default.




Figure 51: Migrate VM window

4. In the **Host** dropdown menu, perform either of the following actions:
 - Retain the *System will automatically select a host* default option if you want to migrate the VM to a host selected by the system.
The system selects a host based on the GPU resources available with the host as appropriate for the VM to be migrated live.
 - Select the target host from the dropdown menu to which you want to migrate the VM.

5. Click **Migrate**.

Prism submits the task and displays the following message:

Successfully submitted migrate operation.
Task details

Task details is a link to the **Tasks** page. Click the **Task details** link to monitor the migration task in the **Tasks** page.

The host name of the VM in the **List** view changes to the host name to which you migrated the VM.

Live Migration of vGPU-enabled VMs

You can live-migrate a vGPU-enabled VM to another host within the cluster.

You can perform live migration of VMs enabled with virtual GPUs (vGPU-enabled VMs) only on commercially reasonable effort, if the destination node is equipped to provide enough resources to the vGPU-enabled VMs. However, if the destination node is not equipped with the enough resources, the vGPU-enabled VMs are shut down and you might experience a downtime.

In a successful migration case, the vGPUs can continue to run while the VMs that are running the vGPUs are seamlessly migrated in the background.

When you perform the LCM update, the vGPU-enabled VMs are listed as Non-HA-protected VMs. LCM also migrates the Non-HA-protected VMs on commercially reasonable effort to the destination node if the following requirements are met:

- Destination node is equipped with the required resources for the VM.
- The VM GPU drivers are compatible with the AHV host GPU drivers.

If the destination node is not equipped with the enough resources or there is any compatibility issue between the VM GPU drivers and AHV host GPU drivers, the LCM forcibly shuts down the Non-HA-protected VMs.

Prerequisites

Ensure the following prerequisites are met before you live migrate the vGPU-enabled VMs:

- The VM is not powered off.
- The host affinity is not set for the VM.

If the host affinity of the VM is set to only one host, you cannot migrate the VM. However, based on the GPU resources required, if the host affinity is set to multiple hosts with the same or similar GPU resources, you can migrate the VM among the hosts with which the affinity is set.

- You have another host in the same cluster that supports the resources necessary to support the vGPU-enabled VM that you want to live migrate.

Limitations

For limitations applicable to live migration of vGPU-enabled VMs, see [Limitations for Live Migration](#) in *AHV Administration Guide*.

On-Demand Cross-Cluster Live Migration

On-demand cross-cluster live migration (OD-CCLM) enables zero-downtime migration of guest VMs and their metadata (such as VM categories) across AHV clusters registered to the same or different Prism Central instances, spanning multiple availability zones (AZs).

To ensure zero downtime, you can live migrate the guest VMs without the need to protect them with synchronous replication schedules or set up Nutanix Disaster Recovery from the Prism Central web console. The guest VMs retain Flow Network Security (FNS) policies when migrated to an AHV cluster registered to the same Prism Central instance.

Note: AHV Metro Availability also supports cross-cluster live migration (CCLM), but only for guest VMs that are protected with synchronous replication schedules. For more information, see [Cross-Cluster Live Migration](#) in the *Nutanix Disaster Recovery Guide*.

You can live migrate a guest VM in the following situations:

- Between overlay subnets in different AZs.
- Between overlay subnets in different virtual private clouds (VPCs) within the same AZ.
- Within the same VPC in an AZ.
- Between NC2 AWS VPCs within the same AZ.
- Between VLAN subnets in different AZs.
- Between VLAN subnets within the same AZ.
- Within the same Prism Central AZ, where guest VMs are part of Flow Network Security policies.

A Layer 2 (L2) network extension creates a tunnel between source and destination subnets, ensuring dependent VMs can communicate with the migrating VM. Configure an L2 network extension when migrating a guest VM in the following scenarios:

- Between overlay subnets in different AZs.
- Between overlay subnets in different virtual private clouds (VPCs) within the same AZ.
- Between VLAN subnets in different AZs.
- Between VLAN subnets within the same AZ.

Tip: An L2 network extension is not required when migrating a VM within the same VPC subnet.

For successful live migration, Nutanix recommends that the round-trip latency between the clusters must be 40 ms or less. The data copy time during live migration is dependent on the data to be copied between the clusters and the available network bandwidth.

On-Demand Cross-Cluster Live Migration (OD-CCLM) Requirements

This section describes the requirements for live migrating guest VMs.

Table 36: Nutanix Software Requirements

Nutanix Software	Software Version Required
AOS	• 6.7 or later. While upgrading the AOS version from 6.6 to 6.7, ensure that the upgrade is completed before performing on-demand migration of the guest VM. • For OD-CCLM of vTPM-enabled guest VMs, 6.8 or later. Both the source and destination clusters must be running AOS 6.8 and later.
Prism Central	• pc.2023.3 or later. • For OD-CCLM of vTPM-enabled guest VMs pc.2024.1 or later.

Table 37: Prism Central User Requirements

Prism Central User	Permissions
Administrator	Only the administrator (local admin user account in Prism Central) can perform on-demand live migration between the clusters registered to different Prism Central deployments.

Prism Central User	Permissions
Non-administrator	Non-administrators can perform on-demand live migration only between the clusters registered to the Prism Central deployment. The user must have the following permissions: • View_Prism_Central • View_Virtual_Machine • View_Subnet • View_Availability_Zone • Allow_Cross_Cluster_VM_Migration • View_Cluster For more information on user management, see User Management in <i>Security Guide</i>.

Cluster and VM Requirements

- The destination cluster subnet must use the same IP network prefix as the source VM subnet.
- If you did not enable the advanced processor compatibility feature for the guest VM, ensure that the destination Nutanix cluster supports all of the CPU feature sets (set of CPU flags) that are supported on the source Nutanix cluster. Otherwise, live migration fails. For more information, see [Advanced Processor Compatibility](#) in *AHV Administration Guide*.

You can use Live migration even when the CPU types of the source and destination clusters do not match exactly. The destination cluster must support the superset of the CPU features of the source cluster.

- If the VM in the source cluster has NGT enabled, then the VM must have at least one empty IDE or SATA CD-ROM to attach the ISO required for reconfiguration on the destination cluster.
- Both the source and the destination clusters must run on the same AHV version.
- Both the source and the destination clusters must have the same storage container name for the guest VMs.

A storage container with the same name as the one on the source cluster must exist on the destination cluster. For example, if a *SelfServiceContainer* storage container exists on the source cluster, the destination cluster must also have a *SelfServiceContainer* storage container.

- The destination cluster must have sufficient storage capacity to host the migrating guest VMs.
- The destination cluster must be reachable from the source cluster.
- The guest VM to be live migrated must be powered on.

The Prism Central web console filters out the guest VMs that are powered off.

Network and Port Requirements

- If HTTP proxy is enabled on any Prism Element within the cluster, ensure that the CVM IP addresses and the virtual IP addresses of the remote cluster (destination Prism Element) are added to the proxy allowlist of both the clusters. You can use either the Prism Element UI or the Prism Element CLI (nCLI) to allow the IP addresses. For UI instructions, see [Configuring HTTP Proxy](#) in *Prism Element Web Console Guide*. For CLI, you can use the same commands as mentioned in [Allowing Prism Central and Its Managed Clusters \(nCLI\)](#) in *Prism Central Admin Center Guide*.

- For clusters running on AOS version 7.3 or later, the system opens the required ports automatically.
- For clusters running on versions earlier than AOS version 7.3, open the ports listed on [Ports and Protocols](#) manually for communication

- To open ports for communication from the source cluster to the destination cluster, run the following command on all CVMs of the source cluster:

```
nutanix@cvm$ allssh 'modify_firewall -f -r remote_cvm_ip,remote_virtual_ip -p 2030,2036,2073,2090 -i eth0'
```

Replace *remote_cvm_ip* with the IP address of the destination cluster CVM. For multiple CVMs, replace *remote_cvm_ip* with the IP addresses of the CVMs separated by comma.

Replace *remote_virtual_ip* with the virtual IP address of the destination cluster.

- To open the ports for communication from the destination cluster to the source cluster, run the following command on all CVMs of the destination cluster:

```
nutanix@cvm$ allssh 'modify_firewall -f -r source_cvm_ip,source_virtual_ip -p 2030,2036,2073,2090 -i eth0'
```

Replace *source_cvm_ip* with the IP address of the source cluster CVM. For multiple CVMs, replace *source_cvm_ip* with the IP addresses of the CVMs separated by comma.

Replace *source_virtual_ip* with the virtual IP address of the source cluster.

- Replace *remote_cvm_ip* or *source_cvm_ip* with the CVM IP addresses, separated by commas for multiple CVMs. Replace *remote_virtual_ip* or *source_virtual_ip* with the cluster virtual IP.
- For standard deployments, use the *eth0* interface only. *eth0* is the default CVM interface that shows up when you install AOS.
- For clusters with Disaster Recovery (DR) network segmentation enabled, use the *ntnx0* interface and run the *modify_firewall* commands.
- For clusters with backplane network segmentation enabled, use the *eth0* interface. The cross cluster communication on clusters that have backplane network segmentation enabled is done through the management interface *eth0*.

On-Demand Cross-Cluster Live Migration Limitations

Consider the following limitations before performing On-Demand Cross-Cluster Live Migration (OD-CCLM) of your guest VMs:

- On-demand live migration of guest VMs is only supported in the IPAM-enabled VPC (overlay) networks and not supported in the IPAM-enabled VLAN networks.
- On-demand live migration to clusters on Nutanix Cloud AZ (DRaaS) is not supported.
- On-demand live migration to Nutanix Cloud Cluster AZ (NC2) is not supported.
- On-demand live migration of up to 60 guest VMs in parallel is supported.

While no limit exists to the number of parallel migrations, the time required for successful migration depends on the guest VM size.

- On-demand live migration of guest VMs fails in the following scenarios:

- Guest VMs are part of Flow Network Security policies in another Prism Central AZ.
- Guest VMs are part of consistency groups.
- Guest VMs have volume groups attached to them.
- Guest VMs have replication factor 1.
- GPU passthrough is enabled on guest VMs.

For information on GPU passthrough verification, see [Checking Live Migration Status of a VM](#).

For information on GPU passthrough verification using acropolis CLI (accli), see [vm](#) in the [Command Reference Guide](#).

- Guest VMs are vNUMA VMs.
- Guest VMs are in a paused state.
- Guest VM has memory overcommit enabled.
- Guest VMs are migrated within an Open Virtual Network (OVN) using VLAN segmentation and a Distributed Virtual Switch (DVS), all managed by a single Prism Central instance in same availability zone (AZ).
- Guest VMs are upgraded before on-demand live migration starts.

When on-demand live migration is in progress, do not perform hotplug upgrades to the memory or CPU. Updates to NIC, disk, memory, and CPU are fatal while data is copied during migration.

- Prism Central upgrade is in progress.
- AHV upgrade is in progress.

Similarly, if on-demand live migration is in progress, upgrading AHV fails.

- LCM upgrade is in progress.
- After migrating a guest VM in an overlay subnet in a virtual private cloud (VPC) to an overlay subnet in another VPC, Prism Central does not retain the floating IP address of the guest VM if the VPCs are attached to different external subnets. Therefore, you must reassign the floating IP address to the migrated guest VM to access the guest VM from the external network.

If the VPCs share the same external subnets, Prism Central retains the floating IP address of the guest VM.

- After migration to the recovery cluster, VM categories are not preserved if those categories do not exist on the recovery cluster.
- Data over the wire is not encrypted during the live migration.
- Automatic defragmentation on the recovery cluster is not supported.
- Live migration of guest VMs running on ESXi is not supported.

On-Demand Cross-Cluster Live Migration (OD-CCLM) Best Practices

Nutanix recommends the following best practices for live migrating your guest VMs (and all of their associated metadata like VM categories):

- Ensure that you set up the relevant user permissions. For more information, see [On-Demand Cross-Cluster Live Migration \(OD-CCLM\) Requirements](#) on page 181.

- Enable the advanced processor compatibility feature for the guest VM and select the oldest CPU generation to maximize the migration capability of the guest VM. For more information, see [Advanced Processor Compatibility in AHV](#).
- Do not protect the guest VM with a DR protection policy when on-demand live migration is in progress.
- Do not perform Nutanix software upgrades when on-demand live migration is in progress. For more information, see [On-Demand Cross-Cluster Live Migration Limitations](#) on page 183.
- Do not perform on-demand live migration when Nutanix software upgrades are in progress. For more information, see [On-Demand Cross-Cluster Live Migration Limitations](#) on page 183.
- Do not change the IP addresses or other network configurations of the guest VMs when on-demand live migration is in progress.
- Ensure that a 10G network link is available in the infrastructure setup.
- When an administrator user live migrates the guest VMs created by a non-administrator to a Nutanix cluster registered to a different Prism Central deployment, the administrator user becomes the owner of the guest VM after the VM migration. The ownership remains the same irrespective of who migrated the guest VM when the live migration happens between the Nutanix clusters registered to the same Prism Central deployment.

Performing On-Demand CCLM

If, due to a failure event (for example, scheduled maintenance of guest VMs) at the primary cluster, you must migrate your applications to another AHV cluster without VM downtime, perform an on-demand live migration to the recovery cluster.

Before you begin

See [On-Demand Cross-Cluster Live Migration \(OD-CCLM\) Requirements](#) on page 181, [On-Demand Cross-Cluster Live Migration Limitations](#) on page 183, and [On-Demand Cross-Cluster Live Migration \(OD-CCLM\) Best Practices](#) on page 184.

About this task

To live migrate the guest VMs, follow these steps at the primary AZ:

Procedure

1. Log in to the Prism Central web console.
2. From **Application Switcher Function**, select the **Infrastructure** application, and navigate to **Compute > VMs** from the **Navigation Bar**.
3. Select the guest VMs to live migrate.

Warning: To see the guest VM configurations that cannot be migrated, see [On-Demand Cross-Cluster Live Migration Limitations](#) on page 183.

4. Click **Migrate Across Clusters** from the **Actions** drop-down menu.

5. Do the following in the **Destination Cluster** tab:
 - a. From the dropdown menu, under **Destination Location**, select the AZ to live migrate your guest VMs.
 - b. From the dropdown menu, under **Destination Cluster**, select the AHV cluster to live migrate your guest VMs.
 - c. Click **Next**.
6. Do the following in the **Network** tab:
 - a. From the dropdown menu under **Destination Subnets**, select the network to map the subnet of the migrated guest VMs.
 - b. Click **Next**.
7. Do the following in the **Migration Checks** tab:
 - a. Review the status of the selected guest VMs after the prechecks run automatically.
If any precheck fails, resolve the issue that is causing the failure and click **check again**.
 - b. When all the **Checks** under the **Selected VMs** show OK, click **Migrate**.
 - (Optional) To cancel a running migration task, go to **Tasks** and select the active task. From the drop-down menu of the **Migrate VM Across Clusters** section, click **Cancel**. When the confirmation window is prompted, click **Confirm**.

The selected guest VMs migrate to the recovery cluster with zero VM downtime.

Creating a VM through Prism Central (ESXi)

You can create virtual machines (VMs) in ESXi clusters through Prism Central.

Before you begin

Ensure that the following prerequisites are met before you create a VM in ESXi cluster:

- All the requirements, rules, and guidelines are considered, and the limitations are observed. For more information, see [External vCenter Server Integration](#) on page 341 and .
- The vCenter Server is registered with your cluster. For information on how to register a vCenter Server, see [Registering a Cluster with a vCenter Server using Prism Central](#) on page 351.

Procedure

To create a VM in ESXi cluster, perform the following steps:

1. Log in to Prism Central.
2. Select **Infrastructure** application from [Application Switcher Function](#), and navigate to **Compute > VMs** from the **Navigation Bar**. For information on the **Navigation Bar**, see [Application-specific Navigation Bar](#). The system displays the **List** tab by default with all the VMs across registered clusters in **Nutanix** environment. For information on how to access the list of non-nutanix VMs managed by an external vCenter, see [VMs Summary View](#) on page 100 and .

3. Click **Create VM**, and enter the following information in the **Configuration** step:
 - a. **Name**: Enter a name for the VM.
 - b. **Description** (optional): Enter a description for the VM.
 - c. **Cluster**: Select the target ESXi cluster from the drop-down list on which you intend to create the VM.
 - d. **Number of VMs**: Enter the number of VMs you intend to create. The created VM names are suffixed sequentially.
 - e. **vCPU(s)**: Enter the number of virtual CPUs to allocate to this VM.
 - f. **Number of Cores per vCPU**: Enter the number of cores assigned to each virtual CPU.
 - g. **Memory**: Enter the amount of memory (in GiBs) to allocate to this VM.
4. In the **Resources** step, perform the following actions to attach a Disk to the VM:

Disks: Click **Attach Disk** , and enter the following information:

 - a. **Type**: Select the type of storage device, **Disk** or **CD-ROM**, from the dropdown list.
 - b. **Operation**: Specify the device contents from the drop-down list.
 - Select **Empty CD-ROM** to create a blank CD-ROM device. A CD-ROM device is needed when you intend to provide a system image from CD-ROM. The **Empty CD-ROM** option is available only when *CD-ROM* is selected as the storage device in the **Type** field.
 - Select **Allocate on Storage Container** to allocate space without specifying an image. Selecting this option means you are allocating space only. You have to provide a system image later from a CD-ROM or other source. The **Allocate on Storage Container** option is available only when *Disk* is selected as the storage device in the **Type** field.
 - c. If you select **Allocate on Storage Container** in the **Operation** field, the system prompts you to specify the **Storage Container**. Select the appropriate storage container.
 - d. **Bus Type**: Select the bus type from the drop-down list.

The options displayed in the drop-down list varies based on the storage device **Type** selected in the **Type** field.

If the storage device **Type** is:

 - **Disk** - The available choices are **SCSI**, **SATA**, **PCI**, or **SATA**.
 - **CD-ROM** - The available choices are **IDE**, or **SATA**
 - e. **Path**: Enter the path to the desired system image.
 - f. **Capacity**: Enter the disk size in GiB.
 - g. When all the field entries are correct, click **Save** to attach the disk to the VM and return to the **Create VM** window.

Repeat this step to attach additional devices to the VM.

5. In the **Resources** step, perform the following actions to create a network interface for the VM:

Networks: Click **Attach to Subnet**. The **Attach to Subnet** window appears.

- a. **Subnet:** Select the target subnet from the drop-down list.

The list includes all the defined networks. For information on how to configure the network, see [Network Configuration for VM Interfaces](#) in *Prism Element Web Console Guide*.

- b. **Network Adapter Type:** Select the network adapter type from the drop-down list.

For information on the list of supported adapter types, see [External vCenter Server Integration](#) on page 341.

- c. **Network Connection State:** Select the state for the network after VM creation:

- *Connected* - If the VM needs to be connected to the network to operate.
- *Disconnected* - If the VM needs to be in disconnected state after creation.

- d. Click **Save** to create a network interface for the VM, and return to the **Create VM** window.

Repeat this step to create additional network interfaces for the VM.

6. In the **Management** step, perform the following actions to define categories and timezone:

- a. **Categories:** Search for the category to be assigned to the VM. The policies associated with the category value are assigned to the VM.

- b. **Guest OS:** Type and select the guest operating system.

The guest operating system that you select affects the supported devices and number of virtual CPUs available for the virtual machine. The Create VM wizard does not install the guest operating system. For information on the list of supported operating systems, see [External vCenter Server Integration](#) on page 341.

7. In the **Review** step, when all the field entries are correct, click **Create VM** to create the VM, and close the **Create VM** window.

The new VM appears in the VMs **Summary** page and **List** page.

Managing a VM through Prism Central (ESXi)

This section describes how to manage VMs in an ESXi cluster through Prism Central in **Nutanix** environment.

About this task

The procedure to manage VMs in an ESXi cluster in Nutanix environment is same as for AHV cluster, however the fields can vary when you create or manage a VM in an ESXi cluster. For information on the fields, see [Creating a VM through Prism Central \(ESXi\)](#) on page 186 and .

About this task

To manage a non-nutanix VM on an external vCenter, you can use playbooks. For more information, see [VMs Summary View](#) on page 100.

Note: Do not install any external software on the Prism Central VM.

Procedure

To manage a VM of an ESXi cluster in **Nutanix** environment:

- For information on how to manage a VM of an ESXi cluster in **Nutanix** environment, see [Managing a VM through Prism Central \(AHV\)](#) on page 161 and .

You can perform only those operations for which you have permissions from the admin.

Nutanix Guest Tools

Nutanix Guest Tools (NGT) is a software package that comes bundled with AOS. You can install NGT in a guest virtual machine (VM) to enable advanced VM management functionalities provided by Nutanix.

Nutanix Guest Tools Overview

The NGT software package for Linux and Windows includes the following components:

- **NGT Installer:** Allows you to install NGT in a guest VM.
- **Nutanix Guest Agent (NGA) Service:** Maintains a communication channel between the Controller VM (CVM) and the guest VMs.
- **Nutanix VirtIO Package:** Includes the drivers required to run the VM on the AHV hypervisor. For more information about VirtIO drivers, see [Nutanix VirtIO for Windows](#) in *AHV Administration Guide*.
- **Python:** The NGA service is written in Python and is shipped with a dedicated installation of Python so it does not interact with the other components in the system.
- **Nutanix VSS package:** Enables Nutanix native in-guest Volume Snapshot Service (VSS) agent to take application-consistent snapshots for all the VMs that support VSS.

In addition, the NGT software package for Windows also includes the following components:

- A custom module that performs checks and fulfills all the prerequisites.
- Microsoft Visual C++ 2015 x64 redistributable component
- Integration with Windows Performance Monitor for hypervisor statistics.

Installing NGT in a guest VM allows you to use the following advanced VM management features:

- **File Level Restore CLI**

Performs self-service file-level recovery from the VM snapshots. For more information about self-service restore, see [Self-Service Restore](#) in the *Data Protection and Recovery with Prism Element* guide.

- **Nutanix VM Mobility**

Facilitates VM migration between ESXi and AHV, in-place hypervisor conversion, and cross-hypervisor disaster recovery (CHDR) features. For more information about cross-hypervisor disaster recovery, see [Nutanix Cross Hypervisor Disaster Recovery](#) in the *Data Protection and Recovery with Prism Element* guide. For more information about in-place hypervisor conversion, see [In-Place Hypervisor Conversion](#) in the *Prism Element Web Console Guide*.

- **Application-consistent snapshots for Windows VMs**

Enables Nutanix native in-guest Volume Snapshot Service (VSS) agent to take application-consistent snapshots for all the VMs that support VSS. This mechanism takes application-consistent snapshots without any VM stuns (temporary unresponsive VMs) and also enables third-party backup providers like CommVault and Rubrik to take application-consistent snapshots on the Nutanix platform regardless of which hypervisor is used in the cluster. For more information about Nutanix VSS-based snapshots for the Windows VMs, see the *Application-consistent Snapshots with Microsoft Volume Shadow Copy Service (VSS)* section in the [Conditions for Application-consistent Snapshots](#) topic of the *Data Protection and Recovery with Prism Element* guide.

- **Application-consistent snapshots for Linux VMs**

Supports application-consistent snapshots for Linux VMs by running specific scripts on VM quiesce. For more information, see the [Conditions for Application-consistent Snapshots](#) information of the *Data Protection and Recovery with Prism Element* guide.

- **Static IP address preservation support after failover for Nutanix Disaster Recovery (DR)**

This feature allows the preservation of the IP address of a guest VM (with a static IP address) for its failover (DR) to an IPAM network. For more information, see the *Networking Requirements* section in the [Nutanix Disaster Recovery Requirements](#) topic of the *Nutanix Disaster Recovery Guide*.

- **In-guest scripts execution support for Nutanix Disaster Recovery (DR)**

In-guest scripts automate various task executions upon recovery of the VMs. For example, you can automate the task of changing the domain controller or performing the reassigning of the DNS IP address and reconnection to the database server at the recovery AZ. For more information about the tasks that can be automated using In-guest scripts, see [Creating a Recovery Plan](#) in the *Nutanix Disaster Recovery Guide*.

Nutanix Guest Tools Requirements

Ensure that the following requirements are met so that you can successfully install NGT, and use all the NGT features.

General Requirements

- NGT can be installed on VMs running specific versions of Windows or Linux only. For information about the supported OS for specific NGT features, see the [NGT](#) section in *Compatibility and Interoperability Matrix*.
- VMs must have at least one empty IDE or SATA CD-ROM to attach the ISO required for the configuration and installation of NGT, and for NGT functionalities such as IP-remapping during DR failover, enabling NGT on a remote cluster after DR failover, or CCLM.

Requirements for IP-based connectivity between NGT and CVM

- A virtual IP address must be configured for the Nutanix cluster. If you change the virtual IP address of the cluster, the reconfiguration impacts all the VMs in your cluster on which NGT is installed. For more information, see [Virtual IP Address Impact](#) in the *Prism Element Web Console Guide*.
- Network access from the guest VM to the virtual IP address of the cluster where IP-based connectivity is used for NGT.
- If IP-based communication is enabled, the TCP port 2074 of the cluster virtual IP address must be accessible from the guest VMs so that the guest VMs can communicate with the CVM. For more information, see [Nutanix Guest Agent and Controller VM Communication](#). For the complete list of required ports, see [Ports and Protocols](#).

Requirements for installing NGT through Prism Central

In addition to the [General Requirements](#), ensure that the following requirements are met so that you can successfully install NGT through Prism Central:

- Network access from the virtual IP address of the cluster to the guest VMs where Prism Central-based management of NGT is required.
- To install NGT on a Windows VM, you need a local user account in the VM with administrative privileges. Domain accounts are currently not supported.

- To install NGT on a Windows VM, Windows Remote Manager Service (winrm) must be running and configured to allow SSL-based connections. The following commands are one of the many methods to configure winrm with SSL. These commands must be executed from an administrative Windows PowerShell console.

```
> $certificate = New-SelfSignedCertificate -DnsName $env:computername -
CertStoreLocation cert:\LocalMachine\My

> winrm create winrm/config/Listener?Address=*&Transport=HTTPS
"@{Hostname=`"$env:computername`";CertificateThumbprint=`"$($certificate.ThumbPrint)`"}"

> cmd /c 'winrm set winrm/config/service/auth @{Basic="true"}'

> netsh advfirewall firewall add rule name="WinRM-HTTPS" dir=in localport=5986
protocol=TCP action=allow
```

Note: Windows Remote Manager Service (winrm) is required to install NGT from Prism Central only. It is not required when you upgrade NGT or install NGT manually by logging in to the guest VM.

- To install NGT on a Linux VM, you must have a guest VM user account with password-based SSH, and root or passwordless SUDO access enabled.

Windows Requirements

In addition to the [General Requirements](#), ensure that the following requirements are met so that you can successfully install NGT in a Windows VM:

- For Windows Server Edition VMs, ensure that Microsoft Volume Shadow Copy Service (VSS) is enabled before you install NGT with the VSS application enabled. For more information about how to verify if Microsoft VSS is enabled on your server, see Microsoft documentation.

Note:

- Before you enable the VSS feature, upgrade the NGT version on the Windows Server Edition VMs to 2.3.2 or later. If not, the VSS snapshots will be marked as crash-consistent.
- NGT's VSS feature is not supported on Windows Client Edition VMs or Linux VMs.

- During NGT installation a few PowerShell commands are executed. To run these commands, the PowerShell executable must be present in the PATH system variable of the environment. Otherwise, the installation of NGT might fail. For more information, see [KB-7284](#).
- The PowerShell version on the Windows VMs must be 3.0 or later.
- Antivirus scanners might need exclusions added for the C:\Program Files\Nutanix directory for successful NGT installation.
- The TCP port 23578 in the guest VM must be accessible if you want to use the VSS service. For the complete list of required ports, see [Ports and Protocols](#).

Linux Requirements

In addition to the [General Requirements](#), ensure that the following requirements are met so that you can successfully install NGT in a Linux VM:

- Ensure that any package manager configuration within the VMs works properly, as old or invalid configurations might prevent NGT package installation.

Note: The guest OS type and version determines the package manager. Depending on which OS type and version you are using, the package manager can be yum, dnf or apt. To confirm which package manager is installed, see your guest OS documentation.

- Ensure that the user has write permissions for the /usr/local directory as NGT is installed at this location.
- Ensure that the user has write permissions for the /mnt/nutanix/ngt directory to mount the NGT ISO at this location.
- Ensure that the user has write access to /tmp directory for the duration of the installation.

General Notes

- NGT includes VirtIO drivers when you install it on a Windows VM. If these drivers are already installed in your guest VM, they might be replaced during the installation of NGT, depending on the version that is bundled with NGT. If the drivers were downgraded, you might need to manually reinstall VirtIO and restart the VM if prompted. For more information about VirtIO drivers, see [Nutanix VirtIO for Windows](#) in *AHV Administration Guide*. For information about the compatible VirtIO driver versions, see the *AOS - NGT - VM Mobility Matrix* section in the *NGT* tab of the [Compatibility and Interoperability Matrix](#).
- VirtIO driver installation updates the NIC driver in the guest VM. During the update, the guest VM might lose network connectivity. The network connectivity is restored after the update is complete.
- If a VM is connected to a volume group (VG), NGT captures the iSCSI Qualified Name (IQN) of the VM and stores the information. If you change the IQN of the VM and take a snapshot of the VM before the NGT refresh cycle (currently at 5 minutes) occurs, NGT does not provide the auto restore capability because the snapshot operation cannot capture the VM-VG connection.

Workaround:

- **Windows VM:** In the command prompt, run the net stop "Nutanix Guest Tools Agent" && net start "Nutanix Guest Tools Agent" command.
- **Linux VM:** Restart the Nutanix Guest Agent (NGA) service by running one of the following commands:
 - For guest VMs that support the **systemd** OS, run the \$ systemctl restart ngt_guest_agent command.
 - For guest VMs that do not support the **systemd** OS, run the \$ sudo service ngt_guest_agent restart command.
- By default, NGT client certificates expire every 1,000 days. Regenerate the certificates before they expire. For more information, see [Regenerating NGT Certificates for Guest VMs](#) information in the *Prism Element Web Console Guide*.

NGT Management in Prism Central

You can use Prism Central to perform the following operations on a guest VM or multiple guest VMs:

- [Setting up NGT](#) on page 193\
- [Enabling NGT and Mounting the NGT Installer on Cloned VMs](#) on page 195
- [Managing NGT Applications](#) on page 196
- [Upgrading NGT](#) on page 196
- [Removing NGT](#) on page 198

In addition, NGT installation through Prism Central provides the following capabilities:

- Enable Nutanix Volume Shadow Copy Service (VSS) and Self Service Restore (SSR) applications as part of the install workflow
- Select multiple VMs in Prism Central from the VM Entity browser and upgrade NGT on these VMs
- Defer restarting of the VMs to a specified later time and date after installing or upgrading NGT

- Define policies that allow you to defer restarting of the VMs to a specified later time and date

Setting up NGT

Set up NGT in a guest VM using Prism Central.

Before you begin

Ensure that all the requirements specified in [Nutanix Guest Tools Requirements](#) are met before you set up NGT.

About this task

NGT is fully automated using Prism Central. You do not need to log in to the VM manually to set up NGT

Note:

- You can install NGT on up to 1,500 VMs in a cluster. If you exceed this limit, migrate or delete the extra VMs to return to the supported limits.
- The guest VM must be powered on to install NGT in the VM.
- You cannot install NGT on VMs created on storage containers with replication factor 1.
- You cannot set up NGT on Nutanix infrastructure VMs.

To set up NGT on a guest VM or multiple guest VMs at the same time in Prism Central, follow these steps :

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and go to **Compute > VMs** from the **Navigation Bar**.

For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

The system displays the **List** tab by default with VMs across all the registered clusters.

Tip: The **NGT** status column associated with a VM displays whether the NGT is installed in the VM, whether it is the latest version, and whether an upgrade is available or not.

3. Select the checkbox associated with the VM, and from the **Actions** dropdown menu, click **Guest Tools > Set Up NGT**.

The system displays the **Set Up Nutanix Guest Tools (NGT)** window.

Note: If you select multiple VMs, the dialog box displays the number of VMs eligible for NGT setup and the VMs that are ineligible. The reason for rejection also appears next to each ineligible VM. For example, if you select a VM with NGT already set up, the system skips it during this operation.

4. In the **Choose Preference** tab, choose one of the following:
 - » Select **New NGT Installation** to automatically or manually set up a new NGT instance.
 - » Select **Configure Existing NGT Installations** to enable NGT and configure applications and certificates in guest VMs with NGT already set up.
5. (Optional) To view the prerequisites for setting up NGT, click **View Prerequisites**.
6. (Optional) To see the guest VMs where Prism Central sets up NGT, in the **Selected VMs** section, click the arrow icon.

7. Click **Next**.

The system displays the **Configure Applications** tab.

8. (Optional) To enable Volume Snapshot Service (VSS) and Self Service Restore (SSR) applications, select the **Volume Shadow Copy Services (VSS)** and **Self Service Restore (SSR)** checkboxes.

If you choose not to enable these features when you set up NGT, you can do it later by following the instructions in [Managing NGT Applications](#).

9. Click **Next**.

The system displays the **Install** tab.

10. In the **Installation Type** section, choose one of the following:

- » Select **Install Automatically** to automatically set up NGT in the selected guest VMs.

The system displays the **VM Credentials** section.

- » Select **Mount Installer** if you do not want to set up NGT automatically or cannot provide the username and password of an account that meets the requirements.

This option allows you to complete the setup manually.

Prism Central enables the NGT feature in the VM, mounts the NGT installer, and attaches an ISO to the virtual CD drive with the volume label NUTANIX_TOOLS to the selected VMs. You can then install NGT by logging in to the VM. For more information, see [Installing NGT in a Windows VM](#) or [Installing NGT in a Linux VM](#) in the *Prism Element Web Console Guide*.

Note: If the expected version of NGT is already installed in the VM, it detects the NUTANIX_TOOLS CD and reconfigures itself. After NGT is installed and the reconfiguration is successful, the system automatically unmounts NUTANIX_TOOLS. For more information, see [CD-ROM Eject Functionality of NGT](#).

11. In the **VM Credentials** section, choose one of the following:

- » Select **Single credential for all VMs** if the guest VMs use the same username and password.

The system displays the **Username** and **Password** fields.

- » Select **Different credentials for different VMs** if the guest VMs do not use the same username and password.

The system displays the **Enter VM Credentials JSON** text box.

12. In the **Username** and **Password** field, enter a username and password of an account that meets the requirements specified in [Nutanix Guest Tools Requirements](#), and is accessible to all the VMs on which NGT is being set up.

Alternatively, in the **Enter VM Credentials JSON** text box, paste the JSON file containing the credentials. To download a sample of the JSON file, click **Download Template JSON**.

13. In the **Restarting Preference** section, select the restart schedule of the guest VM using one of the following options:
 - **Restart immediately after installation completes:** The VM restarts after the installation completes.
 - **Restart at a preferred time:** The VM restarts at the specific date and time that you configure in the **Date** and **Time** fields.
 - **Skip restart:** The guest VM does not restart after the installation completes..

Note: If you select **Skip restart**, the system checks if any NGT policy is defined, and applies the restart action; **Restart as soon as the install is completed** or **Restart at specific date and time after the install is completed** based on the NGT policy (if defined). For more information about how to create an NGT policy, see [Creating a NGT Policy](#) on page 523.

14. Click **Complete Setup** to start the setup process.

When the setup completes successfully on a guest VM, the **NGT Status** column of the guest VM in the List tab displays the status as **Up to date**. It takes a few minutes for the status to reflect in Prism Central because this is an asynchronous operation.

Note: If the [Nutanix Guest Tools Requirements](#) are not met or you experience any issue with the setup process, the NGT setup fails, and an alert message displays. Ensure that all the requirements are met and perform the setup process again to set up NGT in the VMs.

Enabling NGT and Mounting the NGT Installer on Cloned VMs

Enable NGT and mount the NGT installer on cloned VMs.

About this task

If you cloned one or more VMs from a VM that had NGT installed, follow these steps to enable NGT and mount the NGT installer:

Note: After you perform the following steps, you do not need to separately install NGT on the cloned VMs.

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and go to **Compute > VMs** from the **Navigation Bar**.

For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

The system displays the **List** tab by default with VMs across all the registered clusters.
3. Select the checkbox associated with the VM to enable or disable the SSR and VSS applications, and from the **Actions** dropdown menu, click **Guest Tools > Set Up NGT**.

The system displays the **Set Up Nutanix Guest Tools (NGT)** dialog box.

Note: If you select multiple VMs, the dialog box displays the number of VMs eligible to enable NGT and the VMs that are ineligible. The reason for rejection also appears next to each ineligible VM. For example, if you select a VM with NGT already enabled, the system skips it during this operation.

4. In the **Choose Preference** tab, select **Configure Existing NGT Installations**.
5. (Optional) To view the prerequisites for setting up NGT, click **View Prerequisites**.

6. (Optional) To see the guest VMs on which Prism Central sets up NGT, in the **Selected VMs** section, click the arrow icon.
7. Click **Next**.
The system displays the **Configure Applications** tab.
8. (Optional) To enable Volume Snapshot Service (VSS) and Self Service Restore (SSR) applications, select the **Volume Shadow Copy Services (VSS)** and **Self Service Restore (SSR)** checkboxes.
If you choose not to enable these features when you set up NGT, you can do it later by following the instructions in [Managing NGT Applications](#).
9. Click **Next**.
The system displays the **Install** tab.
10. Click **Complete Setup**.
NGT automatically identifies the ISO, updates the configuration, and unmounts the CD.

Managing NGT Applications

Using Prism Central, manage the Nutanix Volume Shadow Copy Service (VSS) and Self-Service Restore (SSR) applications in a guest VM.

About this task

To enable the VSS and SSR applications if you did not enable them during the NGT installation workflow, follow these steps. You can also use this procedure to disable the applications after installing NGT.

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and go to **Compute > VMs** from the **Navigation Bar**.
For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.
The system displays the **List** tab by default with VMs across all the registered clusters.
3. Select the checkbox associated with the VM to enable or disable the SSR and VSS applications, and from the **Actions** dropdown menu, click **Guest Tools > Manage NGT Applications**.
The system displays the **Manage Nutanix Guest Tools (NGT) Applications** dialog box.
4. Select one of the following options to enable or disable the applications:
 - **Keep Existing Settings**: Continue using the current settings.
 - **Enable All**: Enable the applications in the guest VM.
 - **Disable All**: Disable the applications in the guest VM.
5. Click **Save** to save the changes.
When you enable an application, the **NGT Applications Enabled** field in the **Guest Tools** widget of the **Summary** tab of the guest VM displays the NGT services enabled for the VM. For example, if you enable the SSR application, the **NGT Applications Enabled** field displays **Self-Service Restore (SSR)**.

Upgrading NGT

Upgrade NGT on a guest VM using Prism Central.

Before you begin

- Ensure that all the requirements specified in [Nutanix Guest Tools Requirements](#) are met.
- Unless you upgrade AOS, you cannot upgrade NGT.
- Upgrading NGT using Prism Central is supported on VMs configured with NGT version 1.2.3 or later. If the VM is configured with an earlier version of NGT, upgrade NGT using the procedure in the [Upgrading NGT](#) topic in the *Prism Element Web Console Guide*.

About this task

To upgrade NGT on a guest VM, follow these steps:

Note:

- You can install NGT on up to 1,500 VMs in a cluster. If you exceed this limit, migrate or delete the extra VMs to return to the supported limits.
- You can upgrade NGT on a maximum of 60 VMs at the same time. For more information, see [Manage Bulk Operations for NGT](#) in the *Prism Central Infrastructure Guide*.
- For Linux VMs, upgrading a cluster running any NGT version prior to NGT 4.0 to NGT 4.0 or later versions, might fail. For more information, see [KB 15396](#).

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and go to **Compute > VMs** from the **Navigation Bar**.

For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

The system displays the **List** tab by default with VMs across all the registered clusters.

3. Select the checkbox associated with the VM to enable or disable the SSR and VSS applications, and from the **Actions** dropdown menu, click **Guest Tools > Manage NGT Applications**.
The system displays the **Upgrade Nutanix Guest Tools (NGT)** dialog box. The dialog box displays the number of VMs eligible for NGT upgrade and the VMs that are ineligible. The reason for rejection also appears next to each ineligible VM. For example, if you select a VM on which the latest version of NGT is already installed, that VM is skipped during this operation.
4. Select the restart schedule of the VMs using one of the following options:
 - » **Restart immediately after installation completes:** The VM restarts after the installation completes.
 - » **Restart at a preferred time:** The VM restarts at the specific date and time that you configure in the **Date** and **Time** fields.
 - » **Skip restart:** The VM does not restart after the installation completes.

Note: Nutanix recommends that you restart the VM after upgrading NGT.

5. Click **Upgrade** to start the upgrade process.

When the upgrade is complete, the **NGT Status** status column associated with the VMs displays the status as **Up to date**. Because this is an asynchronous operation, it takes a few minutes for the status to reflect in Prism Central.

Note: If the [Nutanix Guest Tools Requirements](#) are not met or if you experience any issue with the upgrade process, the NGT upgrade fails and an alert message displays. Ensure that all the prerequisites are met and perform the upgrade process again to upgrade NGT.

Removing NGT

Remove NGT from a guest VM using Prism Central.

About this task

To remove NGT from a guest VM, follow these steps:

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and go to **Compute > VMs** from the **Navigation Bar**.
For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.
The system displays the **Remove Nutanix Guest Tools (NGT)** dialog box.
3. Click **Remove**.
NGT is removed from the guest VM.

Manage Bulk Operations for NGT

You can install or upgrade NGT in bulk on multiple guest virtual machines (VMs) using third-party endpoint management tools, such as Microsoft Intune or HCL BigFix, or automation tools, such as Ansible. Using these management tools, you can also remove NGT in bulk from guest VMs.

Bulk installation of NGT is supported on guest VMs running AOS 6.7 or later versions only.

The [Nutanix Support portal](#) lists the following NGT installer files:

- EXE file for Windows OS
- TGZ file for RPM-based Linux OS
- TGZ file for DEB-based Linux OS
- Nutanix NGT GnuPG public key

Download the latest installer files to install or upgrade NGT in bulk.

Note: Bulk installation of NGT using third-party endpoint management tools does not require Prism Central. However, you must enable and mount NGT in guest VMs using Prism Central after installing NGT in VMs. For more information, see [Enable and Configure NGT](#).

Prepare the NGT installation files for Distribution

You can install NGT using the installer files available in the [Nutanix Support portal](#) in a single VM or, in bulk, on multiple VMs. When you prepare to install NGT in bulk on multiple VMs, download the installer files and make them available for use by the third-party endpoint management tool deployed at your site. The instructions in this document

assume that the third-party endpoint management tool deployed at your site requires that you host the NGT installer files on a web server.

Hosting the files for Windows VMs

Before you begin

Ensure that the NGT version is compatible with the AOS version installed in your cluster.

Procedure

1. Go to the Nutanix Support portal, select **Downloads > Nutanix Guest Tools (NGT)**, and download the *nutanix-guest-agent-<version>.exe* installer file for Windows, which matches the AOS version installed in your clusters.
2. Host this installation file directly on the web server.

Hosting the files for Red Hat Package Manager (RPM) based VMs

Before you begin

Ensure that the NGT version is compatible with the AOS version installed in your cluster.

Procedure

1. Go to the Nutanix Support portal, select **Downloads > Nutanix Guest Tools (NGT)**, and download the *nutanix-guest-agent-rpm-<version>.tar.gz* installer file for RPM-based distributions, which matches the AOS version installed in your clusters.
2. Extract the *nutanix-guest-agent-rpm-<version>.tar.gz* file, and host the *NUTANIX-NGT-GPG-KEY* file and the *ngt_repo* directory on the web server.

Hosting the files for Debian (DEB) based VMs

Before you begin

Ensure that the NGT version is compatible with the AOS version installed in your cluster.

Procedure

1. Go to the Nutanix Support portal, select **Downloads > Nutanix Guest Tools (NGT)**, and download the *nutanix-guest-agent-deb-<version>.tar.gz* installer file for DEB-based distributions, which matches the AOS version installed in your clusters.
2. Extract the *nutanix-guest-agent-deb-<version>.tar.gz* file, and host the *i386* and *amd64* directories on the web server.
3. (Optional) Perform the following steps to verify the DEB installer packages against the detached signatures using the *NUTANIX-NGT-GPG-KEY* file:
 - a. Run the following command to import the public key:

```
$ gpg --import NUTANIX-NGT-GPG-KEY
```

The following is an example.

```
$ gpg --import NUTANIX-NGT-GPG-KEY
gpg: key 42DBF8BB: public key "Nutanix, Inc. (NGT Packaging)
<security@nutanix.com>" imported
```



```
gpg: Total number processed: 1
gpg:             imported: 1 (RSA: 1)
```

- b. Run the following command to verify the DEB installer packages against the detached signatures:

```
$ gpg --verify os-arch/nutanix-guest-agent.deb.asc os-arch/nutanix-guest-agent_version-1_os-arch.de
```

Replace *os-arch* with the architecture of the OS of guest VM, and *version* with the NGA version.

The following is an example.

```
$ gpg --verify i386/nutanix-guest-agent.deb.asc i386/nutanix-guest-agent_4.0-1_i386.deb
gpg: Signature made Wed 24 May 2023 01:46:29 PM UTC using RSA key ID 42DBF8BB
gpg: Good signature from "Nutanix, Inc. (NGT Packaging) <security@nutanix.com>"
gpg: WARNING: This key is not certified with a trusted signature!
gpg:             There is no indication that the signature belongs to the owner.
Primary key fingerprint: D8B0 18BD CFEB 774C D157  F0A5 11B1 600F 42DB F8BB
```

Install NGT in Bulk on Multiple VMs

This section explains how to install Nutanix Guest Tools (NGT) in bulk across multiple virtual machines using Prism Element. It provides step-by-step instructions for automating the installation process, reducing manual effort, and ensuring consistency. The guidance also includes best practices for managing NGT components to enhance VM performance, security, and operational efficiency. You can install NGT in bulk on multiple Windows and Linux guest VMs using the NGT installer files available in the [Nutanix Support portal](#).

Installing NGT in Bulk on Windows VMs

Before you begin

- Ensure that the cluster meets all NGT requirements. For more information, see [Nutanix Guest Tools Requirements](#) on page 190.
- Ensure that you registered the VMs with the third-party endpoint management tool deployed at your site. For more information, see the tool-specific documentation.
- Review the end user license agreement (EULA) for Nutanix Guest Tools using a manual installation because the following installation procedure requires you to accept the EULA automatically.

Procedure

1. Configure the third-party endpoint management tool deployed at your site to distribute the *nutanix-guest-agent-<version>.exe* file to the VMs where NGT is installed.

For more information, see the tool-specific documentation.

2. Configure the third-party endpoint management tool to install NGT by running one of the following commands:

» `C:\ngtinstaller> nutanix-guest-agent-version.exe /quiet ACCEPTTEULA=yes /norestart`

Replace *version* with the NGA version.

Use this command to ensure that VMs do not restart after you install NGT.

Note: This command might update the Nutanix VirtIO drivers if no Nutanix VirtIO drivers are installed or if a newer version is available, but the updated functionality of the VirtIO drivers is available only after a VM restart.

» `C:\ngtinstaller> nutanix-guest-agent-version.exe /quiet ACCEPTTEULA=yes`

Replace *version* with the NGA version.

Use this command to automatically restart the VM and to make all the updated VirtIO driver functionalities to be available in the VM.

3. (Optional) To generate the NGT logs in a location other than the %TEMP% directory, install NGT by running the following command.

`C:\ngtinstaller> nutanix-guest-agent-version.exe /quiet ACCEPTTEULA=yes /log log_file`

Replace *version* with the NGA version, and *log_file* with the filename to write the logs.

Ensure that the directory containing the filename that you provide has the necessary write permissions. Also, the installation process adds some events to the Windows application event log.

What to do next

After successful installation, enable and configure NGT in guest VMs. For more information, see [Enable and Configure NGT](#).

Installing NGT in Bulk on Linux VMs

About this task

This sample procedure provides steps to install NGT on Red Hat Package Manager (RPM) based operating systems. Use this procedure as a template to install NGT on Debian-based operating systems.

Before you begin

- Ensure that the cluster meets all NGT requirements. For more information, see [Nutanix Guest Tools Requirements](#) on page 190.
- Ensure that you registered the VMs with the third-party endpoint management tool deployed at your site. For more information, see the tool-specific documentation.

Procedure

1. To verify the package signatures, configure the third-party endpoint management tool deployed at your site to install the NUTANIX-NGT-GPG-KEY file on RPM-based operating systems.

For more information, see the tool-specific documentation.

2. Configure the third-party endpoint management tool to set up the repositories.

For more information, see the tool-specific documentation.

For example, guest VMs running RedHat 8.x distribution might need a repository configuration in `/etc/yum.repos.d/nutanix-guest-tools.repo` similar to the following configuration:

```
[NutanixGuestTools]
name=Nutanix Guest Tools
baseurl=http://local-web-server/ngt_repo
enabled=1
gpgcheck=1
gpgkey=http://local-web-server/RPM-GPG-PUBLIC-KEY
repo_gpgcheck=0
```

Replace `local-web-server` with the name of the web server used to distribute the NGT installer files as described in [Preparing the NGT installation files for the Distribution](#).

3. Configure the third-party endpoint management tool to install the Nutanix guest agent package by running the package manager-specific install command.

For example, the `yum install -y nutanix-guest-agent` command installs Nutanix guest agent package using the yum package manager for RedHat-based distributions.

```
[nutanix@localhost ~]$ yum install -y nutanix-guest-agent
```

For information about the install command specific to the package manager at your site, see the package manager-specific documentation.

What to do next

After successful installation, enable and configure NGT in guest VMs. For more information, see [Enable and Configure NGT](#).

Enable and Configure NGT

Use Prism Central to enable and configure NGT in VMs so that you can use the NGT applications such as self-service restore, volume snapshot service, and application-consistent snapshots. The initial configuration of NGT requires mounting an ISO into the CD-ROM drive of guest VMs. The NGT agent detects the ISO and performs the initial configuration, making connections to the CVM.

Enabling and Configuring NGT using Prism Central

Before you begin

- Ensure that the cluster meets all NGT requirements. For more information, see [Nutanix Guest Tools Requirements](#) on page 190.
- Ensure that you install NGT in the VMs.

Note: Because NGT is installed, you do not require WinRM or SSH access to the VM or a local user account in the VM with administrative privileges.

Procedure

1. Log in to Prism Central.

2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and go to **Compute > VMs** from the **Navigation Bar**.

The system displays the **List** tab by default with VMs across all the registered clusters.

Tip: The **NGT** status column associated with a VM displays whether the NGT is installed in the VM, whether it is the latest version, and whether an upgrade is available or not.

3. To enable and configure NGT in a VM, select the target VM checkbox and choose **Install NGT** from the **Actions** dropdown menu.

The system displays the **Install NGT** window.

Note: The Install NGT workflow is the same for installing NGT and enabling and configuring NGT if NGT is already installed.

4. To enable SSR and VSS applications, select the **Enable Self Service Restore (SSR)** and **Enable Volume Snapshot Service (VSS)** checkboxes.

If you choose not to enable these features during NGT configuration, you can do it later by following the instructions in [Managing NGT Applications](#).

5. Select **Skip restart** as the restart schedule.

Note: Because NGT is already installed in the VM, a restart is not required during NGT configuration.

6. Click **Confirm and Enter Password**.

7. Click **Skip and Mount**.

Prism Central enables the NGT feature in the VM, mounts the NGT installer, and attaches an ISO to the virtual CD drive with the volume label NUTANIX_TOOLS to the selected VMs.

Note: Because NGT is already installed in the VM, it detects the NUTANIX_TOOLS CD and reconfigures itself. After NGT is installed and the reconfiguration is successful, the system automatically unmounts NUTANIX_TOOLS. For more information, see [Automatic CD-ROM Ejection](#) on page 209.

What to do next

The Nutanix Guest Agent (NGA) service in the VMs starts periodic communication with the CVM. To verify whether the NGA service is communicating with the CVM, log in to the CVM and run the following command:

```
nutanix@cvm$ nutanix_guest_tools_cli list_vm_tools_entities include_vm_info=true  
vm_name=vm-name
```

Replace `vm-name` with the name of one of the VMs.

Upgrade NGT in Bulk on Multiple VMs

This section explains how to upgrade Nutanix Guest Tools (NGT) in bulk across multiple virtual machines using Prism Element. It provides step-by-step instructions for streamlining the upgrade process, reducing manual intervention, and ensuring consistency across environments. The guidance also includes best practices for maintaining compatibility, improving VM performance, and enhancing operational efficiency. You can upgrade NGT in bulk on multiple Windows and Linux guest VMs using the NGT installer file available in the [Nutanix Support portal](#).

Upgrading NGT in Bulk on Windows VMs

Before you begin

- Ensure that the NGT version is compatible with the AOS version installed in your cluster. For more information, see the [NGT](#) section in the *Compatibility and Interoperability Matrix*.
- Ensure that the cluster meets all NGT requirements. For more information, see [Nutanix Guest Tools Requirements](#) on page 190.
- Ensure that you registered the VMs with the third-party endpoint management tool deployed at your site. For more information, see the tool-specific documentation.

Procedure

- Perform Step 1 to Step 3 in [Installing NGT in Bulk on Windows VMs](#).
The installer detects that an existing NGT version is already present on the VM and performs an upgrade.

Upgrading NGT in Bulk on Linux VMs

About this task

This sample procedure provides steps to upgrade NGT on Red Hat Package Manager (RPM) based operating systems. Use this procedure as a template to upgrade NGT on Debian-based operating systems.

Before you begin

- Ensure that the NGT version is compatible with the AOS version installed in your cluster. For more information, see the [NGT](#) section in the *Compatibility and Interoperability Matrix*.
- Ensure that the cluster meets all NGT requirements. For more information, see [Nutanix Guest Tools Requirements](#) on page 190.
- Ensure that you registered the VMs with the third-party endpoint management tool deployed at your site. For more information, see the tool-specific documentation.

Procedure

1. To verify the package signatures, configure the third-party endpoint management tool deployed at your site to install the NUTANIX-NGT-GPG-KEY file on RPM-based operating systems.
For more information, see the tool-specific documentation.

2. Configure the third-party endpoint management tool to set up the repositories.

For more information, see the tool-specific documentation.

For example, guest VMs running RedHat 8.x distribution might need a repository configuration in `/etc/yum.repos.d/nutanix-guest-tools.repo` similar to the following configuration:

```
[NutanixGuestTools]
name=Nutanix Guest Tools
baseurl=http://local-web-server/ngt_repo
enabled=1
gpgcheck=1
gpgkey=http://local-web-server/RPM-GPG-PUBLIC-KEY
repo_gpgcheck=0
```

Replace `local-web-server` with the name of the web server used to distribute the NGT installer files as described in [Preparing the NGT installation files for the Distribution](#).

3. Configure the third-party endpoint management tool to upgrade the Nutanix guest agent package by running the package manager-specific upgrade command.

For example, the `yum update -y nutanix-guest-agent` command upgrades Nutanix guest agent package using the yum package manager for RedHat-based distributions.

```
[nutanix@localhost ~]$ yum update -y nutanix-guest-agent
```

For information about the upgrade command specific to the package manager at your site, see the package manager-specific documentation.

Uninstall NGT in Bulk from Multiple VMs

You can uninstall NGT in bulk from guest VMs using the third-party management tool deployed at your site.

Note: Before uninstalling NGT from a guest VM, ensure the communication link between guest VM and CVM remains active. The CVM displays that NGT is uninstalled only if the communication between CVM and guest VM is active during the uninstallation. If the communication link is down, CVM continues to display that NGT is installed on the guest VM even after the successful uninstallation of NGT. For information about how to verify that the communication link is active, see Step 4 in [Enabling NGT and Mounting the NGT Installer on Cloned VMs](#) of the *Prism Element Web Console Guide*.

Uninstalling NGT from Windows VMs

About this task

When you install NGT in a Windows VM, NGT registers an uninstaller in the VM that can be accessed manually using Add/Remove Programs in the Windows Control Panel. This uninstaller can also be used by a third-party endpoint management tool to remove NGT from the VM.

Before you begin

Ensure that you registered the VMs with the third-party endpoint management tool deployed at your site. For more information, see the tool-specific documentation.

Procedure

- Configure the third-party endpoint management tool to uninstall the Nutanix Guest Agent package using the uninstaller registered with Add/Remove Programs in the Windows Control Panel.
For more information, see the tool-specific documentation or Microsoft Windows documentation.

Uninstalling NGT from Linux VMs

About this task

This sample procedure provides steps to remove NGT on Red Hat Package Manager (RPM) based operating systems. Use this procedure as a template to remove NGT on Debian-based operating systems.

Before you begin

Ensure that you registered the VMs with the third-party endpoint management tool deployed at your site. For more information, see the tool-specific documentation.

Procedure

- Configure the third-party endpoint management tool to uninstall the Nutanix guest agent package in bulk by running the package manager-specific removal command.
For example, the `yum remove -y nutanix-guest-agent` command uninstalls Nutanix guest agent package using the yum package manager for RedHat-based distributions.

```
[nutanix@localhost ~]$ yum remove -y nutanix-guest-agent
```

For information about the removal command specific to the package manager at your site, see the package manager-specific documentation.

Nutanix Guest Agent and Controller VM Communication

You can set up connectivity between the controller virtual machine (CVM) and the guest VMs using the Nutanix Guest Tools (NGT) service that runs on the CVM and the Nutanix Guest Agent (NGA) service that runs on the guest VMs in the Nutanix cluster. After you install NGT in a guest VM, the NGA service in the guest VM starts periodic communication with the CVM based on the type of communication in a VM:

- **IP-less:** Communication between NGA and CVM happens only over a serial port connected to the guest VM.
- **IP-based:** NGA communicates with the CVM over SSL by connecting to the virtual IP of the CVM on port 2074.

By default on AHV clusters, the NGA service communicates with the CVM over IP-less connections using the serial port with index 1 on the guest VM. If the CVM does not respond, NGA reverts to IP-based communication.

For more information about how to check the communication type in a guest VM, to verify whether the NGA is communicating with the CVM, or to verify whether NGA is communicating with the CVM using IP or without an IP, see [Communication Type Verification](#).

The communication type used between NGA and the CVM is determined when the guest VM starts. NGA might select an alternative communication type following a failover event or an event that requires a VM restart, such as disaster recovery or high availability.

IP-less Communication

In IP-less communication, the NGA service in the guest VM starts periodic communication with the CVM over the serial port with index 1 on the guest VM. The NGA service communicates with the CVM by sending remote procedure calls (RPCs) over the serial port to the CVM. The communication link over the serial port becomes active when NGA receives a response from the CVM. If the CVM does not respond, NGA reverts to IP-based communication.

Note:

- The serial port with index 1 on the guest VM is reserved for IP-less communication using NGT.
- Only AHV clusters support IP-less communication.
- Nutanix recommends that the virtual IP port 2074 on the CVM is accessible so that NGA can revert to IP-based communication in case of any failures.

IP-less communication enhances the security of the communication link because it does not require the CVM to be accessible by the VM over the network, and identification of the VM is guaranteed by the hypervisor without in-guest certificates.

Prerequisites and Requirements

IP-less communication between the CVM and the guest VMs is established automatically when the following prerequisites and requirements are met:

- The guest VM must be running on AOS 6.6 or later, and NGT 3.0 or later versions.

- NGT is installed. For information about how to install NGT, see:
 - [NGT Installation](#) to install using Prism Element web console.
 - [Setting up NGT](#) on page 193 to install using Prism Central.
- If an older version of NGT is already installed on the guest VM, you must upgrade it to the latest version. For information about how to upgrade NGT, see:
 - [Upgrading NGT](#) to upgrade using Prism Element web console.
 - [Upgrading NGT](#) on page 196 to upgrade using Prism Central.
- To switch an existing guest VM from IP-based communication to IP-less communication, Nutanix recommends that you upgrade AHV and AOS to the latest version, then restart the VM when installing or upgrading NGT on the VM. If you do not restart the VM, the serial port remains deactivated and the NGA service continues to communicate with the CVM using the virtual IP port 2074 on the CVM. After NGT is installed or upgraded, you must power cycle the VM to activate the serial port that is attached to it before IP-less communication can be used. An in-guest reboot might not be sufficient.

IP-based Communication

With IP-based communication, after you install NGT in a guest VM, the NGA service in the guest VM starts periodic communication with the CVM over SSL connections.

Note: AHV, ESXi, and mixed clusters support IP-based communication.

Each Nutanix cluster is configured as a certificate authority (CA). When NGT is enabled in a VM, a certificate pair is generated specifically for that VM and it is embedded in an ISO that is configured for that VM. The security certificates are installed inside the VM as part of the installation process.

The NGA service inside the guest VM initiates an SSL connection to port 2074 of the cluster virtual IP address to communicate with the CVM. Any firewall must be configured to allow the guest VM to reach port 2074 on the cluster virtual IP address. For the complete list of required ports, see [Port Reference](#).

NGT provides three levels of security for VM communication:

- SSL certificates that ensure a secure TCP connection between the CVM and the guest VM
- Capability-based authorization that ensures only NGT features that are supported by the cluster are enabled on the guest VM.
- NGT ensures that the communication is established only if the system or BIOS UUID of the guest VM and the UUID of the guest VM (provided by the hypervisor) stored in the CVM are the same.

Communication is successful only if all the three conditions are met. For example, if a VM that has NGT installed is cloned, the new VM cannot communicate with the CVM. You must separately enable NGT, and mount the NGT installer on the cloned VM. The NGT installer includes configuration specific to the VM, and when the VM is re-configured and is able to communicate with the CVM, it automatically unmounts the installation media. You do not need to separately install NGT again on the cloned VMs.

NGA also publishes information about the VM to the CVM, for example, guest OS type, status of VM mobility, and VSS services.

Communication Type Verification

This section explains how to verify the communication type between Controller VMs (CVMs) and guest VMs using Prism Central or CLI. It covers IP-based and IP-less communication methods, verification steps, filtering options, and operational requirements to ensure secure, reliable connectivity for Nutanix Guest Tools (NGT) across supported environments.

Verifying Communication Type Using Prism Central

You can verify the type of communication established between the CVM and the guest VM using Prism Central.

About this task

To verify the communication type using Prism Central, perform the following steps:

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Compute > VMs** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.
The system displays the **List** tab by default with VMs across all the registered clusters.
3. Click **Modify Filters** option to access the list of available filters.
4. In the **NGT COMMUNICATION TYPE** filter, select the **IP-based** or **IP-less** checkbox to filter the VMs based on the NGT communication type.

Note:

- The **NGT COMMUNICATION TYPE** filter on a Prism Central hosted on an AOS 6.5.x cluster might in-correctly display the **IP-less** checkbox even though IP-less communication is not supported on AOS 6.5.x.
- To check the VMs on which NGT is installed, select the **Installed** checkbox in the **NUTANIX GUEST TOOLS** filter. The system displays the list of VMs on which NGT is installed.

The system displays the list of VMs based on the selected options.




Figure 52: IP-Based and IP-Less Filter

Verifying Communication Type Using CLI

You can verify the type of communication established between CVM and guest VM using CVM SSH (CLI) console.

About this task

To verify the communication type using the CVM SSH console, perform the following steps:

Procedure

1. Log in to any CVM in the cluster as a Nutanix user with SSH.
2. Run the following command:

```
nutanix@cvm $ nutanix_guest_tools_cli list_vm_tools_entities include_vm_info=true  
vm_name=vm-name
```

Replace *vm-name* with the name of the guest VM.

Observe the following attributes in the command output:

- `communication_link_active = true` - Indicates that the NGA is communicating with the CVM.
- `communication_link_over_serial_port_active = true` - Indicates that the communication between NGA and the CVM is IP-less.
- `communication_link_over_serial_port_active = false` - Indicates that the communication between NGA and the CVM is over IP.

Automatic CD-ROM Ejection

After NGT is installed or re-configured in a guest VM, the CD-ROM is auto-ejected from the VM in the following stages.

1. CD-ROM is automatically ejected from the VM immediately after the installation of NGT is successfully completed.
2. NGA indicates to the controller VM (CVM) that the installation is complete and the CD-ROM is detached from the VM through the hypervisor. This operation depends on the communication between the guest VM and the CVM and may take up to 10 minutes.

Note:

- These operations can occur at different times. It is possible that the CD-ROM shows up as empty inside a VM, but the status in the NGT-Controller VM service Tools Mounted might display as true. If this situation occurs, unmount the CD-ROM using the Prism Element web console.
- If the NGT software version in the ISO is more recent than the installed version in the guest VM, the CD-ROM does not eject automatically. This functionality enables the upgrade of the NGT software inside the VM.

NGT Usage in Disaster Recovery

If NGT is enabled on a guest VM that is protected by Nutanix Disaster Recovery, the snapshots of the VM include NGT relevant information along with its capabilities.

The NGT information is preserved for a restored guest VM in the following scenarios:

- Migrating the VM to the remote cluster (Planned Failover)
- Performing an in-place restore of the VM on the source cluster
- Creating a clone of the VM on the source cluster

In these scenarios, a new NGT ISO image containing only the relevant configuration information (SSL certificates, Controller VM IP address, and so on) is created for the recovered guest VM and the image is automatically attached to the VM. When the VM is powered on, the NGA service running on the VM copies the relevant configuration information and detaches the NGT ISO CD-ROM automatically.

The NGT information is not preserved for a restored guest VM in the following scenarios:

- Performing an out-of-place restore of the VM on the remote cluster.
- Retrieving and restoring the snapshot from the remote cluster to the source cluster.

In these scenarios, you must enable NGT again in the restored VM.

Storage Quality of Service (QoS)

Storage QoS provides administrators granular control to manage the performance of virtual machines and ensure that the system delivers consistent performance for all workloads. You can use a controllable knob to limit the IOPS that the storage layer would serve for individual virtual machines. IOPS is the number of requests the storage layer can serve in a second. You can set throttle limits on a VM to prevent noisy VMs from over-utilizing the system resources.

Note:

- Storage QoS is applicable only if you have an AOS Pro or above license. See the [License Manager Guide](#) for more information on AOS licenses.
- Setting storage QoS is not allowed when the AOS upgrade is in progress.

Terminology

Throttled IOPS

The maximum IOPS that the storage layer tries to admit over a sustained period. Applications that are running on the VM are not permitted to exceed this level. This prevents VMs from affecting a system beyond the set limits. However, if the system is in a state of flux, your application might not be throttled. To prevent application crashing due to misconfiguration, the throttled IOPS setting must not be less than one hundred. Otherwise, an error is raised.

Setting maximum IOPS on a VM may affect the cluster latency for that VM and may result in higher latency being displayed in the cluster latency charts. Performance of any other VM is not affected by this setting. You can view individual VM latency graph to ensure that the other VMs are not affected. You can also create a custom focus view to display the throttled IOPS for all the VMs.

Throughput

Throughput is the amount of data that an application sees from the storage layer underneath per second. The maximum throughput is calculated by multiplying the number of IOPS with the block size.

Calculation of the IOPS and Throughput

The default block size used for IOPS calculation is 32 KiB. If the block size is less than or equal to 32 KiB, then the system is rate limited by the number of IOPS and the bandwidth is calculated by multiplying the number of IOPS with the block size of the application. Otherwise, if the block size is greater than 32 KiB, the system is rate limited by the throughput and the number of IOPS is calculated by dividing the throughput by block size of the application.

For example, if you have set a throttle limit of 800 IOPS on a VM and the application is admitting IOPS with 8 KiB block size, then the application would see 800 IOPS and throughput of 6.25 MBps. However, if application is admitting IOPS with 64 KiB block size, then the application would see 400 IOPS and a throughput of 25 MBps.

The following table provides the information about the block size, IOPS, and throughput relationship when throttle limit is set to 800, 400, or 200 IOPS:

Table 38: Relationship between Block Size, IOPS, and Throughput

Block size	IOPS	Throughput
8 KiB	800	6.25 MBps
16 KiB	800	12.50 MBps
32 KiB	800	25 MBps
64 KiB	400	25 MBps
128 KiB	200	25 MBps

Limitations

Storage QoS has the following limitations:

- Storage QoS is not supported for an in-place restore, out-of-place restore, and snapshot operation on VM.
- While creating a clone of a VM, any QoS attributes throttle limit set on the original VM might not be applied to the new cloned VM. QoS throttle limit needs to be set on the new cloned VM separately.
- For linked clones, when you are sharing the vDisks across multiple VMs, shared disks are not a part of QoS policy and shared disks would not be throttled. Each VM might exceed the maximum limit depending on the shared vDisk usage. For example, during bootstorm in a VDI setup.
- If a VM has volume groups attached, then QoS is not applicable. This limitation of Storage QoS applies irrespective of whether QoS is defined directly on VMs or defined through a Storage Policy.
- Storage QoS is not supported for Metro Availability, synchronous replication, and Nutanix Files virtual machines.

Setting QoS for an Individual VM

About this task

Storage QoS allows you to set throttle limits on a VM to prevent noisy VMs from over-utilizing the system resources. Perform the following procedure to set the QoS for an individual VM.

To set QoS for an individual VM, perform the following steps:

Procedure

1. Log in to Prism Central.
 2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Compute > VMs** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.
The system displays the **List** tab by default with VMs across all the registered clusters.
 3. Perform either of the following actions:
 - » Select the target VM checkbox, and choose **Set QoS Attributes** from the **Actions** dropdown menu.
- Note:** This action is not permitted (greyed out) for the VM on which Prism Central is hosted.
- » Click the target VM to access the [Viewing VM Details](#) on page 113, and choose **Set QoS Attributes** from the **More** dropdown menu.

4. In **Input Type** field, select either of the following options:

- » IOPS - For more information about IOPS, see *Throttled IOPS* information in [Storage Quality of Service \(QoS\)](#) on page 210.
- » Throughput - For more information about throughput, see *Throughput* information in [Storage Quality of Service \(QoS\)](#) on page 210.

If you select:

- **IOPS**- Enter a value for the throttled IOPS in the **Throttled IOPS** field.
 - **Throughput** - Enter a value for the throttled throughput in the **Throttled Throughput** field.
5. (Optional) Click **Show Details** to populate a table that describes the relationship between block size, IOPS, and throughput.
6. Click **Confirm** to set the defined QoS values for the target VM.
- Click **Clear Attributes** to delete all the values.
- The system stores the number of IOPS or throughput to be served by the storage layer.

Note: To edit the storage QoS setting, follow the same steps as described in this section, and enter a new value for **IOPS** or **Throughput**.

Memory Overcommit Management

The Memory Overcommit feature enables you to optimize the physical memory utilization of a VM. It allows the host to automatically detect whether the memory is under-utilized or over-utilized for a VM. Using Prism Central, you can enable the Memory Overcommit feature.

For more information about Memory Overcommit feature, and how to enable or disable it using Prism Central, see [Memory Overcommit](#) information in *AHV Administration Guide*.

Policies for VM Management

You can define VM-Host Affinity, VM-VM Anti-Affinity, and NGT policies for managing a VM. For information about Affinity and NGT policies, see [VM Policy Management](#) on page 501.

VM Template Management

In Prism Central, you can create VM templates to manage the golden image of a VM.

A VM template can be considered as a golden image or a master copy of a virtual machine. It captures the virtual machine configuration and the contents of the virtual machine including the guest operating system, and the applications installed on the virtual machine. You can use this single template to deploy multiple, consistent VMs across various clusters in your environment. You can create template placement policies to manage which categories of templates can be deployed on which clusters, providing more granular control over resource usage. For more information on template placement policies, see .

Note: You can create or manage a VM template only as an admin user or a user with admin privileges.

VM Template Summary View

The **Templates** page displays the information about all the VM templates across registered clusters.

To access the summary view of all VM templates, perform the following steps:

1. Log in to Prism Central.

2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Compute > Templates** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

The following is an example showing the **Templates** page with all the VM templates:



Figure 53: Summary View - All VM Templates

Table 39: Templates Page - Field Description

Field	Description	Values
Name	Displays the template name. Click the template name to view the detailed information about that template. For more information, see VM Template Details View on page 214.	<Template name>
Compute	Displays the virtual CPU count of the VM that is deployed using this template.	(vCPU count) x vCPU, where x is numeric digit.
Memory	Displays the total amount of memory available to the VM that is deployed using this template.	(Memory) x [MB GiB], where x is numeric digit.
Storage	Displays total amount of storage available to the VM that is deployed using this template.	x [MB GiB], where x is numeric digit.
Network	Displays the NIC count of the VM that is deployed using this template.	x NIC, where x is numeric digit.
Updated On	Displays the last date and time when the template has been updated.	(date and time), Timestamp in MM DD, YYYY, hh:mm AM/PM format. For example, Dec 27, 2022, 06:51 AM
Updated By	Displays the user who most recently updated the template.	<user>. For example admin
Active Version	Displays the name of the active version of the template. An active version is the version of the template that gets deployed by default when you click Deploy VMs .	(version). For example, Initial Version

You can perform the following actions for the VM templates in the **Templates** page:

- Access the detailed information about an individual VM template. For more information, see [VM Template Details View](#) on page 214.

- Deploy a VM. For information about how to deploy a VM using VM template, see [Deploying VM from a Template](#) on page 220.
- Filter the VM templates list based on the template name using **Filters** pane. For more information about **Filters** pane, see [Filters Pane](#) on page 50.
- Sort the list of templates from the **Name** and **Updated On** columns. For information about sorts function, see [Sort Function](#) on page 50.
- Use the **Actions** dropdown menu to perform the available operations as described in [Managing a VM Template](#) on page 223.

VM Template Details View

The VM template details view includes two pages: **Summary** page and **Versions** page.

Summary Page

The **Summary** page of an individual VM template consists of a dashboard that provides the detailed information about the VM.

To access the **Summary** page of an individual VM template:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Compute > Templates** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.
3. Click the `<VM_Template_Name>` to view the **Summary** page of an individual VM template.

Note: Replace `<VM_Template_Name>` with the actual VM name at your site.

The following is an example showing the **Summary** page of an individual VM template:




Figure 54: Summary Page - Individual VM Template

The **Summary** page provides you the options to deploy a VM using the selected VM template, and a horizontal ellipses dropdown menu for actions.

- For information about how to deploy a VM using the VM template, see [Deploying VM from a Template](#) on page 220.
- For information about the list of actions available under **Actions** dropdown menu, see [Managing a VM Template](#) on page 223.

The **Summary** page of an individual VM template provides the following widgets:

- **Template:** This widget displays the following parameters:
 - **Available Versions:** Number of versions available
 - **Description:** User-provided description of the template
 - **Last Updated by:** User who has last updated the template
 - **Last Updated on:** Date and time of last update to the template
- **Active Version Configuration:** This widget displays the following parameters:
 - **Version Name:** Name of the currently active version of the template
 - **Change Notes:** User-provided change notes for the active version
- **VM Properties:** This widget displays the following parameters:
 - **Compute:** Compute resources of the VM such as CPU, core per CPU, and memory
 - **Boot Type:** Boot type of the VM such as legacy or UEFI
 - **Guest OS:** Guest OS of the VM. If the source VM of the template does not have NGT installed on it, then this field is displayed as unknown.
 - **Guest Customization:** Status of the guest customization application (applied/not applied)
 - **NGT Status:** Status of the NGT installation
- **VM resources:** This widget displays the following parameters:
 - **GPU:** Number of GPU assigned to the VM
 - **Disks:** Number of disks and total disk capacity assigned to the VM
 - **Network:** Number of NICs assigned to the VM

Versions Page

To view the **Versions** page of an individual VM template, click **Versions**.

The following is an example showing the **Versions** page of an individual VM template:




Figure 55: Versions Page - VM Template

You can perform the following actions for the VM template version in the **Versions** page:

- Deploy a VM using the selected VM template version. For information about how to deploy a VM using the selected VM template version, see [Deploying VM from a Template](#) on page 220.

Note: By default the active version of the template always gets deployed. In the **Versions** page, you can select the specific VM template version, and make it Active for VM deployment.

- Select a version and set it as active, or delete a selected version from the **Actions** dropdown menu. For more information, see [Managing a VM Template](#) on page 223.
- Filter the VM templates list based on the template name using **Filters** pane. For more information about **Filters** pane, see [Filters Pane](#) on page 50.
- Sort the list of templates from the **Name** and **Updated On** columns. For information on sorts function, see [Sort Function](#) on page 50.

Table 40: Versions Fields

Parameter	Description	Values
Name	Displays the VM Template version name. This column also indicates if the version is an active version.	<Template _Version_name>
Compute	Displays the virtual CPU count of the VM that gets deployed using this version.	(vCPU count) x vCPU, where x is numeric digit.
Memory	Displays the total amount of memory available to the VM that gets deployed using this version.	(Memory) x [MB GiB], where x is numeric digit.
Storage	Displays total amount of storage available to the VM that gets deployed using this version.	x [MB GiB], where x is numeric digit.
Network	Displays the NIC count of the VM that gets deployed using this version.	x NIC, where x is numeric digit.
Added On	Displays the last date and time when the version was added.	(date and time), Timestamp in MM DD, YYYY, hh:mm AM/PM format. For example, Dec 27, 2022, 06:51 AM

Parameter	Description	Values
Added By	Displays the user who added the version.	<user>. For example admin
Notes	Displays the user-specified change note for the version.	<Version-specific-text> that you specify in Change Notes field when you update a VM template version.

You can sort the list of versions by clicking on **Name** and **Added On** columns.

Limitations of VM Template Feature

The current implementation of the VM template feature has the following limitations:

- You cannot create a VM template if any of the following conditions is applicable at your site:
 - VM is not on AHV.
 - VM is an agent or a PC VM.
 - Any volume group is attached to the VM.
 - VM is undergoing vDisk migration.
 - VM has disks located on containers with replication factor 1.
 - VM is protected by PD-based DR.
- VM templates do not copy the following attributes from the source VMs:
 - Host affinity attributes
 - Nutanix Guest Tools (NGT) installation
 - Quality of service (QoS) configuration

You must reconfigure the above-listed attributes at the deployed VMs.

- VM Template does not protect the underlying VM recovery point from deletion by a user having delete permission. VM Template deployment fails if you delete the associated VM recovery point.

Creating a VM Template

This section describes how to create a VM template in Prism Central.

About this task

To create a VM template, perform the following steps:

Procedure

- Log in to Prism Central.
- Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Compute > VMs** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.
The system displays the **List** tab as default

3. In **List** tab, select the VM to create a VM template.

Note: Before you select a VM to create a VM template, ensure that the VM is powered off.

4. In the **Actions** dropdown menu, select **Create VM Template**.

You can also click **Create VM Template** in the **Summary** page of an individual VM. For information about how to access the **Summary** page of an individual VM, see [Viewing VM Details](#) on page 113.

The system displays the **Create template from VM** window.

The screenshot shows the 'Create Template from VM' window. It includes a title bar, a close button, and several input fields. The 'VM Template Name' field contains 'win10-template'. The 'Description' field is empty. The 'Guest Customization' section has a sub-header 'Specify guest OS customization during VM Deployment'. Below this are two dropdown menus: 'Sysprep (Windows)' and 'Custom Script'. There is an 'Upload Script' button and a text area for 'Upload or paste script here'. At the bottom, there is a toggle for 'Allow users to override at VM Deployment?' set to 'Yes'. A 'Next' button is located at the bottom right.

Figure 56: Create Template from VM window

5. In the **Create Template from VM** window, enter the following information:
 - a. **Name** - Name of the VM Template.
 - b. **Description** - Description for the VM template. Description is an optional field.
 - c. For the **Guest Customization** fields, select the following options for guest operating system (OS) for the VMs that you deploy using this VM template:
 - In the **Script Type** field, select **Sysprep (Windows)** to customize the Windows OS, and **Cloud-init (Linux)** to customize the Linux OS.

Note: If you select **No Customization** at the time of creating the template and allow the users to override the guest customization settings using **Allow users to override at VM Deployment?** toggle field, it

gives the maximum customization control to the users. In this case, the users can customize the script type and the configuration method.

- In the **Configuration Method** field, for each of these script types selected in **Script Type** field, select either upload a custom script or opt for a guided setup in the field.

Note:

- If you select **Custom Script**, you can either upload a script to customize the guest OS of the VMs, or you can copy-paste the script in the text box.
- If you select **Guided Script**, enter the following information:
 - **Authentication Type:** Select one of the radio button to set the authentication type:
 - **Password:** Set a username and password for the user who uses this template to deploy the VM.
 - **SSH Key** [Cloudinit (Linux) only]: Set the SSH key for the user who use this template to deploy the VM.
 - **Locale:** Select the locale (language) from the dropdown list.
 - **Hostname:** Enter the hostname for the user who uses this template to deploy a VM
 - **License Key:** Enter the license key.

Important: The information that you enter is used to customize the OS of the VMs that are deployed using this template.

Note: If you opt for **Guest Customization** setup with script (custom or guided script), ensure that the script is in a valid format. The system does not validate the guest customization scripts. In the script is not in valid format, the VM deployment might succeed but the guest customization script may not work as expected. You can observe the discrepancies (if any) only after the VM gets deployed.

- d. Set the **Allow users to override at VM Deployment?** toggle field as per your requirement. This toggle field is used to enable or disable the override permission for guest customization settings.

- If this toggle field is:
 - *Enabled:* The system allows the users (who use this VM template to deploy VM) to modify the guest OS customization settings in VM template. The users can modify the settings only for the **Configuration Method** field. For example, the users can change the authentication information at the time of deploying a VM from VM template, or they can change from a guided setup to a custom script.
 - *Disabled:* The settings that are already provided in the VM template, can be used for VM deployment.

6. Click Next.

On the next page, you can review the template details.

The system prompts you to review the configuration details, resource details, network details, and management details.

7. Click **Save** to save the inputs and create a VM template.

The new template appears in the **Templates** page list.

Note: Once you create a VM template, the template metadata is available in Prism Central. The template data is stored as a VM recovery point and is co-located with the source VM. If you use the VM template to deploy a VM to another cluster, the recovery point is copied to the destination cluster before deployment.

Deploying VM from a Template

Before you begin

- Ensure that the VM template is available. For information on how to create a VM template, see [Creating a VM Template](#) on page 217.
- To deploy a VM on a remote cluster, ensure that the ports and protocols requirements listed for Disaster Recovery (DR) are met. For more information, see [Ports and Protocols for DR](#).

About this task

After you create a VM template, you can use the VM template to deploy any number of VMs across clusters.

To deploy a VM using a VM template, perform the following steps:

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Compute > Templates** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61 .
3. Select the target VM template to deploy VMs using either of the following methods:
 - In the **Templates** page, select the target template checkbox, and click **Deploy VMs**. For more information about **Templates** page, see [VM Template Summary View](#) on page 212.
 - In the **Summary** page of an individual VM template, click **Deploy VMs**. For more information about **Summary** page of an individual VM template, see [VM Template Details View](#) on page 214.
 - In the **Versions** page of an individual VM template, select the target VM template version, and click **Deploy VMs**. For more information on **Versions** page of an individual VM template, see [VM Template Details View](#) on page 214.

Note: In the **Versions** page, you can select any active or non-active VM template version for VM deployment.

Note: By default, the active version of the VM template is used for VM deployment.

The **Deploy VM from Template** window appears. By default, you see a **Quick Deploy** method. You can click **Advanced Deploy** to access the **Advanced Deploy** method in **Deploy VM from Template** window . In **Advanced Deploy** method, you can view and modify some VM properties and network settings.

4. Deploy VM using either of the following deployment methods:

- **Quick Deploy** method:

There is an example showing the **Quick Deploy** method:




Figure 57: Deploy VM from VM Template using the Quick Deploy Method

1. Enter the following information:

- **Name:** Enter a name for the VM.
- **Cluster:** Select the cluster where you want to deploy the VM.
- **Number of VMs:** Enter the number of VMs that you want to deploy.
- **Starting Index Number:** Enter the starting index number for the VMs when you are deploying multiple VMs simultaneously. These index numbers are used in the VM names. For example, if you are deploying two VMs and specify the starting index number as 5, the VMs are named as `vm_name-5` and `vm_name-6`.
- **Guest Customization:** The template can have any one of the following options for guest OS customization:
 - **No Customization**
 - **Sysprep (Windows), or Cloud-init (Linux).** For **Sysprep (Windows)**, or **Cloud-init (Linux)**, you can choose to either upload a custom script or opt for a guided setup.

Note: These fields are enabled for modification only if the VM template allows you to override its guest customization settings while deploying the VM. The **Allow users to override at VM Deployment?** toggle field is used to enable or disable the override for guest customization settings. For information on how to set this toggle field during VM template creation, see [Creating](#)

[a VM Template](#) on page 217. If the override permission for the guest OS customization is disabled, the settings that are already provided in the template can be used for VM deployment.

2. Click **Next** to verify the configuration details of the VMs to be deployed.
3. Click **Deploy** to deploy the VMs.

- **Advanced Deploy** method:

The following is an example showing the **Advanced Deploy** method:

The screenshot shows the 'Deploy VM from Template' wizard with the 'Configuration' tab selected. The wizard has four steps: Configuration, Resources, Management, and Review. The Configuration tab contains the following fields and options:

- Name:** A text input field.
- Description:** A text input field with the placeholder text '(Optional)'.
- Cluster:** A dropdown menu showing 'auto_cluster_nested_61afc380572f30dec6788bc'.
- Number of VMs:** A text input field with the value '1'.
- Starting Index Number:** A text input field with the value '0'.
- VM Properties:** A section with three columns: CPU, Cores Per CPU, and Memory.
 - CPU:** A text input field with the value '1' and a dropdown menu showing 'vCPU'.
 - Cores Per CPU:** A text input field with the value '1' and a dropdown menu showing 'Cores'.
 - Memory:** A text input field with the value '2' and a dropdown menu showing 'GB'.
- Enable Memory Overcommit:** A checkbox that is currently unchecked.
- Buttons:** 'Back to Quick Deploy', 'Cancel', and 'Next'.

Figure 58: Deploy VM from VM Template using the Advanced Deploy Method

1. Enter the following information:

- **Configuration:** Provide inputs for name and description (optional) of the VM, cluster where you want to deploy the VM, Number of VMs to be deployed, and starting index number (only if deploying multiple VMs). In this tab, you can also view and modify the VM properties such as CPU, core per CPU, and memory.
- **Resources:** Review the configuration settings for the VM resources such as disks, networks, and boot configuration. Here, you can modify the network settings but cannot modify any other settings.
- **Management:** The fields in this tab are enabled for modification only if the VM template allows you to override its guest customization settings while deploying the VM. The **Allow users to override at VM Deployment?** toggle field is used to enable or disable the override for guest customization settings. For information on how to set this toggle field during VM template creation, see [Creating a](#)

[VM Template](#) on page 217. If the override permission for the guest OS customization is disabled, the settings that are already provided in the template can be used for VM deployment.

If guest customization has been enabled, provide inputs for authentication type, username, password, locale, and Hostname.

Note: Hostname of the VM is automatically generated based on the VM names that you provide. If you are deploying a single VM, you can override the automatic generation of the hostname by specifying a hostname. If you are deploying multiple VMs, you cannot override the automatic hostname generation.

- **Review:** Review the information displayed in this tab.
2. Click **Deploy** to deploy the VMs.

Managing a VM Template

About this task

After you create a VM template, you can perform the following actions to manage the VM Template:

- Update the guest OS of the source VM specified in the VM template
- Complete guest OS update
- Cancel guest OS update
- Update the configuration of the template to create a new VM template version
- Delete the VM template.

For information on how to create a VM template, see [Creating a VM Template](#) on page 217.

To manage a VM template, perform the following steps:

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Compute > Templates** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61 .

Select the target VM template using either of the following methods:

- Select the target VM template checkbox in the **Templates** page. For more information on **Templates** page, see [VM Template Summary View](#) on page 212.
 - Click the target template to view the **Summary** page of an individual VM template. For more information about **Summary** page of an individual VM template, see [VM Template Details View](#) on page 214.
3. Click the horizontal ellipsis (three-dot icon), and select the required action from the dropdown menu:

Note: The available actions appear in bold and the unavailable actions are greyed out. The available actions depend on the current state of the template and user permissions.

The following actions are available in horizontal ellipsis (three-dot icon) drop down menu:

- **Update Guest OS**

To update the guest OS in a VM Template, perform the following steps:

1. Select the target VM template version on which the guest OS is to be updated in **Select a version to Update** window, and click **Proceed**.

Note: By default the active version of the VM Template is selected for VM Template update, however the system provides you an option to select the VM Template version on which the guest OS is to be updated.

2. Review the information displayed in the **Update Guest OS** window, and click **Proceed**.

At this stage, the system deploys a temporary VM from the selected VM template version, and provides you an option to access the VM. You can also access the new VM from VMs **List** tab. For information on VMs **List** tab, see [VMs Summary View](#) on page 100.




Figure 59: Update Guest OS

3. Start the temporary VM, log on to the VM, and update the guest OS of the temporary VM. For information on how to update a VM, see [Managing a VM through Prism Central \(AHV\)](#) on page 161.
 4. Complete the guest OS update using the **Complete Guest OS Update** option from the **Actions** dropdown menu.
- **Complete Guest OS Update** to complete the process initiated for guest OS update. You must select this option only after successful update of the guest OS of the temporary VM. If you have opted for guest customization for

the deployed VMs, ensure to run Sysprep (for Windows) or Cloud-init (for Linux) before completing the guest OS update.

Note: At this stage, the system prompts you to create a new version of the VM Template with updated guest OS. Specify the details for the new version, and click **Complete Update**.

The temporary VM automatically gets deleted after completion of the guest OS upgrade process.

- **Cancel Guest OS Update** to cancel the process initiated for guest OS update.

The temporary VM automatically gets deleted after cancellation of the guest OS upgrade process.

- **Update Configuration** to modify the VM template configuration. In **Select a version to Update** window, select the VM template version that you want to modify, and create a version on top of it.

Perform the following steps to update the VM template configuration in **Update Template Configuration** window:



Figure 60: Update Template Configuration

1. In the **Configuration** step, view the name of the base version that you want to update, change notes for that version, cluster name, VM properties (CPU, cores per CPU), and memory). In this section, you can modify only VM properties.
2. In the **Resources** step, view the information about disks, networks, and boot configuration. In this section, you can modify only network resources.
3. In the **Management** step, modify the guest customization settings.
4. In the **Review** step, review and modify the configuration settings that you are allowed to modify. You must provide a name and change notes for the new version. You can also choose to set this new version as active version.

Note: An active version is the version of the template gets deployed by default when you click **Deploy VMs** after VM template configuration update.

5. Click **Save** to save the settings and create a new VM template version.
- **Delete Template** to delete a template. The system prompts you to confirm the delete action. Click **OK** to delete the VM template.

4. To only manage a VM template version:

- a. Select the VM template version from the **Versions** page. For information on how to access the **Versions** page of an individual VM template, see [VM Template Details View](#) on page 214.
- b. Select either of the following options from the **Actions** dropdown menu:
 - **Set as Active** to make the selected VM template version as the active version.
 - **Delete** to delete the selected VM template version.

Note: You cannot delete the active version of a VM template.

Kubernetes Clusters Management

The **Kubernetes Clusters** entity in the **Infrastructure** application allows you to see the Kubernetes clusters that you deploy on the Nutanix clusters.

To see the Kubernetes clusters in the **Kubernetes Clusters** entity in the **Infrastructure** application, you must onboard the Kubernetes clusters to Prism Central. See [Onboarding the Kubernetes Cluster to Prism Central](#) on page 227 for more information.

Before you onboard any Kubernetes clusters to Prism Central, the **Kubernetes Clusters** entity page is displayed as follows:




Figure 61: Kubernetes Clusters entity page before onboarding Kubernetes clusters

If you do not want to manage the Kubernetes clusters in Prism Central **Infrastructure** application, off-board the Kubernetes clusters. To off-board the cluster, use the `helm uninstall <helm-chart-name>` command. See the *What to do next* section in [Onboarding Kubernetes Cluster to Prism Central](#) topic in *Nutanix Data Services for Kubernetes*.

Kubernetes Cluster Entity Views

Prism Central provides the following views for the **Kubernetes Clusters**.

- [Kubernetes Clusters Summary View](#) on page 227

- [Kubernetes Cluster Details View](#) on page 228

Onboarding the Kubernetes Cluster to Prism Central

After creating the Kubernetes clusters with NKE Next Generation Workload Clusters (In Technical Preview) or Red Hat OpenShift, you need to onboard the clusters to Prism Central. After onboarding, you can view the Kubernetes clusters on Prism Central.

For information on NKE Next Generation Kubernetes cluster (In Technical Preview), see [Next Generation Workload Cluster \(Technical Preview\)](#) in the *Nutanix Kubernetes Engine Guide*.

For information on onboarding the Kubernetes cluster to Prism Central, see [Onboarding Kubernetes Cluster to Prism Central](#) in the *Nutanix Data Services for Kubernetes*.

Kubernetes Clusters Summary View

The **List** tab in the **Kubernetes Clusters** summary view provides a list of the Kubernetes clusters that you onboard to the Prism Central instance. To access the **Kubernetes Clusters** summary view, perform the following steps:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Compute > Kubernetes Clusters** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

After you onboard the Kubernetes clusters, the **List** tab displays the onboarded Kubernetes clusters as follows.

Name	Size	Status	Version	Distribution	Node VMs	vCPU Allocated	Memory Allocated	Storage Usage (GiB)	Project
nke-s-4	+	Connected	v1.20.6	Nutanix Kubernetes Engine	3	20	28 GB	0 B of 0 B	system
nke-s-3	+	Connected	v1.27.6	Nutanix Kubernetes Engine	3	20	28 GB	0 B of 0 B	system
onboarding-4	+	Connected	v1.28.3	CNCF Kubernetes	4	8	16 GB	0 B of 0 B	internal
onboarding-1	+	Connected	v1.28.3	CNCF Kubernetes	4	8	16 GB	17 GB of 18 GB	system
onboarding-5	+	Connected	v1.28.3	CNCF Kubernetes	4	8	16 GB	0 B of 0 B	system
nke-s-1	+	Connected	v1.27.6	Nutanix Kubernetes Engine	3	20	28 GB	0 B of 0 B	system
h-ocp-2	+	Connected	v1.26.9+cnf05e7	Red Hat OpenShift	6	36	96 GB	0 B of 0 B	system
onboarding-3	+	Connected	v1.28.3	CNCF Kubernetes	4	8	16 GB	17 GB of 18 GB	system
nke-s-mgmt	+	Connected	v1.26.8	Nutanix Kubernetes Engine	0	0	0 B	0 B of 0 B	-
nke-s-2	+	Connected	v1.27.6	Nutanix Kubernetes Engine	3	20	28 GB	0 B of 0 B	system
h-ocp-1	+	Connected	v1.26.9+cnf05e7	Red Hat OpenShift	6	36	96 GB	0 B of 0 B	system-px-c

Figure 62: Kubernetes Clusters entity page after onboarding Kubernetes clusters

The following table describes the fields that appear in the **List** tab of the **Kubernetes Clusters** page. A dash (-) is displayed in a field when a value is not available or applicable.

Table 41: Kubernetes Clusters - List Tab Field Description

Field	Description	Values
Name	Displays the name of the Kubernetes cluster.	String

Field	Description	Values
State	Displays the state of the connection to the Kubernetes cluster onboarding agent.	Connected or Disconnected
Version	Displays the version of Kubernetes bundle used to deploy the Kubernetes cluster.	<version>. For example, 1.27.6
Distribution	Displays the name of the tool used to create the Kubernetes.	CNCF Kubernetes Red Hat OpenShift Nutanix Kubernetes Engine
Node VMs	Displays the number of Node VMs in the Kubernetes cluster.	Integer
vCPU Allocated	Displays the number of vCPUs allocated to the Kubernetes cluster.	Integer
Memory Allocated	Displays the amount of memory in GB allocated to the Kubernetes cluster.	(Decimal number) GB
Storage Usage (VGs)	Displays the amount of Volume Group storage in bytes used by the Kubernetes cluster.	(Decimal number) B of (Decimal number) B
Project	Displays the name of the Project that manages the Kubernetes cluster. Displays Internal if the Kubernetes cluster infrastructure resources are assigned to the default project VM.	String

Kubernetes Cluster Details View

The Kubernetes cluster details view displays the properties of individual Kubernetes cluster listed on the **List** tab of the **Kubernetes Clusters** summary view.

The Kubernetes cluster details view page provides information about the Kubernetes cluster on the following tabs:

- **Summary** tab—This tab is the default landing tab of the details view page.
- **Nodes** tab—This tab provides details of all the Kubernetes nodes in the Kubernetes cluster.
- **Storage** tab—This tab provides the details of the storage (Volume Groups or VGs) associated with the Kubernetes cluster.
- **Alerts** tab—This tab provides the details of the infrastructure alerts raised for the Kubernetes cluster.

To access the details view of an Kubernetes cluster, perform the following steps:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Compute > Kubernetes Clusters** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.
3. Click the **Name** of the Kubernetes cluster to view the **Summary** tab of the details view.

Summary Tab

The following is an example showing the **Summary** tab of an individual Kubernetes cluster.




Figure 63: Kubernetes Cluster - Summary Tab

The **Summary** tab of an individual Kubernetes cluster provides the following widgets:

- **Properties:** This widget provides the following parameters.
 - **Connection Status:** The state of the connection to the Kubernetes cluster onboarding agent.
 - **Kubernetes Version:** The version of Kubernetes bundle used to deploy the Kubernetes cluster.
 - **Kubernetes Distribution:** The name of the Kubernetes distribution used to create the Kubernetes cluster.
 - **OS Image:** The distribution and version of the operating system underlying the Kubernetes cluster. For example, Ubuntu 22.04.3 LTS.
 - **Project:** The name of the project that is configured in the **Prism Central Admin Center** application, to manage the infrastructure resources utilized by the Kubernetes cluster.
 - **Nutanix Add-Ons:** The number of the Nutanix add-on packages such as Nutanix CSI or any Container Network Interface (CNI) packages.
- **Infrastructure Alerts:** This widget displays the purpose of this widget in a banner. This widget provides the following parameters for the alerts from the underlying VMs and VGs for this Kubernetes cluster.
 - **Critical Alerts:** The number of critical alerts generated in the AHV host for the VMs and VGs associated with the Kubernetes cluster.
 - **Warning Alerts:** The number of warning alerts generated in the AHV host for the VMs and VGs associated with the Kubernetes cluster.
 - **Info Alerts:** The number of informational alerts generated in the AHV host for the VMs and VGs associated with the Kubernetes cluster.

For more information about alerts, see the [Prism Central Alerts and Events Reference Guide](#).

- **Storage:** This widget provides the **Number of Volume Groups** associated with the Kubernetes cluster. It also provides a clickable link that opens the **Storage** tab in the Kubernetes cluster details view.

- **Compute Resources:** This widget provides the following parameters about the VMs associated with the Kubernetes cluster.
 - **Total Number of Node VMs:** The total number of VMs on the AHV host associated with the Kubernetes cluster.
 - **vCPUs Allocated:** The number of vCPUs allocated to the Kubernetes cluster.
 - **Memory Allocated:** The amount of memory allocated (in GB) to the Kubernetes cluster.

The widget also provides a clickable link that opens the **Storage** tab in the Kubernetes cluster details view.

Nodes Tab

The **Nodes** tab provides information about the AHV hosts associated with the individual Kubernetes cluster. The following is an example showing the **Nodes** tab of the Kubernetes cluster.

VM Name	IP Addresses	vCPUs	CPU Usage	Memory Usage	Disk Usage	Prism Element Cluster
k8s-john-icp-dgfm	10.47.30.43, 10.47.30.45	2	11.99%	0 of 4 GB	5 of 40 GB	prism-element-cluster
k8s-john-wend-z6uik-gtazl	10.47.30.56	2	4.95%	0 of 4 GB	5 of 40 GB	prism-element-cluster
k8s-john-wend-z6uik-kazjd	10.47.30.58	2	3.85%	0 of 4 GB	5 of 40 GB	prism-element-cluster
k8s-john-wend-z6uik-zzp0k	10.47.30.57	2	3.84%	0 of 4 GB	5 of 40 GB	prism-element-cluster

Figure 64: Kubernetes Cluster - Nodes Tab

Table 42: Nodes Tab Field Descriptions

Parameter	Description	Values
VM Name	Displays the name of the VM in the AHV host, that is associated with the individual Kubernetes cluster.	(String)
IP Addresses	Displays the IP addresses of the VM.	Comma-separated list of IP addresses in XXX.XXX.XXX.XXX without CIDR prefix.
vCPUs	Displays the number of vCPUs associated with the AHV host.	Integer
CPU Usage	Displays the CPU usage in percentage.	Decimal number with two decimal places suffixed with %.

Parameter	Description	Values
Memory Usage	Displays the usage of the allocated memory.	<Usage> of <Allocated> memory suffixed with GB.
Disk Usage	Displays the usage of the allocated disk space.	<Usage> of <Allocated> disk space suffixed with GB.
Prism Element Cluster	Displays the name of the Prism Element cluster.	(String)

Storage Tab

The **Storage** tab provides the storage allocation and usage information associated with the individual Kubernetes cluster. The following is an example showing the **Storage** tab for the Kubernetes cluster.

Name	Space Usage	Disks	Connections	IOPS	IO Bandwidth	IO Latency	Cluster Name
pvc-01d2f57b-7ada-4c38-bc01-27cd3d320...	1.6 MB of 3.0 GB	1	1				auto_cluster_prod...
pvc-08250326-7a4a-410b-8a0c-e2778a19c...	1.6 MB of 3.0 GB	1	1				auto_cluster_prod...
pvc-2c0ab663-311f-48ff-9c63-11433a14ac96	1.6 MB of 3.0 GB	1	1				auto_cluster_prod...
pvc-61c5f1c9-lac9-4197-93a6-4e0e7eaahd17	1.6 MB of 3.0 GB	1	1				auto_cluster_prod...
pvc-6c34a042-5b1c-423c-9bda-e2b908c5...	1.6 MB of 3.0 GB	1	1				auto_cluster_prod...
pvc-9020660b-f6b3-4bf1-b574-b1eccc700a...	1.6 MB of 3.0 GB	1	1				auto_cluster_prod...
pvc-bbd97855-007a-4d67-3263-0a86f591...	1.6 MB of 3.0 GB	1	1				auto_cluster_prod...
pvc-be780efb-c5ac-406f-83ab-d0624e009...	3.1 MB of 6.0 GB	1	1				auto_cluster_prod...
pvc-c94a3663-2c74-44b0-8a9b-cbe05c58...	1.6 MB of 3.0 GB	1	1				auto_cluster_prod...
pvc-d88c5db9-4253-4299-bc90-af02a026...	3.1 MB of 6.0 GB	1	1				auto_cluster_prod...
pvc-d34a7582-227e-4e5b-a65d-60a65b16...	1.6 MB of 3.0 GB	1	1				auto_cluster_prod...
pvc-e72f938d-ccc5-43e1-a18a-44fb3b6406...	1.6 MB of 3.0 GB	1	1				auto_cluster_prod...

Figure 65: Kubernetes Cluster - Storage Tab


Table 43: Storage tab Field Descriptions

Parameter	Description	Values
Name	Displays the name of the Volume Group (VG).	Volume Group name
Storage usage	Displays the usage of the VG storage.	XX GB of <total> capacity GB
Disks	Displays the number of disks associated with the VG.	(Integer number)
Connections	Displays the number of connections to the VG.	(Integer number)
IOPS	Displays the IOPS of the VG.	(Integer number)
IO Bandwidth	Displays the Input Output (I/O) Bandwidth in MBps.	(Integer number)

Parameter	Description	Values
IO Latency	Displays the I/O latency.	Time in ms
Cluster Name	Displays the name of the cluster on which the VG is located.	(Cluster name)

Alerts Tab

The **Alerts** tab provides the alerts information associated with the individual Kubernetes cluster. The following is an example showing the **Alerts** tab for the Kubernetes cluster.



Details	Entity	Impact Type	Severity	Status	Created Time	Last Occurred
nutest-capx-10141220-kcp-wwzmz - VM CP...	nutest-capx-10141220-kcp-wwzmz	Performance	Critical	-	10/24/2023	10/26/2023
test-alert - VM CPU Usage for 'entity_name...	test-alert	Performance	Critical	-	10/24/2023	10/26/2023

Figure 66: Kubernetes Cluster - Alerts Tab

The Alerts tab shows an extract of the entity's alert tab. For example, the alert page shows the alerts for the VM that the workload nodes are created on. This tab provides the same features and options as the Alerts dashboard, however it is filtered to display the alerts only for the selected Kubernetes cluster. For more information about alerts, see [Prism Central Alerts and Events Reference Guide](#).

Table 44:

Parameter	Description	Values
Details	Displays the name of the entity and type of alert.	Concatenated value of entity name (for example, VM name) and the alert name.
Entity	Displays the name of the entity (VM or VG) that the alert applies to.	<entity-name>
Impact Type	Displays the category in which the alert is classified.	Availability, Capacity, Configuration, Performance, System Indicator

Parameter	Description	Values
Severity	Displays the severity level of this condition. There are three levels: Critical An actionable critical situation has been detected, and action is required immediately. The cluster can have the potential issues and can stop running, or run into irreparable issues. Warning An actionable issue has been detected, and user intervention is required. A more serious issue can occur if the alert is not resolved soon. Info (Informational) An actionable minor problem has been detected. It should be resolved relatively soon and not ignored.	Critical, Warning, Informational

Parameter	Description	Values
Status	Indicates whether the alert is resolved. Resolving an error means you set that error as fixed. The alert can return if the condition is scanned again at a future point. If you do not want to be notified about the condition again, turn off the alert for this condition. For more information about how to turn off the alert using alert policies, see Alert Policies (Prism Central). • A blank value means the alert is not resolved. • An Acknowledged By value means this alert is acknowledged manually by the specified user at the specified date and time. To manually acknowledge an alert, select the target alert check box and click Acknowledge. • An Auto Resolved (date_time) value means this alert was resolved automatically at the specified date and time. This functionality requires the auto resolve setting to be enabled for this alert type. For more information about how to set the auto resolve using alert policies, see Alert Policies (Prism Central). • A Resolved By user (date_time)" value means this alert is resolved manually by the specified user at the specified date and time. To manually resolve an alert, check the box for that alert and click Resolve.	(blank), Acknowledged By user (date_time), Auto Resolved (date_time), Resolved By user (date_time)
Created Time	Displays the date and time when the alert occurred.	(date and time)
Last Occurred	Displays the date and time when the alert is last generated (before this occurrence). If this is the first occurrence, both Created Time and Last Occurred display the same date and time.	(date and time)

OVA Management

An Open Virtual Appliance (OVA) file is a tar archive file created by converting a virtual machine (VM) into an Open Virtualization Format (OVF) package for easy distribution and deployment. OVA helps you to quickly create, move or deploy VMs on different hypervisors.

The minimum supported versions to perform OVA operations are AOS 5.18, Prism Central 2020.8, and AHV-20190916.253.

You can perform the following OVA operations in Prism Central:

Note: All the following operations are supported from the **OVA**s page except exporting a VM as an OVA. You can export a VM as an OVA from the **List** tab in [VMs Summary View](#) on page 100 or from the **Summary** page in [Viewing VM Details](#) on page 113.

- Export a VM as an OVA. For more information, see [Exporting a VM as an OVA](#) on page 238.
- Upload an OVA. For more information, see [Uploading an OVA](#) on page 239.
- Deploy an OVA file as a VM. For more information, see [Deploying an OVA as VM](#) on page 249.
- Download an OVA file to your local machine. For more information, see [Downloading an OVA](#) on page 256.
- Rename an OVA file. For more information, see [Renaming an OVA](#) on page 256.
- Delete an OVA file. For more information, see [Deleting an OVA](#) on page 257.
- Resume upload when an upload is interrupted and shows an error.

Note: The resume upload action is not available in case the concatenate process of the upload is interrupted. In such a case, you need to run the Concatenate API from the **REST API Explorer** to resume the upload.

To perform any action on an OVA, select the target OVA in the **OVA**s page, and choose the required action from the **Actions** dropdown menu.

Note:

The **Actions** dropdown menu appears only if you select an OVA in the **OVA**s page. If the **OVA**s page is empty, then the **Actions** dropdown menu is available only after you create an OVA using **Upload OVA** or export an existing VM as an OVA.

The following is an example showing the **Actions** dropdown menu in **OVA**s page.




Figure 67: Actions Dropdown Menu - OVAs Page

You can check the task progress for these actions in **Tasks** page or from the [Tasks icon](#). For information about **Tasks** page, see [Tasks View](#) on page 496.

OVAs Summary View

This section provides the information about the OVAs summary view, OVA-specific alerts, and the fields that appear in the **OVAs** page.

Note: Only the **super admin** role has the permission to view OVAs.

To access the summary view of all OVAs across registered clusters, perform the following steps:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Compute > OVAs** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

Prism Central displays **OVAs** page that enables you to view the information about OVAs available in the cluster.




Figure 68: OVAs Page

The following table describes the fields that appear in the **OVAs** page. A dash (-) is displayed in a field when a value is not available or applicable.

Table 45: OVAs Page - Field Description

Field	Description	Values
OVA Name	Displays the name of the OVA file.	<OVA_File Name>
Source VM	Displays the name of source VM that was exported as an OVA. The alert indicators for the task status are displayed beside the source VM for which the task is run. See the Alert indicator table for more information.	<Source_VM_File Name>
Disk Format	Displays the format of the disks in the OVA.	[QCOW2 VMDK]

Field	Description	Values
Added by	Indicates the user who added this OVA.	<User_Name>. For example, admin
Date added	Displays the date and time when the OVA was added.	(Timestamp in mm/dd/yyyy format, time in hr:min:sec [AM] PM] format) Example: 01/07/23, 3:00:20 PM

OVA Alerts

This section provides the information about the OVA specific alerts generated in Prism Central.

Table 46: OVA Alerts

OVA Alerts	Status Message sample	Description
 Figure 69: Failure Alert	Incorrect file uploaded. Multiple OVA files present in OVA.	Indicates the upload process failed during validation. You must fix the issue in the OVA file or upload the correct, compatible OVA file.
 Figure 70: Error Alert	Upload for OVA file has not started. Please start OVA file upload using Resume Upload option. Uploaded OVA file has not been verified. Please start verification using Resume Upload option.	Indicates there is an interruption in one of the three sub-tasks in the upload process. You can resume the upload. See the <i>Resume Upload</i> action procedure.
 Figure 71: Progress Alert	OVA file upload is in progress.	Indicates the OVA upload is in progress without any error or interrupts. Check the Tasks page after some time to verify the successful upload. For information on how to access Tasks page, see Icons in Prism Central UI on page 36 (using the Tasks icon) or Tasks View on page 496.

A sample of the error hover message is as follows:




Figure 72: Sample error hover message

Filters Pane - OVAs page

You can filter the information in the **OVAs** page based on the following fields that are available in the **Filter** pane. For information about how to use **Filters** option in Prism Central, see [Prism Central GUI Organization](#) on page 49.

Table 47: Filter Pane Field Description - OVAs page

Field	Description	Values
OVA Name	Filters based on the OVA file name. It returns a list of OVAs that matches the name condition/string.	<OVA file name >
Source VM	Filters based on the Source VM name. It returns a list of OVAs for which the Source VM field matches the condition/string.	<Source VM name>
Disk Format	Filters based on the disk format. Select one or more checkbox for the OVAs of the required disk format.	VMDK, QCOW2

Exporting a VM as an OVA

Before you begin

If you intend to run the OVA in a non-Nutanix environment, ensure that you install the necessary drivers as per the requirements of the destination hypervisor.

About this task

This section describes how to export a VM as an OVA.

You can select the **Export as OVA** from the **Actions** dropdown menu in the VMs **List** tab or click **Export as OVA** action in the **Summary** page of an individual VM. For information about how to access the **List** tab or **Summary** page of an individual VM, see [VMs Summary View](#) on page 100 and [Viewing VM Details](#) on page 113.

Note: You can export a VM as an OVA with a specified disk format for the disks configured on the VM.

To export a VM as an OVA, perform the following steps:

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Compute > VMs** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.
The system displays the **List** tab by default with all the VMs across registered clusters in **Nutanix** environment. For information about how to access the list of VMs managed by external vCenter, see [VMs Summary View](#) on page 100.
3. Select the target VM checkbox, and choose **Export as OVA** from the **Actions** dropdown menu.
You can also click the target VM to access [Viewing VM Details](#) on page 113, and select **Export as OVA** from the **More** dropdown menu.

The system displays the **Export as OVA** window:




Figure 73: Export as OVA window

4. Specify the following information:
 - a. Enter the name for the OVA file in the **Name** field.
 - b. Select the disk format in which you want to export the OVA file in the **Export Disk Format** field.
The default format is QCOW2. If you want to use the OVA to deploy a VM with disks formatted as VMDK, then select VMDK as the disk format.
5. Click **Export**.
You can check the progress of the *Export as OVA* task in the **Tasks** page or from the [Tasks icon](#). For information about **Tasks** page, see [Tasks View](#) on page 496.
The exported OVA is available on the OVA dashboard.

Uploading an OVA

About this task

This section describes how to upload an OVA in Prism Central.

You can upload an OVA from the following sources:

- *From a local folder:* When you upload an OVA from a local source, you can upload it to only one target cluster.
- *From a URL:* When you upload an OVA from a URL, you can upload it to multiple clusters. The upload operation runs concurrently on all the selected clusters.

Select the appropriate upload source based on your requirement; a single cluster upload or multiple-cluster upload option.

Note:

- Prism Central supports the upload of up to three OVA files concurrently from local folders. Prism Central queues other OVA files that you upload until one or more of the concurrent uploads are completed.
- There is no restriction on the number of concurrent uploads for OVA uploads from URL.
- You need admin privileges for uploading an OVA.
- For OVA upload from local folders, the upload process can be interrupted due to the following reasons:
 - Network issues when the upload is running.
 - You close the browser window or tab in which the upload is running.
- The concatenate and validate operations of the OVA upload operation fail when an error occurs in the upload services because the checksum provided for the OVA file when it was created is incorrect.

If the upload process is interrupted due to any reason, you can resume the process. If the validation operation is interrupted due to a concatenation error, you must run the concatenate API to resume concatenation and validation. See [Concatenating Upload using APIs](#) on page 246.

Prism Central uses indicators to display the status of OVAs that you may have tried to upload. The status message is displayed when you hover the mouse or pointing device cursor on the indicator. See *OVA Alerts* section in [OVAs Summary View](#) on page 236.

To upload an OVA, perform the following steps:

Procedure

1. Log in to Prism Central as a user with admin privileges.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Compute > OVAs** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.
The system displays the **OVAs** page

3. Click **Upload OVA**.

The system displays the **Upload OVA** window.

The screenshot shows the 'Upload OVA' window with the following elements and annotations:

- Upload OVA**: Title of the window.
- OVA Source**: Radio buttons for **OVA File** (selected) and **URL**.
- Cluster**: A dropdown menu showing **auto_cluster_prod_**. An annotation box labeled **Select the clusters** with an arrow points to this dropdown.
- You can upload the OVA to multiple clusters together using URL upload**: A note below the cluster selection.
- Name**: An empty text input field.
- Checksum**: A dropdown menu showing **Optional** and a button labeled **SHA-256**.
- OVA File**: A section with the text **No file selected** and a blue **Upload** link. An annotation box labeled **Browse to the OVA file in your local folder and select** with an arrow points to the **Upload** link.
- Buttons**: **Cancel** and **Upload** buttons at the bottom right.

Figure 74: Upload OVA window

4. Specify the following field information in the **Upload OVA** window:

- a. **OVA Source:** Select the OVA source. For example, **OVA File** to upload from a local folder or **URL** to upload from a URL.

If you choose **URL** in the **OVA Source** field, you can select multiple clusters in the **Select AHV Cluster** field. One cluster gets selected by default in the **Select AHV Cluster** field.

- b. **Name:** Enter the name for the OVA file.

By default, Prism Central uses the file name of the OVA file that you upload, if you do not specify OVA file name in the **Name** field.

- c. (Optional) **Checksum:** Enter the checksum value of the OVA file based on the selected encryption as SHA-1 or SHA-256 hash algorithm.

- d. Based on your selection in the **OVA Source** field, perform the following relevant action:

- **OVA File:** Click **Select File** to navigate to the location of the OVA file in your local folder and open it.
- **OVA URL:** Enter the source URL from where you want to upload the OVA file.

Note: Only NFS and HTTP protocols are supported.

To specify one or more URLs, enter the appropriate URL address in the field using the following syntax for either NFS or HTTP:

```
nfs://[hostname|IP_addr]/path
http://[hostname|IP_addr]/path
```

Enter either the name of the host (*hostname*) or the host IP address (*IP_addr*) and the path to the file. If you use a *hostname*, the cluster must be configured to point at a DNS server that can resolve that

name (see [Configuring Name Servers for Prism Central](#) in *Prism Central Admin Center Guide*). A file uploaded through NFS must have 644 permissions.

The following is an example showing the OVA URL field:



OVA URL

Required

Figure 75: OVA URL field

Note:

Ensure that the URL that you provide is not redirected to another URL. The upload fails for redirected URLs.

Prism Central adds the URL to a list on the page and clears the **OVA URL** text box for another URL.

e. Click **Upload**.

The OVA upload starts. Wait until the **Continue in Background** option is displayed.

Note: The **Continue in Background** option is not displayed for uploads from URL.

Upload OVA

OVA Source

☒ OVA File ☐ URL

Cluster

auto_cluster_prod

You can upload the OVA to multiple clusters together using URL upload

Name

UP-Test-qcow-2.ova

Checksum

Optional

SHA-256

OVA File

UP-Test-qcow-2.ova

Wait

Waiting...

Cancel

Upload

Figure 76: Upload OVA - Waiting

Upload OVA

OVA Source

☒ OVA File

☐ URL

Cluster

auto_cluster_prod_deepanshu_singhal_1aeeed3012f9e

You can upload the OVA to multiple clusters together using URL upload

Name

Testing

Checksum

OptionalSHA-256

OVA File

UP-Test-qcow-2.ova

0%

Cancel

Continue in Background

Figure 77: Continue in Background button

When the upload task is successfully completed, the system displays the following message:

Upload of the file is successfully completed. Verification of file is progressing in background. Please track it from Tasks page.

5. To resume an upload process that was interrupted, select the target OVA in the **OVA**s page, and choose **Resume Upload** from the **Actions** dropdown menu.

The status message appears when you hover the mouse over the status icons for each OVA. The status message provides the indication about whether you can use the **Resume Upload** option to resume the upload.

Note:

Check the reason for failure of **OVA Validate** task in the **Tasks** page. If the upload process is interrupted due to concatenation and validation failure resulted from an incorrect checksum, the **Resume Upload** option is not available.

Upload the OVA again in case the upload process was interrupted due to a checksum error.

Concatenating Upload using APIs

About this task

This section describes how to run the concatenation API using the **REST API Explorer** in Prism Central.

During an OVA upload, the OVA file is chunked and the chunks are uploaded. After upload, during the validation phase, the chunks are concatenated and validated.

When the OVA upload process is interrupted due to the concatenation error, you need to run the concatenation API to resume the upload process.

To run the concatenation API successfully, you need to perform the following actions:

- Derive the UUID of the OVA for which the upload is interrupted. The UUID is derived using the name of the OVA. For more information, see *Step 1 to Step 4* in the following procedure.
- Run the concatenation API. For more information, see *Step 5 to Step 6* in the following procedure.

Procedure

To run the concatenation API, perform the following steps:

1. Log in to Prism Central.
2. Select **REST API Explorer** from the `<user>` dropdown menu in the upper-right corner.

The **REST API Explorer** opens in a new browser tab and displays a list of the objects that can be managed by the API.

3. In the **REST API Explorer** browser tab, navigate to **ovas > Post /ovas/list**, and click **Try it Out**. The **Post /ovas/list** helps you filter the OVAs and get the details of the details of the interrupted OVA.




Figure 78: REST API Explorer




Figure 79: REST API Explorer -Try it Out

4. Enter the name of the OVA in the **"filter"** field of **get_entities_request** API in the following format, and click **Execute**.

POST /ovas/list

Get a list of existing OVAs ^

This operation gets a list of OVAs, allowing for sorting and pagination. Note: Entities that have not been created successfully are not listed.

Parameters

Cancel

Name	Description
get_entities_request * required	
object	
(body)	<pre>{ "kind": "ova", "sort_attribute": "ova", "filter": "name:: OVA_Site_A", "length": 1, "sort_order": "ASCENDING", "offset": 0 }</pre>

Cancel

Parameter content type

application/json

Execute

Clear

Figure 80: Parameter Filter

In the **Response Body** box, the UUID of the queried OVA is listed. Note or copy the UUID.

```
"kind": "ova",
"creation_time": "2023-01-02T06:06:08Z",
"uuid": "1b3adf92-878d-40e1-a327-008f90603d68"
```

Figure 81: OVA Upload UUID

5. In the **REST API Explorer** browser tab, navigate to **ovas > Post /ovas/{uuid}/chunks/concatenate**, and click **Try it Out**.

6. Enter the UUID of the OVA in the **uuid** field, and click **Execute**.

Name	Description
uuid * required	The UUID of the entity.
string(\$UUID)	1b3adf92-878d-40e1-a327-008f90603d68
(path)	

Figure 82: OVA Upload UUID

The **Response Code** field displays 202 indicating that the concatenation and validation request has been accepted.

Deploying an OVA as VM

About this task

This section describes how to deploy an OVA as VM.

Note:

For the OVAs that are not created from AHV, ensure that the pre-installed VirtIO drivers are available for them. For more information about VirtIO drivers, see [Nutanix VirtIO for Windows](#) in *AHV Administration Guide*.

To deploy an OVA as a VM, perform the following steps:

Note: You can navigate back and forth between the steps using **Back** and **Next**.

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Compute > OVAs** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.
The system displays the **OVAs** page

3. Select the target OVA, and choose **Deploy as VM** from the **Actions** dropdown menu. The system displays the **Deploy as VM** window.

You need to use the steps in the **Deploy as VM** window to configure the deployment.

Deploy as VM

1 Configuration

2 Resources

3 Management

4 Review

Name

test

Description

(Optional)

Cluster

auto_cluster_prod_deepanshu_singhal_1aeed3012f9e

VM Properties

CPU

1

▼▲

vCPU

Cores Per CPU

1

▼▲

Cores

Memory

2

▼▲

GB

Next

Figure 83: Deploy as VM - Configuration Step

4. In the **Deploy as VM** window, specify the following information in the **Configuration** step:

- a. **Name:** Enter the name of the VM that needs to be deployed.
- b. **Description:** Enter that description such as Backup VM for Prism.
- c. **Cluster:** Select one of the following options from the dropdown menu:
 - Target cluster on which you intend to place the guest VM.
Only the clusters for which you have access and where OVAs are available, are displayed in the **Cluster** dropdown menu. By default, Prism Central selects a cluster in the dropdown menu for deployment.
 - **Automatic cluster selection**, if you intend to place the guest VM on a system-selected cluster.
For more information, see [Automatic Cluster Selection for VM Placement](#) section in the *AHV Administration Guide*.

Note: Automatic cluster selection configuration for VMs with GPU specifications is not supported.

- d. **CPU, Cores Per CPU and Memory in VM Properties:** Specify the CPU and memory requirements as required.
- e. Click **Next**.
The system displays the **Resources** step.

5. Specify the following information in the **Resources** step:

Note:

You can add new disks, NICs and GPUs.

- a. **Boot Configuration:** Select **UEFI Mode** only if your hardware supports UEFI boot mode, else select **Legacy BIOS Mode**.

- b. **Set Boot Priority:** The default priority is already set.

- c. **Attach Disk:** Used to attach additional disks in addition to the disks already available in the OVA.

In the **Attach Disk** window, specify the required information, and click **Save**.

For more information about how to attach a disk, see [Creating a VM through Prism Central \(AHV\)](#) on page 149.

- d. **Attach to Network:** Used to attach the network.

In the **Attach to Network** window, select **Subnet** and **Network Connection State** from the dropdown menus.

For more information about how to configure network, see [Creating a VM through Prism Central \(AHV\)](#) on page 149.

- e. Click **Add GPU** to add GPU resources.

Note:

- Automatic cluster selection configuration for VMs with GPU specifications is not supported. If you have selected the **Automatic cluster selection** option in step 4.c on page 251, the system does not provide you the option to add GPU.
- If the selected cluster has GPU resources, then you can add multiple vGPUs to the same VM. Addition of multiple vGPUs is based on the installed GPU resources.

- For information about how to add multiple GPUs, see [Adding Multiple vGPUs to the Same VM](#) on page 172.

- For information about multiple vGPU support, see [Multiple Virtual GPU Support](#) information in *AHV Admin Guide*.

- f. Click **Next**.

The system displays the **Management** step.

Deploy as VM

1 Configuration
2 Resources
3 Management
4 Review

Categories

Type to search...

Tag the VM with Category: Value to assign policies associated with value

Timezone

(UTC) UTC

☐ Use this VM as an Agent VM

Guest Customization

Script Type

No Customization

Configuration Method

Custom Script

Back

Next

Figure 84: Deploy as VM - Management Step

- g. **Categories:** Select the category to be mapped to the VM. Type in this field to display a list of available categories.
For more information on categories, see [Category Management](#) in the *Prism Central Admin Center Guide*.
- h. **Timezone:** Select the appropriate timezone as per your requirement.
- i. **Use this VM as an Agent VM-** Select this checkbox if you want to use the deployed VM as an agent VM.
- j. **Guest Customizations:** Used to customize the VM. Select Cloud-init (for Linux VMs) or Sysprep (for Windows VMs).

Script Type: Select the VM script customization type from the dropdown menu. The options are *No Customization*, *Sysprep(Windows)*, or *Cloud-init (Linux)*.
Configuration Method: Select the VM script customization method from the dropdown menu. This field is activated when the **Script Type** is either *Sysprep(Windows)* or *Cloud-init (Linux)*. The options are *Custom Script* or *Guided Script*.

The system displays the options required to configure Cloud-init and Sysprep, such as options to specify a configuration script and option to upload script.

- k. To specify a user data file (Linux VMs) or answer file (Windows VMs) script for unattended provisioning, perform either of the following actions:

- » If the file is available on your local computer, click **Upload Script**, choose and upload the file.
- » Create or paste the contents of the file in the text box below the **Upload Script**.

Note: The script type supports the following file formats.

- Sysprep: XML
- Cloud-init (Linux): YAML, JSON, or Shell.

- l. Click **Next**.
The system displays the **Review** step.

Deploy as VM

- 1 Configuration 2 Resources 3 Management 4 Review

Configuration

Name	test
CPU	1 vCPU
Cores Per CPU	1 Cores
Memory	2 GB

Resources

Disks

#	Type	Image	Size	Bus Type
1	Disk	-	0.94 GiB	SCSI
2	Disk	-	0.94 GiB	SCSI

Boot Configuration Default Boot Order (CD-ROM, Disk, Network)

Networks

Subnet	Private IP
pc_network	Auto-Assign

Management

Categories	None
Timezone	UTC
Guest Customization	No Customization

6. Review the deployment configuration in the **Review** step, and Click **Launch** .
You can check the progress of the deployment task in the **Tasks** page or from the [Tasks icon](#). For information about **Tasks** page, see [Tasks View](#) on page 496.

Downloading an OVA

This section describes how to download an OVA on your local machine from Prism Central.

About this task

To download an OVA, perform the following steps:

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Compute > OVAs** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.
The system displays the **OVAs** page.
3. Select the OVA to be downloaded, and choose **Download** from the **Actions** dropdown menu.
The system displays the **Download** window.

The following is an example showing the **Download** window:




Figure 86: Download window

Note: You can download up to five OVAs in parallel.

4. To confirm the download action, click **Continue**. To cancel the download action, click **Cancel**.

The downloaded file is saved in a local folder based on your browser's download settings.

Once the download starts, it continues even after you sign out from the console, as the TCP connection does not break. However, if you close the browser or the internet connection is interrupted, the download fails as the TCP connection breaks.

Renaming an OVA

This section describes how to rename an OVA in Prism Central.


About this task

To rename an OVA, perform the following steps:

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Compute > OVAs** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.
The system displays the **OVAs** page
3. Select the OVA to be renamed, and choose **Rename** from the **Actions** dropdown menu.
The system prompts you to update the new name for the OVA file.

The following is an example showing the **Rename** window:



The screenshot shows a 'Rename' dialog box. The title bar contains the word 'Rename', a question mark icon, and a close 'X' icon. The main area has a label 'Name' above a text input field. The input field contains the text 'Test-OVA-vGPU-UploadBack.ova'. At the bottom right of the dialog are two buttons: 'Cancel' and 'Save'.

Figure 87: Rename window

4. Delete the old name and enter the new name in the **Name** field.
5. To confirm the rename action, select **Save**.
To cancel the rename action, select **Cancel**.

Deleting an OVA

You can permanently delete an OVA in Prism Central.

About this task

To delete an OVA, perform the following steps:

Procedure

1. Log in to Prism Central.
2. From the [Application Switcher Function](#) on page 40, select the **Infrastructure** application, and from the **Navigation Bar**, navigate to **Compute > OVAs**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.
The system displays the **OVAs** page.

3. Select the OVA, and from the **Actions** dropdown menu, choose **Delete**.
The system prompts you to confirm the delete action.




Figure 88: Delete OVA window

4. To confirm the delete action, select **Delete**.
To cancel the delete action, select **Cancel**.

Note: The **Delete** action deletes the OVA permanently. You cannot recover a deleted OVA.

Image Management

Prism Central provides a centralized location to manage the images you require on registered AHV clusters. You can upload images to the clusters and maintains an inventory of the images on it.

Note:

- You can configure policies that govern the image placement process. These policies allow you to specify the clusters on which the images need to be uploaded. For more information about the image placement policies, see [Image Placement Policies](#) on page 527.
- You can also update the uploaded images. For more information, see [Modifying an Image](#) on page 280
- Cluster scoping is not supported for images.

Image Source

In Prism Central, you can add images from the following sources:

- Workstation. For more information, see [Adding Images from a Workstation](#) on page 264.

- URLs to images on a remote server. For more information, see [Adding Images from a Remote Server](#) on page 269.
- VM disk. For more information, see [Adding Images from a VM Disk](#) on page 273.

For information about these sources, see [Adding an Image](#) on page 264.

Image Placement Methods

When you add images, the system provides you the following image placement methods:

- Select target clusters during images upload, and place image directly on the target clusters.
- Specify target clusters using image placement policies, and place images on the target clusters based on image placement policies.

In this case, you can select the image categories, and the system uploads images to all the registered clusters mapped to the image categories in image placement policy. For information about image placement policies, see [Image Placement Policies](#) on page 527.

Note: You can also suspend or resume the image placement policies based on your requirement. For more information, see [Suspending or Resuming Enforcement of an Image Placement Policy](#) on page 278

For information about how to add images, see [Adding Images from a Workstation](#) on page 264 and [Adding Images from a Remote Server](#) on page 269.

Images Import from Registered Clusters to Prism Central

In addition to the sources defined in [Image Source](#) on page 258, you can also import images from registered AHV clusters. You can import images only from registered clusters that are running AOS 5.8 or later. For more information, see [Importing Images to Prism Central](#) on page 281.

Images Summary View

The **List** tab in **Images** page displays the information about all the images across registered clusters.

Note: This section describes the information and options that appear in the **Images** page.

- See [Prism Central GUI Organization](#) on page 49 for instructions on how to view and organize that information in various ways.
- See [Image Management](#) on page 258 for information about how to add and manage images through Prism Central.
- See [Image Placement Policies](#) on page 527 for information about image placement policies.
- See [Bandwidth Throttling Policies](#) on page 535 for information about bandwidth throttling policies.

To access the summary view of all images, perform the following steps:

1. Log in to Prism Central.

2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Compute > Images** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

The system displays the **List** tab by default.

The following is an example showing the **List** tab in **Images** page:




Figure 89: List tab - Images

The following table describes the fields that appear in the **List** tab in **Images** page:

Note: A dash (-) is displayed in a field when a value is not available or not applicable.

Table 48: Images Page - Field Description

Parameter	Description	Values
Name	Displays the image name. Click the image name to view the detailed information about that image. For more information, see, Image Details View on page 262	<Image name>
Description	Indicates which user uploaded this image.	(text string)
Type	Displays the image type.	ISO, Disk
Size	Displays the image size.	x [MB GB], where x is numeric digit.
Creator	Displays who created the image.	<user>. For example admin

You can perform the following actions for the images in the **List** tab:

- Access the detailed information about an individual image. For more information, see [Image Details View](#) on page 262.
- Filter the images list based on available parameter values using **Filters** pane. For more information about **Filters** pane, see [Filters Pane - Images](#) on page 261.
- Export the table that contains the list of images and their information to a file in a CSV format. For more information about **Export** option, see [Export](#) on page 54.

- Group the images based on pre-defined criteria. For information about how to group the images, see [Group by](#) on page 51.
- View images based on pre-defined criteria or create a custom view. For information about available views and how to create a custom view, see [View by](#) on page 51.
- Perform the following image-specific actions on a single or multiple images using the **Actions** dropdown menu:
 - **Update** - Used to update the image file attributes. For details, see [Modifying an Image](#) on page 280.
 - **Delete** - Used to delete the image. Select the image(s), and choose **Delete** from the **Actions** dropdown menu.
 - **Add Image to Catalog** - Used to add an image to the catalog. For more information, see [Modifying an Image](#) on page 280.

For information on how to manage the catalogs, see [Catalog Management](#) on page 283.

- **Manage Categories** - Used to assign a category to the images. For information about how to manage image categories, see [Modifying an Image](#) on page 280.
- **Download Image** - Used to download the image. Select the image, and choose **Download Image** from the **Actions** dropdown menu.

The downloaded image does not have any file extension. If you download an .ISO image, you must rename the file to append .ISO extension before using the image. A downloaded disk image is in raw format and you must manually convert it using *qemu-img* or similar tool before using the disk image.

The Image download speed is dependent on the following factors:

- Bandwidth between Prism Element where image is stored and Prism Central
- Bandwidth between Prism Central and client system

Note: You can also perform these actions from the **Summary** page of an individual image. For more information, see [Image Details View](#) on page 262.

Filters Pane - Images

You can filter the information in the **Images** page based on the following fields that are available in the **Filters** pane. For information about how to use **Filters** option in Prism Central, see [Prism Central GUI Organization](#) on page 49.

Table 49: Filter Pane Field Description - Images page

Field	Description	Values
Name	Filters based on the image name. It returns a list of storage containers that satisfy the name condition/string.	<Image name >
Description	Filters based on the image description.	(description string)
Type	Filters based on the image type. Check the box(es) for the desired image types.	Disk, ISO

Policies Tab

The **Policies** tab displays **Placement Policies** and **Bandwidth Throttling Policies**.

- For information about how to configure **Placement Policies**, see [Image Placement Policies](#) on page 527.
- For information about how to configure **Bandwidth Throttling Policies**, see [Bandwidth Throttling Policies](#) on page 535.

Image Details View

To access the details view of an individual image, perform the following steps:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Compute > Images** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

The system displays the **List** tab by default.

3. Click the `<Image_Name>` to view the **Summary** tab of an individual image.

Note: Replace `<Image_Name>` with the actual image name at your site.




Figure 90: Image Details View (Summary Tab)

For information about the fields available in **Info** widget, see [Images Summary View](#) on page 259.

Location Tab

The following table describes the information displayed in the **Location** tab

Table 50: Location Tab -Field Description

Field	Description	Values
Name	Displays the name of the cluster in which the image is stored. Each row represents another cluster that contains the image.	<Cluster name>
AOS Version	Displays the AOS version running in the cluster.	AOS version number
Hypervisor	Displays the name of the hypervisor running in the cluster.	AHV, ESXi, Hyper-V
Host Count	Displays the number of hosts in the cluster.	(integer)
VM Count	Displays the number of VMs in the cluster	(integer)

Policies Tab

The **Policies** tab lists the policies that apply to the selected image. For information about the image-related policies, see [Image Policy Management](#) on page 527.

- For information about how to configure **Placement Policies** for images, see [Image Placement Policies](#) on page 527.
- For information about how to configure **Bandwidth Throttling Policies** for images, see [Bandwidth Throttling Policies](#) on page 535.

Requirements

You must meet the following requirements to access the image management feature in Prism Central:

- The version of both Prism Central and AOS on registered clusters is 5.8 or later.
- Clusters is registered with Prism Central with AHV installed in it.
- The port 2007 must be open for the image service.

Limitations

The following limitations apply to the image management feature in Prism Central:

- You cannot update images that another Prism Central instance managed earlier. However, the images are displayed in Prism Central, and you can use the images to create VMs or delete the images you no longer need.
- You cannot choose a container when uploading images from Prism Central. Prism Central uploads images to the container named **SelfServiceContainer**.
- When you upload images from a workstation and place the images on multiple clusters during the initial image upload, the uploaded image becomes active (in Prism Element) on only one of the selected clusters and remains inactive on the other clusters registered to Prism Central.
- If you add a new cluster that contains images to an existing Prism Central, you must first import the images to Prism Central and then use the image placement policies to push the images to any other cluster already registered with Prism Central.
- If an image is not active on a cluster, as an admin user it is possible that you are unable to use Prism Central to create that VM on that cluster. The cluster on which you want to create the VM cannot check out the image from

a remote cluster. In this scenario, you must upload the image manually to the cluster on which you want to create the VM.

When you create a VM, Prism Central uses API v2 workflows. API v2 does not have the checkout workflow. The image checkout workflow is used to copy an image on demand from cluster B to cluster A if the image does not exist or is inactive on cluster A.

Note:

You can remove this limitation using an image placement policy. Assign a category to the image that is available on multiple remote clusters, and apply the image placement policy to the image category. When you apply the image placement policy to an image, Prism Central propagates the image to all the clusters that are included in the image placement policy. Prism Central can then checkout the image from the remote cluster where the image is active and create the VM.

The system can take up to 15 minutes to switch the image propagation from an alternative cluster. For example, if an image is located on clusters A and B. As a result of an image placement policy you applied to that image, cluster A starts copying the image to cluster C. During the copy operation, if cluster A becomes unavailable, the system takes up to 15 minutes to switch to cluster B, and starts copying the image from cluster B to cluster C.

Adding an Image

This section describes how to add an image in Prism Central.

About this task

You can add the images in Prism Central from any of the following sources:

Procedure

- Workstation, see [Adding Images from a Workstation](#) on page 264.
- Remote Server, see [Adding Images from a Remote Server](#) on page 269.
- VM disk, see [Adding Images from a VM Disk](#) on page 273.

Adding Images from a Workstation

This section describes how to add (or upload) the images in Prism Central from the workstation that you use to access Prism Central.

About this task

Note:

- By default, the system selects all the clusters. The system also provides you the option to clear the cluster as per your requirement.
- The image becomes active (in Prism Element) on the selected cluster and inactive on other clusters registered to Prism Central.
- Image placement policies, if applied, take effect after a short duration.
- Most modern browsers impose file size limitations that affect this upload method. If you intend to upload images larger than 2 GB, Nutanix recommends you to upload the images from a remote server.
- The browser type, and CPU and RAM utilization on the workstation limit the number of concurrent uploads. Concurrent uploads that exceeds the default limit of the browser are queued or throttled by the browser and can consume more time.

- Large file uploads, and high CPU and memory utilization can slow down the browser.

To upload an image from a workstation, perform the following steps:

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Compute > Images** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.
The system displays the **List** tab by default.

3. Click **Add Image**.
The system displays the **Add Images** window.

Add Images

1

Select Image

2

Select Location

Image Source

☒ Image File

☐ URL

☐ VM Disk

+ Add File

Source: [LOCAL]\pc.2020.11.0.1.ova

Remove

Name

test_image.qcow2

Type

Disk

Description

Checksum

SHA-256

Cancel

Next

Figure 91: Add Images: File

4. Specify the following information in the **Select Image** step:

- a. **Image Source** - Select **Image File** radio button.
- b. Click **Add File**. The system displays the file attributes.
- c. Browse to the location of the image file, and then click **Open**.
- d. Specify the following attributes for the image file:
 - **Image Name** - Enter the image name. By default, the system pre-fills the name of the file you selected, however you can change the image name as per your requirement.

Note: Ensure that the name of the image is unique across all the images in Prism Central.

- **Image Type** - Select the type of image.
 - **Image Description** - Enter the description for the image file.
 - **Checksum** - Select the hashing algorithm.
- e. Repeat step b to step d if you want to add multiple image files.

To remove an image file entry, locate the entry and click **Remove**.
 - f. Click **Next** after you add all the image files.

The system displays the **Select Location** step.

Add Images

1

Select Image

2

Select Location

Placement Method

- ☒ Place image directly on clusters

This option is good for smaller environments. The image will be placed on all selected clusters below.

- ☐ Place image using Image Placement policies

This option is good for larger environments. It requires you to first set up Image Placement policies between categories assigned to clusters and categories assigned to images. From then on, you only need to associate a relevant category to an image while uploading it here.

Select Clusters

Select the set of clusters to use for placement

- ☒ All clusters

NAME



[Redacted]

Back

Save

Figure 92: Add Images: Select Location

5. Specify the following information in the **Select Location** step:

a. **Placement Method** - Select one of the following placement methods:.

- **Place image directly on clusters**- Used to place the images directly on the selected clusters.
- **Place image using Image Placement policies** - Used to delegate image placement decisions to configured policies and assign categories to the images

b. If you select:

- **Place image directly on clusters** - Specify the following information:

Select Clusters - Select the clusters in which you want to add the image file in the **Name** column.

Note:

- By default, the system selects all the registered clusters.
- If you want to add the images to all registered clusters, ensure that you select all the clusters in the **Name** column.
- If you want to upload to only a subset of the registered clusters, clear all clusters in the **Name** column, and only select the clusters you want from the list.

- **Place image using Image Placement policies** - Specify the following information:

Select Image Categories - Click inside the search box and select the category you want from the list. You can also type the name of the category to reduce the list to matching names.

Note:

- To select multiple categories, click [Add icon](#) beside each category name that you want to include in the categories.
- To remove a category, click [Remove icon](#) beside each category name that you want to remove from the selected categories.
- For information about how to create a category and assign a category value to an image, see [Creating a Category](#) in the *Prism Central Admin Center Guide* and [Associating Images with Categories](#) on page 530.

6. Click **Save**.

The system adds the image files in batches and takes some time to enforce the image placement policies (if selected).

Note: You can also suspend or resume the image placement policies based on your requirement. For details, see [Suspending or Resuming Enforcement of an Image Placement Policy](#) on page 278.

Adding Images from a Remote Server

You can add an image from a remote server to registered Nutanix clusters. To add an image from a remote server, you need the URL of the image. You can also specify URLs to multiple images as part of a single operation. When you specify image URLs, Prism Central adds the images to all registered clusters. Adding the image from a remote server also allows you to overcome file size limitations imposed by modern browsers. The file size limitation is usually 2 GB.

About this task

To add an image from a remote server, perform the following steps:

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Compute > Images** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.
The system displays the **List** tab by default.
3. Click **Add Image**.
The system displays the **Add Images** window.

1 Select Image 2 Select Location

Image Source

☐ Image File ☒ URL ☐ VM Disk

Image URL

http://<remote-image-location>

Authentication (optional) ^

Username Password

john_doe

Show

Note: Authentication will take place after the image upload has been submitted.

+ Add URL

Cancel Next

Figure 93: Add Images: URL

4. Specify the following information in the **Select Image** step:
 - a. **Image Source** - Select **URL** radio button.
 - b. **Enter Image URL** - Enter the appropriate URL address in the field using the following syntax for NFS, HTTP, or HTTP, and then click **Upload File**.

Note: The system supports only NFS, HTTP, and HTTPS protocols.

```
nfs://[hostname|IP_addr]/path
```

`http://[hostname|IP_addr]/path`

Enter either the name of the host (*hostname*) or the host IP address (*IP_addr*) and the path to the file. If you use a *hostname*, the cluster must be configured to point at a DNS server that can resolve that name. For more information, see [Configuring Name Servers for Prism Central](#) in *Prism Central Admin Center Guide*.

Note: A file uploaded through NFS must have 644 permissions.

- c. Optionally, if the remote image location requires authentication credentials, enter the **Username** and **Password**, then click **Next**.

Note: The authentication is validated after the image upload request is submitted.

- d. Specify the following attributes for the image file:

- **Image Name** - Enter the image name. By default, the system pre-fills the name of the file you selected, however you can change the image name as per your requirement.

Note: Ensure that the name of the image is unique across all the images in Prism Central.

- **Image Type** - Select the type of image.
- **Image Description** - Enter the description for the image file.

The screenshot shows a web form for adding image attributes. At the top, it says 'Source: CentOS' and has a 'Remove Image' link. Below this is a 'General' section with fields for 'Name' (pre-filled with 'CentOS'), 'Type' (a dropdown menu showing 'ISO'), 'Description' (pre-filled with 'Generic CentOS image'), and 'Checksum' (a text field with a dropdown menu showing 'SHA-1'). Below the 'General' section is an 'Authentication (optional)' section with fields for 'Username' (pre-filled with 'john_doe') and 'Password' (a masked field with a 'Show' link). At the bottom right are 'Cancel' and 'Next' buttons.

Figure 94: Add Image Attributes: URL

- e. Repeat step b to step c if you want to add multiple image files.
- To remove an image file entry, locate the entry and click **Remove**.
- f. Click **Next** after you add all the image files.
- The system displays the **Select Location** step.

Add Images

1 Select Image

2 Select Location

Placement Method

☒ Place image directly on clusters

This option is good for smaller environments. The image will be placed on all selected clusters below.

☐ Place image using Image Placement policies

This option is good for larger environments. It requires you to first set up Image Placement policies between categories assigned to clusters and categories assigned to images. From then on, you only need to associate a relevant category to an image while uploading it here.

Select Clusters

Select the set of clusters to use for placement

☒ All clusters

NAME	
☐	

Back

Save

Figure 95: Add Images: Select Location

5. Specify the following information in the **Select Location** step:

a. **Placement Method** - Select one of the following placement methods:

- **Place image directly on clusters**- Used to place the images directly on the selected clusters.
- **Place image using Image Placement policies** - Used to delegate image placement decisions to configured policies and assign categories to the images

b. If you select:

- **Place image directly on clusters** - Specify the following information:

Select Clusters - Select the clusters in which you want to add the image file in the **Name** column.

Note:

- By default, the system selects all the registered clusters.
- If you want to add the images to all registered clusters, ensure that you select all the clusters in the **Name** column.
- If you want to upload to only a subset of the registered clusters, clear all clusters in the **Name** column, and only select the clusters you want from the list.

- **Place image using Image Placement policies** - Specify the following information:

Select Image Categories - Click inside the search box and select the category you want from the list. You can also type the name of the category to reduce the list to matching names.

Note:

- To select multiple categories, click [Add icon](#) beside each category name that you want to include in the categories.
- To remove a category, click [Remove icon](#) beside each category name that you want to remove from the selected categories.
- For information about how to create a category and assign a category value to an image, see [Creating a Category](#) in the *Prism Central Admin Center Guide* and [Associating Images with Categories](#) on page 530.

6. Click **Save**.

The system adds the image files in batches and takes some time to enforce the image placement policies.

Note: You can also suspend or resume the image placement policies based on your requirement. For more information, see [Suspending or Resuming Enforcement of an Image Placement Policy](#) on page 278.

Adding Images from a VM Disk

This section describes how to add an image from a VM disk on the registered cluster.

Before you begin

Ensure that the following prerequisites are met before you add an image from a VM disk.

- The source VM belongs to an AHV cluster.
- The source VM on which the disk resides is in the powered-off state.

- To place the image on clusters associated with relevant categories, you must first set up the image placement policies between categories assigned to clusters and categories assigned to images.

For information about how to create a category and assign a category value to an image, see [Creating a Category](#) in the *Prism Central Admin Center Guide* and [Associating Images with Categories](#) on page 530.

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Compute > Images** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.
The system displays the **List** tab by default.

3. Click **Add Image**.
The system displays the **Add Images** window.

Add Images

Image Source

☐ Image File ☐ URL ☒ **VM Disk**

Search by VM Name ?

Search by VM Name... ▾

Disk ?

Select Disk ▾

select a vm disk from pull-down list

click to add the vm disk

vm name entered in the seach field

Source: SCSI.0 (1 GIB) VM **VM-test** [Remove](#)

Cluster: auto_cluster_nested_609b6c4573101aefe244297b

Image Name

VM-test---SCSI.0-1

Image Type

Disk ▾

Image Description

This image was created from VM disk from 'VM-test' VM

[Cancel](#) [Next](#)

Figure 96: Add Images: VM Disk

4. Specify the following information in the **Select Image** step:

- a. **Image Source** - Select **VM Disk** radio button.
- b. **Search by VM Name** - Click inside this field text box, and select the target VM from the displayed list. The list displayed consists of all the powered-off VMs in the cluster.
The system adds the VM to a list and populates the list of disks attached to this VM in the **Disks** dropdown list.
- c. Select the target disk from the **Disks** dropdown list, and click the [Add icon](#),
The system displays the disk, and cluster information of the VM source.
- d. Specify the following attributes for the image file:

- **Image Name** - Enter the image name. By default, the system pre-fills the name of the file you selected, however you can change the image name as per your requirement.

Note: Ensure that the name of the image is unique across all the images in Prism Central.

- **Image Type** - Select the type of image.
 - **Image Description** - Enter the description for the image file.
- e. Repeat step b to step d if you want to add multiple image files.
To remove an image file entry, locate the entry and click **Remove**.
 - f. Click **Next** after you add all the image files.
The system displays the **Select Location** step.

Add Images

1 Select Image

2 Select Location

Placement Method

☒ Place image directly on clusters

This option is good for smaller environments. The image will be placed on all selected clusters below.

☐ Place image using Image Placement policies

This option is good for larger environments. It requires you to first set up Image Placement policies between categories assigned to clusters and categories assigned to images. From then on, you only need to associate a relevant category to an image while uploading it here.

Select Clusters

Select the set of clusters to use for placement

☒ All clusters

NAME	
☐	

Back

Save

Figure 97: Add Images: Select Location

5. Specify the following information in the **Select Location** step:

a. **Placement Method** - Select one of the following placement methods:

- **Place image directly on clusters**- Used to place the images directly on the selected clusters.
- **Place image using Image Placement policies** - Used to delegate image placement decisions to configured policies and assign categories to the images

b. If you select:

- **Place image directly on clusters** - Specify the following information:

Select Clusters - Select the clusters in which you want to add the image file in the **Name** column.

Note:

- By default, the system selects all the registered clusters.
- If you want to add the images to all registered clusters, ensure that you select all the clusters in the **Name** column.
- If you want to upload to only a subset of the registered clusters, clear all clusters in the **Name** column, and only select the clusters you want from the list.

- **Place image using Image Placement policies** - Specify the following information:

Select Image Categories - Click inside the search box and select the category you want from the list. You can also type the name of the category to reduce the list to matching names.

Note:

- To select multiple categories, click [Add icon](#) beside each category name that you want to include in the categories.
- To remove a category, click [Remove icon](#) beside each category name that you want to remove from the selected categories.
- For information about how to create a category and assign a category value to an image, see [Creating a Category](#) in the *Prism Central Admin Center Guide* and [Associating Images with Categories](#) on page 530.

6. Click **Save**.

The system adds the image files in batches and takes some time to enforce the image placement policies.

Note: You can also suspend or resume the image placement policies based on your requirement. For more information, see [Suspending or Resuming Enforcement of an Image Placement Policy](#) on page 278.

Suspending or Resuming Enforcement of an Image Placement Policy

This section describes how to suspend or resume the enforcement of an Image Placement Policy in Prism Central.

Before you begin

The minimum supported versions to suspend or resume the enforcement of an image placement policy in Prism Central are AOS 6.6 for Prism Element (with bundled AHV version) and PC 2022.9 for Prism Central.

About this task

Enforcement of placement policy involves copying images between clusters using image data transfer tasks that are typically long-running tasks and consume a considerable amount of network bandwidth. If you need to manage your local bandwidth and resources while the enforcement of the image placement policy is still going on, you can suspend the enforcement of the image placement policy, and can resume the enforcement at a later point of time.

To suspend or resume the enforcement of an image placement policy, perform the following steps:

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Compute > Images** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.
The system displays the **List** tab by default.
3. Select **Placement Policies** in the **Policies** tab.

The following is an example showing the placement policies view:



Figure 98: Placement Policy View

4. Select the target policy checkbox for which you want to suspend or resume enforcement, and choose **Suspend Enforcement** or **Resume Enforcement** from the **Actions** dropdown menu.
The policy enforcement is suspended or resumed based on the selected action.

Image Upload Verification

This section describes how to verify whether an image upload to multiple clusters succeeded.

When you add an image and select multiple clusters for initial placement of the image using the **Place image directly on clusters**, the status message box shows that the task is successful. The status message displays the successful completion of the task even if the image is successfully uploaded to only one cluster. Nutanix recommends you to go to the **Tasks** page and check whether the Image upload and Image update subtasks succeeded for all the clusters that you selected. It is possible that the Image upload or update subtasks for some of the clusters are failed.

Nutanix recommends that you use **Place image using Image Placement policies** option when you add images to multiple clusters. Image Placement policy based placement ensures that images propagate successfully to multiple clusters in a single operation.

For information about Image Placement policy, see [Image Placement Policies](#) on page 527 section.

For information about how to place the image on the registered clusters using image placement policies, see [Adding an Image](#) on page 264 section.

Add Images

1 Select Image

2 Select Location

Placement Method

☐ Place image directly on clusters
This option is good for smaller environments. The image will be placed on all selected clusters below.

☒ Place image using Image Placement policies
This option is good for larger environments. It requires you to first set up Image Placement policies between categories assigned to clusters and categories assigned to images. From then on, you only need to associate a relevant category to an image while uploading it here.

Select Image Categories

Clusters will be selected based on categories and associated placement policies.

Search for a category

+

Back

Save

Figure 99: Image Placement Method selection for multiple clusters

Modifying an Image

This section describes how to modify images in Prism Central.

About this task

You can perform the following actions to modify the images in Prism Central:

- Update the image attributes.

- Delete the image
- Add image to Catalog.
- Manage image categories

To modify an image (delete, update, or add to catalog), perform the following steps:

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Compute > Images** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.
The system displays the **List** tab by default.
3. Select the target image checkbox, and choose any one of following actions from the **Actions** dropdown menu:
 - **Delete** to delete the image file.
 - **Update** to edit the image file. In the **Edit Image** window, update the image name, image description, or image type as desired, and click **Save**. For field details, see [Adding an Image](#) on page 264.
 - **Add Image to Catalog** to add an image to the catalog. For more information on how to add an image to the catalog, see [Catalog Management](#) on page 283.
 - **Manage Categories** to manage the image categories. In **Manage Categories** window, perform the following steps:
 1. Use the [Add icon](#) and [Remove icon](#) available under **Set Categories** field to add or remove the categories for the selected image.
 2. Click **Save**.

For information about how to create an image category, see [Creating a Category](#) in the *Prism Central Admin Center Guide*.

Importing Images to Prism Central

This section describes how to import images from registered clusters and manage the images centrally from Prism Central.

About this task

An image imported to Prism Central continues to reside on the cluster that owns it. Prism Central only creates and stores image metadata locally, and it uses that metadata when you perform an action on the image. After you import an image, the image remains visible on the cluster from which you imported it, however the system never allows you to update the image on the cluster. You can update the image only from Prism Central.

There is no impact on the following images when you import images in Prism Central:

- Any images that you choose not to import.
- Any images that you add subsequently to the cluster from its web console. These images remain editable on the cluster until you import them to Prism Central.

You can import the images in Prism Central using any of the following criteria in a single operation:

- All images from all registered clusters
- All images from a selection of clusters

- Selection of images from some of the clusters.

To import images from registered clusters, perform the following steps:

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Compute > Images** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.
The system displays the **List** tab by default.
3. Click **Import Image**.
The system displays the **Import Image** window.
4. In the **Import Images** window, select one of the following options:
 - » **All Images** - To import all images from all registered clusters.
 - » **Images On a Cluster** - To import all images from a selection of registered clusters.
 - » **Select Images** - To import specific images from a given cluster. Click the link provided for the cluster, select the images that you want to import, and click **Done**.
Repeat this step for all the clusters from which you want to import the specific images.
5. Click **Save** to begin import.
Prism Central imports the metadata of the selected images and marks the images as read-only entities on the clusters.
The image continues to stay in the container that hosted the disk before importing the image.

Uploading Images to Objects

You can use the Object Lite service to upload images to Objects. The Objects Lite service is a compact version of core Object product, specifically designed to operate within the Controller Microservices Platform (CMSP).

Before you begin

Ensure that you meet all the prerequisites mentioned in the [Accessing Objects Lite Using AWS CLI- Put object](#) on page 613 topic or [Uploading Files to Objects Lite Using AWS CLI - Multipart Upload](#) on page 614.

About this task

To upload an image to Objects using Objects Lite and create an image through API, follow these steps:

Procedure

1. Upload the file using the PutObject API with AWS CLI:

```
> aws s3api put-object --bucket vmm-images --body <file-path> --key <object-key> --no-verify-ssl
```

You must use vmm-images as bucket name for images. For more information about command parameters, see [Accessing Objects Lite Using AWS CLI- Put object](#) on page 613. After the file is uploaded, the command displays the ETag of the uploaded object.

For multipart upload, use the commands mentioned in [Uploading Files to Objects Lite Using AWS CLI - Multipart Upload](#) on page 614.

2. After uploading the image to Objects Lite, run the POST request API with the following details:

- **Endpoint URL:**

```
https://<pc-ip>:9440/api/vmm/v4.0/content/images
```

- **Payload:**

```
{
  "name": "CirrosImage",
  "type": "DISK_IMAGE",
  "source": {
    "$objectType": "vmm.v4.content.ObjectsLiteSource",
    "key": "<object-key>"
  }
}
```

- The following list describes the parameters in the endpoint URL and payload:
 - *pc-ip*: The IP address of the Prism Central instance.
 - *name*: The user-defined name of the image.
 - *type*: The type of the image.
 - *key*: Key that identifies the source object in the bucket. Use the same value that you used as key in the put-object CLI.
 - *objectType*: Source object type.

Policies for Image Management

You can define the Image Placement and Bandwidth Throttling policies for managing images. For information about Image Placement and Bandwidth Throttling policies, see [Image Policy Management](#) on page 527.

Catalog Management

Prism Central provides a catalog service to store VM snapshots and images. A Prism Central administrator or project manager creates this catalog of objects for self-service users (who have permissions to create a VM) to use them to deploy VMs in a project.

Note:

- The catalog service is a self-service feature that appears in Prism Central only when Prism Self-Service is enabled. For information about Prism Self-Service, see [Prism Self Service Setup](#) on page 569.
- You can create a custom role and enable or disable the Marketplace Item permissions to manage the Catalog Item entity permissions. For example, if you select **No Access** permission for the Marketplace Item entity for the custom role, the user with the custom role cannot access the **Catalog Item** page. For information about how to create a custom role with access to catalog items entity, see [Creating A Custom Role](#) information in the *AOS Security Guide*.

Catalog Items Summary View

To access the summary view of all the catalog items, perform the following steps:

1. Log in to Prism Central.

2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Compute > Catalog Items** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

The system displays **Catalog Items** page that enables you to view the catalog items available for all the projects in the cluster.




Figure 100: Catalog Items Page

The following table describes the fields that appear in the **Catalog Items** page. A dash (-) is displayed in a field when a value is not available or not applicable.

Table 51: Catalog Items Page - Field Description

Parameter	Description	Values
Name	Displays the catalog item name.	<Catalog_Item_Name>
Type	Displays the type of catalog item (Image or VM).	Image, VM
Description	Indicates the description for the catalog item	(description string)

Filters Pane - Catalog Items page

You can filter the information in the **Catalog Items** page based on the following fields that are available in the **Filter** pane. For information about how to use **Filters** option in Prism Central, see [Prism Central GUI Organization](#) on page 49.

Table 52: Filter Pane Field Description - Catalog Items page

Parameter	Description	Values
Name	Filters based on the catalog item name. It returns a list of catalog items that satisfy the name condition/string.	<Catalog_Item_Name>
Description	Filters based on the description. It returns a list of catalog items whose description field satisfy the condition/string.	(description string)
Type	Filters based on the item type. Select one or more checkbox for the catalog items of the required types.	VM, Image

Adding a Catalog Item

This section describes how to add a VM snapshot or Image as a catalog item in Prism Central.

About this task

You can add the following entities as catalog items:

- *VM* - When you add a VM as the catalog item, a VM snapshot is created. The VM snapshot is available to users across all the self-service projects in the cluster. Users who have the requisite permissions can create VMs from the snapshot.
- *Image* - When you add an image as the catalog item, the image is available to users across all the self-service projects in the cluster. Users who have the requisite permissions can create VMs from the image.

Note: After you add a VM or Image as the catalog item, a copy of these entities (VM snapshot or image) is added to the catalog. There is no impact to the catalog items (VM snapshot or image) even if the VM or image is deleted in Prism Central.

To add a VM snapshot or image as the catalog item, perform the following steps:

Note: Nutanix recommends to power off the VM before you add a VM snapshot to the catalog.

Procedure

1. Log in to Prism Central.
2. Perform any of the following actions based on the type of catalog item you want to add:
 - » *Image* - Navigate to the **Images** page. For information about how to access the **Images** page, see [Images Summary View](#) on page 259.
 - » *VM* - Navigate to the **List** tab of the **VMs** page. For information about how to access the VMs **List** tab, see [VMs Summary View](#) on page 100.

3. To add an image, perform the following actions:

- a. Select the target image checkbox, and choose **Add Image to Catalog** from the **Actions** drop-down menu. The system displays the **Add Image to Catalog** window.




Figure 101: Add Image to Catalog Window

b. Enter the following information in the **Add Image to Catalog** window:

- **Name** - Enter a name for the image.
- **Description** - Enter a description of the image.

c. Click **Save**.

The image is added as the catalog item.

4. To add a VM, perform the following actions:

- a. Select the target VM checkbox, and choose **Add to Catalog** from the **Actions** drop-down menu. The system displays the **Add VM to Catalog** window.

b. Enter the following information in the **Add VM to Catalog** window:

- **Name** - Enter a name for the VM.
- **Description** (optional) - Enter a description of the VM.
- **Guest Customization** - Select the following options for guest operating system (OS) for the VMs that you deploy using this VM template:.
 - From the **Script Type** drop-down list, select the setup script as **Sysprep (Windows)** to customize the Windows OS or **Cloud-init (Linux)** to customize the Linux OS.

Note: If you select **No Customization** at the time of creating the template and allow the users to override the guest customization settings using **Allow users to override at VM Deployment?**

toggle field, it gives the maximum customization control to the users. In this case, the users can customize the script type and the configuration method.

- From the **Configuration Method** drop-down list, for each of these script types selected in **Script Type** field, select either upload a custom script or opt for a guided setup in the field.

Important: The information that you enter is used to customize the OS of the VMs that are deployed using this template.

- c. (Optional) If you select the **Script Type** as **Sysprep (Windows)** and the **Configuration Method** as **Guided Setup**, enter the following information:
 - **Authentication:** Select the checkbox to allow the user who deploys the VM to set a username and password.
 - **Locale:** Select the checkbox to allow the end user to specify the locale (language).
 - **Hostname:** Select the appropriate radio button to specify the host name source: the deployed VM name, a name provided by the person deploying the VM, or restricted hostname access.
 - **License Key:** Specify the license key source. Check the **Enter License Key** radio button and enter the key in the field to set the license key, check the **Allow end user to input License Key** radio button to let the user do it, or check the **No License Key** radio button to not require a license key.
- d. (Optional) If you select the **Script Type** as **Cloud-init (Linux)** and the **Configuration Method** as **Guided Setup**, enter the following information:
 - **Authentication:** Select the checkbox to allow the user who deploys the VM to set a username and password.
 - **SSH Key** Select the checkbox to allow the user who deploys the VM to provide an SSH key.
 - **Locale:** Select the checkbox to allow the end user to specify the locale (language).
 - **Hostname:** Select the appropriate radio button to specify the host name source: the deployed VM name, a name provided by the person deploying the VM, or restricted hostname access.
- e. (Optional) If you select the **Script Type** as **Sysprep (Windows)** and the **Configuration Method** as **Custom Script**, do one of the following:
 - » Click **Upload Script** to upload a script to customize the guest OS of the VMs.
 - » Copy the script and paste the script in the text box.
- f. (Optional) If you select the **Script Type** as **Cloud-init (Linux)** and the **Configuration Method** as **Custom Script**, do one of the following:
 - » Click **Upload Script** to upload a script to customize the guest OS of the VMs.
 - » Copy the script and paste the script in the text box.
- g. Click **Save**.
The VM snapshot is added as the catalog item.

Deleting a Catalog Item

This section describes how to delete a catalog item in Prism Central.

About this task

VM snapshots and images deleted from a catalog makes those items unavailable to project members.

To delete a VM snapshot or image from the catalog items, perform the following steps:

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Compute > Catalog Items** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61
The system displays **Catalog Items** page that enables you to view the catalog items available for all the projects across the registered clusters.
3. Select the target VM snapshot or Image catalog item, and click **Delete Catalog Item**.
The selected catalog item is deleted from all the projects in the cluster.

STORAGE ENTITIES

You can access the following entity items from the **Storage** entity of the **Infrastructure** application:

- Storage Containers (see [Limitations for Storage Containers](#) on page 311)
- Volume Groups (see [Volume Group Management](#) on page 318)
- vCenter Datastores (see [External vCenter Server Integration](#) on page 341)

For information about how to access the entities items available in **Storage** entities, see [Application-specific Navigation Bar](#) on page 61.

Storage Container Management

Storage Management (Prism Central)

Storage in a Nutanix cluster is organized into several components that allow you to manage capacity and performance. Nutanix clusters include storage pool, storage container, volume group, and virtual disk components. Prism Element allows you to manage all aspects of storage for a cluster, but you can also manage storage components directly from Prism Central.

A storage pool is created during cluster configuration. There is one storage pool per cluster. Storage efficiency features that include compression, deduplication, erasure coding, and replication factor are enabled at the storage container level.

The **Storage Containers** page displays information about all the storage containers across the registered clusters.

Note:

- A storage pool and three storage container are created automatically when the cluster is created.
- A storage container is created only if you have configured the Controller VMs (CVMs) with enough memory. The CVM allocation requirements differ depending on the models and features that are used at your site. For more information about the minimum CVM requirements, see [Controller VM \(CVM\) Field Specifications](#) information in *Acropolis Advanced Administration Guide*.
- The NutanixManagementShare, SelfServiceContainer, and a default-container are created by default.
 - The NutanixManagementShare storage container is a built-in storage container for Nutanix clusters for use with the Nutanix Files and Self-Service Portal (SSP) features. This storage container is used by Nutanix Files and SSP for file storage, feature upgrades, and other feature operations. The NutanixManagementShare storage container is not intended to be used as storage for vDisks, including Nutanix Volumes.
 - SelfServiceContainer is a built-in storage container within a Nutanix cluster that is used for storage by VMs created using Image service features such as Self-Service and OpenShift. SelfServiceContainer can also be used like any other container for regular VMs, volume groups, and images. General requirements for using container-level configurations such as compression, deduplication, erasure coding, and replication factor in SelfServiceContainer are the same as any other container.

Important: To ensure proper operation of these features, do not delete these storage containers. Nutanix also recommends that you should not delete these storage containers even if you are not using these features.

- The **Summary** tab in the **Storage Containers** page displays information about storage containers across the registered clusters. For details, see [Storage Containers Summary View](#) on page 290.
- The **Summary** tab in the individual storage container page displays the detailed information about the selected storage container. For more information, see [Storage Container Details View](#) on page 296.
- For information about storage management tasks that are performed through Prism Element web console, see the [Storage Management](#) chapter in the *Prism Element Web Console Guide*.

You can perform the following actions to manage a storage container from Prism Central:

- Create and manage storage directly from Prism Central when the hypervisor is either ESXi, Hyper-V or AHV. For more information, see the following sections:
 - [Creating a Storage Container](#) on page 311
 - [Modifying a Storage Container](#) on page 316
 - [Deleting a Storage Container](#) on page 318
- Enable inline erasure coding. For more information, see [Enabling Inline Erasure Coding](#) on page 309

Storage Containers Summary View

The **Summary** tab in **Storage Containers** page provides the information about all the storage containers across the registered clusters.

To access the summary view of all storage containers, perform the following steps:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and go to **Storage > Storage Containers** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.
3. Click the **Summary** tab.

The system displays **Summary** page for all storage containers.




Figure 102: Summary Page - All Storage Containers

The **Storage Container** page includes five tabs on the left (**Summary**, **List**, **Alerts**, **Events**, and **Metrics**) with a display area below the selected tab.

The **Summary** tab of all storage containers displays the following widgets:

- **Highlighted Entities:** Displays a list of the storage containers with the highest usage of the *<parameter>* you select from the dropdown menu on the right of the widget. The *<parameter>* involves **IO Latency**, **IOPS**, and **Bandwidth**. Click **View All XX Storage Containers** link at the bottom to display the **List** tab.

- **Alerts:** Displays a list of alerts related to storage container that are generated during the specified *<interval>* you select from the dropdown menu on the right of the widget. The *<interval>* involves **Last week**(default), **Last 24 hours**, and **Last 1 hour**. When an alert appears, you can click the graph to view a list of those alerts. Click any alert to display the details page for that alert.
- **Anomalies:** Displays a graph of memory, I/O, CPU, networking, or disk anomalies that occurred during the specified *<interval>* you select from the dropdown menu on the right of the widget. The *<interval>* involves **Last week**(default), **Last 24 hours**, and **Last 1 hour**. When an anomaly appears, you can click the graph to view a list of those anomalies. Click any anomaly to display the event page for that anomaly.

Lists Tab

The **List** tab displays the list of storage containers across all clusters.

Name	Physical Usage	Cluster Fault Tolerance	Replication Factor	Compression	Capacity Deduplication	Erasure Coding	Cluster Name
default-container-701407...	17.27 TiB	1N&1D	3	Off	Off	Off	TestClusterUp...
NutanixManagementShare	67.43 GiB	1N&1D	3	On	Off	Off	TestClusterUp...
objectsif3e7f2f267104df...	0 GiB	1N&1D	3	On	Off	Off	TestClusterUp...
SelfServiceContainer	419.95 GiB	1N&1D	3	Off	Off	Off	TestClusterUp...

Figure 103: Storage Containers List Tab

The following table describes the fields that appear in the **List** tab of **Storage Containers** page:

Note: The fields vary based on the **View by** and **Group by** options. A dash (-) is displayed in a field when a value is not available or applicable. For more information about **View by** and **Group by** options, see [Prism Central GUI Organization](#) on page 49.

Table 53: Storage Containers List Tab - Field Description

Field	Description	Values
Select General from View by option.		
Name	Displays the name of the storage container. You can click the storage container name to view the detailed information for that storage container. For information about storage container details view, see Storage Container Details View on page 296.	<i><Storage_Container_Name></i>
Physical Usage	Displays the amount of used physical storage space. A bar also appears that graphically indicates the amount of used and free storage space available in the storage container. When you hover the cursor over the bar, the system displays a window that lists the used space, free space, and total space in the storage container.	xxx [GiB TiB]

Field	Description	Values
Cluster Fault Tolerance	Displays the fault tolerance, which is the number of failures the cluster can withstand if a node or a disk in the cluster fails. The cluster fault tolerance is specified when the cluster is created.	1N/1D, 1N&1D, 2N/2D
Replication Factor	Displays the replication factor, which is the number of maintained data copies. The replication factor is specified when the storage container is created.	Integer
Compression	Displays whether compression is enabled.	[Off On]
Capacity Deduplication	Displays whether capacity deduplication is enabled, which allows data duplication compression on hard disks (HDD). Performance tier deduplication is a prerequisite for capacity deduplication.	[Off On]
Erasure Coding	Displays whether erasure coding is enabled or not. For more information about Erasure Coding, see Erasure Coding on page 305.	[Off On]
Cluster Name	Displays the name of the cluster in which the storage container resides. You can click the cluster name to view the detailed information for that cluster. For information about cluster details view, see Cluster Details View on page 431.	<Cluster_Name>
Select Performance from View by option.		
Name	Displays the name of the storage container. You can click the storage container name to view the detailed information for that storage container. For information about storage container details view, see Storage Container Details View on page 296.	<Storage_Container_Name>
Free Space (Physical)	Displays the amount of free physical space available in the storage container.	xxx [GiB TiB]
Used Space (Physical)	Displays the amount of used physical space in the storage container.	xxx [GiB TiB]
Total Space (Physical)	Displays the total amount of physical storage space in the storage container.	xxx [TiB]
IOPS	Displays the current I/O operations per second (IOPS) for the storage container. The controller IOPS, I/O bandwidth, and I/O latency fields record the I/O requests serviced by the Controller VM (CVM). The I/O can be served from memory, cache (SSD), or disk.	[0 - unlimited]
IO Bandwidth	Displays I/O bandwidth used per second for Controller VM-serviced requests in this storage container.	xxx [MBps KBps]
IO Latency	Displays the average I/O latency for Controller VM-serviced requests in this storage container.	xxx [ms]
Select Optimization from View by option.		

Field	Description	Values
Name	Displays the name of the storage container. You can click the storage container name to view the detailed information for that storage container. For information about storage container details view, see Storage Container Details View on page 296.	<Storage_Container_Name>
Data Reduction Ratio	Displays the capacity optimization (as a ratio) that results from the combined effects of deduplication, compression, and erasure coding.	xx:1
Data Reduction Savings	Displays the amount of storage capacity saved from the combined effects of deduplication, compression, and erasure coding.	xxx [GiB TiB]
Compression Delay	Displays the time delay to perform compression on the data.	xx [m]
Effective Free Space	Displays the amount of physical free space after data reduction.	xxx [GiB TiB]
Overall Efficiency	Displays the capacity optimization (as a ratio) that results from the combined effects of data reduction (deduplication, compression, and erasure coding), cloning, and thin provisioning.	xx:1

You can perform the following actions for the storage containers in the **Lists** tab:

- Access the detailed information about an individual storage container. For more information, see [Storage Container Details View](#) on page 296.
- Filter the storage containers list based on available parameter values using **Filters** pane. For more information about **Filters** pane, see [Filters Pane - Storage Containers Page](#) on page 293.
- Export the table that contains the list of storage containers and their information to a file in a CSV format. For more information about **Export** option, see [Export](#) on page 54.
- Group the storage containers based on pre-defined criteria. For information about how to group the storage containers, see [Group by](#) on page 51.
- View storage containers based on pre-defined criteria or create a custom view. For information about available views and how to create a custom view, see [View by](#) on page 51.
- Create a storage container. For more information about how to create a storage container, see [Creating a Storage Container](#) on page 311.
- Perform the actions specific to storage container on a single or multiple storage containers using the **Actions** dropdown menu. For instructions on how to perform these actions, see [Modifying a Storage Container](#) on page 316 or [Deleting a Storage Container](#) on page 318.

Note: The **Actions** dropdown menu is grayed out unless you select a storage container or if you select multiple storage containers.

Filters Pane - Storage Containers Page

The following table describes the fields available in the **Filters** pane:

Table 54: Filter Pane Fields

Parameter	Description	Values
Name	Filters based on the storage container name. Select a condition from the dropdown list and enter a string in the field. The system returns a list of storage containers that satisfy the storage container name condition/string.	(storage container name string)
	Note: In this field, the condition menu options are Contains, Doesn't contain, Starts with, Ends with, and Equal to.	
Cluster Name	Filters based on the cluster name. Select a condition from the dropdown list and enter a string in the field. The system returns a list of storage containers that satisfy the cluster name condition/string.	(cluster name string)
Replication Factor	Filters based on the replication factor. Select the checkbox(es) for the desired replication factor (1, 2, 3). Only existing RF values appear. For example, if all containers are at RF 2, just a single checkbox for RF 2 appears. The number of storage containers currently in each setting are displayed in the List tab.	1, 2, 3
Capacity Deduplication	Filters based on the capacity deduplication setting. Select the checkbox(es) for the desired settings to return a list of storage containers with those settings. The number of storage containers currently in each setting are displayed in the List tab.	On, Off, None, Inline, Post Process
Erasure Coding	Filters based on the erasure coding setting. Select the checkbox(es) for the desired settings to return a list of storage containers with those settings. The number of storage containers currently in each setting are displayed in the List tab.	On, Off
Free Space (Physical)	Filters based on the available storage space for a storage container. Select the checkbox(es) for the desired range or enter an amount range in the " From <low> to <high> GiB TiB " field. It returns a list of storage containers with available capacity in that range.	([xx] to [yy] GiB TiB range)
Health	Filters based on the storage container health state. Select one or more states checkboxes to return a list of storage containers in those states. The number of storage containers currently in each state is displayed in the List tab.	Critical, Warning, Good

Parameter	Description	Values
Used Space (Physical)	Filters based on the used storage space for a storage container. Select the checkbox(es) for the desired range or enter an amount range in the " From <low> to <high> GiB TiB " field. It returns a list of storage containers with used storage in that range.	([xx] to [yy] GiB TiB range)
Total Space (Physical)	Filters based on the total available capacity for a storage container. Select the checkbox(es) for the desired range or enter an amount range in the " From <low> to <high> GiB TiB " field. It returns a list of storage containers with maximum capacity in that range.	([xx] to [yy] GiB TiB range)
IOPS	Filters based on the current IOPS. Check the box for the desired range or enter a range in the " From <low> to <high> iops " field. It will return a list of storage containers with IOPS in that range.	([xx] to [yy] IOPS range)
IO Bandwidth	Filters based on the I/O bandwidth used. Check the box for the desired range or enter a range in the " From <low> to <high> MBps " field. It will return a list of storage containers with I/O bandwidth usage in that range.	([xx] to [yy] MBps range)
IO Latency	Filters based on the average I/O latency. Check the box for the desired range or enter a range in the " from <low> to <high> ms " field. It returns a list of storage containers with average I/O latency in that range.	([xx] to [yy] ms range)

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options that are available on the **Activity > Alerts** page, however it only displays the alerts related to storage containers across the registered clusters. For more information about alerts, see *Prism Central Alerts and Events Reference Guide*.

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options that are available on the **Activity > Events** page, however it only displays the events related to storage containers across the registered clusters. For more information about events, see *Prism Central Alerts and Events Reference Guide*.

Metrics Tab

The **Metrics** tab allows you to view performance metrics across the storage containers. Click the **Metrics** tab to display dropdown menu of available metrics, and select the metric name to display the relevant performance information.

Note: The **Metrics** dropdown menu is hypervisor-specific, and might vary based on the hypervisors used in the cluster.

The following is an example showing the dropdown menu of the **Metrics** tab:




Figure 104: Storage Containers Metrics Tab

Table 55: Metrics Tab Fields

Metric	Description
IOPS	Displays total, read, and write IOPS graphs listing current values and total containers (number). The current values are split into intervals (for example, less than 700, 700-1400, 1400-2000, more than 2000).
Note: The same format also applies to the other metrics in this table.	
IO Latency	Displays total, read, and write I/O latency rate graphs.
IO Bandwidth	Displays total, read, and write I/O bandwidth rate graphs.
Storage Usage	Displays storage usage graph.
Optimization	Displays replication factor, data reduction ratio, and overall efficiency graphs.

Storage Container Details View

To access the details view of an individual storage container, perform the following steps:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Storage > Storage Container** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

By default, the system displays the **List** tab with storage containers across all the registered clusters.

3. Click the target `<Storage_Container_Name>` to view the **Summary** tab of an individual storage container.

Note: Replace `<Storage_Container_Name>` with the actual storage container name at your site.

The following is an example showing the **Summary** tab of an individual storage container:




Figure 105: Summary Tab - Individual Storage Container

The **Summary** tab of an individual storage container provides the following widgets:

- **Properties** - Displays summary information about the storage container.
- **Alert**- Displays a list of related alerts that occurred during the specified interval. Select either **Last week** (default), **Last 1 hour** or **Last 24 hours** from the dropdown menu on the top right corner of the widget.
- **Anomalies** - Displays a graph of memory, I/O, CPU, networking, or disk anomalies that occurred during the specified interval. Select either **Last week** (default), **Last 1 hour** or **Last 24 hours** from the dropdown menu on the top right corner of the widget. When an anomaly appears, you can click the graph to display a list of those anomalies. If you click an individual anomaly, the system displays the event page for that anomaly.
- **Usage** - Displays usage data for the storage container.
- **Optimization & Performance** - Displays optimization and performance data for the storage container.

<Action> available above the widgets. Click the appropriate <Action> to run that administrative action on the storage container. For more information about how to perform any <Action>, see [Modifying a Storage Container](#) on page 316 and [Deleting a Storage Container](#) on page 318.

Storage Container Widgets - Field Details

The following table describes the fields in the **Properties**, **Usage**, and **Optimization & Performance** widgets. A dash (-) in a field indicates that there is not enough data to evaluate or a value is not assigned. The displayed fields vary by hypervisor.

Table 56: Storage Container Widgets - Field Description

Field	Description	Values
Properties widget		
Cluster	Displays the name of the cluster in which the storage container resides. You can click the name to view the detailed information about the cluster. For more information, see Cluster Details View on page 431.	(cluster name)
Replication Factor	Displays the replication factor, which is the number of maintained data copies. The replication factor is specified when the storage container is created.	[1-3]
Compression	Displays whether compression is enabled.	[Off On]

Field	Description	Values
Capacity Deduplication	Displays whether capacity deduplication is enabled, which allows data duplication compression on hard disks (HDD). Performance tier deduplication is a prerequisite for capacity deduplication.	[On, Off]
Erasure Coding	Displays whether erasure coding is enabled or not. For more information about Erasure Coding, see Erasure Coding on page 305.	[Off On]
Filesystem allowlists	Displays the IP addresses for file systems that are allowed.	[IP addresses]
Compression Delay	Displays the delay (number of minutes) before data changes are compressed. A zero value indicates that compression is immediate (not delayed).	xx min
Thick Provision	Displays the reserved storage capacity for thick provisioned VMs.	xxx [GiB TiB]
Effective Free Space	Displays the amount of logical free space after data reduction (logical free space x data reduction ratio).	xxx [GiB TiB]
Usage widget		
Free Space (Physical)	Displays the amount of physical free space available in the storage container.	xxx [GiB TiB]
Used Space (Physical)	Displays the amount of physical used space in the storage container.	xxx [GiB TiB]
Reserved Capacity (Physical)	Displays the amount of physical reserved space in the storage container.	xxx [GiB TiB]
Total Space (Physical)	Displays the amount of physical total space in the storage container.	xxx [GiB TiB]
Optimization & Performance widget		
Data Reduction Ratio	Displays the capacity optimization (as a ratio) that results from the combined effects of deduplication, compression, and erasure coding.	xx:1
Data Reduction Savings	Displays the amount of storage capacity saved from the combined effects of deduplication, compression, and erasure coding.	xxx [GiB TiB]
Overall Efficiency	Displays the capacity optimization (as a ratio) that results from the combined effects of data reduction (deduplication, compression, and erasure coding), cloning, and thin provisioning.	xx:1
IOPS	Displays the current I/O operations per second (IOPS) for the storage container. The controller IOPS, I/O bandwidth, and I/O latency fields record the I/O requests serviced by the Controller VM (CVM). The I/O can be served from memory, cache (SSD), or disk.	[0 - unlimited]

Field	Description	Values
IO Bandwidth	Displays I/O bandwidth used per second for Controller VM-served requests in this storage container.	xxx [MBps KBps]
IO Latency	Displays the average I/O latency for Controller VM-served requests in this storage container.	xxx [ms]
Erasure Coding	Displays whether erasure coding is enabled or not. For more information about Erasure Coding, see Erasure Coding on page 305.	[Off On]

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, however it is filtered to display the alerts only for the selected storage container (individual storage container). For more information about alerts, see [Prism Central Alerts and Events Reference Guide](#).

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options as the Events dashboard, except it is filtered to display just events for this storage container. For more information about events, see [Prism Central Alerts and Events Reference Guide](#).

Metrics Tab

The **Metrics** tab allows you to view usage metrics for the storage container. Click the **Metrics** tab to view graphs for all the metrics. The graph is a rolling time interval performance or usage monitor. The baseline range appears as a blue band in the graph.

Note: The baseline range and identified anomalies are based on sophisticated machine-learning capabilities. For more information, see [Behavioral Learning Tools](#) in *Intelligent Operations Guide*. The machine-learning algorithm uses 21 days of data to monitor and predict performance. A graph or baseline band may not appear if less than 21 days of data is available.

You can perform the following actions in the **Metrics** tab:

- Place the cursor anywhere on the horizontal axis to display the value at that time.
- Select the duration (time interval) from the dropdown list on the right (last 1 hour, last 24 hours, last week, last 21 days).
- From the **Actions** dropdown menu that appears on the top-right corner of each metrics widget, you can perform the following operations:
 - Select **Add to Analysis** to add the selected metric to **Analysis** dashboard. For more information, see [Analysis Dashboard](#) in *Intelligent Operation Guide*.
 - Choose **Select Analysis Session** to assign the metric to a target session.
- Click the **Filters** option to select one or more appropriate metric checkboxes to display the selected metrics in the **Metrics** page.

The following table describes the metrics available in **Metrics** tab:

Note: Metrics are hypervisor-dependant, and might not be available on all hypervisors.

Table 57: Metrics Tab Fields

Metric	Description
Storage Controller IOPS	Displays the graph for total I/O operations per second (IOPS) for the storage controller.
Storage Controller Read IOPS	Displays the graphs for Read I/O operations per second (IOPS) for the storage controller.
Storage Controller Write IOPS	Displays the graph for Write I/O operations per second (IOPS) for the storage controller.
Storage Controller Latency	Displays the graph for total latency (in milliseconds) for storage controller
Storage Controller Read Latency	Displays the graph for read I/O latency (in milliseconds) for storage controller.
Storage Controller Write Latency	Displays the graph for write I/O latency (in milliseconds) for storage controller.
Storage Controller I/O Bandwidth	Displays the graph for total I/O bandwidth used per second (MBps or KBps) by storage controller.
Storage Controller Read I/O Bandwidth	Displays the graph of bandwidth used per second (MBps or KBps) for Read I/O operations by storage controller
Storage Controller Write I/O Bandwidth	Displays the graph of bandwidth used per second ((MBps or KBps) for Write I/O operations by storage controller.




Figure 106: Metrics Tab

Storage Usage Tab

The **Storage Usage** tab displays the following graphs:

- **Usage Summary** graph - It displays a rolling time interval monitor of storage container storage usage that can vary from one to several hours depending on activity moving from right to left. Placing the cursor anywhere on

the horizontal axis displays the value at that time. For more in-depth analysis, you can add the monitor to the analysis page by clicking the blue link in the upper right of the graph.

- **Tier-wise Usage** graph - It displays a pie chart divided into the percentage of container storage space used by each disk tier (SSD and DAS-SATA).




Figure 107: Storage Container Usage Tab

Storage Components

Storage in a Nutanix cluster is organized into the following components.

Storage Tiers

Each type of storage hardware (SSD-PCIe (NVMe), SSD (SATA SSD), and HDD) is placed in a storage tier. You can determine the tier breakdown for disks in a storage pool through the web console. For more information, see [Storage Table View](#) information in *Prism Element Web Console Guide*.

Storage Pools

Storage pools consist of groups of physical disks from one or more tiers providing physical separation since each storage device belongs to only one storage pool at a time. Nutanix recommends creating a single storage pool per cluster to enable dynamic optimization of capacity and performance.

While isolating disks into separate storage pools ensures physical separation, it might lead to resource imbalances if the disks remain underutilized. When you expand a cluster by adding new nodes, their disks can also be seamlessly integrated into the existing storage pool. This scale-out architecture allows you to build a cluster that grows with your needs.

By default, when a cluster is created, a predefined storage pool is available. This pool includes the total capacity of all the disks on all the hosts in the cluster.

Storage Containers

A storage container is a subset of available storage within a storage pool. Storage containers are created within a storage pool to hold virtual disks (vDisks) used by virtual machines. For more information, see [Creating a Storage Container](#). By default, storage is thinly provisioned, which means that the physical storage is allocated to a storage container as needed when data is written, rather than allocating the predefined capacity when the storage container is created. Storage efficiency features such as compression, deduplication, and erasure coding are enabled at the container level.

When you create a Nutanix cluster, the following storage containers are created by default:

- **NutanixManagementShare:** The NutanixManagementShare storage container is a built-in storage container designed for Nutanix clusters, specifically supporting Nutanix Files and the Self-Service Portal (SSP) features. This storage container is used for file storage, feature upgrades, and various operational tasks related to these features. To ensure proper functionality, do not delete this storage container. Nutanix recommends that you do not

delete this storage container even if you are not using Nutanix Files or SSP. This storage container is not intended for vDisk storage, including Nutanix Volumes.

- **SelfServiceContainer:** SelfServiceContainer is a built-in storage container within a Nutanix cluster that is used for storage by VMs created using Image service features such as Self-Service and OpenShift. SelfServiceContainer can also be used like any other container for regular VMs, volume groups, and images. General requirements for using container-level configurations such as compression, deduplication, erasure coding, and replication factor in SelfServiceContainer are the same as any other container. Nutanix recommends that you do not delete this storage container.
- **Default-Container-XXXX:** Default-Container-XXXX is a built-in storage container used by VMs to store vDisks for guest VMs and applications. You can rename the Default-Container or delete it and create a new one according to your naming convention.

Volume Groups

A volume group is a collection of logically related virtual disks (or volumes). A volume group is attached to a VM either directly or using iSCSI. You can add vDisks to a volume group, attach them to one or more consumers, include them in disaster recovery policies, and perform other management tasks. You can also detach a volume group from one VM and attach it to another, possibly at a remote location to which the volume group is replicated.

You manage a volume group as a single unit. When a volume group is attached to a VM, the VM can access all the vDisks in the volume group. You can add, remove, and resize the vDisks in a volume group at any time.

Each volume group is identified by a UUID, a name, and an iSCSI target name. Each disk in the volume group also has a UUID and an iSCSI index that specifies ordering within the volume group. A volume group can be configured for either exclusive or shared access.

You can backup, protect, restore, and migrate volume groups. You can include volume groups in a protection domains configured for asynchronous data replication (Async DR), either exclusively or with VMs. However, volume groups cannot be included in a protection domain configured for metro availability, in a protected vStore, or in a consistency group for which application consistent snapshots are enabled.

vDisks

A vDisk is created within a storage container or volume group to provide storage to the virtual machines. A vDisk shows up as a SCSI device when it is mapped to a VM.

Containers for VMware and Hyper-V (Datastores/SMB Shares)

In vSphere, a datastore is a logical container for files necessary for VM operations. Nutanix supports both iSCSI and NFS protocols when mounting a storage volume as a datastore within vSphere. NFS has many performance and scalability advantages over iSCSI, and is the recommended datastore type.

In Hyper-V environments, storage containers are mounted as an SMB share.

Note: Nutanix does not recommend using an AOS Storage container as a general-purpose NFS or SMB share. For NFS and SMB file service, use Nutanix Files.

NFS Datastores. AOS Storage reduces unnecessary network chatter by localizing the data path of guest VM traffic to its host. This boosts performance by eliminating unnecessary hops between remote storage devices that is common with the pairing of iSCSI and VMFS. To enable vMotion and related vSphere features (when using ESX as the hypervisor), each host in the cluster must mount an NFS volume using the same datastore name. The Nutanix web console and nCLI both have a function to create an NFS datastore on multiple hosts in a Nutanix cluster.

To correctly map the local ESXi datastore to the Nutanix container:

- Map the NFS share with 192.168.5.2 (internal IP address) and not the Controller VM IP address or cluster virtual IP address.
- The name of the datastore should be same as the name of the container.




Figure 108: vSphere Configuration of NFS Datastore

SMB Library Share. The Nutanix SMB share implementation is the Hyper-V equivalent of an NFS Datastore, offering feature and performance parity with a vSphere configuration. The registration of an AOS Storage container as an SMB Library share can be accomplished through a single powershell script, or through the Virtual Machine Manager GUI.




Figure 109: Hyper-V Configuration of an SMB Share

Storage Efficiency

Nutanix provides storage efficient capabilities that allows strategic management of storage resources to maximize capacity utilization, enhance performance, and ensure reliability.

In a Nutanix cluster, where data is distributed across multiple nodes, achieving efficiency is critical to reducing costs, optimizing resource usage, and maintaining scalability.

To minimize redundancy, conserve space, and enhance fault tolerance, a Nutanix cluster supports the following storage efficiency capabilities:

- [Deduplication](#) on page 304
- [Compression](#) on page 303
- [Erasure Coding](#) on page 305

Compression

You can enable compression on a storage container to save physical storage space, improve I/O bandwidth, and optimize memory usage thereby enhancing overall system performance.

Note: If the metadata usage is high, compression is automatically disabled. If compression is automatically disabled, an alert is generated.

A storage container in a Nutanix cluster supports the following types of compression:

Post-process compression

Data is compressed after it is written. The delay time between write and compression is configurable, and Nutanix recommends a delay of 60 minutes.

Inline compression

Data is compressed as it is written. When you create a new storage container, inline compression is enabled by default for all license tiers. It is set to a delay of 0, compressing data immediately as it is written.

Data Reduction Ratios and Data Reduction Savings

You can view data reduction ratios and data reduction savings in Prism Central for any of the following entities:

- *For the entire cluster*

Navigate to **Summary** tab in the **Clusters** page, and observe the **Data Reduction** column entry for the target cluster in **Storage Usage** widget. For information about how to access the **Clusters** page, see [Clusters Summary View](#) on page 424.

Note: The **Data Reduction Savings** for the entire cluster is visible only in the **Storage** dashboard in Prism Element. For more information, see [Storage Overview View](#) information in *Prism Element Web Console Guide*.

- *For an individual Storage Container*

Navigate to the **Summary** tab of an individual storage container, and observe the **Data Reduction Ratio** and **Data Reduction Savings** fields in **Optimization & Performance** widget. For information about how to navigate to the **Summary** tab of an individual storage container, see [Storage Container Details View](#) on page 296.

Deduplication

Deduplication reduces space usage by consolidating duplicate data blocks on AOS Storage when you enable capacity deduplication on a storage container.

Important:

- Deduplication is only supported on clusters with a minimum of three nodes.
- If you enable deduplication on storage containers with protected VMs, the system lowers the replication speed.
- Nutanix recommends that you do not enable deduplication for VAAI clone or linked clone environments.

Capacity Deduplication

Capacity deduplication refers to deduplication performed on data stored in hard disk drives (HDDs). Enabling capacity deduplication for persistent data helps reduce storage usage.

Note:

- Capacity deduplication is not enabled by default.

- Capacity deduplication is available if you have purchased a Nutanix Cloud Infrastructure (NCI) Starter or higher license.

Important: Nutanix recommends that you configure the Controller VMs with at least 32 GiB of RAM and 300 GiB SSDs for the metadata disk to enable capacity deduplication.

How to enable Deduplication

Deduplication is enabled at the storage container level in Prism Central. For information about how to enable deduplication, see [Creating a Storage Container \(AHV\)](#) for AHV, [Creating a Storage Container \(ESXi\)](#) for ESXi or [Creating a Storage Container \(Hyper-V\)](#) for Hyper-V.

Deduplication Best Practices

The following table provides the scenarios where deduplication is recommended and where it is not recommended:

Enable deduplication	Do not enable deduplication
• Full clones • Physical-to-virtual (P2V) migration • Persistent desktops	• <i>Linked clones or Nutanix VAAI clones:</i> Duplicate data is managed efficiently by AOS Storage so deduplication has no additional benefit • <i>Server workloads:</i> Redundant data is minimal so may not see significant benefit from deduplication

Erasure Coding

Erasure coding increases the usable capacity of a cluster. Instead of replicating data, erasure coding uses parity information to rebuild data in the event of a disk failure. The capacity savings of erasure coding are in addition to deduplication and compression savings.

Important: Erasure coding is supported on clusters with a minimum of 4 nodes when using replication factor 2 and a minimum of 6 nodes when using replication factor 3.

If you have configured 1N/1D cluster fault tolerance, two data copies are maintained. For example, consider a 6-node cluster with 4 data blocks (a b c d). In this example, we start with 4 data blocks (a b c d) configured with 1N/1D cluster fault tolerance.

The white text represents the data blocks and the green text represents the copies.




Figure 110: Data copies before Erasure Coding

When the data becomes cold, the erasure code engine performs an exclusive OR operation to compute parity “P” for the data.

$$a \quad b \quad c \quad d \quad = \quad P \text{ (parity)}$$

Figure 111: During Computing Parity

After parity is computed, the data block copies are removed and replaced with the parity information. Redundancy through parity results in data reduction because the total data on the system is now $a+b+c+d+P$ instead of $2 \times (a+b+c+d)$.

Note: Each block in the stripe is placed on a separate node to protect from a single node failure.




Figure 112: After Computation of Parity

If the node that contains data block c fails, block c is rebuilt using the rest of the erasure coded stripe (a b d and P) as displayed in the following example:




Figure 113: Post Node Failure

Block c is then placed on a node that does not have any other members of this erasure coded stripe.

Note: When the cluster is configured with 2N/2D cluster fault tolerance, two parity blocks are maintained so that the erasure coded data has the same resiliency as the replicated data. An erasure coded stripe with two parity blocks can handle the failure of two nodes.

Example of Data Reduction Savings from Erasure Coding

The space savings from the erasure coding depends on the cluster size, cluster fault tolerance setting, and percentage of cold data.

You can view the data reduction savings from Erasure coding using any of the following methods:

- Access the [Storage Container Details View](#), and observe the following fields:
 - **Erasure Coding** value as *On* in the **Properties** widget.
 - **Data Reduction Savings** in the **Optimization & Performance** widget.

- Access the [Storage Containers Summary View](#), and observe the **Data Reduction Savings** column for the target storage container using any of the following methods:

- Click the **Lists** tab, and filter the storage containers based on the **Erasure Coding** value as *On*. The system displays the results in the **Lists** page.

For information about how to use filters, see [Filters Pane](#).

- Select **Optimization** from the **View by** option and **Erasure Coding** from the **Group by** option. The system displays the results in the **Lists** page with **Erasure Coding: Off** and **Erasure Coding: On**.

In a 6-node cluster configured with replication factor 2, erasure coding uses a stripe size of 5 where 4 nodes are for data and 1 node is for parity. The sixth node in the cluster ensures that if a node fails, another node is available for rebuild. You can view the erasure coding usage savings from the storage container summary.




Figure 114: Storage Container Summary: Usage Savings Screen

Erasure coding stripe size adapts to the size of the cluster, starting with a minimum 4 nodes and a maximum stripe width of 5 nodes. The following is an example displaying the various configurations of cluster size, possible stripe widths, and approximate savings that might occur when erasure coding is enabled.




Figure 115: Example of Space Saving from Erasure Coding on 20 TiB Nodes

Erasure Coding Best Practices and Requirements

Nutanix recommends the following best practices and requirements to implement Erasure Coding:

- A cluster must have at least four nodes/blocks/racks to enable erasure coding. The cluster can have all four flash nodes or a combination of flash and hybrid nodes, or all hybrid nodes. If erasure coding is enabled on a storage container, a minimum of four blocks for replication factor 2 or six blocks for replication factor 3 is required to maintain block awareness.
- The following table provides the information about the recommended minimum configuration for multiple node removal operations:

Table 58: Minimum Recommended Configuration for Erasure Coding

Desired Awareness Type	Cluster Fault Tolerance	Min. Units	Simultaneous Failure Tolerance
Node	1N/1D	4 nodes	1 node
Node	2N/2D	6 nodes	2 nodes
Block	1N/1D	4 blocks	1 block
Block	2N/2D	6 blocks	2 blocks
Rack	1N/1D	4 racks	1 rack
Rack	2N/2D	6 racks	2 racks

Note: Ensure that you maintain a cluster size that is at least one node greater than the combined strip size (data + parity) to allow space to rebuild the strips if a node fails.

- AOS dynamically calculates the erasure coding strip sizes depending on the number of nodes, blocks, and racks. The maximum supported and recommended strip sizes are (4,1) or (4,2), depending on the nodes, blocks, and racks. Nutanix recommends that you do not change the strip size. Greater strip sizes increase the space savings; however, they also increase the cost of rebuild.
- Erasure coding effectiveness (data reduction savings) might reduce on workloads that have many overwrites outside of the erasure coding window. The default value for erasure coding window is seven days for write cold.
- Read performance is affected during rebuild and the amount depends on cluster strip size and read load on the system.
- Erasure coding is an asynchronous process, and hence the time taken to calculate and display space savings depends on the type and coldness of data. A minimum of two full Curator scans are required to calculate the data savings.
- Ensure that you have replication factor+1 storage heavy or storage only nodes for all-flash clusters with storage heavy nodes. For example, if you have a four-node replication factor 2 enabled cluster, then you must add a minimum of three storage heavy nodes for optimum performance.

Inline Erasure Coding

Inline erasure coding creates erasure coding strips by erasure coding data without waiting for the data to become write cold.

There are two types of inline erasure coding:

- *Same vDisk strips*: Strips that are created using the data blocks from the same vDisk. Nutanix recommends that you configure inline erasure coding type as same vDisk strips for workloads that do not require data locality.
- *Cross vDisk strips*: Strips that are created using the data blocks across multiple vDisks. Nutanix recommends that you configure inline erasure coding type as cross vDisk strips for workloads that require data locality.

By default, same vDisk strips are created when you enable inline erasure coding.

Note: Inline erasure coding with same vDisk strips can be enabled for clusters running AOS version 5.18 or later; and with cross vDisk strips can be enabled for clusters running AOS version 6.6 or later versions.

Enabling Inline Erasure Coding

Inline erasure coding can be enabled only using nCLI. Inline erasure coding is added as a storage container parameter in Zeus.

Before you begin

Caution:

- Nutanix recommends that you enable inline erasure coding for Object storage containers only. To enable inline erasure coding for any other type of storage container, contact Nutanix Support.
- Erasure coding must be enabled on the container to enable inline erasure coding. For information about how to enable erasure coding, see [Creating a Storage Container](#) on page 311.

Procedure

To enable inline erasure coding, perform the following actions:

- Run the following nCLI command:

```
ncli> container create name=container_name sp-id=storage_pool_id erasure-code=on
inline-ec-enabled=true
```

Replace *container_name* and *storage_pool_id* with the storage container name and storage pool ID on which you want to enable erasure coding.

- To explicitly configure inline erasure coding type, run the following nCLI commands:

- For inline erasure coding type: *Same vDisk strips*

```
ncli> container create name=container_name sp-id=storage_pool_id erasure-
code=on inline-ec-enabled=true inline-ec-type=same-vdisk-strips
```

- For inline erasure coding type: *Cross vDisk strips*

```
ncli> container create name=container_name sp-id=storage_pool_id erasure-
code=on inline-ec-enabled=true inline-ec-type=cross-vdisk-strips
```

Replace *container_name* and *storage_pool_id* with the storage container name and storage pool ID on which you want to enable erasure coding.

- To change an existing inline erasure coding type, run the following ncli commands:

- To change to *Same vDisk strips*:

```
ncli> container edit inline-ec-enabled=true inline-ec-type=same-vdisk-strips
id=container_id
```

- To change to *Cross vDisk strips*:

```
ncli> container edit inline-ec-enabled=true inline-ec-type=cross-vdisk-strips
id=container_id
```

Replace *container_id* with the ID of the storage container.

- To verify if inline erasure coding is enabled, run the following nCLI command:

```
ncli> container ls name=container_name
```

Replace *container_name* with the name of the storage container on which you enabled inline erasure coding.

The system displays Inline EC Enabled : true if inline erasure coding is enabled.

Capacity Reservation Best Practices

Capacity reservation ensures that a storage container has a minimum allocated capacity, preventing other storage containers from using that reserved space.

By default, each storage container has access to all of the unused storage in the storage pool. If a storage pool consists of multiple storage containers, one storage container might take all the remaining storage space and leave others with no available space. To ensure that there is space available for a storage container, you can enable capacity reservation.

Nutanix recommends the following best practices to enable capacity reservation:

- Reserve capacity for a storage container only if the storage pool consists of multiple storage containers. Unless there is a specific reason to have multiple storage containers, Nutanix recommends you to configure a single storage pool with a single storage container.
- Do not reserve more than 90% of the total space in the storage pool.
- When you set an advertised capacity for a storage container, be aware that some extra space must be allocated beyond the projected size of any VMs placed in the container. This extra space is to allow room for data that is not

yet garbage collected. The extra space is based on the workload and can be substantial, for example, 10% or more of the storage capacity in some cases.

Limitations for Storage Containers

Containers with replication factor 1 are not supported in Prism Central. You can create containers with replication factor 1 only from Prism Element.

Note: You cannot update or delete the containers with replication factor 1 from Prism Central.

Creating a Storage Container

This section describes how to create a storage container in a cluster.

Before you begin

Ensure that the following prerequisites are met before you create a storage container in AHV cluster:

- The AHV cluster is configured to synchronize time with NTP servers. For information about how to configure the NTP server, see [Configuring NTP Servers for Prism Central](#) in *Prism Central Admin Center Guide*.
- The time on the Controller VMs and Files is synchronization to the current time. If the time on the Controller VMs and Files is ahead of the current time, cluster services might fail to start.

About this task

AOS automatically creates the following type of access to the storage container for each hypervisor:

- *AHV* - Storage container is accessible transparently for AHV.
- *ESXi* - Storage container is accessible as an NFS datastore. In this case, the access to the storage container requires access to the vSphere APIs. Ensure that you have appropriate license of vSphere to access the APIs.
- *Hyper-V* - Storage container is accessible as an SMB share.

Important: To use cross cluster live migration (CCLM) and on-demand cross cluster live migration (OD-CCLM), ensure that both the source and the destination clusters have the same storage container name for the guest VMs. For example, if a SelfServiceContainer storage container exists on the source cluster, the destination cluster must also have a SelfServiceContainer storage container.

For more information on CCLM, see [Cross Cluster Live Migration](#) in *Nutanix Disaster Recovery Guide*.

For more information on OD-CCLM, see [On-Demand Cross-Cluster Live Migration](#) on page 180.

Procedure

To create a storage container, follow these steps:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and go to **Storage > Storage Containers** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.
The system displays the **List** tab as default.

3. Click **Create Storage Container**, and enter the following information:

- a. **Name:** Enter a name for the storage container.

Note: This entity has the following naming restrictions.

- **Container Name Length:** Maximum length is 75 characters.
- **Supported Characters:** Uppercase and lowercase standard Latin letters (A-Z and a-z), Simplified Chinese, decimal digits (0-9), dots (.), hyphens (-), hash (#), and underscores (_).
- **Case Sensitivity:** Container names are case sensitive.

- b. **Clusters:** Select a cluster from the dropdown list.

- If you select an ESXi cluster in the **Clusters** dropdown list, the system prompts you to specify the **NFS Datastore**.
- If you select a Hyper-V cluster in the **Clusters** dropdown list, the system prompts you to specify any of the following mechanisms to make this storage container a default store for VM configuration and virtual hard disk files on Hyper-V:
 - **Make default on all Hyper-V hosts** - Used to make this storage container as a default store location for all the Hyper-V hosts.
 - **Make default on particular Hyper-V hosts** - Used to make this storage container as default store for selected Hyper-V hosts.

- c. **Max Capacity** (*Read-Only*): Displays the amount of free physical space available in the selected cluster. For example, 3.48 TiB (Physical Unreserved Capacity).

4. Click **Advanced Settings** to configure the following additional parameters:

- a. **Cluster Fault Tolerance** (*Read-Only*): Displays the fault tolerance of the cluster that you configured when you created the cluster.

- b. **Replication Factor:** Displays the number of copies of data in the container.

The default value is based on the cluster fault tolerance that you select when you create the cluster. For more information on cluster fault tolerance, see [Cluster Fault Tolerance](#) in the *Prism Web Console Guide*.

To modify the replication factor, select the value from the dropdown menu. You can set the replication factor to 1, 2, or 3 depending on the fault tolerance of the cluster. For more information on replication factor, see [Replication Factor](#) in the *Prism Web Console Guide*.

Note:

- You can modify the replication factor of a container in a cluster with 1N&1D or 2N/2D cluster fault tolerance. You cannot modify the replication factor of a container in a cluster with 1N/1D cluster fault tolerance; however, you can modify the replication factor of the container after you modify the fault tolerance of the cluster from 1N/1D to 2N/2D. For more information, see [Managing Cluster Fault Tolerance](#) on page 439.
- Nutanix supports a replication factor of 1, 2, or 3. Nutanix supports a replication factor of 1 if you enable the replication factor 1 setting only. For more information, see [Enabling Replication Factor 1](#) in the *Prism Web Console Guide*. If you do not enable the replication factor 1 setting, Nutanix supports a replication factor of 2 or 3 depending on the fault tolerance configured in

the cluster. Setting the replication factor to 3 adds an extra layer of data protection at the cost of storing an additional copy of the data.

Create Storage Container



Name

Storage Container Name



Cluster

TestClusterUpgrade



Max Capacity

10.26 TiB (Physical Unreserved Capacity)

Advanced Settings

Cluster Fault Tolerance

1N&1D (Tolerates failure of any 1 Node AND 1 Disk)

Replication Factor

3

Recommended . Original + 2 copies



2

Original + 1 copy

3

Recommended . Original + 2 copies



Reserved Capacity (Physical) -

Advertised Capacity (Logical)



GiB

Advertised Capacity (Physical) -

Compression

Type

Inline Compression



Delay

0

min

☐ Erasure Coding

- c. **Reserved Capacity (Logical):** Enter the capacity (in GiB) to reserve storage space for this storage container.

You can reserve space for a storage container to ensure a minimum storage capacity is available. If you reserving space for a storage container, the reserved space is no longer available to other storage containers even if it is unused. For more information, see [Capacity Reservation Best Practices](#) on page 310.

Reserved Capacity (Physical) (Read-Only): Displays the amount of physical capacity that is reserved based on the logical reserved capacity value.

- d. **Advertised Capacity (Logical):** Enter the capacity (in GiB) to reserve the maximum storage space for this storage container.

This field is used to set an advertised capacity, which is the maximum storage size that the storage container can use. This can be set to any value, but if you configure **Reserved Capacity**, ensure that you set the advertised capacity greater than or equal to the **Reserved Capacity** for the storage container. The hypervisor ensures that the storage does not go beyond the advertised capacity of the storage container.

Advertised Capacity (Physical) (Read-Only): Displays the amount of physical capacity that is advertised based on the logical advertised capacity value.

- e. **Deduplication:** Select the **Capacity** checkbox under **Deduplication** to optimize performance by performing post-process deduplication of persistent data.

Capacity deduplication is primarily recommended for full clone, persistent desktops, and physical to virtual migration use cases that need both storage capacity savings and performance savings from deduplication. For more information, see [Capacity Deduplication](#) on page 304 information.

- f. **Compression:** Select the **Compression** checkbox to compress all the data in the storage container.

For more information about data compression, see [Compression](#) on page 303.

By default, **Inline compression** is selected in the **Type** field and **0** is selected in the **Delay** field (In Minutes) which indicates that the data is compressed immediately as it is written.

To configure the delay time between the write and compression, under **Type** field, select **Post Process Compression** from the dropdown menu. For post-process compression, where data is compressed after it is written, Nutanix recommends that you set a delay of 60 minutes to delay the compression activity for 60 minutes after the initial write operation.

- g. **Erasure Coding:** Select the **Erasure Coding** checkbox to enable erasure coding. Erasure coding increases the effective or usable capacity on a cluster. For more information about erasure coding, see [Erasure Coding](#) on page 305.

- h. **Filesystem allowlists:** Enter the comma-separated IP address and netmask value in the form `ip_address/netmask`.

An allowlist is a set of addresses that are allowed access to this storage container. Allowlists are used to allow appropriate traffic when unauthorized access from other sources is denied. If you set an allowlist at storage container level, the system overrides any global allowlist for this storage container.

Setting an allowlist helps you provide access to the container via NFS. Some manual data migration workflows might require the allowlist to be configured temporally, while some third-party backup vendors might require the allowlist to be configured permanently to access the container via NFS.

Caution:

- User authentication is not available for NFS access, and the IP address in the allowlist has full read or write access to the data on the container.
- Nutanix recommends to allow single IP addresses (with net mask such as 255.255.255.255) instead of allowing subnets (with netmask such as 255.255.255.0).

5. Click **Create**.

The system creates a storage container in the cluster.

Modifying a Storage Container

A storage container is a defined subset of available storage within a storage pool that can be modified as conditions change.

About this task

Storage containers can be modified to change how the data in that storage container is handled. For example, apply compression to the storage container.

Note: You cannot rename a storage container if any of the following conditions is met:

- The cluster is the AHV cluster.
- Storage container include vdisks.

To modify a storage container, perform the following steps:

Procedure

1. Log in to Prism Central.
2. Select **Infrastructure** application from [Application Switcher Function](#) on page 40, and go to **Storage > Storage Containers** from the **Navigation Bar**.

For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

The system displays the **List** tab as default.

3. Select the target storage container checkbox, and choose **Update** from the **Actions** dropdown menu.
The system displays the **Update Storage Container** window.

Update Storage Container

?

×

General

Name

default-container-17475

Cluster

auto_cluster_prod_shubham_shrivastava_f62bbdc9bfa4

Max Capacity

8.25 TiB (Physical)

Advanced Settings

Cluster Fault Tolerance

2N/2D (Tolerates failure of any 2 Nodes OR 2 Disks)

Replication Factor

3

Recommended . Original + 2 copies

Reserved Capacity (Logical)

0

GiB

Reserved Capacity (Physical) -

Advertised Capacity (Logical)

GiB

Advertised Capacity (Physical) -

Compression

Type

None

Delay

0

min

☐ Erasure Coding

Cancel

Save

Figure 116: Update Storage Container

4. Update the storage container configuration as per your requirement, and click **Save**.

The fields that appear in **Update Storage Container** and **Create Storage Container** windows are same. For field details, see [Creating a Storage Container](#) on page 311.

Note:

- If the compression policy is changed from compressed to uncompressed (or vice versa), the existing compressed (uncompressed) data in the storage container will be uncompressed (compressed) as a background process when the next data scan detects the data that needs this change.
- You can modify the replication factor of a container in a cluster with 1N&1D or 2N/2D cluster fault tolerance. You cannot modify the replication factor of a container in a cluster with 1N/1D cluster fault tolerance; however, you can modify the replication factor of the container after you modify the fault tolerance of the cluster from 1N/1D to 2N/2D. For more information, see [Managing Cluster Fault Tolerance](#) on page 439.

Deleting a Storage Container

A storage container is a defined subset of available storage within a storage pool that can be deleted.

About this task

Note: A SelfServiceContainer storage container is created on the target cluster and used by Prism Self Service for storage and other feature operations. To ensure proper operation of these features, do not delete this storage container. For more information about Prism Self Service, see [Prism Self Service Administration](#) on page 569.

To delete a storage container, perform the following steps:

Procedure

1. Log in to Prism Central.
2. Select **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Storage > Storage Containers** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.
The system displays the **List** tab as default.
3. Select the target storage container checkbox, and choose **Delete** from the **Actions** dropdown menu.
The system prompts you to confirm the delete action.
4. Click **Delete**.
The selected storage container is deleted from the cluster.

Volume Group Management

Volumes act as one or more targets for client Windows or Linux operating systems running on a bare metal server or as guest VMs using iSCSI initiators. You can use any storage available in any new or existing Nutanix cluster for Volumes.

Note: Do not use volumes to create an iSCSI datastore to Hyper-V or ESXi hosts. This configuration is not supported.

- The **Summary** tab on the **Volume Groups** page displays information about volume groups across the registered clusters. For more information, see [Volume Groups Summary View](#) on page 319.
- The **Summary** tab on the individual volume group page displays the detailed information about the selected volume group. For more information, see [Volume Group Details View](#) on page 324.

- For more information about volumes and volume groups, see [Volumes Guide](#).

You can create and update volume groups directly from Prism Central web console.

- For information about how to create a volume group, see [Creating a Volume Group](#) on page 332.
- For information about how to modify a volume group, see [Modifying a Volume Group](#) on page 336.

You can configure CHAP secret password from the Prism Central web console. For more information, see [Configuring Mutual CHAP Authentication](#) on page 339

You can attach volume groups to guest VMs from the Prism Central web console. For more information, see [Attaching Volume Groups to Guest VMs](#) on page 340

You can also define the Cluster Role Based Access Control (RBAC) for a volume group. For details, see [Cluster RBAC for Volume Group](#) on page 341

Volume Groups Summary View

The **Summary** tab in **Volume Groups** page displays information about volume groups across the registered clusters and allows you to access detailed information about each volume group.

To access the summary view of all volume groups, perform the following steps:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Storage > Volume Groups** from the **Navigation Bar**. For information on the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

The **Volume Groups** page includes these tabs on the top (**Summary**, **List**, **Alerts**, and **Metrics**) with a display area below the selected tab.

The system displays the **List** tab by default with all the volume groups across registered clusters.

3. Click the **Summary** tab.

The system displays **Summary** page.




Figure 117: Summary Page - All volume Groups

The **Summary** tab of all volume groups displays the following widgets:

- **Highlighted Entities:** Displays a list of the volume groups with the highest usage of the parameter you select from the pull-down menu on the right of the widget. The options are **IO Latency**, **IOPS**, and **IO Bandwidth**. Click the name to display the details page for that volume group. Click the **View all XX Volume Groups** link at the bottom to display the **List** tab.
- **Alerts:** Displays a list of volume group-related alerts that occurred during the specified interval. Select **Last 24 hours** (default), **Last week**, or **Last Hour** from the drop-down menu.

List Tab

The **List** tab displays the list of volume groups across all clusters.




Figure 118: Volume Groups List Tab

The following table describes the fields that appear in the **List** tab of **Volume Groups** page:

Note: The fields vary based on the **View by** and **Group by** options. A dash (-) is displayed in a field when a value is not available or applicable. For more information on **View by** and **Group by** options, see [Prism Central GUI Organization](#) on page 49.

Table 59: Volume Groups List Tab - Field Description

Parameter	Description	Values
Select General from View by option.		
Name	Displays the name of the volume group. You can click the volume group name to view the details page for that volume group. For information on volume group details, see Volume Group Details View on page 324.	<Volume_Group_Name>
Note: An exclamation point appears if the cluster is running a pre-6.0 AOS version, which indicates the available information and actions are limited.		
Usage	Displays the amount of used storage space. A bar also appears that graphically indicates the amount of used and free storage space available in the volume group. When you hover the cursor over the bar, the system displays a window that lists the used space, free space, and total space in the volume group.	xxx [GiB TiB]

Parameter	Description	Values
Disks	Displays the number of virtual disks in the volume group. You can click the number to view the Virtual Disks tab in the Volume Group Details View on page 324.	(integer)
Connections	Displays the number of external client connections to the volume group. You can click the number to view the Connections tab in the Volume Group Details View on page 324.	(integer)
IOPS	Displays the current I/O operations per second (IOPS) for the volume group. The controller IOPS, I/O bandwidth, and I/O latency fields record the I/O requests serviced by the Controller VM. The I/O can be served from memory, cache (SSD), or disk.	[0 - unlimited]
IO Bandwidth	Displays I/O bandwidth used per second for Controller VM-serviced requests in this volume group.	xxx [MBps KBps]
IO Latency	Displays the average I/O latency for Controller VM-serviced requests in this volume group.	xxx [ms]
Cluster Name	Displays the name of the cluster in which the volume group resides. You can click the cluster name to view the cluster details. For information on cluster details, see Cluster Details View on page 431.	(cluster name)
Select Data Protection from View by option.		
Name	Displays the name of the volume group. You can click the volume group name to view the details page for that volume group. For information on volume group details, see Volume Group Details View on page 324.	<Volume_Group_Name>
	Note: An exclamation point appears if the cluster is running a pre-6.0 AOS version, which indicates the available information and actions are limited.	
Category	Displays the category of the volume group. For information on categories in Prism Central, see Category Management in the <i>Prism Central Admin Center Guide</i> .	String (Category name)
Consistency Group	Displays the consistency group of the volume group. For information on Consistency Groups in Prism Central, see Consistency Groups information in <i>Nutanix Disaster Recovery Guide</i> .	Yes/No
Protection Status	Displays the Protection Status of the volume group.	Protected/Unprotected
Protection Type	Displays the type of protection defined for the volume group.	Protection Policy

Parameter	Description	Values
Protection Policy	Displays the protection policy that is applicable for the VM. For more information on protection policies, see Nutanix Disaster Recovery Guide .	(String) Protection Policy Name
Recovery Plan	Displays the recovery plans that is applicable for the volume group. For more information on VG recovery plans, see Nutanix Disaster Recovery Guide .	(String) Recovery Plans Name

You can perform the following actions for the storage containers in the **Lists** tab:

- Access the detailed information about an individual volume group. For more information, see [Volume Group Details View](#) on page 324.
- Filter the volume groups list based on available parameter values using **Filters** pane. For more information on **Filters** pane, see [Filters Pane - Volume Groups Page](#) on page 322.
- Export the table that contains the list of volume groups and their information to a file in a CSV format. For more information on **Export** option, see [Export](#) on page 54.
- Group the volume groups based on pre-defined criteria. For information on how to group the storage containers, see [Group by](#) on page 51.
- View volume groups based on pre-defined criteria or create a custom view. For information on available views and how to create a custom view, see [View by](#) on page 51.
- Create a volume group. For more information on how to create a storage container, see [Creating a Volume Group](#) on page 332.
- Perform the actions specific to storage container on a single or multiple storage containers using the **Actions** dropdown menu. For instructions on how to perform these actions, see [Modifying a Volume Group](#) on page 336.

Note: The **Actions** dropdown menu is grayed out unless you select a volume group.

Filters Pane - Volume Groups Page

The following table describes the fields available in the **Filters** pane:

Table 60: Filter Pane Fields - Volume Groups

Parameter	Description	Values
Name	Filters based on the volume group name. Select a condition from the dropdown menu and enter a string in the field. The system returns a list of storage containers that satisfy the storage container name condition/string.	(volume group name string)

Note: In this field, the condition menu options are **Contains**, **Doesn't contain**, **Starts with**, **Ends with**, and **Equal to**.

Parameter	Description	Values
Usage	Filters on the used storage space for a volume group. Check the box for the desired range or enter an amount range in the " from <low> to <high> GiB " field. It returns a list of volume groups with used storage in that range.	([xx] to [yy] GiB range)
Cluster Name	Filters based on the cluster name. Select a condition from the dropdown menu and enter a string in the field. The system returns a list of storage containers that satisfy the cluster name condition/string.	(cluster name string)
Category	Filters based on category names. Enter a category name in the field and then check the box. As you type a dropdown menu appear to help you select the correct category. A new field appears where you can add more categories to the filter. The number of volume groups tagged to each selected category are displayed in the List tab.	(category name)
IOPS	Filters based on the current IOPS. Check the box for the desired range or enter a range in the From <low> to <high> iops field. It returns a list of volume groups with IOPS in that range.	([xx] to [yy] IOPS range)
IO Bandwidth	Filters based on the I/O bandwidth used. Check the box for the desired range or enter a range in the From <low> to <high> MBps field. It returns a list of volume groups with I/O bandwidth usage in that range.	([xx] to [yy] MBpS range)
IO Latency	Filters based on the average I/O latency. Check the box for the desired range or enter a range in the from <low> to <high> ms field. It returns a list of volume groups with average I/O latency in that range.	([xx] to [yy] ms range)
Internal Volume Group(s)	Filters based on the volume group names specified	Yes No

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options that are available on the **Activity > Alerts** page, however it only displays the alerts related to storage containers across the registered clusters. For more information on alerts, see *Prism Central Alerts and Events Reference Guide*.

Metrics Tab

The **Metrics** tab allows you to view performance metrics across the volume groups. Click the **Metrics** tab to display dropdown menu of available metrics, and select the metric name to display the relevant performance information.

Note: The **Metrics** dropdown menu is hypervisor-specific, and might vary based on the hypervisors used in the cluster.




Figure 119: Volume Groups Metrics Tab

Table 61: Metrics Tab Fields

Metric	Description
Usage	Displays volume group usage graph.
IOPS	Displays total, read, and write IOPS graphs listing current values and total volume groups (number). The current values are split into intervals (for example, less than 700, 700-1400, 1400-2000, more than 2000).
Note: The same format also applies to the other metrics in this table.	
IO Latency	Displays total, read, and write I/O latency rate graphs.
IO Bandwidth	Displays total, read, and write I/O bandwidth rate graphs.

Volume Group Details View

To access the details view of an individual volume group, perform the following steps:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Storage > Volume groups** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

The system displays the **List** tab by default with all the volume groups across registered clusters.

3. Click the target `<Volume_Group_Name>` to view the **Summary** tab of an individual volume group.

Note: Replace `<Volume_Group_Name>` with the actual storage container name at your site.

The following is an example showing the **Summary** tab of an individual volume group:




Figure 120: Summary Tab - Individual Volume Group

The **Summary** tab of an individual volume group provides the following widgets:

- **Properties** - Displays summary information about the volume group. For information about the fields available in **Properties** widget, see [Table 62: Volume Group Widgets - Field Description](#) on page 325.
- **Usage and Performance** - Displays the usage and performance data for the volume group. For information about the fields available in **Usage and Performance** widget, see [Table 62: Volume Group Widgets - Field Description](#) on page 325.
- **Alert-** Displays a list of related alerts that occurred during the specified interval. Select either **Last week** (default), **Last 1 hour** or **Last 24 hours** from the dropdown menu on the top right corner of the widget.

<Action> available above the widgets. Click the appropriate <Action> to run that administrative action on the volume group. For more information about how to perform any <Action>, see [Modifying a Volume Group](#) on page 336.

Volume Group Widgets - Parameter Details

The following table describes the fields available in the **Properties** and **Usage and Performance** widgets. A dash (-) in a field indicates that there is not enough data to evaluate or a value is not assigned. The displayed fields vary by hypervisor.

Table 62: Volume Group Widgets - Field Description

Parameter	Description	Values
Properties widget		
IP Addresses	Displays the IP address assigned to the volume group.	(IP address)
Cluster	Displays the name of the cluster in which the volume group resides. You can click the name to view the detailed information about the cluster. For more information, see Cluster Details View on page 431).	(cluster name)
Flash Mode	Displays whether Flash Mode is enabled for the volume group.	[Enabled, Disabled]
Virtual Disks	Displays the number of virtual disks in the volume group. You can click the numeric value to view the Virtual Disks tab.	(integer)
Connections - External Clients	Displays the number of external client connections to the volume group. You can click the numeric value to view the Connections tab.	(integer)

Parameter	Description	Values
iSCSI Target Name	Displays the name of the iSCSI target.	(iSCSI name)
Cluster	Displays the name of the cluster in which the volume group resides. You can click the cluster name to view the cluster details. For more information, see Cluster Details View on page 431.	(cluster name)
Usage and Performance Widget		
Free Space (Logical)	Displays the amount of logical (effective) free space available in the volume group.	xxx [GiB TiB]
	Note: Logical space accounts for the volume group replication factor. Replication factor 1 means the logical and physical spaces are the same, replication factor 2 means the logical space is half the physical space, and replication factor 3 means the logical space is a third of the physical space.	
Used Space (Logical)	Displays the amount of logical used space in the volume group.	xxx [GiB TiB]
Total Space (Logical)	Displays the amount of logical total space in the volume group.	xxx [GiB TiB]
IOPS	Displays the current I/O operations per second (IOPS) for the volume group. The controller IOPS, I/O bandwidth, and I/O latency fields record the I/O requests serviced by the Controller VM (CVM). The I/O can be served from memory, cache (SSD), or disk.	[0 - unlimited]
IO Bandwidth	Displays I/O bandwidth used per second for Controller VM-serviced requests in this volume group.	xxx [MBps KBps]
IO Latency	Displays the average I/O latency for Controller VM-serviced requests in this volume group.	xxx [ms]
Storage		
Virtual Disks	Displays the number of virtual disks associated with the VG.	(integer)
Defined via	Displays the entity that defines the storage attributes of the VG. The storage attributes of the VG could be inherited from the storage properties of the cluster or applied by a Storage Policy. See Storage Policy Management on page 545 for more information.	(System Default, <Name of Storage Policy>)
Replication Factor	Displays the Replication Factor selected.	Inherit from Container or 2 or 3

Parameter	Description	Values
Encryption	Displays the encryption mode selected.	Enabled or Inherit from Cluster
Compression	Displays the compression mode selected.	Inline or Post Process (if On) Off Inherit from Cluster
QOS Metric (Throttled IOPS)	Displays the throttled throughput value in terms of IOPS.	(Integer number)
QOS Metric (Throttled Throughput)	Displays the throttled throughput value as throughput in MBps.	(Integer number)

For information about Storage Policies, see [Storage Policy Management](#) on page 545.

Virtual Disks Tab

Click the **Virtual Disks** tab to view a list of virtual disks for the volume group.




Figure 121: Virtual Disks Tab

You can perform the following actions for the virtual disks in the **Virtual Disks** tab:

- Filter the Virtual Disks list based on available field values using **Filters** pane. For more information about how to use the **Filters** pane, see [Filters Pane](#) on page 50.
- View Virtual Disks based on pre-defined criteria or create a custom view. For information about available views and how to create a custom view, see [View by](#) on page 51.
- Add, update, and delete a virtual disk. For more information about how to manage a virtual disk, see [Modifying a Volume Group](#) on page 336.

Note: The available actions appear in bold; other actions are grayed out. For grayed out options, a tool tip explaining the reason is provided. The available actions depend on the current state of the selected VM(s).

The following table describes the fields that appear in the list. A dash (-) is displayed in a field when a value is not available or applicable:

Table 63: Virtual Disks - Field Description

Fields	Description	Values
Select General from View by option.		
Index	Displays the index number assigned to the virtual disk.	(integer)
Storage Container	Displays the name of the storage container in which the virtual disk is located.	(container name)
Used Space	Displays the amount of used space in the virtual disk.	xxx [GiB TiB]
Total Space	Displays the total amount of storage space in the virtual disk.	xxx [GiB TiB]
Read IOPS	Displays the current read I/O operations per second (IOPS) for the virtual disk.	[0 - unlimited]
Read Latency	Displays the average read I/O latency for the virtual disk.	xxx [ms]
Read Bandwidth	Displays read I/O bandwidth used per second for the virtual disk.	xxx [MBps KBps]
Write IOPS	Displays the current write I/O operations per second for the virtual disk.	[0 - unlimited]
Write Latency	Displays the average write I/O latency for the virtual disk.	xxx [ms]
Write Bandwidth	Displays write I/O bandwidth used per second for the virtual disk.	xxx [MBps KBps]
Select +Add Custom from View by option to create a custom view with the following additional fields:		
IOPS	Displays the current total I/O operations per second for the virtual disk.	
IO Bandwidth	Displays the total I/O bandwidth used per second for the virtual disk.	
IO Latency	Displays the average I/O latency for the virtual disk.	
Nutanix File Path	Displays the path of the storage file in Nutanix File server. For information about Nutanix Files, see Nutanix Files User Guide .	File path. For example /home/nutanix/

The following table describes the **Filters** pane of **Virtual Disks** tab:

Table 64: Filters Pane - Virtual Disks Tab

Field	Description	Values
Storage container	Filters based on the Storage Container name. Select a condition from the dropdown menu and enter a string in the field. The system returns a list of virtual disks that satisfy the Storage container name condition/string.	(storage container name string)
Note: In this field, the condition menu options are Contains, Doesn't contain, Starts with, Ends with, and Equal to.		
[Total Used] Space	Filters based on the total or used storage space for a virtual disk. Enter an amount range in the From : <low> To: <high> GiB field. The system returns a list of virtual disks with total or used storage in that range.	(([xx] to [yy] GiB range)
[Read Write] IOPS	Filters based on the current read or write IOPS. Enter a range in the From : <low> To: <high> IOPS field. It returns a list of virtual disks with IOPS in that range.	(([xx] to [yy] range)
[Read Write] Bandwidth	Filters based on the read or write I/O bandwidth used. Enter a range in the From : <low> To: <high> MBps field. It returns a list of virtual disks with I/O bandwidth usage in that range.	(([xx] to [yy] range)
[Read Write] Latency	Filters on the average I/O latency. Enter a range in the From : <low> To: <high> ms field. It returns a list of virtual disks with average I/O latency in that range.	(([xx] to [yy] range)

Connections Tab

Click the **Connections** tab to view a list of external client connections to the volume group.




Figure 122: Connections Tab

You can add, update, and delete an external connection for the volume group. For more information about how to manage a volume group connection, see [Modifying a Volume Group](#) on page 336.

Note: The available actions appear in bold; other actions are grayed out. For grayed out options, a tool tip explaining the reason is provided. The available actions depend on the current state of the selected VM(s).

The following table describes the fields that appear in the list. A dash (-) is displayed in a field when a value is not available or applicable:

Table 65: Connections - Field Description

Field	Description	Values
Connection	Displays the IP address or IQN of the external client.	(IP or IQN address)
Type	Displays the type of connection.	external client

Recovery Point Tab

Click the **Recovery Point** tab to view a list of recovery points created with the selected volume group.




Figure 123: Recovery Points Tab

You can clone, revert, replicate, and delete a recovery point of the volume group. For more information about how to manage a volume group recovery points, see [Nutanix Disaster Recovery Guide](#).

The following table describes the fields that appear in the list. A dash (-) is displayed in a field when a value is not available or applicable:

Table 66: Recovery Point - Field Description

Field	Description	Values
Create Time (Name)	Displays the name of the VG recovery point. By default, the system displays the VG recovery point creation time as the name when you create a VG recovery point, however you can change it as per your requirement.	(Name)
Location	Displays the location where the VG recovery point is created in <i>AZ location:Cluster name</i> format.	Local AZ : auto_cluster_prod_f3829eb9649
Replicated From	Displays the location from where the VG recovery point is replicated in <i>AZ location:Cluster name</i> format.	
Reclaimable Space	Displays the space that can be reclaimed for the VG group recovery point.	(Integer) B. For example, 0 B

Field	Description	Values
Expiry	Displays the expiry time of the VG group recovery point.	Timestamp in hh:mm AM/PM, DD MM YYYY. For example, 7:57 PM, 18 Feb 2091
Owner	Displays the owner name.	-
Action	Displays the applicable actions for the VG recovery point.	-

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, however it is filtered to display the alerts only for the selected volume group (individual volume group). For more information about alerts, see [Prism Central Alerts and Events Reference Guide](#).

Metrics Tab

The **Metrics** tab allows you to view usage metrics for the storage container. Click the **Metrics** tab to view graphs for all the metrics. The graph is a rolling time interval performance or usage monitor. The baseline range appears as a blue band in the graph.

Note: The baseline range and identified anomalies are based on sophisticated machine-learning capabilities. For more information, see [Behavioral Learning Tools](#) in *Intelligent Operations Guide*. The machine-learning algorithm uses 21 days of data to monitor and predict performance. A graph or baseline band may not appear if less than 21 days of data is available.

You can perform the following actions in the **Metrics** tab:

- Place the cursor anywhere on the horizontal axis to display the value at that time.
- Select the duration (time interval) from the dropdown menu on the right (Last 1 hour, Last 24 hours, and Last week).
- From the **Actions** dropdown menu that appears on the top-right corner of each metrics widget, you can perform the following operations:
 - Select **Add to Analysis** to add the selected metric to **Analysis** dashboard. For more information, see [Analysis Dashboard](#) in *Intelligent Operation Guide*.
 - Choose **Select Analysis Session** to assign the metric to a target session.
- Click the **Filters** option to select one or more appropriate metric checkboxes to display the selected metrics in the **Metrics** page.

The following table describes the metrics available in **Metrics** tab:

Note: Metrics are hypervisor-dependant, and might not be available on all hypervisors.

Table 67: Metrics Tab Fields

Metric	Description
Usage	

Metric	Description
Controller IOPS	Displays the graph for total I/O operations per second (IOPS) for the controller.
Controller Read IOPS	Displays the graph for Read I/O operations per second (IOPS) for the controller.
Controller Write IOPS	Displays the graphs for Write I/O operations per second (IOPS) for the controller.
Controller AVG I/O Latency	Displays the graph for average latency of total I/O operations (in milliseconds) for controller
Controller AVG Read I/O Latency	Displays the graph for average latency of Read I/O operations (in milliseconds) for controller
Controller AVG Write I/O Latency	Displays the graph for average latency of Write I/O operations (in milliseconds) for controller
Controller I/O Bandwidth	Displays the graph for total I/O bandwidth used per second (MBps or KBps) by controller.
Controller Read I/O Bandwidth	Displays the graph of bandwidth used per second (MBps or KBps) for Read I/O operations by controller
Controller Write I/O Bandwidth	Displays the graph of bandwidth used per second (MBps or KBps) for Write I/O operations by controller.




Figure 124: Metrics Tab

Creating a Volume Group

Create a volume group in Prism Central.

Before you begin

To ensure seamless management and persistent connectivity when using NVMe-oF/TCP with an external client, ensure that the cluster meets the following requirements:

- Configure a persistent discovery connection using the `-p` option in `nvme CLI`. This ensures continued discovery, even when the cluster topology changes.

- Set the autoconnect-on-reboot option to ensure that NVMe devices reconnect automatically after the host reboots. This prevents connection loss across reboots.
- To avoid compatibility issues and ensure stability:
 - Use Rocky Linux 9.4 or later.
 - Use libnvme version 1.7 or later.

About this task

The following collection of tasks describes how to create a volume group in Prism Central:

- [Configuring Volume Group Specifications](#)
- [Configuring the Client Connections](#)
- [Managing Storage Configurations](#)
- [Review](#)

Complete all of the tasks to create a volume group.

To create a volume group, follow these steps:

Procedure

To create a volume group, perform the following steps:

1. Log in to Prism Central.
2. Select **Infrastructure** application from [Application Switcher Function](#) on page 40, and go to **Storage > Volume Groups** from the **Navigation Bar**.
The system displays the **List** tab that lists all the volume groups across registered clusters.
3. Click **Create Volume Group**.
The system displays the **Create Volume Group** window.
4. Complete each of the following configuration procedures:
 - [Configuring Volume Group Specifications](#)
 - [Configuring the Client Connections](#)
 - [Managing Storage Configurations](#)
 - [Review](#)

Configuring Volume Group Specifications

Before you begin

Complete the previous task before configuring the volume group parameters.

About this task

To configure the volume group specifications, in the **Configuration** tab of the **Create Volume Group** window in Prism Central, follow these steps:

Procedure

1. In the **Name** field, enter a name for the volume group.

2. (Optional) In the **Description** field, enter a description for the volume group.
3. From the **Cluster** dropdown menu, select the cluster to create the volume group.
The list includes all clusters registered with the Prism Central.
4. In the **Target Name Prefix** field, enter the target name prefix of the volume group.
5. To add a virtual disk, click **Add Disk**.
In the **Add Virtual Disk** dialog box that opens, perform the following steps:
 - a. From the **Storage Container** dropdown menu, select the target storage container.
 - b. In the **Size** field, enter the disk size (in GiBs).
 - c. Click **Add**.Repeat this step to add additional disks.
To edit or delete a disk, click the Edit icon or Delete icon available under the **Actions** column.
6. (Optional) To view the **Flash Mode** checkbox, click **Advanced Configuration**.
To ensure that no down migration of data occurs from the flash tier, select the **Flash Mode** checkbox.
7. Click **Next**.

Configuring the Client Connections

Before you begin

Complete the previous task before configuring the client connections.

About this task

In the **Connections** tab of the **Create Volume Group** window in Prism Central, follow these steps:

Procedure

1. Depending on the type of client access, do one of the following:
 - » To configure direct attachment to VMs on the same cluster, select **Virtual Machines**.
 - » To configure access to external clients or to clients that are not on the same cluster, select **External Clients**.
2. If you select **Virtual Machines** in Step 1, follow these steps:
 - a. To attach only one VM to the volume group, select the **Non-Shared** option.
To attach multiple VMs to the volume group, select the **Shared** option.
 - b. Select the VMs from the **Select Virtual Machines** dropdown menu.
 - c. (Optional) To enable load balancing between the selected VMs, select the **Load Balance Volume Group** checkbox.
Selecting this checkbox ensures that the virtual disks in the volume group are evenly distributed between the CVMs instead of being hosted locally.

Note: You cannot change the load balance configuration once you attach the VMs.

3. If you select **External Clients** in Step 1, follow these steps:

a. Depending on the type of protocol to connect to the external clients, select one of the following in the **Select Protocol** dropdown menu:

- » To connect the volume group to an external client using the iSCSI protocol, select **iSCSI**.
- » To connect the volume group to an external client using the NVMeoF/TCP protocol, select **NVMeoF/TCP**.

For limitations that apply when you access a volume group using NVMeoF/TCP protocol, [Limitations of Accessing a Volume Group using NVMeoF/TCP](#) on page 340.

b. If you select **iSCSI** in Step 3a, follow these steps:

- To enable the Challenge Handshake Authentication Protocol (CHAP) on the external client, select the **CHAP Authentication** checkbox.
- Enter the password in the **Target Password** field.
- To attach only one external client to the volume group, select the **Non-Shared** option.
To attach multiple external clients to the volume group, select the **Shared** option.
- To attach a new external client, click **Attach New External Client**.
- In the **Add External Client** dialog box that appears, enter the client IP address or IQN designator name in the **Client IQN/IP Address** field, select the **CHAP Authentication** checkbox, enter the password in the **Client Password** field.
- Click **Add**.

The **Attached Clients** table lists the iSCSI external clients that you added.

Note: The system does not allow you to change the client IQN/IP address.

c. If you select **NVMeoF/TCP** in Step 3a, follow these steps:

- To attach a new external client, click **Attach New External Client**.
- In the **Add External Client** dialog box that appears, enter the client NQN in the **Client NQN field**, and click **Add**.

The **Attached Clients** table lists the NVMeoF/TCP external clients that you added.

By default, the system selects the **Non-Shared** option because NVMeoF/TCP does not support attaching multiple external clients to the volume group.

4. Click **Next**.

Managing Storage Configurations

Before you begin

Complete the previous task before managing the storage configurations across the volume group.

About this task

In the **Management** tab of the **Create Volume Group** window in Prism Central, follow these steps:

Procedure

1. To apply the default storage policy to the volume group that you create, toggle the **Enable 'Default-Storage' policy** slider.

The **Enable 'Default-Storage' policy** slider is turned off by default. For more information on the default storage policy, see [Default Storage Policy](#).

If you enable the **Enable 'Default-Storage' policy** slider for the volume group, the system displays the **Storage:\$Default** category by default in the **Categories** dropdown menu. To assign more categories to the volume group, select the categories from the dropdown menu.

2. To assign a category to the volume group, use the **Categories** dropdown menu to search for the category, then select the checkboxes for the categories to assign to the volume group.

Note: To enable the **Enable 'Default-Storage' policy** slider for an entity such as a VM or volume group, do not assign any other category that is already linked to a different storage policy. If the entity is already associated with another category and you enable the **Enable 'Default-Storage' policy** slider, the system adds the **Storage:\$Default** category, but the storage policy linked to the other category overrides the default policy.

3. Click **Next**.

Review

Before you begin

Complete the previous task before reviewing the steps to create a volume group.

About this task

In the **Review** tab of the **Create Volume Group** window in Prism Central, follow these steps:

Procedure

1. Review and verify if all the field entries are correct.
2. (Optional) To modify an entry, click the Edit option associated with the entry.
3. Click **Create**.
The system creates the new volume group, which appears on the **Summary** and **List** pages of the **Volume Groups** window.

Modifying a Volume Group

This section describes how to modify volume groups from the Prism Central web console.

About this task

You can perform the following actions to modify the volume groups from the Prism Central web console.

- Update the volume group settings.
- Delete the volume group
- Manage volume group connections.
- Manage volume group virtual disks.
- Manage volume group categories
- Configure mutual CHAP authentication.

- Attach volume groups to guest VMs.
- Manage the following operations related to Nutanix Disaster Recovery setup.
 - Protect the volume group.
 - Create recovery points for the volume group.
 - Add volume group or volume group category to a recovery plan.

For more information about the operations related to Nutanix Disaster Recovery setup, see [Nutanix Disaster Recovery Guide](#).

Procedure

To modify a volume group, perform the following steps:

1. Log in to the Prism Central web console.
2. Select **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Storage > Volume Groups** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.
The system displays the **List** tab by default with all the volume groups across registered clusters.
3. Select the target volume group checkbox, and click **Update**. The system displays the **Update Volume Group** window - **Configuration** step.

4. Modify the volume group settings as required in **Configuration**, **Connections**, and **Management** steps. In the **Review** step, verify if all the field entries are correct and click **Save**. For field details, see [Creating a Volume Group](#) on page 332.

To delete a volume group, you can select the target volume group and choose **Delete** from the **Actions** dropdown menu.

Note: If there are active connections, you must first remove the connections before deleting the volume group.

You can perform the following actions to manage connections, virtual disks, and categories associated with the volume group:

- To only update volume group connections, you can select the target volume group checkbox, and choose **Manage Connections** from the **Actions** dropdown menu. For field details, see [Creating a Volume Group](#) on page 332.

The following is an example showing the **Manage Connections** window:

Manage Connections

Configure External Clients

☒ CHAP Authentication

Target Password Show

Target Password must be set to enable CHAP Authentication on the Client

iSCSI Target Name: calm-DB0a6584mutanix

Select iSCSI clients or add external IP that you want to allow access to the Volume Group.

Attached Client: 1 Add Client

Client iQN/IP Address	Actions
☒ iqn.1994-05-01.rdg:00000000000000000000000000000000	
☐ iqn.1994-05-01.rdg:00000000000000000000000000000000	

To prevent data loss, only allow multiple concurrent access with filesystem clustering solutions, such as Microsoft Failover Clusters, or Oracle RAC.

Cancel Save

Figure 125: Manage Connections Window - Volume Group

- To only add, update and remove volume group connections, navigate to **Connections** tab of the target volume group from [Volume Group Details View](#) on page 324 , and perform the following action:
 - Click **Add Connection** to add a volume group connection.
 - Select the target volume group connection, and click **Update** to update the volume group connection.
 - Select the target volume group connection, and click **Remove** to delete the volume group connection.
- For field details, see [Creating a Volume Group](#) on page 332.
- To only update, add, and remove virtual disks attached to a volume group, navigate to **Virtual Disks** tab of the target volume group from [Volume Group Details View](#) on page 324, and perform the following actions:
 - Click **Add Virtual Disk** to add a virtual disk to the volume group.

- Select the target volume group virtual disk, and click **Update** to update the volume group virtual disk.
- Select the target volume group virtual disk, and click **Delete** to delete the volume group virtual disk.

For field details, see [Creating a Volume Group](#) on page 332.

- To update volume group categories, you can select the target volume group checkbox, and choose **Manage Categories** from the **Actions** dropdown menu.

In **Manage Categories** window, perform the following steps:

1. Use the [Add icon](#) and [Remove icon](#) available under **Set Categories** field to add or remove the categories for the selected volume group
2. Click **Save**.

For information about how to create a volume group category, see [Creating a Category](#) in the *Prism Central Admin Center Guide*.

Configuring Mutual CHAP Authentication

About this task

You must set the CHAP secret password in the Prism Central web console and then set the same password in the guest VMs hosting your initiators. It is important to use the same password as the target CHAP password; otherwise, the target CHAP password changes, and VM clients connecting to the target experience a loss of connectivity (in both clustered and non-clustered applications).

To set up mutual CHAP authentication in the Prism Central web console, do the following.

Procedure

1. Log in to the Prism Central web console.
2. Select the **Infrastructure** application from Application Switcher Function, and navigate to **Storage > Volume Groups** from the **Navigation Bar**.
3. Select a volume group to configure CHAP authentication and click **Update**.
4. In the **Connections** tab, select **CHAP Authentication**, enter the **Target Password** and then click **Save**.
5. Go to the **List** tab under **Storage > Volume Groups** and from the **View by** drop-down menu, select **General**.
6. Click the volume group and do the following under the **Connections** tab.
 - a. Select the iSCSI client and click **Update** on the top.
 - b. In the **Edit Connections** window, select **CHAP authentication** and enter the **Client Password**.

Note: The client password must be different from the target password, and length should be between 12 and 16 characters.

What to do next

Log in to the guest VM and configure CHAP on it. To set up mutual CHAP authentication in the Windows guest VMs, see [Configuring CHAP Authentication](#) in the *Volumes Guide*.

When you update the target password on the volume group, ensure that you update the associated recovery plan also with the new target password. If you do not update the new target password in the recovery plan, the new target password on the volume group is not retained after the failover.

Attaching Volume Groups to Guest VMs

Procedure

1. Log in to the Prism Central web console.
2. Select the **Infrastructure** application from Application Switcher Function, and navigate to **Storage > Volume Groups** from the **Navigation Bar**.
3. Select the volume group and click **Manage Connections** from the **Actions** drop-down menu.
4. To connect different clients for clustered applications to synchronously replicated volume groups, click **Connect different clients**.

Note: Clustered applications generally span over Nutanix clusters. Therefore, for volume groups protected with synchronous replication, you should have more external clients on the primary or the replicated volume groups.

5. Select the IQN or the IP Address of the volume groups you want to attach.
6. Click **Save**.

Note: Alternatively, you can select the volume group, click **Update** from the **Actions** drop-down menu, and then click **Connect different clients** to attach the clients to the replicated volume groups.

You can use the **Primary** and **Replicated** tabs in the **Update Volume Groups** page to connect an external client either with the primary instance or the replicated instance of the Volume Group. Thus, you can switch between **Primary** and **Replicated** tabs for client connections.

Attaching a Guest VM to a volume group protected with a synchronous replication schedule.

Select external clients and attach them using the add client option.

In the **Add External Client** window, add the Client IQ or IP address and, if necessary, the CHAP authentication client password.

The volume groups are connected to the guest VM. To view the connected volume groups, go to the **List** tab under **Storage > Volume Groups** and then select General from the **View by** drop-down menu on the right.

For both the primary and dormant volume group attachment, log on to the guest VM, discover the iSCSI targets, and establish connections. To complete the iSCSI connections, refer to [Windows Clients](#) or [Linux Clients](#).

Limitations of Accessing a Volume Group using NVMeoF/TCP

The following limitations apply when you access a volume group using the NVMeoF/TCP protocol:

- You cannot access a volume group over the NVMeoF/TCP protocol if the cluster has network segmentation enabled.
- You cannot access a volume group over the NVMeoF/TCP protocol if CHAP authentication is required for client connections.
- You cannot access a volume group over the NVMeoF/TCP protocol from multiple external clients.
- You cannot recover a volume group over the NVMeoF/TCP protocol during a failover.
- You cannot access a volume group over the NVMeoF/TCP protocol if hugepages are not enabled in the cluster.
- You can configure up to 4096 VGs as NVMe endpoints. However, Nutanix recommends that you do not create more than 2500 VGs per cluster to ensure optimal performance and manageability.
- When using VGs as NVMe endpoints, the system supports a maximum of 1024 disks only.

- You cannot use Application Data Services (ADS) dynamic load balancing with NVMeoF-attached volume groups.
- You cannot boot from SAN using the NVMeoF/TCP protocol.

Cluster RBAC for Volume Group

Cluster role-based access control (RBAC) for Volume Group feature enables a super-admin user to provide Prism Admin and Prism Viewer roles access to one or more clusters registered with Prism Central. A user with a Prism Admin role can view and update the entities like volume groups, virtual disks, and storage containers from the allowed clusters. However, a user with a Prism Viewer role can only view the entities.

Cluster RBAC is currently supported on an on-prem Prism Central instance hosted in a Prism Element cluster running AHV. After you enable the Micro Services Infrastructure feature on Prism Central, the Cluster RBAC feature is then automatically enabled.

Cluster RBAC for Volume Group feature is supported on AHV and ESXi clusters.

Note: Prism Central supports Cluster RBAC for VG feature from PC.2022.6 release.

Table 68: List of Permissions for Prism Admin and Prism Viewer Roles

Role	Privileges
Prism Admin	Full administrator privileges except for creating or modifying the user accounts
Prism Viewer	View-only privileges

For information about how to configure cluster RBAC for volume group, see [Configuring Cluster RBAC for Volume Group](#) information in *Security Guide*.

External vCenter Server Integration

The VM management through Prism Central for ESXi provides a unified management interface for all of the external vCenter server instances. All the external vCenter server instances that are registered or not registered with any of the clusters are auto-discovered and displayed. If you have not registered the external vCenter server in Prism Element, you can register the clusters to the external vCenter server instances directly from Prism Central.

During the registration process, you have an option to select the ESXi clusters that you want to manage using Prism Central. For more information about registering external vCenter server, see [Registering a Cluster with a vCenter Server using Prism Central](#) on page 351. After you successfully register external vCenter server instances, you can perform the following operations directly from Prism Central.

- Create, clone, update, and delete VMs.
- Create and delete NICs.
- Attach and delete disks.
- Power operations: Power on or off, reset, suspend, resume.
- Open and launch VM console.
- Enable and disable NGT.

Note: Managing VMware guest tools is not supported through Prism Central.

Note:

- You can perform the power operations and launching of VM console even when external vCenter server is not registered.
- If you are creating VM through Prism, configuration changes to the VM when it is powered on is enabled by default and it depends on the guest operating system that is deployed on the VM.

Rules and Guidelines

- Ensure that all the hosts in the cluster are managed by a single external vCenter server.
- Ensure that DRS is enabled on all the external vCenter server instances.
- Ensure that you are running ESXi and external vCenter server 5.5 or later releases.
- Ensure that you have homogeneous network configuration for all the external vCenter server instances. For example, network should have either 1G or 10G NICs.
- Ensure that you unregister the external vCenter server from the cluster before changing the IP address of the external vCenter server. After you change the IP address of the vCenter Sever, you must register the external vCenter server again with the new IP address.
- The **vCenter Registration** page displays the registered vCenter Server. If for some reason the **Host Connection** field changes to *Not Connected*, it implies that the hosts are being managed by a different external vCenter server. In this case, there will be new vCenter entry with host connection status as Connected and you need to register to this external vCenter server. For more information about registering external vCenter server again, see [Migrating a Nutanix Cluster between Two vCenter Servers using Prism Central](#) on page 352.

Caution: If multiple external vCenter servers are managing the hosts of a single Nutanix cluster, you will not be able to perform the VM management operations. Move all the hosts into one external vCenter server.

Requirements and Limitations

- The E1000, E1000e, PCnet32, VMXNET, VMXNET 2, VMXNET 3 network adapter types (NICs) are supported.
- Only SCSI and IDE disks are supported. SATA and PCI disks are not supported.
- Creating a VM by using a template is not supported.
- Creating a VM by using image service is not supported.
- If a VM is deleted, all the disks that are attached to the VM get deleted.
- Network configuration (creation of port groups or VLANs) is not supported.

External vCenter Datastores Summary View

The **Summary** tab on the **External vCenter Datastores** page provides a dashboard of external vCenter datastores across registered vCenter instances.

To access the summary view of all external vCenter datastores, perform the following steps:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Storage > External vCenter Datastores** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

The system displays the **Lists** tab by default.

3. Click the **Summary** tab.

The system displays the **Summary** view of all the external vCenter datastores in the registered external vCenter server:




Figure 126: Summary View - All External vCenter Datastores

The **External vCenter Datastores** page includes five tabs on the left (**Summary**, **List**, **Alerts**, **Events**, and **Metrics**) with a display area below the selected tab.

Note: This section describes the information and options that appear in the **Summary** page for all external vCenter datastores. For instructions about how to view and organize that information in various ways, see [Prism Central GUI Organization](#) on page 49.

The **Summary** tab for all external vCenter datastores displays the following three widgets:

- **Highlighted Entities:** Displays a list of the external vCenter datastores with the highest usage of the *<parameter>* you select from the dropdown menu on the right of the widget. The *<parameter>* involves only **vSphere Datastore Numberreadaveraged Average**. Click **View All XX External vCenter Datastores** link at the bottom to display the **List** tab.
- **Alerts:** Displays a list of external vCenter datastore-related alerts that are generated during the specified *<interval>* you select from the dropdown menu on the right of the widget. The *<interval>* involves **Last week**(default), **Last 24 hours**, and **Last 1 hour**. When an alert appears, you can click the graph to view a list of those alerts. Click any alert to display the details page for that alert.
- **Anomalies:** Displays a graph of memory, I/O, CPU, networking, or disk anomalies that occurred during the specified *<interval>* you select from the dropdown menu on the right of the widget. The *<interval>* involves **Last week**(default), **Last 24 hours**, and **Last 1 hour**. When an anomaly appears, you can click the graph to view a list of those anomalies. Click any anomaly to display the event page for that anomaly.

List Tab

The **List** tab displays the list of external vCenter datastores across all clusters.

To access the **List** tab of all external vCenter datastore:

1. Log in to Prism Central.

2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Storage > External vCenter Datastores** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

The system displays the **List** tab by default with external vCenter datastore across all the registered clusters.

Note: You can use the **View by** and **Group by** options to create your own customized view and add the necessary columns to that view. For more information on how to **View by** and **Group by** options, see [Prism Central GUI Organization](#) on page 49.

Container Name	Type	Uncommitted	Storage Usage	Free space	Capacity	Cluster
Datastore-1	VMFS	289.55 GiB	136.66 GiB	12.34 GiB	149 GiB	Vcenter_Collector-2
Datastore-2	VMFS	472.45 GiB	1.92 TiB	824.91 GiB	2.73 TiB	Vcenter_Collector-2
Datastore-vcenter	VMFS	117 TiB	481.03 GiB	450.22 GiB	931.25 GiB	Vcenter_Collector-2

Figure 127: External vCenter Datastores - List Tab

The following table describes the fields that appear in the external vCenter datastore **List** tab.

Note: The fields vary based on the **View by** and **Group by** options. A dash (-) is displayed in a field when a value is not available or applicable. For more information about **View by** and **Group by** options, see [Prism Central GUI Organization](#) on page 49.

Table 69: External vCenter Datastores List Tab - Field Description

Field	Description	Values
Container Name	Displays the name of the datastore container. You can click the container name to view the detailed information for that container. For more information about datastore container details view, see External vCenter Datastore Details View on page 347.	(container name)
Type	Displays the file system type.	VMFS, NFS
Uncommitted	Displays the amount of uncommitted disk space.	xx [GiB TiB]
Storage Usage	Displays the amount of used storage space.	xx [GiB TiB]
Free Space	Displays the amount of free storage space.	xx [GiB TiB]
Capacity	Displays the total amount of storage capacity.	xx [GiB TiB]
Cluster	Displays the name of the cluster in which the datastore resides.	(cluster name)

You can perform the following actions for the external vCenter datastores in the **Lists** tab:

- Access the detailed information about an individual external vCenter Datastore. For details, see [External vCenter Datastore Details View](#) on page 347.
- Filter the external vCenter datastores list based on available parameter values using **Filters** pane. For more information about **Filters** pane, see [Filters Pane - External vCenter Datastores Page](#) on page 345.
- Export the table that contains the list of external vCenter datastores and their information to a file in a CSV format. For more information about **Export** option, see [Export](#) on page 54.
- Group the external vCenter datastores based on pre-defined criteria. For information about how to group the external vCenter datastores, see [Group by](#) on page 51.
- View external vCenter datastores based on pre-defined criteria or create a custom view. For information about available views and how to create a custom view, see [View by](#) on page 51.

Filters Pane - External vCenter Datastores Page

The following table describes the fields available in the **Filters** pane:

Table 70: Filter Pane Fields

Parameter	Description	Values
Container Name	Filters based on the container name. Select a condition from the dropdown menu and enter a string in the field. The system returns a list of containers that satisfy the container name condition/string.	(name string)
Note: In this field, the condition menu options are Contains , Doesn't contain , Starts with , Ends with , and Equal to .		
Cluster	Filters on the cluster name. Enter a string in the field to display a list of clusters that satisfy the name string.	(cluster name)
Storage Usage	Filters on the amount of used storage space. Check the box(es) for the desired range or enter a range in the " from <low> to <high> TiB " field. The number of containers in each range is displayed on the right of the line.	([xx] to [yy] range)
Capacity	Filters based on the amount of total storage capacity.	([xx] to [yy] range)
Free Space	Filters based on the amount of free storage space.	([xx] to [yy] range)
Uncommitted	Filters based on the amount of uncommitted storage space.	([xx] to [yy] range)
Type	Filters based on file system type.	VMFS, NFS

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options that are available on the **Activity > Alerts** page, however it only displays the vCenter Datastore-related alerts across the registered clusters. For more information about alerts, see *Prism Central Alerts and Events Reference Guide*.

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options that are available on the **Activity > Events** page, however it only displays the vCenter Datastore-related events across the registered clusters. For more information about events, see *Prism Central Alerts and Events Reference Guide*.

Metrics Tab

The **Metrics** tab allows you to view performance metrics across the external vCenter datastores. Click the **Metrics** tab to display dropdown menu of available metrics, and select the metric name to display the relevant performance information.

Note: The **Metrics** dropdown menu is hypervisor-specific, and might vary based on the hypervisors used in the cluster.

The following is an example showing the dropdown menu of the **Metrics** tab:




Figure 128: External vCenter Datastores Metrics Tab

The following table describes the dropdown menu of the **Metrics** tab:

Table 71: Metrics Tab Dropdown Menu

Metric	Description
IOPS	Displays the number of read and write IOPS for the listed number of containers. You can click the IOPS value to view the Summary tab with a filtered list of containers for the selected IOPS value. The Total aggregate column summaries field displays the number of containers for which the IOPS value is applicable.
Disk Usage	Displays the disk usage amount for the listed number of containers. You can click the disk usage amount to view the Summary tab with a filtered list of containers for the selected disk usage amount. The Total aggregate column summaries field displays the number of containers for which the disk usage value is applicable.

Metric	Description
Disk Capacity	Displays the disk capacity for the listed number of containers. You can click the disk capacity to view the Summary tab with a filtered list of containers for the selected disk capacity. The Total aggregate column summaries field displays the number of containers for which the disk capacity value is applicable.
Disk Provisioned	Displays the provisioned disk space for the listed number of containers. You can click the provisioned value to view the Summary tab with a filtered list of containers for the selected provisioned value. The Total aggregate column summaries field displays the number of containers for which the disk provisioned value is applicable.

External vCenter Datastore Details View

Summary Tab

The **Summary** tab of an individual external vCenter datastore consists of a dashboard that provides the detailed information about the external vCenter datastore.

To access the **Summary** tab of an individual external vCenter datastore:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Storage > Storage Container** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

By default, the system displays the **List** tab with external vCenter datastores across all the registered clusters.

3. Click the target <vCenter_Datastore_Name> to view the **Summary** tab of an individual external vCenter datastore.

Note: Replace <vCenter_Datastore_Name> with the actual external vCenter datastore name at your site.

The following is an example showing the **Summary** tab of an individual external vCenter datastore:




Figure 129: Summary Tab - Individual External vCenter Datastore

The **Summary** tab of an individual external vCenter datastore provides the following widgets:

- **Properties** - Displays summary information about the external vCenter datastore. For information about the fields available in **Properties** widget, see [External vCenter Datastore Widgets - Parameter Details](#) on page 348.

- **Alerts**- Displays a list of related alerts that occurred during the specified interval. Select either **Last week** (default), **Last 1 hour** or **Last 24 hours** from the dropdown menu on the top right corner of the widget.
- **Anomalies** - Displays a graph of memory, I/O, CPU, networking, or disk anomalies that occurred during the specified interval. Select either **Last week** (default), **Last 1 hour** or **Last 24 hours** from the dropdown menu on the top right corner of the widget. When an anomaly appears, you can click the graph to display a list of those anomalies. If you click an individual anomaly, the system displays the event page for that anomaly.
- **Metrics** - Displays the average metrics information for **Last 15 minutes avg** for the external vCenter datastore (container) that includes **IOPS Read**, **IOPS Writes**, **Disk Usage**, **Disk Capacity**, and **Disk Provisioned** fields. For information about the fields available in **Metrics** widget, see [External vCenter Datastore Widgets - Parameter Details](#) on page 348.

External vCenter Datastore Widgets - Parameter Details

The following table describes the fields in the **Properties** and **Metrics** widgets. A dash (-) in a field indicates that there is not enough data to evaluate or a value is not assigned. The displayed fields vary by hypervisor.

Table 72: External vCenter Datastore Widgets - Field Description

Field	Description	Values
Properties widget		
Capacity	Displays the total amount of storage capacity.	xx [GiB TiB]
Storage Usage	Displays the amount of used storage space.	xx [GiB TiB]
Free Space	Displays the amount of free storage space.	xx [GiB TiB]
Uncommitted	Displays the amount of uncommitted disk space.	xx [GiB TiB]
Cluster	Displays the name of the cluster in which the datastore resides. Click on the cluster name to displays details about that cluster.	(cluster name)
Type	Displays the file system type.	VMFS, NFS
Metrics widget		
IOPS Reads	Displays the average number of read IOPS for the selected container in last 15 minutes.	Integer (IOPS)
IOPS Writes	Displays the average number of write IOPS for the selected container in last 15 minutes.	Integer (IOPS)
Disk Usage	Displays the average disk usage amount for the selected container in last 15 minutes.	Integer (MiB)
Disk Capacity	Displays the total disk capacity for the selected container.	Integer (GiB)
Disk Provisioned	Displays the average provisioned disk space for the selected container in last 15 minutes.	Integer (MiB)

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, however it is filtered to display the alerts only for the selected external vCenter datastore (individual external vCenter datastore). For more information about alerts, see *Prism Central Alerts and Events Reference Guide*.

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options as the Events dashboard, except it is filtered to display just events for this external vCenter datastore. For more information about events, see [Prism Central Alerts and Events Reference Guide](#).

Metrics Tab

The **Metrics** tab allows you to view usage metrics for the external vCenter datastore. Click the **Metrics** tab to view graphs for all the metrics. The graph is a rolling time interval performance or usage monitor. The baseline range appears as a blue band in the graph.




Figure 130: Metrics Tab

Note: The baseline range and identified anomalies are based on sophisticated machine-learning capabilities. For more information, see [Behavioral Learning Tools](#) in *Intelligent Operations Guide*. The machine-learning algorithm uses 21 days of data to monitor and predict performance. A graph or baseline band may not appear if less than 21 days of data is available.

You can perform the following actions in the **Metrics** tab:

- Place the cursor anywhere on the horizontal axis to display the value at that time.
- Select the duration (time interval) from the dropdown menu on the right (last 1 hour, last 24 hours, last week, last 21 days).

- From the **Actions** dropdown menu that appears on the top-right corner of each metrics widget, you can perform the following operations:
 - Select **Add to Analysis** to add the selected metric to **Analysis** dashboard. For more information, see [Analysis Dashboard](#) in *Intelligent Operation Guide*.
 - Choose **Select Analysis Session** to add the metric to an existing or new analysis session. The system displays the **Select A Session** window.

Figure 131: Session Window

Select the target session from the **Session** dropdown menu, and then click **Add to Session**.

- Select **Alert Settings** to create a custom alert policy for the metric. For more information about how to create a custom alert policy, see [Prism Central Alerts and Events Reference Guide](#).
- Click the **Filters** option to select one or more appropriate metric checkboxes to display the selected metrics in the **Metrics** page.

The following table describes the metrics available in **Metrics** tab:

Note: Metrics are hypervisor-dependant, and might not be available on all hypervisors.

Table 73: Metrics Tab Fields

Metric	Description
Average read requests per second	Average number of read commands issued per second to the external vCenter datastore during the collection interval.
Average write requests per second	Average number of write commands issued per second to the external vCenter datastore during the collection interval.
Space actually used	Amount of space actually used by the virtual machine or the external vCenter datastore. It might be less than the amount provisioned at any given time, depending on whether the virtual machine is powered-off, whether snapshots have been created or not, and other such factors. Available from a datastore and virtual machine target entities.
Capacity	Configured size of the external vCenter datastore. Available from an external vCenter datastore entity only.

Metric	Description
Space potentially used	Amount of storage set aside for use by an external vCenter datastore or a virtual machine. Files on the external vCenter datastore and the virtual machine can expand to this size, but not beyond it. Available from an external vCenter datastore and virtual machine target entities.

Registering a Cluster with a vCenter Server using Prism Central

Register your Nutanix cluster with an external vCenter Server using Prism Central.

Before you begin

- Ensure that your cluster is connected to the external vCenter Server.
- If you choose to register a cluster with an external vCenter Server using extension-based authentication, ensure that you have the required extension privileges. These privileges provide permission to register the Nutanix cluster with vCenter Server using extension-based authentication.

About this task

You can register a cluster with an external vCenter Server directly from Prism Central.

Important:

- Nutanix stores external vCenter Server credentials only when you register your cluster with vCenter Server using SSO-based authentication. Nutanix does not store credentials when you use extension-based authentication for registration.
- Whenever a new node is added to a cluster, vCenter Server registration for the new node is automatically performed.

To register a cluster with an external vCenter Server using Prism Central, follow these steps:

Procedure

1. Log in to Prism Central.
2. Select **Infrastructure** application from the [Application Switcher Function](#) on page 40.
3. From the left navigation menu, select **Prism Central Settings**.
The system displays the **Prism Central Settings** page.
4. From the **Setup** section in the left navigation menu, click **vCenter Integration**.
The system displays the **vCenter Integration** page that lists all the clusters that are connected to an external vCenter Server.
5. Select the checkbox associated with a cluster to register with vCenter Server and click **Connect Integration**.
The system displays the **Connect a Cluster** dialog box. The system auto-discovers the external vCenter Server that manages the hosts in the cluster and displays the IP address of the vCenter Server in the vCenter IP Address field, along with the port number in the PORT field.

Note: Do not change the port number. For the complete list of required ports, see [Ports and Protocols](#).

6. In the **Username** and **Password** fields, enter the user name and password of the external vCenter Server.

7. From the Authentication Method dropdown menu, select one of the following:

- » **Extension:** Choose this option to register using the extension-based authentication method.
- » **SSO:** Choose this option to register using the SSO-based authentication method.

8. Click **Save**.

The system starts registering the cluster with the vCenter Server and generates a certificate to communicate with the external vCenter Server during the registration process.

After registration is complete, the **Integration Status** column for the cluster on the **vCenter Integration** page shows **Active**, and the **Authentication Method** column displays the method you used to register the cluster.

Also, a relevant message is displayed in the **Tasks** page. For more information on how to access the Tasks page, see [Tasks View](#) on page 496.

Unregistering a Cluster from the External vCenter Server using Prism Central

Unregister your cluster from an external vCenter Server using Prism Central.

About this task

To unregister a cluster from an external vCenter Server, follow these steps:

Procedure

To unregister the cluster from the external vCenter server, perform the following steps:

1. Log in to Prism Central.
2. Select **Infrastructure** application from the [Application Switcher Function](#) on page 40.
3. From the left navigation menu, select **Prism Central Settings**.
The system displays the **Prism Central Settings** page.
4. From the **Setup** section in the left navigation menu, click **vCenter Integration**.
The system displays the **vCenter Integration** page that lists all the clusters that are registered or connected to an external vCenter Server.
5. Select the checkbox associated with a cluster to unregister from vCenter Server and click **Disconnect Integration**.
The system prompts you to confirm the disconnect action.
6. Perform one of the following:
 - » If the cluster was registered to the vCenter Server using SSO-based authentication, click **Disconnect Integration**.
 - » If the cluster was registered to the vCenter Server using extension-based authentication, enter the administrator user name and password of the vCenter Server in the **Username** and **Password** fields, and click **Disconnect Integration**.

The system starts unregistering the cluster from the vCenter Server.

After unregistration is complete, the **Integration Status** column for the cluster on the **vCenter Integration** page displays **Not Configured**. Also, a relevant message is displayed in the **Tasks** page. For more information on how to access the Tasks page, see [Tasks View](#) on page 496.

Migrating a Nutanix Cluster between Two vCenter Servers using Prism Central

Migrate a cluster from an existing external vCenter Server instance to a new vCenter Server instance using Prism Central.

About this task

The vCenter Integration page displays the clusters registered to an external vCenter Server. If the system displays **Not Configured** in the **Integration Status** column for a cluster, it implies that the clusters (hosts) are managed from a different external vCenter server. In this case, the system also displays a new external vCenter server instance with **Integration Status** field as **Configured** and you must register the cluster to this external vCenter server instance.

To migrate a cluster from an existing external vCenter Server instance to a new vCenter Server instance, follow these steps:

Procedure

To re-register the external vCenter server with Prism Central, perform the following steps:

1. Log in to Prism Central.
2. Select **Infrastructure** application from the [Application Switcher Function](#) on page 40.
3. From the left navigation menu, select **Prism Central Settings**.
The system displays the **Prism Central Settings** page.
4. From the **Setup** section in the left navigation menu, click **vCenter Integration**.
The system displays the **vCenter Integration** page that lists all the clusters that are connected to an external vCenter Server.
5. Unregister the cluster from the existing vCenter Server instance by clicking **Disconnect Integration**.
For more information on how to unregister a cluster from a vCenter Server, see [Unregistering a Cluster from the External vCenter Server using Prism Central](#) on page 352.
6. Register the cluster to a new vCenter Server instance by clicking **Connect Integration**.
For more information on how to register a cluster with a vCenter Server, see [Registering a Cluster with a vCenter Server using Prism Central](#) on page 351.

Updating the vCenter Service Account Credentials in Prism Central

Update the vCenter Server service account credentials used to register a cluster with vCenter Server in Prism Central.

Before you begin

Ensure that you register the cluster with vCenter Server using SSO-based authentication.

About this task

Each time you update your user credentials in vCenter Server, you must also update them in Prism Central. If you do not, Prism Central cannot perform VM operations, and the **Status** column for a cluster on the vCenter Integration page displays **Inactive**.

Procedure

1. Log in to Prism Central.
2. Select **Infrastructure** application from the [Application Switcher Function](#) on page 40.
3. From the left navigation menu, select **Prism Central Settings**.
The system displays the **Prism Central Settings** page.

4. From the **Setup** section in the left navigation menu, click **vCenter Integration**.
The system displays the **vCenter Integration** page that lists all the clusters that are connected to an external vCenter Server.
5. To update the credentials, select the checkbox associated with the cluster registered to the vCenter Server and click **Update Integration**.
The system displays the **Update Integration** dialog box.
6. In the **Username** and **Password** fields, enter the new administrator user name and password of the vCenter Server.
7. Click **Save**.

NETWORK AND SECURITY ENTITIES

You can access the following networking and security entity items from the **Network and Security** entity of the **Infrastructure** application. For information on how to access the entity items available in **Network and Security** entity, see [Application-specific Navigation Bar](#) on page 61.

- **Subnets:** This page displays the subnets and the operations you can perform on subnets. For more information, see [Subnets](#) on page 365.
- **Virtual Private Clouds:** This page displays the VPCs and the operations you can perform on VPCs. For more information, see [Virtual Private Clouds Summary View](#) on page 370.
- **Floating IPs:** This page displays a list of floating IP addresses that you are using in the network. It allows you to request for floating IP addresses from the free pool of IP addresses available to the clusters managed by the Prism Central instance. For more information, see [Floating IPs Summary View](#) on page 403.
- **Connectivity:** This page allows you to manage the following networking capabilities. For more information, see [Connectivity](#) on page 404.
 - **Gateways:** This page provides a list of network Gateways you have created and configured, and the operations you can perform on the network Gateways. For more information, see [Gateways Summary View](#) on page 405.
 - **VPN Connections:** This page provides a list of VPN connections you have created and configured, and the operations you can perform on the VPN connections. For more information, see [VPN Connections Summary View](#) on page 408.
 - **Subnet Extensions:** This page provides a list of subnets that you have extended at the Layer 2 level using VPN (point-to-point over Nutanix VPN) or VTEP (point-to-multi-point including third party). For more information, see [Subnet Extensions Summary View](#) on page 412.
 - **BGP Sessions:** This page provides a list of BGP sessions you have created and configured, and the operations you can perform on the BGP sessions. For more information, see [BGP Sessions Summary View](#) on page 416.
- **Security Policies:** This page provides a list of security policies you configured using Flow Segmentation. For more information, see [Security Policies](#) on page 420.
- **Security Dashboard:** This page provides dynamic summary of the security posture across all registered clusters. For more information, see [Security Dashboard](#) on page 420.

For information on how to configure network connections, see [Network Configuration](#).

Subnets (Overlay IP subnets), Virtual private clouds, floating IPs, and Connectivity are Flow virtual networking features. These features support flexible app-driven networking that focuses on VMs and applications instead of virtual LANs and network addresses. Flow virtual networking powers network virtualization to offer a seamless network experience with enhanced security. It is disabled by default. It is a software-defined network virtualization solution providing overlay capabilities for the on-premises AHV clusters.

Security policies drives the Flow Segmentation features for secure communications. For more information, see [Flow Microsegmentation Guide](#).

Virtual Switches

You can create and manage multicluster virtual switches only in Prism Central.

The virtual switches that you create and manage in Prism Central span multiple clusters managed by the Prism Central instance. Deploying and using these virtual switches, also called multicluster virtual switches, in Prism Central requires minimum Nutanix software versions as follows:

- Prism Central pc.7.3
- Network Controller 6.0.0
- AOS 7.3
- AHV 10.3

A multicluster virtual switch that you create on Prism Central is a software layer that represents the uplink ports that are connected to the same switching domain, across hosts and clusters, and share the same attributes. You can map subnets to such a multicluster virtual switch to span across the clusters that the virtual switch represents. You can also create VM vNICs for such subnets that provide multicluster reachability. You can configure a multicluster virtual switch to manage all the hosts and uplink ports connected to the same Ethernet broadcast domain across all the clusters managed by Prism Central.

Note: A multicluster virtual switch is different from a single-cluster default virtual switch created in a Prism Element cluster. A single-cluster default virtual switch that the cluster deployment process creates in the Prism Element web console physically connects all the hosts and ports on the cluster in which it is created.

You can map only a (Network Controller-based) VLAN Subnet to a multicluster virtual switch. For more information on VLAN Subnet association, see [Assign a Subnet to a Virtual Switch](#) on page 365. You cannot map a (AHV-based) VLAN Basic Subnet or an individual Overlay subnet to a virtual switch that you create in Prism Central.

Note: While you create a VLAN Subnet, the multicluster virtual switches created in Prism Central are not displayed in the **Virtual Switch** dropdown list in the **Create Subnet** page. Ensure that you clear the **VLAN Basic Networking** checkbox in **Advanced Configuration** to display the multicluster virtual switches in the **Virtual switch** dropdown list.

For information on single-cluster virtual switches, see [About Virtual Switch](#) in the *AHV Administration Guide* and [Creating or Updating a Virtual Switch](#) in the *Prism Element Web Console Guide*.

Limitations

The limitations of Prism Central-based multicluster default virtual switches are as follows:

- You cannot attach VLAN Basic Subnets or individual Overlay subnets.
- A virtual switch created on a Prism Central instance is a multicluster virtual switch even if the Prism Central instance manages only one cluster.

Only a virtual switch created in the Prism Element web console is a single-cluster virtual switch.

- You cannot migrate an existing single-cluster virtual switch created in a Prism Element cluster, to a multicluster (Prism Central-based) virtual switch.

You can perform the following actions to manage a multicluster virtual switch in the **Virtual Switches** dashboard of Prism Central:

- Create a multicluster virtual switch. For more information, see [Creating a Virtual Switch](#) on page 362.
- View the list of multicluster virtual switches that were created in the Prism Central instance. For more information, see [Viewing Virtual Switches Summary Page](#) on page 357.
- View a specific multicluster virtual switch that was created in the Prism Central instance. For more information, see [Viewing Virtual Switch Details](#) on page 360.
- Update a multicluster virtual switch. For more information, see [Updating a Virtual Switch](#) on page 363.
- Delete a multicluster virtual switch. For more information, see [Deleting a Virtual Switch](#) on page 364.

Viewing Virtual Switches Summary Page

You can view and manage virtual switches in the **Virtual Switches** dashboard.

About this task

The **Virtual Switches** dashboard displays the list of virtual switches in all the clusters managed by Prism Central AZ. You can also update or delete virtual switches that you create on the dashboard.

Procedure

1. Log in to Prism Central.
2. From [Application Switcher Function](#) on page 40, select the **Infrastructure** application and navigate to **Network & Security > Virtual Switches** from the **Navigation Bar**.

The **Virtual Switches** page appears, displaying the **List** tab. This tab lists all the virtual switches configured in all the clusters managed by Prism Central.

For information on the information and configurations in the **List** tab, see [Virtual Switches Summary View Attributes](#) on page 357.

Virtual Switches Summary View Attributes

The **List** tab of the **Virtual Switches** dashboard provides a list of all the virtual switches in all the clusters managed by Prism Central. It provides the details of the listed virtual switches in a table.

Table 74: Attribute Description and Value

Attribute	Description	Values
Name	Displays the name of the virtual switch. If the virtual switch is the default virtual switch created on the Prism Element cluster, the name is suffixed with Single Cluster. If you created the virtual switch on Prism Central, the name is not suffixed.	(Virtual switch name string)
Scope	Displays whether the virtual switch was created on Prism Element web console or Prism Central. The scope of the virtual switch is displayed as Single Cluster or Multi Cluster respectively. You can view this attribute only by creating a custom view and adding this attribute to the view. For more information on configuring a custom view, see <i>View by</i> in Prism Central GUI Organization on page 49.	Multi Cluster Single Cluster

Attribute	Description	Values
Cluster	Displays the name or names of the cluster or clusters that the virtual switch uplinks.	(Cluster name)
Bond Type	Displays the bond type associated with the virtual switch. For more information on bond types, see the <i>Bond Types</i> table.	(Bond type name)
MTU	Displays the MTU set on the virtual switch. The default MTU value is 1500.	(MTU value)

Table 75: Bond Types

Bond Type	Use Case	Maximum VM NIC Throughput	Maximum Host Throughput
Active-Backup	Recommended. Default configuration, which transmits all traffic over a single active adapter.	10 GB	10 GB
Active-Active with MAC pinning Also known as balance-slb	Works with caveats for multicast traffic. Increases host bandwidth utilization beyond a single 10 Gb adapter. Places each VM NIC on a single adapter at a time. Do not use this bond type with link aggregation protocols such as LACP.	10 GB	20 GB
Active-Active Also known as LACP with balance-tcp	LACP and link aggregation required. Increases host and VM bandwidth utilization beyond a single 10 Gb adapter by balancing VM NIC TCP and UDP sessions among adapters. Also used when network switches require LACP negotiation. The default LACP settings are: • Speed—Fast (1s) • Mode—Active fallback-active-backup • Priority—Default. This is not configurable.	20 GB	20 GB

Bond Type	Use Case	Maximum VM NIC Throughput	Maximum Host Throughput
No Uplink Bond	No uplink or a single uplink on each host. Virtual switch configured with the No uplink bond uplink bond type has 0 or 1 uplinks. When you configure a virtual switch with any other bond type, you must select at least two uplink ports on every node.	-	-
Note: The Maximum VM NIC Throughput and Maximum Host Throughput values are not restricted to the values provided in this table. The values in the table are indicated for an assumption of 2 x 10 Gb adapters.			

Actions

You can perform the following actions on the **List** tab of the **Virtual Switch** dashboard:

- Click the name of a virtual switch to open the virtual switch details page, which displays detailed information about the subnet. For more information, see Viewing Virtual Switch Details.
- Create a virtual switch by clicking Create Virtual Switch. For more information, see Creating a Virtual Switch.
- Filter the subnets list based on a variety of parameter values using the **Modify Filters** pane. For more information, see Virtual Switches Summary View Filters Pane.
- Perform the following virtual switch-specific actions on a single virtual switch using the **Actions** dropdown menu. The **Actions** dropdown appears when a virtual switch is selected.

Table 76: Actions

Action	Description
Update	Updates a virtual switch. For more information, see Updating a Virtual Switch.
Delete	Deletes a virtual switch. For more information, see Deleting a Virtual Switch.

Virtual Switches Summary View Filters pane

You can view the **Filters** pane and set the filters for the virtual switch summary view by clicking the **Modify Filters** link on the **List** tab.

Table 77: Summary View Filters Pane

Filter	Description	Value
Name	Filters based on the virtual switch name. It returns a list of virtual switches that satisfy the name condition or string.	(Virtual switch name string)

Filter	Description	Value
Scope	Filters based on the scope of the virtual switches. Select the checkboxes of one or both of the scope options.	(Checkbox selection) Single Cluster Multi Cluster
MTU	Filters based on the MTU range provided. Provide the starting MTU value and the ending MTU value in the From and To fields, respectively.	(MTU values)
Bond Type	Filters based on the bond type of the virtual switches. Select checkboxes of any or all of the listed bond types.	(Checkbox selection) • Active-Backup • Active-Active with MAC pinning • Active-Active • No Bond

Viewing Virtual Switch Details

You can view the details of a virtual switch in its details page.

About this task

To view the details page of a virtual switch, follow these steps:

Procedure

1. Log in to Prism Central.
2. From [Application Switcher Function](#) on page 40, select the **Infrastructure** application and navigate to **Network & Security > Virtual Switches** from the **Navigation Bar**.

The **Virtual Switches** page appears, displaying the **List** tab.

3. Click the name of a virtual switch to open the details page.

The **Summary** tab appears, displaying the details of the virtual switch in widgets.

The virtual switch page has the following tabs:

- Summary
- Alerts

Summary Tab Attributes

The **Summary** tab summarizes the details of the virtual switch in widgets.

Table 78: Summary Tab Widgets and Attributes

Attribute	Description	Values
Properties widget		

Attribute	Description	Values
Name	Displays the name of the virtual switch. If the virtual switch is the default virtual switch created on the Prism Element cluster, the name is suffixed with Single Cluster. If you created the virtual switch on Prism Central, the name is not suffixed.	(Virtual switch name string)
Scope	Displays whether the virtual switch was created on Prism Element web console or Prism Central. The scope of the virtual switch is displayed as Single Cluster or Multi Cluster respectively. You can view this attribute only by creating a custom view and adding this attribute to the view. For more information on configuring a custom view, see <i>View by</i> in Prism Central GUI Organization on page 49.	Multi Cluster Single Cluster
Cluster	Displays the name or names of the cluster or clusters that the virtual switch uplinks.	(Cluster name)
Bond Type	Displays the bond type associated with the virtual switch. For more information on bond types, see the <i>Bond Types</i> table.	(Bond type name)
MTU	Displays the MTU set on the virtual switch. The default MTU value is 1500.	(MTU value)
Associated Subnets widget		
Subnets	Displays the names of the subnets associated with the virtual switch.	(Subnet-name)
Type	Displays the type of subnet. The default subnet type is VLAN Subnet.	VLAN
Uplink Ports widget		
Clusters	Displays the names of the clusters uplinked in the virtual switch.	(Cluster names)

Attribute	Description	Values
25G NICs	Displays the number of 25G NIC ports connected in the virtual switch.	(Number) Port(s)
10G NICs	Displays the number of 10G NIC ports connected in the virtual switch.	(Number) Port(s)

Alerts Tab Attributes

The **Alerts** tab displays all the alerts generated specifically for the virtual switches managed by Prism Central.

The tab shows both active and resolved alerts, with the status of each alert visible in its respective column. A virtual switch alert typically indicates a problem with a virtual switch (VS) or its associated Open vSwitch (OVS) bridges on AHV hosts. For more information on virtual switch alerts, see the [Prism Central Alerts and Events Reference Guide](#).

Creating a Virtual Switch

You can create a multicluster virtual switch in the **Virtual Switches** dashboard.

About this task

To create a multicluster virtual switch, follow these steps:

Procedure

1. Log in to Prism Central.
2. From [Application Switcher Function](#) on page 40, select the **Infrastructure** application and navigate to **Network & Security > Virtual Switches** from the **Navigation Bar**.

The **Virtual Switches** page appears, displaying the **List** tab.

3. Click **Create Virtual Switch**.

The **Create Virtual Switch** page appears, opening in the **General** tab.

4. Provide the following information in the fields in the **General** tab:

- **Name:** Provide a name for the multicluster virtual switch.
- **Description:** (Optional) Provide a suitable description for the multicluster virtual switch.
- **Physical NIC MTU (bytes):** Provide a MTU in bytes. This field displays the default MTU of 1500 bytes. Delete it first and then enter the required MTU value.

You can enter any value between 1280 and 9000 bytes.

Nutanix recommends that you enter a value higher than 1500 bytes.

5. Click **Next**.

6. On the **Uplink Ports** tab, follow these steps:

- **Bond Type:** Select a bond type from the dropdown list in this field.

Each bond type in the list displays an [information icon](#). Hover the cursor over the information icon to see more information on the bond type and its selection.

All bond types except **No Bond** require the selection of a minimum of two ports on each host in the cluster.

- **Uplink Ports:** To select the uplink ports, follow these steps:

- Click **+Add Uplink Ports**.

The **Select Uplink Ports** dialog box appears.

- Click the dropdown list ([chevron down](#)) button displayed for each cluster to view the hosts in the cluster.
- Select the checkbox of each cluster that must be uplinked with this virtual switch.

A multicluster virtual switch uplinks multiple clusters managed by the Prism Central. The **Viewing Clusters** pane lists all the clusters managed by the Prism Central. You can select multiple clusters in the list.

- Click **Select Uplink Ports>** to display the **Viewing Ports** pane.
- Click the dropdown list ([chevron down](#)) button displayed for each of the selected clusters to view the ports on the hosts of the cluster.
- Select the necessary ports on each host of the selected clusters.
- Click **Save** to return to the **Uplink Ports** tab.

7. Click **Save** on the **Uplink Ports** tab.

Updating a Virtual Switch

You can update the attribute values of an existing virtual switch.

About this task

To update a virtual switch, follow these steps:

Procedure

1. Log in to Prism Central.
2. From [Application Switcher Function](#) on page 40, select the **Infrastructure** application and navigate to **Network & Security > Virtual Switches** from the **Navigation Bar**.
The **Virtual Switches** page appears, displaying the **List** tab.
3. You can open the **Update Virtual Switch** page for the virtual switch that needs an update in one of the following methods:
 - » Select the virtual switch in the **List** tab and click **Actions > Update**.
 - » Click the name of the virtual switch in the **List** tab and click **Update** on the virtual switch details page.
4. In the **General** tab, you can update any attribute and click **Next**.

5. In the **Uplink Ports** tab, you can update the **Bond Type** selection.

Note: Updating the **Bond Type** selection resets the selected uplink ports. A confirmation dialog box appears with this message. Click **Confirm** or **Cancel** to update or reject the **Bond Type** selection. If you clicked **Confirm** to change the **Bond Type**, you must select the uplink ports.

You can also update the uplink ports without changing the **Bond Type**.

6. Click the **Manage** link displayed for the **Uplink Ports** pane.

The **Select Uplink Ports** dialog box appears. To update the uplink port selection, follow these steps:

- a. Click the dropdown list (**chevron down**) button displayed for each cluster to view the hosts in the cluster.
- b. Select the checkbox of each cluster that must be uplinked with this virtual switch.

A multicluster virtual switch uplinks multiple clusters managed by the Prism Central. The **Viewing Clusters** pane lists all the clusters managed by the Prism Central. You can select multiple clusters in the list.

- c. Click **Select Uplink Ports** to display the **Viewing Ports** pane.
- d. Click the dropdown list (**chevron down**) button displayed for each of the selected clusters to view the ports on the hosts of the cluster.
- e. Select the necessary ports on each host of the selected clusters.
- f. Click **Save** to return to the **Uplink Ports** tab.

7. Click **Save**.

The **Save Configuration Update** dialog box appears.

8. Select the method to apply the virtual switch configuration updates.

- » **Standard (Recommended):** This method ensures no disruptions occur to the workloads by putting the hosts in maintenance mode and migrating the VMs out of the host before applying the configuration. This process requires a longer duration of time to complete. The time required depends on the number and configuration of VMs.

In this method, the virtual switch configuration is deployed in the rolling update process.

- » **Quick:** This method interrupts the workloads running on the hosts. Use this method only if you are not running production workloads because it may result in network interruptions. In this method, the virtual switch configuration is deployed in a rolling update process but the nodes are not put in maintenance mode before modifying the virtual switch configuration on the node.

9. Click **Save and Apply**.

What to do next

Verify the changes on the virtual switch details page.

Deleting a Virtual Switch

You can delete a virtual switch after you remove all the networking resources such as VLANs, Overlay subnets, and Externally Routable Prefixes.

About this task

To delete a virtual switch, follow these steps:

Procedure

1. Log in to Prism Central.
2. From [Application Switcher Function](#) on page 40, select the **Infrastructure** application and navigate to **Network & Security > Virtual Switches** from the **Navigation Bar**.
The **Virtual Switches** page appears, displaying the **List** tab.
3. You can open the **Delete Virtual Switch** dialog box for the virtual switch to be deleted in one of the following methods:
 - » Select the virtual switch in the **List** tab and click **Actions > Delete**.
 - » Click the name of the virtual switch in the **List** tab and click **Delete** on the virtual switch details page.
4. Click **Delete** to confirm the deletion of the virtual switch.
Click **Cancel** to close the dialog box without deleting the virtual switch.

What to do next

Verify whether the virtual switch is deleted on the **Virtual Switches** page.

Assign a Subnet to a Virtual Switch

You can associate a multicluster virtual switch to a (Network Controller-based) VLAN Subnet or a Virtual Private Cloud (VPC).

Associate a VLAN Subnet

Associate the multicluster virtual switch to the new (Network Controller-based) VLAN Subnet when you create a new VLAN Subnet or an existing VLAN Subnet when you update the existing VLAN Subnet. You cannot assign a VLAN Basic Subnet or an Overlay subnet to a multicluster virtual switch.

For information on creating a VLAN Subnet, see [Creating a Subnet](#). For information on updating a VLAN Subnet, see [Updating a Subnet](#).

Associate a VPC

You can associate a VPC to a multicluster virtual switch only using v4 APIs. For more information on the **Vpc Virtual Switch Mappings** v4 API, see the [API Reference](#).

Subnets

Manage a subnet from Prism Central.

- [Creating a Subnet](#)
- [Updating a Subnet](#)
- [Deleting a subnet](#)
- [Creating a subnet extension](#)
- [Assigning a Category Value to a Subnet](#)
- [Migrating VMs between VLAN and VPC networks](#)

Subnets Summary View

The **Subnets** page displays the list of subnets across all the registered clusters.

To access the **Subnets** page:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Network & Security > Subnets** from the **Navigation Bar**.

The **Subnets** page opens displaying the **List** tab. This tab provides information about all the subnets configured for the registered clusters.

Note: For information on the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

The following table describes the fields that appear in the **Subnets** page.

Note: The fields vary based on the **View by** and **Group by** options. A dash (-) is displayed in a field when a value is not available or applicable. For more information, see the [View by](#) on page 51 and [Group by](#) information.

Table 79: Subnets – Field Description

Field	Description	Values
Name	Displays the subnet name.	(subnet name)
External Connectivity	Displays whether or not the subnet has external connectivity configured.	(Yes/No)
Type	Displays the subnet type.	VLAN or VLAN Basic or Overlay
VLAN ID	Displays the VLAN identification number.	(ID number)
VPC	Displays the name of the VPC in which the subnet is used.	(Name of VPC)
Virtual Switch	Displays the virtual switch that is configured for the VLAN you selected. The default value is the default virtual switch <code>vs0</code> .	(virtual switch name)
Note: The virtual switch name is displayed only if you add a VLAN ID in the VLAN ID field.		
IP Prefix	Displays the IPv4 address of the network with the prefix.	(IPv4 Address/Prefix)
Cluster	Displays the name of the cluster for which this subnet is configured.	(cluster name)
Hypervisor	Displays the hypervisor that the subnet is hosted on.	(Hypervisor)

You can perform the following actions from the **Subnets** page:

- Click the name of a subnet to open the subnet details page, which displays the detailed information about the subnet. For more information, see [Subnet Details View](#) on page 368.
- Create a subnet by clicking **Create Subnet**. For more information, see [Creating a Subnet](#) in the *Flow Virtual Networking Guide*.

- Migrate VMs between VLAN network and VPC network by clicking **Migrate**. For more information, see [Migrating VMs between VLAN Backed and VPC Subnets](#) in the *Flow Virtual Networking Guide*.
- Configure network connections for a cluster by clicking **Network Config**. For more information, see [Network Configuration](#).
- Group the subnets based on pre-defined criteria. For more information, see [Group by](#) on page 51.
- View subnets based on a pre-defined criteria or create a custom view. For more information, see [View by](#) on page 51.
- Filter the subnets list based on a variety of parameter values using the **Filters** pane. For more information, see [Filters Pane - Subnets page](#).
- Perform the following subnet-specific actions on a single or multiple subnets using the **Actions** dropdown menu. The **Actions** dropdown appears when one or more subnets are selected.

Table 80: Subnet Actions

Action	Description
Update	Click this action to update the subnet. For more information, see Updating a Subnet in the <i>Flow Virtual Networking Guide</i> .
Extend	Click this action to create a subnet extension. For more information, see Layer 2 Virtual Subnet Extension Over VTEP in the <i>Flow Virtual Networking Guide</i> .
Manage Categories	Click this action to associate the subnet with a category or change the categories that the subnet is associated with. For more information, see Assigning a Category to an Entity .
Delete	Click this action to delete the subnet. For more information, see Deleting Subnets, Policies, or Routes in the <i>Flow Virtual Networking Guide</i> .

Filters Pane - Subnets page

You can filter the information in the **Subnets** page based on the following fields that are available in the **Filters** pane. For information on how to use the Filters pane, see the [Filters Pane](#) information.

Table 81: Filter Pane Field Description - Subnets page

Field	Description	Values
Name	Filters based on the subnet name. It returns a list of subnets that satisfy the name condition/string.	(Subnet name string)
External Connectivity	Filters based on whether the subnet has external connectivity configured or not.	(Yes/No)
Type	Filters based on the subnet type.	(VLAN/VLAN (External)/Overlay)
VLAN ID	Filters based on VLAN identification number.	(ID number)

Field	Description	Values
VPC	Filters based on the name of the VPC in which the subnet is used.	(Name of VPC)
Cluster	Filters based on the name of the cluster for which this subnet is configured.	(cluster name)
Hypervisor	Filters based on the hypervisor that the subnet is hosted on.	ESXi/AHV/Hyper-V/XenServer/Mixed Hypervisor/Null Hypervisor

Subnet Details View

The Subnet details page consists of a dashboard that provides the detailed information about the subnet.

The details page has the **Summary**, and **Throughput** tabs.

To access the details page of an individual subnet:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Network & Security > Subnets** from the **Navigation Bar**.

Prism Central displays the **Subnets** page that contains information about all the subnets configured for the registered clusters.

3. Click a subnet to open the details page of the subnet.

The **Summary** tab opens displaying the detailed information about the subnet in widgets.

Note: For information on the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

Summary Tab

The **Summary** tab provides detailed information about the subnet in widgets. A dash (-) is displayed in a field when a value is not available or applicable.

The **Summary** tab has the following widgets:

Widget Name	Information provided
Properties	Provides the following: • Type — Displays the type of network like VLAN or Overlay. • VLAN ID — Displays the VLAN ID. This parameter is displayed only for VLAN networks. • VPC — Displays the VPC name. This parameter is displayed only for Overlay networks. • Cluster — Displays the cluster that the VLAN network is configured on. This parameter is displayed only for VLAN networks. • IP Address Prefix — Displays the IP address prefix configured for the network. This parameter is displayed for both VLAN and Overlay networks.

Widget Name	Information provided
IP Address Pools	Provides the following: • The IP address Pool Range assigned to the network. • The total number of used and available IPs in the cluster. • Used IPs in Subnet — Displays the number of used IPs in the subnet. • Used IPs in Pools — Displays the number of used IPs in the pool. • Free IPs in Pools — Displays the number of free IPs in the pool. • Free IPs in Subnet — Displays the number of free IPs in the subnet.
Domain Settings	Provides the following DHCP settings configured for a VM in a subnet: • Domain Name Servers — Displays the total number of DNS IP addresses. • Domain Search — Displays the VLAN domain name. • Domain Name — Displays the domain name. • TFTP Server Name — Displays the name of the TFTP server where you host the host boot file. • Boot File Name — Displays the name of the boot file that the VMs need to download from the TFTP host server.

The **Summary** tab provides the following options, at the top of the page. For more information, see the *Subnet Actions* table in [Subnets Summary View](#) on page 365.

- **Update**
- **Extend**
- **Manage Categories**
- **Delete**

Throughput Tab

The **Throughput** tab provides a graphical representation of the throughput of the subnet.

Network Configuration

Each VM network interface is bound to a virtual network, and each virtual network is bound to a single single-cluster VLAN. The **Network Configuration** window displays information about the configured virtual networks.

For information to access the **Network Configuration** window, see [Network Configuration for VM Interfaces](#) in the *Prism Element Web Console Guide*.

LAG and LACP on the ToR Switch

For more information, see [Enabling LACP and LAG \(AHV Only\)](#).

Virtual Private Clouds

You can create, view, configure and manage the virtual private clouds (VPCs) in the **Virtual Private Clouds** page.

You can deploy and manage virtual private clouds (VPCs) on a Nutanix cluster infrastructure using the Network Controller.

You can view the list of the VPCs in the **Summary** view. For more information, see [Viewing the Virtual Private Cloud Summary](#) on page 370.

You can view the details of the VPCs in the VPC details view. For more information, see [Viewing the Virtual Private Cloud Details](#) on page 371.

Viewing the Virtual Private Cloud Summary

Access the **Virtual Private Clouds** page in **Network & Security**.

About this task

To access the **Virtual Private Clouds** page, follow these steps.

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Network & Security > Virtual Private Clouds** from the **Navigation Bar**.

Note: For information on the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

The **Virtual Private Clouds** page opens displaying the **List** tab. The tab provides a list of virtual private clouds you have created and configured, and the operations you can perform on them.

Virtual Private Clouds Summary View

The **Virtual Private Clouds** page displays the list of virtual private clouds (VPCs) across all the registered clusters.

The following table describes the fields that appear in the **Virtual Private Clouds** page.

Note: The fields vary based on the **View by** and **Group by** options. A dash (-) is displayed in a field when a value is not available or applicable. For more information, see the [View by](#) on page 51 and [Group by](#) information.

Table 82: Virtual Private Clouds – Field Description

Field	Description
Name	Displays the name of the VPC. The Name of a VPC is suffixed with Transit VPC when you configure the VPC as a transit VPC.
Associated External Subnets	Displays the external subnet that the VPC is assigned to.
Categories	Displays the number of categories associated with the VPC.
Externally Routable IP Addresses	Displays the externally routable IP address.
Hypervisor	Displays the hypervisor that the VPC is hosted on.
Inter VN Traffic	Displays the traffic flowing between the virtual networks or VPCs.
Internet Traffic	Displays the traffic flowing to and from the Internet.

Field	Description
IPv4 Gateway	Displays the IPv4 gateway IP address.
IPv4/Subnet	Displays the IPv4 network IP with subnet prefix. For example, 10.20.30.0/24.
On-Prem Traffic	Displays the traffic flowing in the on-premises network.
VLAN ID	Displays the VLAN identification number. VLAN ID is a parameter used for Transit VPC networking in Nutanix Cloud Cluster with Microsoft Azure.

You can perform the following actions for the VPCs from the **Virtual Private Clouds** page:

- Click the name of a VPC to open the VPC details page, which displays the detailed information about the VPC. For more information, see [Virtual Private Cloud Details View](#) on page 372.
- Create a VPC by clicking **Create VPC**. For more information, see [Creating Virtual Private Cloud](#) in the *Flow Virtual Networking Guide*.
- Update or delete an existing VPC using the **Actions** dropdown menu. The **Actions** dropdown appears when one or more VPCs are selected. For more information, see [Updating Virtual Private Cloud](#) or [Deleting a Virtual Private Cloud](#) in the *Flow Virtual Networking Guide*.
- View VPCs based on pre-defined criteria or create a custom view. For more information on the **View by** option, see [View By](#) information.
- Filter the VPC list based on a variety of parameter values using **Filters** pane. For more information, see [Filters Pane - Virtual Private Clouds Page](#).

Filters Pane - Virtual Private Clouds Page

You can filter the information in the **Virtual Private Clouds** page based on the following fields that are available in the **Filters** pane. For information on how to use the Filters pane, see the [Filters Pane](#) information.

Table 83: Filter Pane Field Description - Virtual Private Clouds page

Field	Description	Values
Name	Filters based on the VPC name. It returns a list of IP addresses that satisfy the name condition/string.	(Virtual private cloud name string)
Associated External Subnets	Filters based on the external subnet that the VPC is assigned to.	(External Subnet)

Viewing the Virtual Private Cloud Details

You can view the details of each VPC listed in **Virtual Private Clouds**.

About this task

To access the details page of an individual VPC, follow these steps.

Procedure

1. Log in to Prism Central.

2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Network & Security > Virtual Private Clouds** from the **Navigation Bar**.

Note: For information on the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

Prism Central displays the **Virtual Private Clouds** page that contains information about all the VPCs configured for the registered clusters.

3. Click a VPC to open the details page of the VPC.

The **Summary** tab opens displaying the detailed information about the VPC in widgets.

Virtual Private Cloud Details View

The Virtual Private Cloud (VPC) details page consists of a dashboard that provides the detailed information about the VPC.

The details page has the **Summary**, **Subnets**, **Policies**, **Routes**, and **Metrics** tabs.

Summary Tab

The **Summary** tab provides detailed information about the VPC in widgets.

The **Summary** tab has the following widgets:

Widget Name	Information provided
External Connectivity	Provides the following: • Associated External Subnets — Displays the number of external subnets associated with the VPC. • Externally Routable IP Addresses — Displays the external routable IP addresses associated with the VPC.
Transit VPC	Displays Yes if the VPC is a Transit VPC. Displays No if the VPC is not a Transit VPC.
Domain Name Servers (DNS)	Displays the IP address or the FQDN of the DNS servers used by the VPC.
Associations	Provides the following: • Subnets (Overlay) — Displays the number of subnets associated with the VPC. • Policies — Displays the number of policies associated with the VPC. • Routes — Displays the number of routes associated with the VPC.
Floating IP Addresses	Provides the following: • Assigned Floating IPs — Displays the floating IP addresses assigned to the VPC. • Available Floating IPs — Displays the available floating IP addresses that can be assigned to the VPC.

Subnets Tab

The **Subnets** tab displays the list of subnets added to the VPC.

The following table describes the fields that appear in the **Subnets** tab.

Table 84: Subnets Tab – Field Description

Field	Description
Name	Displays the subnet name.
IP Range	Displays the IP address range configured for the subnet.
DHCP IP Pool	Displays the IP address pool range assigned to the subnet.
Default Gateway IP	Displays the IP address used as the default gateway by the entities in the subnet.
Actions	Action link for editing or deleting the subnet.

You can perform the following actions for a subnet from the **Subnets** tab:

- Click the name of the subnet to open the subnet details page, which displays the detailed information about the subnet. For more information, see [Subnet Details View](#) on page 368.
- Create a subnet by clicking **Create Subnet**. For more information, see [Creating a Subnet](#) in the *Flow Virtual Networking Guide*.
- Update an existing subnet using the **Delete** option associated with the subnet. For more information, see [Updating a Subnet](#) in the *Flow Virtual Networking Guide*.
- Delete an existing subnet using the **Delete** option associated with the subnet. For more information, see [Deleting Subnets, Policies, or Routes](#) in the *Flow Virtual Networking Guide*.

Policies Tab

The **Policies** tab displays information about the security-based traffic shaping policies you configured.

The following table describes the fields that appear in the **Policies** tab.

Note: The fields vary based on the **View by** option. A dash (-) is displayed in a field when a value is not available or applicable. For more information, see the [View by](#) on page 51 information.

Table 85: Policies Tab – Field Description

Field	Description
Description	Displays the user-provided description of the policy.
Action	Displays the appropriate action for the implementation of the policy. • Permit: Permits traffic and services based on the parameters set. • Deny: Denies traffic and service based on the parameters set. • Re-route: Sends matching traffic to the next-hop IP address specified by the Reroute IP.
Priority	Displays the traffic priority.
Rule	Displays the Permit or Deny rule set for the priority.
Rule Type	Displays whether the rule is system generated or user defined.
Traffic	Displays the traffic type that the priority and rule must be applied to.

Field	Description
Virtual Network	Displays the ID of the subnet.
Source	Displays the source IP or subnet for which you need to manage traffic.
Destination	Displays the destination IP or subnet for which you need to set the priority.
Source Subnet	Displays the subnet IP and prefix designated as the source for the policy.
Destination Subnet	Displays the subnet IP and prefix designated as the destination for the policy.
Reroute Address	Displays the IP address to which the traffic is rerouted.
Bidirectional Policy	Displays whether the policy is bidirectional or not.
Protocol	Displays the type of protocol for which the policy is configured.
Protocol Number	Displays the protocol number for which the policy is configured.
ICMP Type	Displays the type of ICMP message associated with the policy.
ICMP Code	Displays the ICMP code of the policy.
Byte Count	Displays the total number of traffic bytes that matches the policy. The count is updated periodically.
Packet Count	Displays the total number of traffic packets that matches the policy. The count is updated periodically.

You can perform the following actions for a policy from the **Policies** tab:

- Create a policy by clicking **Create Policy**. For more information, see [Creating a Policy](#) in the *Flow Virtual Networking Guide*.
- Perform the following actions using the **Actions** dropdown menu. The **Actions** dropdown appears when one or more policies are selected.
 - **Update**: Update the policy. For more information, see [Updating a Subnet](#) in the *Flow Virtual Networking Guide*.
 - **Delete**: Delete the policy. For more information, see [Deleting Subnets, Policies, or Routes](#) in the *Flow Virtual Networking Guide*.
 - **Clear Counters**: Reset the counters for the selected policy.
 - **Clear All Counters**: Reset the counters for all the policies.
- View policies based on predefined criteria or create a custom view. For more information on the **View by** option, see [View By](#) information.

Routes Tab

The **Routes** tab displays the list of static routes added to the VPC.

The following table describes the fields that appear in the **Routes** tab.

Note: The fields vary based on the **View by** option. A dash (-) is displayed in a field when a value is not available or applicable. For more information on the View by option, see [View by](#) on page 51 information.

Table 86: Routes Tab – Field Description

Field	Description
Destination Prefix	Displays the IP address and prefix of the destination.
Next Hop	Displays the next hop network or subnet for the traffic exiting the VPC.
Priority	Displays the traffic priority.
Type	Displays the type of route, local or static.
Status	Displays the status of the route, whether it is active or not

You can perform the following actions for a route from the **Routes** tab:

- View routes based on predefined criteria or create a custom view. For more information on the View by option, see [View by](#) on page 51 information.
- Perform the following actions using the **Manage Static Routes** option:
 - **Add Static Route:** Create a static route. For more information, see [Creating Static Routes](#) in the *Flow Virtual Networking Guide*.
 - Update an existing static route. For more information, see [Updating Static Routes](#) in the *Flow Virtual Networking Guide*.
 - Delete a static route. For more information, see [Deleting Subnets, Policies or Routes](#) in the *Flow Virtual Networking Guide*.

Metrics Tab

The **Metrics** tab displays detailed information about the VPC metrics.




Figure 132: Metrics Tab

The following table describes the fields that appear in the **Metrics** tab.

Table 87: Metrics Tab – Field Description

Field	Description
External Traffic Through	Select All External Networks (default) or (name_of_external_network_associated_with_the_VPC) from the dropdown menu. The page displays the metrics based on your selection.
Last (time_period)	Select the period for which you need to display the metrics. The dropdown menu provides the following options: • Last 24 Hours (default) • Last One Hour • Last Week
Direction of traffic	Select the direction of traffic for which you need to display the metrics. The dropdown menu provides the following options: • Both directions (default) — Includes both directions, Ingress and Egress. • Ingress — Traffic entering the externally connected subnet. • Egress — Traffic leaving the externally connected subnet.
Bandwidth	Displays graphically the bandwidth utilization of the VPC on a timeline as set in the Last (time_period) parameter.
Rx and Tx Packets	Displays graphically the received and transmitted packet volume on a timeline as set in the Last (time_period) parameter.

Network Services

You can create and manage **Network Load Balancer** and **Traffic Mirroring** sessions from the **Network Services** dashboard.

The **Network Services** dashboard provides the **Network Load Balancer** and the **Traffic Monitoring** tabs. The **Network Services** dashboard opens in the **Network Load Balancer** tab by default.

The **Network Load Balancer** tab displays information on all the network load-balancer sessions created across the clusters managed by the Prism Central. For information on **Network Load Balancer** views, see [Network Load Balancer](#) on page 376. For information on managing **Network Load Balancer**, see [Network Load Balancer Management](#).

The **Traffic Mirroring** tab displays information on all the Traffic Mirroring sessions created across the clusters managed by Prism Central. For more information, see [Traffic Mirroring](#) in *Prism Central Infrastructure Guide*.

Network Load Balancer

With the **Network Load Balancer** feature, you can create and manage Layer 4 load balancer sessions.

Types of Load Balancers

Load-balancers are of the following types:

- **Layer 4 load balancers:** These load balancers distribute the data traffic flowing through the network (IP) and transport (FTP, TCP and UDP) layer of the network. Layer 4 load balancers distribute the network load. This type is called the network load balancer.

- Layer 7 load balancers: These load balancers distribute the protocol request traffic based on the data flows in the application layer protocols (such as HTTP). Layer 7 load balancers distribute the application load. This type is called the application load balancer.

Nutanix supports Layer 4 Load Balancing natively in Flow Virtual Networking. For Layer 7 Load Balancing integration using policy based routing or the Service Insertion feature, contact Nutanix Support.

Network Load Balancer

With the **Network Load Balancer** feature, you can achieve the following efficiencies:

- Distribute the network traffic load across multiple guest VMs efficiently, by allowing you to add and remove load balancers and their members based on demand for more efficient distribution of network load.
- Distribute the requests only to active VMs for better high availability and reliability.
- Monitor the health of target VMs to mark active and inactive targets.

The **Network Load Balancer** is a distributed load balancer that is implemented in the AHV host to distribute traffic across the network. Network Load Balancers improve the capacity, reliability, and overall fault tolerance of the network and applications using the network. The Network Load Balancer feature implements high availability for network and applications and improves the performance of the network and applications.

For load-balancing, the overlay subnet in the VPC assigns a virtual IP address to the load balancer. The **Network Load Balancer** feature provides the option of requesting an IP address on the **Floating IP** page from the NAT external subnet range and assigning it to the load balancer as virtual IP address to be used as external IP address, for a NAT external subnet.

External Load Balancing

This load balancing involves the distribution of traffic flowing into the VPC from sources external to the VPC.

Internal Load Balancing

This load balancing involves distributing the traffic flowing within the VPC, in other words, intra-VPC traffic, among a set of VMs in the VPC. For such load balancing, the virtual IP address assigned to the load balancer need not be reachable from outside the VPC.

For information on creating, updating or deleting network load balancers, see [Network Load Balancer Management](#).

Health Monitoring

The **Network Load Balancer** monitors the health of target VM NICs for both TCP and UDP protocols.

When the target is configured with TCP, the local AHV host where the target VM runs sends a TCP SYN on the configured port to the target VM NIC and expects a TCP ACK in response to note the healthy state. The AHV host then sends a TCP RST to close the connection. Lack of a response, or an ICMP unreachable message indicates an unhealthy state.

When the target is configured with UDP, the local AHV host sends a UDP message to the target VM on the configured port and expects no response to note the healthy state. An ICMP unreachable message indicates an unhealthy state.

Limitation

Traffic to the floating IP (for example, FIP1) on, for example, port 22 of the guest VM fails when a guest VM has the following configurations:

- A floating IP (for example, FIP1) is configured on, for example, port 22.
- Load balancing is configured to use the same port, for example, port 22.

- A floating IP (for example, FIP2) is configured for load balancing to reach the guest VM from outside the VPC.
- The private IP address that is normally assigned to the guest VM within the VPC.

Viewing Network Load Balancer Summary

Access the **Network Load Balancer** tab in **Network Services**.

About this task

The **Network Load Balancer** tab displays the list of load balancer sessions configured for all the clusters managed by the Prism Central AZ. You can also update or delete the **Network Load Balancer** sessions that you create on the **Network Load Balancer** tab.

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Network & Security > Network Services** from the **Navigation Bar**.

The **Network Services** dashboard opens in the **Network Load Balancer** tab by default. The **Network Load Balancer** tab displays the list of load balancer sessions.

Network Load Balancer Summary View Attributes

This section provides the details of the **Network Load Balancer** summary view attributes.

The fields vary based on the View by and Group by options. A dash (-) is displayed in a field when a value is not available or applicable.

Fields	Description	Values
Name	Displays the name provided to the load balancer session when it was created.	(String)
VPC	Displays the name of the VPC in which the load balancer session is operating.	(String)
Listener Protocol: Port	Displays the transport protocol and the port number configured for the Listener.	TCP or UDP:(Integer) Port number Example: TCP:80
Listener Virtual IP	Displays the IP address assigned to the Listener as virtual IP address of the Listener.	IP address
Listener Floating IP	Displays the Floating IP address assigned to the Listener from the NAT external subnet for the NAT external connectivity enabled for the load balancer session.	IP address
Target Port	Displays the ports configured for the target VM NICs. Traffic destined for the Listener port is directed to the target VMs on these ports.	Port numbers list

Fields	Description	Values
Target VM NICs Health	Displays the health status of the target VM NICs. The Network Controller checks the reachability of the ports using TCP or UDP based on the Listener Protocol selected while creating the load balancer session. If all the NICs have the same status, then the status is displayed as All <number of NICs>;status as Healthy or Unhealthy. For example you have configured three target VMs with one NIC each and all the NICs are Healthy, then the status is displayed as All 3 Healthy.	Status list as Healthy or Unhealthy.
Listener Protocol	Displays the transport protocol.	TCP or UDP

Actions

You can perform the following actions from the Network load balancer tab:

- Click the name of a subnet to open the subnet details page, which displays detailed information about the subnet. For more information, see [Viewing Network Load Balancer Session Details](#) on page 380.
- Create a load balancer session by clicking **Create Load Balancer Session**. For more information, see [Creating a Load Balancer Session](#).
- Filter the subnets list based on a variety of parameter values using the **Modify Filters** pane. For more information, see [Network Load Balancer Summary View Filters Pane](#) on page 379.
- Perform the following load balancer session-specific actions on a single load balancer session using the **Actions** dropdown menu. The **Actions** dropdown appears when a load balancer session is selected.

Action	Description
Update	Click this action to update the subnet.
Delete	Click this action to delete the subnet.

Network Load Balancer Summary View Filters Pane

This section provides the details of the **Network Load Balancer** summary view filter pane.

You can filter the information in the **Network Load Balancer** page based on the following fields that are available in the Filters pane. For information on the regular expressions that you can use for filtering, see the *Filter Expressions* section in the [Prism Central Infrastructure Guide](#).

Fields	Description	Values
Name	Filters based on the load balancer session name. It returns a list of load balancer sessions that satisfy the name condition or string.	(load balancer name string)

Fields	Description	Values
VPC	Filters based on the name of the VPC in which the load balancer sessions are operating.	(VPC name string)
Listener Protocol	Filters based on the Listener Protocol used by the load balancer sessions. Select the check box of the TCP or UDP or both options.	TCP and UDP
Listener Virtual IP	Filters based on the IP address configured as virtual IP for the Listener	

Viewing Network Load Balancer Session Details

This section provides the procedure to access the **Network Load Balancer** session details page.

About this task

The **Network Load Balancer** details page provides detailed information on a selected load balancer session.

To access the details page of a selected load balancer session:

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Network & Security > Network Services** from the **Navigation Bar**.

The **Network Services** dashboard opens in the **Network Load Balancer** tab by default. The **Network Load Balancer** tab displays the list of load balancer sessions.

3. Click a load balancer session to open the details page of the load balancer session.
The **Summary** tab opens displaying the details of the load balancer session in widgets.

The load balancer session page has the following tabs:

- **Summary**
- **Target VM NICs**
- **Alerts**
- **Audits**

Summary Tab Attributes

This section provides the details of the **Summary** tab attributes in the Network Load Balancer details view.

The **Summary** tab provides detailed information in widgets.

A dash (-) is displayed in a field when a value is not available or applicable.

The **Summary** tab provides information in two widgets—**Properties** and **Target VM Status**.

Widgets and Fields	Information provided
Properties This widget provides information in a two-column table divided into different sections	
Basic Configuration	Provides the following: • Name—Name of the load balancer session • Description—Description of the load balancer session • VPC—Name of the VPC that the load balancer session operates on.
Listener Configuration	Provides the following: • Protocol—The transport protocol set for the Listener, is TCP or UDP. • Port(s)—The port that the Listener listens on. • Subnet—The VPC that the Listener is assigned its virtual IP from. • Virtual IP—The IP address assigned as virtual IP to the Listener. • Floating IP—The Floating IP address assigned for external connectivity to the Listener. • Load Balancing Algorithm—Five Tuple Hash which is the default algorithm.
Targets	Provides the following: • Total VM NICs—The number of VM NICs configured for the load balancer. • Port(s)—The ports configured on the VM NICs for load-balancing.

Widgets and Fields	Information provided
VM Health Check Configuration	A VM health check run is a • TCP Syn message sent by the Network Controller. An Ack message is expected in response to the Syn. The connection is then reset. • UDP Syn message sent by the Network Controller. No response is expected. The VM health check configuration provides the following settings: • Check Run Every—The time interval in seconds between two VM health check runs. The VM health check run occurs at the lapse of the time interval. • Timeout After—The timeout interval in seconds that the VM health check runs for before failing. • Marked Healthy After—Number of consecutive successes (time in seconds). Each consecutive success adds 5 seconds to the time. For example, the default number of successes is three, adding a time interval of 15 seconds, represented in the widget as 3 consecutive successes (15 seconds). • Marked Unhealthy After—Number of consecutive failures (time in seconds). Each consecutive failure adds 5 seconds to the time. For example, the default number of failures is three, adding a time interval of 15 seconds, represented in the widget as 3 consecutive failures (15 seconds).
Target VM Status	
This widget provides information as a donut usage chart with the data stacked adjacent to the chart.	
NIC Health	Provides the number of Healthy and Unhealthy NICs
CPU Usage	Provides the number of Target VMs with CPU usage in the following categories: • <50 % • 50 - 75 % • >75 %
Memory Usage	Provides the number of Target VMs with memory usage in the following categories: • <50 % • 50 - 75 % • >75 %

Widgets and Fields	Information provided
IO Latency	Provides the number of Target VMs with IO latency in the following categories: • <2 ms • 2 - 5 ms • >5 ms

The **Summary** tab provides the following actions, at the top of the page. For information on these actions, see the [Network Load Balancer Summary View Attributes](#) on page 378

- **Update:** Select the load balancer session you need to update and select **Update**. For information, see [Updating a Load Balancer Session](#).
- **Delete:** Select the load balancer session you need to delete and select **Delete**. For information, see [Deleting a Load Balancer Session](#).

Target VM NICs Tab Attributes

This section provides the details of the **Target VM NICs** tab attributes in the Network Load Balancer details view.

The **Target VM NICs** tab provides detailed information about the target VM NICs in a table.

A dash (-) is displayed in a field when a value is not available or applicable.

Field	Description	Values
Name	Displays the name of the target VM	(VM name string)
IP Address	Displays the IP address of the target VM.	(IP address)
Health	Displays the health status of the target VM NIC	(Healthy or Unhealthy)
VM Category	Displays the name of the category that the VM is assigned to.	(Category name string)
Subnet	Displays the name of the subnet that the VM is networked in.	(Subnet name string)
Cluster	Displays the name of the cluster hosting the VM.	(Cluster name string)
Port	Displays the port configured on the VMNIC for load balancing.	(Port number integer)
CPU Usage	Displays the CPU usage of the VM.	(Fractional number with two decimal places with percentage sign (%)) Example: 0.57%
Memory Usage	Displays the CPU usage of the VM.	(Fractional number with two decimal places with percentage sign (%)) Example: 9.58%

Field	Description	Values
IO Latency	Displays the storage IO latency of the VM in milliseconds.	(Fractional number with two decimal places suffixed with “ms”) Example: 49.37 ms

Alerts Tab Attributes

This section provides the details of the **Alerts** tab attributes in the **Network Load Balancer** details view.

The **Alerts** displays the list of alerts generated for the load balancer sessions.

Note: The fields vary based on the View by and Group by options. A dash (-) is displayed in a field when a value is not available or applicable.

Field	Description	Values
Title	Displays the title of the alert.	(Alert title string)
Source Entity	Displays the load balancer session for which the alert is raised.	(load balancer session name string)
Impact Type	Displays the impact type of the issue that the alert is raised for.	(Impact type category string)
Severity	Displays the severity category that the alert belongs to.	(Severity category string) Example: Critical, Warning, or Info
Status	Displays the status of alert whether acknowledged or resolved.	(Status category string) Example: Acknowledged or Resolved
Created Time	Displays the time the alert was raised.	(Date and time in MMM DD, YYYY, HH:mm AM/PM)
Last Occurred	Displays the last time the alert was raised.	(Date and time in MMM DD, YYYY, HH:mm AM/PM)
Cluster	Displays the cluster on which the alert was raised such as Prism Central cluster or the name of the Prism Element host cluster.	(Cluster name string)
Provider	Displays the name of the source of the alert, in other words, the name of the entity that raised the alert.	(Entity name string)
Resolved	Indicates whether the alert was resolved or not. Displays the mode of resolution with date and time stamp. Displays bank (dash) if the alert is not resolved.	Auto (MM/DD, HH:mm AM/PM) or Manual (MM/DD, HH:mm AM/PM) or blank (dash)

Field	Description	Values
Resolved By	Indicates whether the name of the user or role which resolved the alert, for manual resolution only. Displays “-” for Auto or No.	User name or role name (string) or blank (dash)
Resolved At	Displays the time stamp for the time of resolution.	(Date and time stamp)
Acknowledged	Indicates whether the alert was acknowledged or not. Displays blank (dash or “-”) if the alert was automatically acknowledged.	Auto (date and time stamp) or Manual (date and time stamp)
Acknowledged By	Indicates whether the alert has been acknowledged. A blank (dash) value means the alert is not acknowledged or automatically acknowledged. Displays the name of the user or role that acknowledged the alert for manual acknowledgement.	Blank (dash) or User name or role name (string)
Acknowledged At	Displays the time stamp for the time of acknowledgement.	(Date and time stamp)
Alert Type	Displays the Alerts type code message. The last word in the message indicated the severity-based type such as WARNING or CRITICAL	(Alert type string) Example: networking_atlas-rules_LoadBalancerSessionTargetsUnhealthy_WAR
Diagnostics		

You can perform the following actions for alerts in the tab.

- Configure the email preferences to receive the alerts on email, using the **Email Configuration** button. For information, see [Configuring Email for Alerts](#) on page 385.
- Resolve an alert by selecting the checkbox for an alert and selecting **Actions > Resolve**.
- Acknowledge an alert by selecting the checkbox for an alert and selecting **Actions > Acknowledge**.

You can select an alert to open the details page of the alert.

For more information, see the [Prism Central Alerts and Events Reference Guide](#).

Configuring Email for Alerts

This section provides the steps to configure the email IDs that the notifications of the alerts must be sent to.

Procedure

1. Log in to Prism Central.

2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Network & Security > Network Services** from the **Navigation Bar**.

The **Network Services** dashboard opens in the **Network Load Balancer** tab by default. The **Network Load Balancer** tab displays the list of load balancer sessions.

3. In the **Network Load Balancer** tab, click the name of the load balancer session that you want to configure the **Alert Email Configuration** for.
4. Click **Alerts > Email Configuration**.
5. On the **Alert Email Configuration** dialog box, provide the necessary information.
For information on the attributes that you need to configure, see [Email Configuration Attributes](#) on page 386.
6. Click **Save**.

Email Configuration Attributes

Click **Save** on each tab to save the configuration made on that tab.

Attribute	Description and Value
Settings tab	
You can configure the following standard preferences for emails in this tab.	
Email Preference	
Every Single Alert	Select the Every Single Alert checkbox to ensure that the email recipients receive an email each for every single alert.
Daily Digest	Select the Daily Digest checkbox to ensure that the email recipients receive a summary email sent every day with all alerts. If you have selected the Daily Digest checkbox, do the following: • Set a time to ensure that the recipients receive the emails at the specified time of the day in A summary email sent every day with all alerts at. • Select the Skip the daily digest email if there are no alerts generated on a given day to skip summary emails if no alerts are generated during the 24 hours preceding the time of sending the email.
Email Recipients	Enter the email IDs of the recipients. For multiple email IDs, click Enter , Tab , Space or Comma after entering each email ID and enter the next email ID.

Attribute	Description and Value
Tunnel Connection	The attributes presented in this section are information. You cannot enter or modify any information in this section. When you open the Alert Email Configuration dialog box after configuring Email Recipients and clicking Save, the following changes occur in this section: The Status changes to Success. The Service Center attribute displays the details of the service domain used to send the email. The Established Since attribute displays the date and time stamp of the time when the tunnel connection was established.
Mode	Displays the name of the tunnel used to send the emails. The default value is Default Nutanix Tunnel .
Status	Displays the status as Enabled or Disabled. The Status is displayed as Disabled, before you configure the Email Recipients attribute. Note: This attribute value changes to Success when you open the Alert Email Configuration dialog box after configuring Email Recipients and clicking Save and the tunnel connection is implemented.
Service Center	Displays the details of the service domain used to send the email. The default service center is nsc01.nutanix.net. Note: This attribute is displayed only when you open the Alert Email Configuration dialog box after configuring Email Recipients and clicking Save.
Established Since	Displays the date and time stamp of the time when the tunnel connection is established. Note: This attribute is displayed only when you open the Alert Email Configuration dialog box after configuring Email Recipients and clicking Save.
Custom Settings tab	Create custom preferences in the Create New Setting section that is displayed when you click New Setting on this tab.
Create New Setting	Enter or select the appropriate values in the following attributes or fields in this section.

Attribute	Description and Value
Alert Severity	Select from this dropdown list, the checkbox for appropriate severity or All for which you want the recipients to receive the email.
Impact Type	Select from this dropdown list, the checkboxes for the appropriate Impact types or All. The Impact Type options are: • All • Performance • Capacity • CPU Capacity • Configuration • Availability • System Indicator • Memory Capacity • Storage Capacity
Cluster	Select from this dropdown list, the checkboxes for the appropriate cluster options or All. The cluster options include All and Prism Central, followed by a list of Prism Element clusters managed by the Prism Central instance.
Alert Contains	Enter the appropriate text content that filters the alerts for inclusion in the email notification.
Email Recipients	Enter the email IDs of the recipients. For multiple email IDs, click Enter , Tab , Space or Comma after entering each email ID and enter the next email ID.
Email Content tab	
This tab lets you provide additional content in the email notifications.	
Email Content	Displays the subject and body of an alert email notification. The subject and body text are derived from the Alert content.
Prepend Subject	Enter the text that you want to prepend to the text in the subject line. As you enter the text, it is displayed in the subject line of the Email Content field above.
Append Body	Enter the text that you want to append to the text in the email body text. As you enter the text, it is displayed in the body section of the Email Content field above.

Audits Tab Attributes

This section provides the details of the **Audit** tab attributes in the Network Load Balancer details view.

The **Audits** tab displays the list of successful and audited actions generated for creation, update and deletion events of the load balancer sessions.

Note: The fields vary based on the View by options. A dash (-) is displayed in a field when a value is not available or applicable.

The audited actions are available in the list only if these actions were successful. Failed actions are not provided in this list.

Field	Description	Values
Action Description	Displays the description of the action task.	(String)
User Name	Displays the user name of the user, such as Admin, who initiated the action.	(String)
User IP	Displays the IP address of the user's Prism Central Virtual IP or Leader VM IP address.	IP address
Affected Entities	Displays the name of load balancer session or other entity that the alert is raised for.	(String)
Entity Affected	Displays the name of the load balancer session that the action impacts.	(String)
Entity Type	Displays the type of entity affected by the action. It displays Network Load Balancer for the load balancer session that is affected.	Network load balancer
Operation Type	Displays the type of action performed on the load balancer session, such as Create or Update .	(Action type string)
Request Time	Displays the time that the action was requested.	Date and time stamp
Cluster	Displays the name of the cluster such as Prism Central or the name of the Prism Element cluster, on which the action was performed.	(Name string)
Status	Displays the status of the action performance. FAILED: Indicates audit failure. SUCCEEDED: Indicates audit success. ABORTED: Indicates aborted audit.	(String)

Audited Action Entity Attributes

The sidebar on the action details page provides the attributes of the action.

For more information, see [Audits Tab Attributes](#) on page 389.

The dashboard provides the attributes of the load balancer session on which the audited action was performed.

Field	Description
Load-balancer Session Name	Displays the name of the load balancer session that the action was performed on.
Load-balancer Session UUID	Displays the UUID of the load balancer session.
Description	Displays the description of the load balancer session.
VPC UUID	Displays the UUID of the VPC that the load balancer session is configured in.
Listener Subnet UUID	Displays the UUID of the Subnet in the VPC that the load balancer session is configured in.
Virtual IP Address	Displays the virtual IP address assigned to the Listener in the subnet.
Listener Protocol	Displays whether the transport protocol configured for the Listener is TCP or UDP.
Listener Port Ranges	Displays the port ranges that are configured on the target VM NICs.
Target VM NICs	Comma-separated list of NIC UUID and associated port of each NIC. For example, d72f1940-xxxx-4exx-axxx-73xx32exxxxx:8080

Audited Action Details Attributes

The **Details** page consists of a dashboard that provides detailed information about the audited **Actions** that are listed on the **Audits** tab.

To access the details page of an audited **Action**, click the name of an audited action to see the details page of the audited action.

Traffic Mirroring

Traffic mirroring replicates traffic from the interfaces of the AHV hosts to the virtual NIC (vNIC) of guest VMs.

Traffic mirroring replicates the packets from a set of source ports to a set of destination ports. You can mirror inbound, outbound, or bidirectional traffic flowing on a set of source ports. You can then use the mirrored traffic for security analysis and to gain visibility of traffic flowing through the set of source ports. Traffic mirroring is a useful tool for troubleshooting packets and necessary for compliance.

Important: If your cluster is not registered to Prism Central, you can use aCLI to configure the Traffic mirroring session on an AHV host. For more information, see [Configuring Traffic Mirroring on an AHV Host](#) section in AHV Administration Guide.

In Prism Central you can select the following entity types for source and destination:

- Source:
 - Individual host ports
 - Bonded host ports
 - Virtual Machines

Note: You can only select up to four entities in total for all the source entities.

- Destination: Virtual Machines

Note: You can select up to two destination entities.

Source Ports

Prism Central supports the following types of source ports in a traffic mirroring session:

- A bond port that is already mapped to a virtual switch (vs) such as vs0, vs1, or any other vs that you created
- A non-bond port that is already mapped to a vs such as vs0, vs1, or any other vs that you created
- An uplink port that is not assigned to any vs or bridge on the host
- A virtual port attached to a VM

Important Considerations

Consider the following before you create traffic mirroring sessions:

- Prism Central supports traffic mirroring only from physical and VM interfaces.
- Prism Central supports traffic mirroring only to destination guest VMs running on the same AHV host (to which the source ports belong) or on a remote AHV host.
- VM interface port or bond supports traffic mirroring to a destination on a remote host.
- Physical source port or bond interface does not support traffic mirroring to a destination on a remote host.
- In a single traffic mirroring session, if both remote and local host destinations are available, the physical source port or bond interface is mirrored only to the local destination, while the VM interface port or bond is mirrored to both remote and local destinations.
- Nutanix does not support traffic mirroring on Nutanix clusters connected to AWS or Azure.
- Prism Central supports different types of source ports in one session. For example, you can create a session with br0-up (bond port) and eth5 (single uplink port) on the same host as two different source ports in the same session. You can have two different bond ports in the same session.
- The traffic mirroring session displays an error status in the following conditions:
 - If you do not delete the traffic mirroring session before you delete the traffic mirroring destination VM or vNIC.
 - If you do not delete the vNIC if you migrate a destination VM, because the traffic mirroring session displays an error status if the physical port is down.
 - If the session has physical NICs (pNICs).
If all vNICs are migrated to the same host, the traffic mirroring session displays active status.
- You cannot configure the direction of the traffic if the ports are not managed by a virtual switch.

- You can select one source VM vNIC for a single traffic mirroring session.
- You can select one destination VM vNIC for a single traffic mirroring session.

Traffic Mirroring Session Scale

Traffic mirroring session supports the following scale:

Table 88: Traffic Mirroring Session Scale

Entities	Scale
Source ports or entities	4 entities per traffic mirroring session
Destination ports or entities	2 entities per traffic mirroring session
Maximum number of Traffic Mirror sessions	1,000 per cluster
Maximum number of active sessions	2 per host

Viewing Traffic Mirroring Sessions

You can view the details of all the traffic mirroring sessions configured across registered clusters on the **Traffic Mirroring** page.

Procedure

To view the detailed summary of all the traffic mirroring sessions, perform the following steps:

1. Log in to Prism Central.
2. From the [Application Switcher](#), select the **Infrastructure** application, and from the navigation bar, select **Network & Security > Network Services**.
3. Click the **Traffic Mirroring** tab.
You can view all the traffic mirroring sessions configured for the registered clusters.

Traffic Mirroring Session Fields

The following traffic mirroring session fields are displayed on the **List** tab:

Note: The fields vary based on the **View by** option. A dash (-) is displayed in a field when a value is not available or applicable. For more information, see the [View by](#) on page 51 information.

Table 89: Traffic Mirroring - Field Description

Field	Description	Values
Name	Traffic mirroring session name	(traffic mirroring session name)
Source (Ingress)	Traffic entering the source	(source VM, MAC address)
Source (Egress)	Traffic leaving the source	(source VM, MAC address)
Destinations	Destination virtual machine with MAC address	(destination VM, MAC address)

Field	Description	Values
State	Current operational status of the traffic mirroring session	(Enabled, Disabled)

Filters Pane - List Tab

You can filter the information in the **List** tab based on the following fields available in the **Filters** pane. For information about how to use the **Filters** pane, see the [Filters Pane](#) information.

Table 90: Filter Pane Field Description - Traffic Mirroring Session

Field	Description	Values
Name	Filters the information based on the application instance name. Select the checkbox, choose a condition from the dropdown menu (Contains, Equal to, Not Equal to, Doesn't contain, Starts with, Ends with, or), and enter a string in the field.	(Traffic Mirroring Session name string)
State	Filters the information based on the state of the traffic mirroring session based on one of the following checkboxes you select: Enabled, Disabled, or both .	Enabled, Disabled

Traffic Mirroring Session Actions

You can perform the following actions on a traffic mirroring session using the **Actions** dropdown menu:

Table 91: Traffic Mirror Session Actions

Action	Description
Disable Session	Click this action to disable the traffic mirroring session. For more information, see Disabling a Traffic Mirroring Session on page 400.
Enable Session	Click this action to enable the traffic mirroring session. For more information, see Enabling a Traffic Mirroring Session on page 401.
Update	Click this action to update the traffic mirroring session. For more information, see Updating a Traffic Mirroring Session on page 401.
Delete	Click this action to delete the traffic mirroring session. For more information, see Deleting a Traffic Mirroring Session on page 401.

Actions from List Tab

You can view and perform the following actions from the **List** tab:

- Click the name of a traffic mirroring session to open the details , which displays detailed information about the traffic mirroring session.
- Create a traffic mirroring session by clicking **Create Mirror Session**. For more information, see [Creating a Traffic Mirroring Session](#) on page 396.
- Download the table of traffic mirroring sessions in CSV format by clicking **Export**.
- View the reports based on pre-defined criteria or create a custom view. For more information, see [View by](#) on page 51.
- Filter the traffic mirroring session list based on a variety of parameter values using the **Modify Filters** pane.
- Perform the following actions on a traffic mirroring session using the **Actions** dropdown menu.

Traffic Mirroring Details View

The Traffic Mirroring details page consists of dashboard widgets that provide detailed information about the traffic mirroring sessions.

To access the details page of an individual traffic mirroring session, perform the following steps:

1. Log in to Prism Central.
2. From the [Application Switcher](#) , select the **Infrastructure** application, and from the navigation bar, select **Network & Security** > **Network Services**.
3. Click the **Traffic Mirroring** tab.

You can view all the traffic mirroring sessions configured for the registered clusters.

4. Click the name of a traffic mirroring session to open the details page.

The **Summary** tab opens displaying the detailed information about the traffic mirroring session in widgets.

The following figure shows the **Summary** page of an individual **Traffic Mirroring Session**:




Figure 133: Traffic Mirroring Session Details View

The **Summary** page has the following widgets:

Table 92: Traffic Mirroring Session Summary Tab Widgets

Widget	Parameters
Entities	- Cluster: Displays the cluster name. - Virtual Switch for Mirrored Traffic: Displays the virtual switch used for traffic mirroring.
Source: Host Ports (Individual)	- Hosts: Displays the host name. - Ports: Displays the host port details. - Direction: Displays the traffic direction of the host port.
Source: Virtual Machines	- VM: Displays the name of the source virtual machine. - MAC Address: Displays the MAC address of the source virtual machine.
Statistics	- Status: Displays the operational status of the traffic mirroring session. - Packet Count: Displays the packet count mirrored and received from the destination. - Packet Byte Count: Displays the packet byte count derived from the packet count.
Destination: Virtual Machines	- VM: Displays the name of the destination virtual machine. - MAC Address: Displays the MAC address of the destination virtual machine.
Source: Host Ports(Bonds)	- Hosts: Displays the host name. - Bonds: Displays the bond name. - Ports: Displays the details of the bonds on the host port. - Direction: Displays the direction of the traffic on the host port bonding.

You can perform the following actions from the Summary tab:

- Disable an existing traffic mirroring session by clicking **Disable Session**.
- Update an existing traffic mirroring session by clicking **Update**.
- Delete an existing traffic mirroring session by clicking **Delete**.

Creating a Traffic Mirroring Session

Create a traffic mirroring session to mirror inbound, outbound, or bidirectional traffic on a set of source ports.

Before you begin

Ensure that you configure the traffic mirror type on the destination VM's vNIC. For more information on adding a traffic mirror type vNIC to the destination VM, see [Creating a VM through Prism Central \(AHV\)](#) on page 149.

About this task

Follow these steps to create a traffic mirroring session that replicates the traffic from a set of source ports to a set of destination ports:

Procedure

1. Log in to Prism Central.
2. From the [Application Switcher](#), select the **Infrastructure** application and from the navigation bar, select **Network & Security > Network Services**.
3. Click the **Traffic Mirroring** tab.
You can view all the traffic mirroring sessions configured for the registered clusters.
4. Click **Create Mirror Session**.
5. In the **General** tab, enter the following information:
 - a. **Name**: Enter a name for the traffic mirroring session.
 - b. (Optional) **Description**: Enter a description for the traffic mirroring session.
 - c. **Cluster**: From the dropdown list, select the cluster.
 - d. **Virtual Switch for mirrored traffic**: From the dropdown list, select the virtual switch where the traffic mirroring session is to be created.

Note: Nutanix recommends utilizing a non-default virtual switch with a Maximum Transmission Unit (MTU) configured between 1600 and 9000 for traffic mirroring to a destination on the remote host. You can also use the default virtual switch **VS0** for traffic mirroring with an MTU set to 1600 bytes. For more information on MTU, see [Requirements and Limitations of Flow Virtual Networking](#) and [Configuring Virtual Switch for VPC Traffic Types](#).

To create traffic mirroring sessions on a non-default virtual switch, run the following `acli` command to configure the IP addresses for the hosts on the non-default virtual switch along with a gateway IP address for that network.

```
nutanix@cvm$ acli net.update_virtual_switch virtual-switch-name  
host_ip_addr_config='{host-uuid1:host_ip_address/prefix;host-  
uuid2:host_ip_address/prefix;host-uuid3:host_ip_address/prefix}'  
gateway_ip_address=IP_address
```

The IP addresses used in this command must be reachable among hosts within the cluster. Ensure that the IP addresses are unique within the same cluster. All the IP addresses used in the command must be in the same subnet. The IP addresses must be in a different subnet than the management IP address of interface `br0` of the Controller VM (CVM).

You can find the uuid required in the previous command by using the `acli host.list` command.

6. Click **Next**.

7. In the **Source & Destination** tab, follow these steps to select the individual host ports:
 - a. From the **Source Type(s)** dropdown menu, select **Host Ports (Individual)**.
 - b. On the **Hosts Ports (Individual)** panel, click **Select Host Ports**.
Select Host Ports (Individual) window appears.
 - c. Click the down arrow icon next to the host name.
 The system displays **NICs** associated with the host and the dropdown menu to select the **Traffic Direction**.
 - d. Under **NICs**, select the checkbox of the Ethernet port to mirror the traffic for.
 - e. Under **Traffic Direction**, from the dropdown menu, select the direction of the traffic:
 - **Both** (default): Includes both directions, Ingress and Egress.
 - **Ingress**: Traffic entering the source port.
 - **Egress**: Traffic leaving the source port.

Note: Traffic direction is unavailable if the ports are not managed by a virtual switch (vs).
 - f. Click **Save**.

Select Host Ports (Individual) ✕

Select desired host ports and traffic direction

Selected 2 out of 21 Ports

Hosts	NICs	Traffic Direction
🔍 Search for ports		
ntnx-nahv-65a62710731...	7	⌵
ntnx-nahv-65a62710731...	7	⌴
☒ eth1 ☐ eth3 Ingress Both		

Cancel
Save

Figure 134: Individual Host Ports - Source

8. In the **Source & Destination** tab, follow these steps to select the bonded host ports:
 - a. From the **Source Type(s)** dropdown menu, select **Host Ports (Bond)**.
 - b. On the **Hosts Ports (Bond)** panel, click **Select Host Ports**.
Select Host Ports (Bond) window appears.
 - c. Click the down arrow icon next to the host name.
 The system displays **Bonds** and **Ports** associated with the host along with a drop-down menu to select the **Traffic Direction**.
 - d. Under **Bonds**, select the checkbox of the bonded port to mirror the traffic for.
 - e. Under **Direction**, from the dropdown menu, select the direction of the traffic:
 - **Both** (default): Includes both directions, Ingress and Egress.
 - **Ingress**: Traffic entering the source port.
 - **Egress**: Traffic leaving the source port.
 - f. Click **Save**.

Select Host Ports (Bond) [X]

Select desired bond and traffic direction

Selected 1 out of 3 Bonds

Hosts	Bonds	Ports	Direction
ntnx-nahv-65...	1		
ntnx-nahv-65...	1		
	☒ vs0 ⓘ	eth0, eth2	Both ▾
ntnx-nahv-65...	1		

Cancel Save

Figure 135: Bonded Host Ports - Source

9. In the **Source & Destination** tab, follow these steps to select the host virtual machines:
- From the **Source Type(s)** dropdown menu, select **Virtual Machines**.
 - On the **Virtual Machines** panel, click **Select Virtual Machines**.
Select Virtual Machines window appears.
 - Under **Virtual Machine**, from the dropdown menu, select the source virtual machine.
 - Under **MAC Address**, from the dropdown menu, select the MAC address.
 - Under **Direction**, from the dropdown menu, select the direction of the traffic:
 - Both** (default): Includes both directions, Ingress and Egress.
 - Ingress**: Traffic entering the source VM.
 - Egress**: Traffic leaving the source VM.
 - (Optional) To add a VM, click **Add VM**.
 - Click **Save**.

Select Virtual Machines

Select desired virtual machine with MAC address and traffic direction

Selected 1 VMs [Add VM](#)

Virtual Machine	MAC Address	Direction	
test-vm-1	50:6b:8d:d6:00:0f	Both	

Cancel Save

Figure 136: Virtual Machine - Source

10. In the **Source & Destination** tab, follow these steps to select the destination virtual machines:
 - a. Under destination, click **Select Virtual Machines**.
Select Virtual Machines window appears.
 - b. Under **Virtual Machine**, from the dropdown menu, select the destination virtual machine.
 - c. Under **MAC Address**, from the dropdown menu, select the MAC address.
 - d. (Optional) To add a VM, click **Add VM**.
 - e. Click **Save**.




Figure 137: Virtual Machine - Destination

11. Click **Next**.
12. In the **Summary** tab, verify the traffic mirroring session details and do one of the following:
 - » To create a session, click **Create Session**.
 - » To create and enable a session, click **Create and Enable Session**.

Disabling a Traffic Mirroring Session

Disable an existing traffic mirroring session to stop mirroring traffic to the destination VM.

Procedure

1. Log in to Prism Central.
2. From the [Application Switcher](#), select the **Infrastructure** application and from the navigation bar, select **Network & Security > Network Services**.
3. Click the **Traffic Mirroring** tab.
 You can view all the traffic mirroring sessions configured for the registered clusters.
4. Select the checkbox of the traffic mirroring session to disable.
5. From the **Actions** dropdown menu, select **Disable Session**.
 Prism Central displays the **Disable Session** window.

6. Click **Disable Session**.

Enabling a Traffic Mirroring Session

Enable a previously created traffic mirroring session to start mirroring traffic to the destination VM.

Procedure

1. Log in to Prism Central.
2. From the [Application Switcher](#), select the **Infrastructure** application and from the navigation bar, select **Network & Security > Network Services**.
3. Click the **Traffic Mirroring** tab.
You can view all the traffic mirroring sessions configured for the registered clusters.
4. Select the checkbox of the traffic mirroring session to enable.
5. From the **Actions** dropdown menu, select **Enable Session**.
Prism Central displays the **Enable Session** window.
6. Click **Enable Session**.

Updating a Traffic Mirroring Session

Update a traffic mirroring session to modify parameters such as name, description, cluster, host, source and destination types.

Procedure

1. Log in to Prism Central.
2. From the [Application Switcher](#), select the **Infrastructure** application and from the navigation bar, select **Network & Security > Network Services**.
3. Click the **Traffic Mirroring** tab.
You can view all the traffic mirroring sessions configured for the registered clusters.
4. Select the checkbox of the traffic mirroring session to update.
5. From the **Actions** dropdown menu, select **Update**.
Prism Central displays the **Update Mirror Session** window.
6. Modify the traffic mirroring session based on your requirements and click **Update Session**.

Deleting a Traffic Mirroring Session

Delete a traffic mirroring session to permanently remove it from the traffic mirroring session page list.

Procedure

1. Log in to Prism Central.
2. From the [Application Switcher](#), select the **Infrastructure** application and from the navigation bar, select **Network & Security > Network Services**.
3. Click the **Traffic Mirroring** tab.
You can view all the traffic mirroring sessions configured for the registered clusters.
4. Select the checkbox of the traffic mirroring session to delete.

5. From the **Actions** dropdown menu, select **Delete**.
Prism Central displays the **Delete Session** window.
6. Click **Delete Session**.

Traffic Mirroring Session Alerts

Prism Central generates the traffic mirroring session alerts for these inconsistent state scenarios.

Note: For more information about the Alerts generated in Prism Central, see [Prism Central Alerts and Events Reference Guide](#).

Table 93: Traffic Mirror Session Alerts

Alert Title	Possible Cause
Inconsistent SPAN Session State Detected	A VM NIC is removed from the destination VM.
Inconsistent SPAN Session State Detected	A destination VM is migrated to another host.
Inconsistent SPAN Session State Detected	A destination VM is powered off.

Role-Based Access Control in Traffic Mirroring

With the Prism Central role-based access control (RBAC), you can provide customized access to users for the traffic mirroring entity.

You need to configure and provide customized access to users based on their assigned roles for the traffic mirroring entity. For more information on configuring RBAC, see [Controlling User Access \(RBAC\)](#).

From the Prism Central roles dashboard, you can define and assign the following roles to users or user groups:

- Predefined built-in roles (system roles)
- Custom roles

Any user with the following predefined built-in roles can configure the traffic mirroring functionality:

- Network Infra Admin
- Prism Admin

For more information on predefined built-in and custom roles, see [Built-in Role Management](#) and [Custom Role Management](#) in the *Security Guide*.

Traffic Mirroring Permissions

You can add the permissions listed to a custom role and assign that custom role to the user to perform various traffic mirroring tasks:

- Create Traffic Mirror
- Delete Traffic Mirror
- Update Traffic Mirror
- View Traffic Mirror
- View Traffic Mirror Stats
- View Cluster

- View Cluster Networking Capabilities
- View Host
- View Uplink Bond
- View VM

Floating IPs

You can access the floating IP addresses you have created and configured, from the **Floating IPs** page.

For information on floating IP addresses and their role in flow virtual networking, see the *SNAT and Floating IP Address* section in the [Essential Concepts](#) topic of the *Flow Virtual Networking Guide*.

Note: Floating IP addresses are not reachable (Pings fail) unless you associate them to primary or secondary IP addresses of VMs. For more information, see [Assigning Secondary IP Addresses to Floating IPs](#) in the *Flow Virtual Networking Guide*.

For information on the limitation on using floating IP on a guest VM with load balancing configuration, see [Network Load Balancer](#) in the *Flow Virtual Networking Guide*.

Floating IPs Summary View

The **Floating IPs** page displays the list of floating IP addresses across all the registered clusters.

To access the **Floating IPs** page:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Network & Security > Floating IPs** from the **Navigation Bar**.

The **Floating IPs** page opens displaying the **List** tab. This tab provides a list of floating IPs you have created and configured, and the operations you can perform on the IPs.

Note: For information on the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

The following table describes the fields that appear in the **Floating IPs** page.

Note: The fields vary based on the **View by** option. A dash (-) is displayed in a field when a value is not available or applicable. For more information, see the [View by](#) on page 51 information.

Table 94: Floating IPs – Field Description

Parameter	Description	Values
Floating IP Address	Displays the floating IP address assigned.	(IP address)
External Subnet	Displays the name of the external subnet that the IP address is assigned to.	(Name of the assigned subnet)
Association Status	Displays the status of association between the IP address and the external subnet and VPC.	Associated
VPC	Displays the name of the VPC associated with the IP address.	(Name of the associated VPC)
VM Name	Displays the name of the VM associated with the IP address.	(Name of the assigned VM)

Parameter	Description	Values
Private IP	Displays the private IP address assigned to the same VM. This private IP address is assigned from the internal private subnet that the network controller creates when you create a network gateway.	(IP address)

You can perform the following actions for the floating IP addresses from the **Floating IPs** page:

- Request a floating IP address by clicking **Request Floating IP**. For more information, see [Requesting Floating IPs](#) in the *Flow Virtual Networking Guide*.
- Update or delete an existing floating IP address using the **Actions** dropdown menu. The **Actions** dropdown appears when one or more addresses are selected.
 - **Update**: Assign or change the assignment of the floating IP address. You can assign the floating IP address to a IP address such as a private IP address in a VPC or the primary IP address of a VM or a secondary IP address created on a VM.
 - **Delete**: Delete the floating IP address. The deleted IP address returns to the IP address pool as unused. Before you delete a floating IP address, ensure that it is not assigned to a private IP address or a VM. Change the assignment to None if it is already assigned, using the **Update** option.
- View floating IP addresses based on pre-defined criteria or create a custom view. For more information on the **View by** option, see [View By](#) information.
- Filter the floating IP addresses list based on a variety of parameter values using **Filters** pane. For more information, see [Filters Pane - Floating IPs Page](#).

Filters Pane - Floating IPs Page

You can filter the information in the **Floating IPs** page based on the following fields that are available in the **Filters** pane. For information on how to use the Filters pane, see the [Filters Pane](#) information.

Table 95: Filter Pane Field Description - Floating IPs page

Field	Description	Values
Floating IP Address	Filters based on the floating IP address assigned. It returns a list of IP addresses that satisfy the string.	(Floating IP address)
External Subnet	Filters based on the external subnet that the IP address is assigned to.	(External Subnet)

Connectivity

You can access network gateways, VPN connections, subnet extensions, and BGP sessions from the **Connectivity** page.

To access the **Connectivity** page:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Network & Security** > **Connectivity** from the **Navigation Bar**.

The **Connectivity** page opens displaying the **Gateways** tab. This tab provides a list of network Gateways you have created and configured, and the operations you can perform on the network Gateways.

- To view the VPN connections, click the **VPN Connections** tab.
- To view the subnets extended across the clusters, click the **Subnet Extensions** tab.
- To view the BGP sessions created for the clusters, click the **BGP Sessions** tab.

Gateways Summary View

The **Gateways** page displays a list of gateways created for the clusters managed by Prism Central.

To access the **Gateways** page:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Network & Security > Connectivity** from the **Navigation Bar**.

The **Gateways** page opens displaying the list of network gateways you have created and configured, and the operations you can perform on the network gateways.

Note: For information on the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

The following table describes the fields that appear in the **Gateways** page.

Table 96: Field Descriptions for the Gateway Page

Parameter	Description	Values
Name	Displays the name of the gateway.	(Name of gateway)
Type	Displays the gateway type.	(Local or Remote)
Service	Displays the service that the gateway uses.	(VPN or VTEP)
Service IP	Displays the IP address used by the service.	(IP address)
Status	Displays the operational status of the gateway.	(Up or Down)
Attachment Type/Vendor	Displays the type of subnet associated with the gateway.	(VLAN or Overlay-VPC name)
Connections	Displays the number of service connections (such as VPN connections) configured and operational on the gateway.	(Number)

You can perform the following actions for a gateway from the **Gateways** page:

- Click the name of a gateway to open the gateway details page, which displays the detailed information about the gateway. For more information, see [Gateway Details View](#) on page 406.
- Create a local or remote gateway with VPN or VTEP service by clicking the **Create Gateway** dropdown menu. For more information, see [Creating a Network Gateway](#) in the *Flow Virtual Networking Guide*.
- Update or delete an existing gateway using the **Actions** dropdown menu. The **Actions** dropdown menu appears when one or more gateways are selected. For more information, see [Updating a Network Gateway](#) or [Deleting a Network Gateway](#) in the *Flow Virtual Networking Guide*.
- Filter the gateway list based on various parameter values using the **Filters** pane. For more information, see [Filters Pane - Gateways Page](#).

Filters Pane on the Gateways Page

You can filter the information in the **Gateways** page based on the following fields that are available in the **Filters** pane. For information on how to use the Filters pane, see the [Filters Pane](#) information.

Table 97: Filter Pane Field Descriptions for the Gateways page

Field	Description	Values
Name	Filters based on the gateway name. It returns a list of gateways that satisfy the name condition/string.	(Gateway name string)
Service IP	Filters based on IP address used by the service.	(IP address)
Status	Filters based on the operational status of the gateway.	(Up or Down)

Gateway Details View

The **Summary** page of an individual gateway consists of a dashboard that provides the detailed information about the gateway.

To access the **Summary** page of an individual gateway:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Network & Security > Connectivity** from the **Navigation Bar**.

The **Gateways** page opens displaying the list of network gateways that you have created and configured.

3. Click a gateway to view the **Summary** page of the gateway.

Note: For information on the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

The gateway **Summary** page has the following widgets:

Table 98: Field Descriptions for the Gateway Widgets

Parameter	Description	Values
Properties widget		
Type	Displays the gateway type.	(Local or Remote)
Attachment Type	Displays the network entity like VLAN or VPC that the gateway is attached to.	(VLAN or VPC)
VPC or Subnet (VLAN)	Displays the name of the attached VPC or VLAN subnet.	(Name of VLAN or VPC)
Vendor (Applicable only if you select remote gateway)	Displays the name of the vendor of the gateway appliance at the remote site.	(Name of Vendor)
Floating or Private IP Address	Displays the Floating (for VPC) or Private (for VLAN) IP address assigned to the gateway.	(IP Address)
External IP (Applicable only if you select remote gateway)	Displays the IP address assigned to the remote gateway.	(IP Address that you assigned to the remote gateway.)

Parameter	Description	Values
Status	Displays the operational status of the gateway.	(Up or Down)
Gateway Version	Displays the version of the Nutanix gateway appliance deployed.	(Version)
Cluster	Displays the name of the cluster on which the gateway is created.	(Cluster name)
Gateway VM	Displays the name of the VM on which the gateway is created.	(Name of VM - actionable link. Click the name-link to open the VM details page of the gateway VM.)
Service Configuration widget		
Service	Displays the service used by the gateway.	(VPN or VTEP or BGP)
VPN Service Configuration		
External Routing	Displays the type of routing associated with the gateway for external traffic routing.	(Static or eBGP with ASN)
Internal Routing	Displays the type of routing associated with the gateway for internal traffic routing.	(Static or eBGP with ASN)
VPN Connections	Displays the total number of VPN connections associated with the gateway.	(Number - actionable link. Click the link to open the VPN connection details page for the associated VPN connection.)
View VPN Connections	Click this link to open the VPN Connections tab.	-
VTEP Service Configuration		
VXLAN (UDP) Port	Displays the VXLAN (UDP) Port for the gateway.	(Number)
Subnet Extensions	Displays the total number of subnet extensions associated with the gateway.	(Number - actionable link. Click the link to open the subnet extensions details page for the associated subnet extension.)
View Subnet Extensions	Click this link to open the Subnet Extensions tab.	-
BGP Service Configuration		
ASN	Displays the ASN of the EBGp route.	(Number)
BGP Sessions	Displays the total number of BGP sessions associated with the gateway.	(Number - actionable link. Click the link to open the BGP sessions details page for the associated BGP session.)
Serviced VPC	Displays VPC service used by the gateway.	(Name of VPC)
View BGP Sessions	Click this link to open the BGP Sessions tab.	-

You can perform the following actions for a gateway from the **Summary** tab:

- Update an existing gateway by clicking **Update**. For more information, see [Updating a Network Gateway](#) in the *Flow Virtual Networking Guide*.
- Delete the gateway by clicking **Delete**. For more information, see [Deleting a Network Gateway](#) in the *Flow Virtual Networking Guide*.

VPN Connections Summary View

The **VPN Connections** page displays a list of VPN connections created for the clusters managed by Prism Central.

A VPN connection represents the VPN IPsec tunnel established between local gateway and remote gateway. When you create a VPN connection, you must select two gateways between which you want to create the VPN connection.

To access the **VPN Connections** page:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Network & Security > Connectivity** from the **Navigation Bar**.

The **Gateways** page opens displaying the list of network gateways.

3. Click the **VPN Connections** tab.

The **VPN Connections** page opens displaying the list of VPN connections created for the clusters.

Note: For information on the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

The following table describes the fields that appear in the **VPN Connections** page.

Table 99: Field Descriptions for the VPN Connections Page

Parameter	Description	Values
Name	Displays the name of the connection.	(gateway name)
IPSec Status	Displays the connection status of IPsec tunnel.	(Connected or Not Connected)
EBGP Status	Displays the status of the EBGP gateway connection.	(Established or Not Established)
Local Gateway	Displays the name of the local gateway used for the connection.	(Name of local gateway)
Remote Gateway	Displays the name of the remote gateway used for the connection.	(Name of remote gateway)
Dynamic Routing Priority	Displays the dynamic routing priority assigned to the connection for throughput management. You can assign any value in the range of 100-1000. Nutanix Flow Virtual Networking assigns the first VPN connection the value 500 by default. Thereafter, subsequent VPN connections are assigned values decremented by 50. For example, the first connections is assigned 500, then the second connection is assigned 450, the third one 400 and so on.	(Number in the range of 100-1000. User assigned.)

You can perform the following actions for a VPN connection from the **VPN Connections** page:

- Click the name of a VPN connection to open the VPN connection details page, which displays the detailed information about the connection. For more information, see [VPN Connection Details View](#) on page 409.
- Create a VPN connection by clicking **Create VPN Connection**. For more information, see [Creating a VPN Connection](#) in the *Flow Virtual Networking Guide*.
- Update or delete an existing VPN connection using the **Actions** dropdown menu. The **Actions** dropdown appears when one or more VPN connections are selected. For more information, see [Updating VPN Connection](#) or [Deleting a VPN Connection](#) in the *Flow Virtual Networking Guide*.
- Filter the VPN connection list based on various parameter values using the **Filters** pane. For more information, see [Filters Pane - VPN Connections Page](#).

Filters Pane on the VPN Connections Page

You can filter the information in the **VPN Connections** page based on the following fields that are available in the **Filters** pane. For information on how to use the Filters pane, see the [Filters Pane](#) information.

Table 100: Filter Pane Field Descriptions for the VPN Connections page

Field	Description	Values
Name	Filters based on the VPN connection name. It returns a list of VPN connections that satisfy the name condition/string.	(VPN connection name string)
EBGP Status	Filters based on the status of the EBGP gateway connection.	(Established or Not Established)
IPSEC Status	Filters based on the connection status of IPsec tunnel.	(Connected or Disconnected)

VPN Connection Details View

The VPN Connection details page provides detailed information about a VPN connection.

The details page has the **Summary**, **Throughput**, **IPSec Logging**, and **Routing Protocol Logging** tabs.

To access the details page of an individual VPN connection:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Network & Security > Connectivity** from the **Navigation Bar**.

The **Gateways** page opens displaying the list of network gateways you have created and configured.

3. Click the **VPN Connections** tab.

The **VPN Connections** page opens displaying the list of VPN connections created for the clusters.

4. Click the name of a VPN connection to open the details page of the connection.

The **Summary** tab opens displaying the detailed information about the VPN connection in widgets.

Note: For information on the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.




Figure 138: VPN Connection Details

Summary Tab

The **Summary** tab provides detailed information about a VPN connection in widgets.

The following table describes the fields that appear in the **Summary** tab.

Table 101: Field Descriptions for the Summary Tab

Parameter	Description	Values
VPN Connection widget		
IPSec Status	Displays the connection status of IPSec tunnel.	(Connected or Not Connected)
EBGP Status	Displays the status of the EBGP gateway connection.	(Established or Not Established)
Dynamic Routing Priority	Displays the dynamic routing priority assigned to the connection for throughput management. You can assign any value in the range of 100-1000. Flow Virtual Networking assigns the first VPN connection the value 500 by default. Thereafter, subsequent VPN connections are assigned values decremented by 50. For example, the first connections is assigned 500, then the second connection is assigned 450, the third one 400 and so on.	(Number in the range of 100-1000. User assigned.)
Local Gateway Properties widget		
Gateway Name	Displays the name of the local gateway used for the connection.	(Name of local gateway)
Type	Displays the type of gateway.	(Local)
Attachment Type	Displays the network entity like VLAN or VPC that the gateway is attached to.	(VLAN or VPC)

Parameter	Description	Values
VPC or Subnet (VLAN)	Displays the name of the attached VPC or VLAN subnet.	(Name of VLAN or VPC)
Tunnel IP	Displays the Tunnel IP address of the local gateway.	(IP Address)
Connection Type	Displays the connection type you selected while creating the VPN connection. The connection type may be Initiator or Acceptor of a VPN connection between the local and remote gateways. T	(Initiator or Acceptor)
External Routing	Displays the type of routing associated with the gateway for external traffic routing.	(Static or eBGP with ASN)
Internal Routing	Displays the type of routing associated with the gateway for internal traffic routing.	(Static or eBGP with ASN)
Floating or Private IP Address	Displays the Floating (for VPC) or Private (for VLAN) IP address assigned to the gateway.	(IP Address that you assigned to the local gateway with /30 prefix when you configured the VPN connection.)
Status	Displays the operational status of the gateway.	(Up or Down)
Cluster	Displays the name of the cluster on which the gateway is created.	(Cluster name)
Gateway VM	Displays the name of the VM on which the gateway is created.	(Name of VM - actionable link. Click the name-link to open the VM details page of the gateway VM.)
Remote Gateway Properties widget		
Gateway Name	Displays the name of the remote gateway used for the connection.	(Name of remote gateway)
Type	Displays the type of gateway.	(Remote)
Tunnel IP	Displays the Tunnel IP address of the remote gateway.	(IP Address)
Connection Type	Displays the connection type you selected while creating the VPN connection. The connection type may be Initiator or Acceptor of a VPN connection between the local and remote gateways. T	(Initiator or Acceptor)
External Routing	Displays the type of routing associated with the gateway for external traffic routing.	(Static or eBGP with ASN)
ASN	Displays the ASN of the EBGP route. This information is only displayed if you configured EBGP as the External Routing protocol.	(Number)
Vendor	Displays the name of the vendor of the gateway appliance at the remote site.	(Name of vendor of gateway appliance)

Parameter	Description	Values
External IP	Displays the IP address assigned to remote the gateway.	(IP Address that you assigned to the remote gateway with /30 prefix when you configured the VPN connection.)
Status	Displays the operational status of the gateway.	-
Protocol Details widget		
Service	Displays the service used by the gateway.	(VPN or VTEP)
Gateway Routes widget	Displays the status of the routes used by the gateways.	(Sent)

You can perform the following actions from the **Summary** tab:

- View the detailed information of a VPN connection. For the list of available parameters, see the *VPN Connection Summary Tab* table above.
- Update an existing VPN connection by clicking **Update**. For more information, see [Updating VPN Connection](#) in the *Flow Virtual Networking Guide*.
- Delete an existing VPN connection by clicking **Delete**. For more information, see [Deleting a VPN Connection](#) in the *Flow Virtual Networking Guide*.

Throughput Tab

The **Throughput** tab provides a graphical representation of the throughput of the VPN connection.

IPSec Logging

The **IPSec Logging** tab provides running logs for the IPSec tunnel of the VPN connection.

Routing Protocol Logging

The **Routing Protocol Logging** tab provides logs for the routing protocol used in the VPN connection.

Subnet Extensions Summary View

The **Subnet Extensions** page displays a list of subnet extensions created for the clusters managed by Prism Central.

To access the **Subnet Extensions** page:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Network & Security** > **Connectivity** from the **Navigation Bar**.

The **Gateways** page opens displaying the list of network gateways.

3. Click the **Subnet Extensions** tab.

The **Subnet Extensions** page opens displaying the list of subnet extensions created for the clusters.

Note: For information on the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

The following table describes the fields that appear in the **Subnet Extensions** page.

Table 102: Field Description for the Subnet Extensions Page

Parameter	Description	Values
Name	Displays the name of the subnet extension.	(Name of subnet extension)
Type	Displays the subnet extension type.	(Across Availability Zones or To a Third Party Data Center)
Extension Over	Displays the service that the subnet extension uses.	(VPN or VTEP)
Extension Uses	Displays the name of the local network gateway that the subnet extension uses.	(Name of local network gateway)
Local Subnet	Displays the name of the local subnet that the subnet extension uses.	(Name of local subnet)
Remote Site	Displays the name of the remote network gateway that the subnet extension uses.	(Name of remote network gateway)
Connection Status	Displays the status of the connection that is created by the subnet extension.	(Not Available, Connected, or Disconnected)
	Note: Not Available status indicates that Prism Central is unable to ascertain the status.	
Interface Status	Displays the status of the interface that is used by the subnet extension.	(Connected or Down)

You can perform the following actions for a subnet extension from the **Subnet Extensions** page:

- Click the name of a subnet extension to open the subnet extension details page, which displays the detailed information about the extension. For more information, see [Subnet Extension Details View](#) on page 414.
- Extend a subnet **Across Availability Zones** or **To a Third Party Data Center** by clicking the **Create Subnet Extension** dropdown menu. You can extend a subnet using **VPN** or **VTEP** service. For more information, see [Layer 2 Virtual Network Extension](#) in the *Flow Virtual Networking Guide*.
- Update or delete existing subnet extension using the **Actions** dropdown menu. The **Actions** dropdown appears when one or more subnet extensions are selected. For more information, see [Updating an Extended Subnet](#) or [Removing an Extended Subnet](#) in the *Flow Virtual Networking Guide*.
- Filter the subnet extension list based on various parameter values using the **Filters** pane. For more information, see [Filters Pane - Subnet Extensions Page](#).

Filters Pane on the Subnet Extensions Page

You can filter the information in the **Subnet Extensions** page based on the following fields that are available in the **Filters** pane. For information on how to use the Filters pane, see the [Filters Pane](#) information.

Table 103: Filter Pane Field Descriptions for the Subnet Extensions page

Field	Description	Values
Name	Filters based on the subnet extension name. It returns a list of subnet extensions that satisfy the name condition/string.	(Subnet extension name string)
Connection Status	Filters based on the status of the connection that is created by the subnet extension.	(Connected or Disconnected)
Interface Status	Filters based on the status of the interface that is used by the subnet extension.	(Connected or Not Available)

Subnet Extension Details View

The Subnet Extension details page provides detailed information about a subnet extension.

The details page has the **Summary**, **Address Table**, and **Throughput** tabs.

To access the details page of an individual subnet extension:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Network & Security > Connectivity** from the **Navigation Bar**.

The **Gateways** page opens displaying the list of network gateways you have created and configured.

Note: For information on the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

3. Click the **Subnet Extensions** tab.

The **Subnet Extensions** page opens displaying the list of subnet extensions created for the clusters.

4. Click a subnet extension to open the details page of the extension.

The **Summary** tab opens displaying the detailed information about the extension in widgets.

Summary Tab

The **Summary** tab provides detailed information about the subnet extension in widgets.

The subnet extension **Summary** tab has the following widgets:

Table 104: Subnet Extension Summary Tab Widgets

Parameter	Description	Values
Properties widget		
Type	Displays the subnet type.	(VLAN or Overlay)
VLAN ID	(For VLAN subnets only) Displays the VLAN ID of the VLAN subnet that is extended.	(VLAN ID number)
VPC	(For Overlay subnets only) Displays the name of the VPC subnet that is extended.	(Name of VPC)
Cluster	(For VLAN subnets only) Displays the cluster that the VLAN subnet belongs to.	(Name of cluster)
IP Address Prefix	Displays the network IP address with prefix, of the VLAN subnet that is extended.	(IP Address with prefix)

Parameter	Description	Values
Virtual Switch	(For VLAN subnets only) Displays the virtual switch on which the VLAN subnet is configured.	(Virtual Switch name such as vs0 or vs1)
IP Address Pools widget		
Pool Range	Displays the range of IP addresses in the pool configured in the subnet that is extended.	(IP address range)
(Interactive Graphic Pie Chart)	Displays a dynamic pie chart that displays the statistic you hover on. Displays the following IP address statistics outside the pie chart, that you can hover on: • Total number of IP addresses available. • Used IP addresses in the subnets • Used IP addresses in the IP address pools • Free IP addresses in the subnets • Free IP addresses in the IP address pools	(IP Address statistics)
Subnet Extension widget		
Subnet Extension (properties) - Common		
Type	Displays the subnet extension type.	(Across Availability Zones or To a Third Party Data Center)
Interface Status	Displays the status of the interface that is used by the subnet extension.	(Connected or Down)
Connection Status	Displays the status of the connection that is created by the subnet extension. Not Available status indicates that Prism Central is unable to ascertain the status.	(Not Available, Connected, or Disconnected)
Local IP Address	Displays the IP address that you entered in the Local IP Address field while creating the subnet extension.	(IP Address)
Local Subnet	Displays the name of the local subnet that the subnet extension uses.	(Name of local subnet)
Subnet Extension (properties) - (Only for Across Availability Zones type)		
Local Availability Zone	(Only for Across Availability Zones type) Displays the name of the local AZ that is hosting the subnet that is extended.	(Name of the local Availability Zone)
Remote Availability Zone	(Only for Across Availability Zones type) Displays the name of the remote AZ that the subnet is extended to.	(Name of the remote Availability Zone)
Remote Subnet	(Only for Across Availability Zones type) Displays the name of the remote subnet that the subnet extension connects to.	(Name of remote subnet)

Parameter	Description	Values
Remote IP Address	(Only for Across Availability Zones type) Displays the IP address that you entered in the Remote IP Address field while creating the subnet extension.	(IP Address)
Subnet Extension (properties) - (Only for To a Third Party Data Center type)		
Local Gateway	(Only for To a Third Party Data Center type) Displays the name of the local gateway used for the subnet extension.	(Name of local gateway)
Remote Gateway	(Only for To a Third Party Data Center type) Displays the name of the remote gateway used for the subnet extension.	(Name of remote gateway)

You can perform the following actions from the **Summary** tab:

- View the detailed information of a subnet extension. For the list of available parameters, see the *Subnet Extension Details - Summary Tab Fields* table above.
- Update an existing subnet extension by clicking **Update**. For more information, see [Updating an Extended Subnet](#) in the *Flow Virtual Networking Guide*.
- Delete an existing subnet extension by clicking **Delete**. For more information, see [Removing an Extended Subnet](#) in the *Flow Virtual Networking Guide*.

Address Table Tab

The **Address Table** tab provides MAC Address information only when the subnet extension uses VTEP service. The tab provides the following information:

- **MAC Address:** This provides the MAC addresses of devices connected to the remote VTEP endpoint in the subnet extension.
- **Remote VTEP Endpoint:** This provides the IP address of the remote VTEP endpoint in the subnet extension.

Throughput Tab

The **Throughput** tab provides a graphical representation of the throughput of the subnet extension.

BGP Sessions Summary View

The **BGP Sessions** page displays a list of BGP sessions created for the clusters managed by Prism Central.

To access the **BGP Sessions** page:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Network & Security > Connectivity** from the **Navigation Bar**.

The **Gateways** page opens displaying the list of network gateways.

Note: For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

3. Click the **BGP Sessions** tab.

The **BGP Sessions** page opens displaying the list of BGP sessions created for the clusters.

The following table describes the fields that appear in the **BGP Sessions** page.

Table 105: BGP Sessions – Field Description

Parameter	Description	Values
Name	Displays the name of the BGP session.	(Name of BGP session)
Serviced VPC	Displays the name of the VPC that the BGP session services.	(Name of VPC)
Local Gateway	Displays the name of the local BGP gateway that the BGP session uses.	(Name of local BGP gateway)
Remote Gateway	Displays the name of the remote BGP gateway that the BGP session uses.	(Name of remote BGP gateway)
Session Status	Displays the status of the eBGP session. - Displays Established if the session is Up. - Displays Active when the network controller is attempting to establish the session.	Established or Active
Route Priority	Displays an integer number that denotes the route priority. When the route priority is assigned dynamically, then the network controller assigns integer numbers (usually between 600 and 800 starting with 700) in descending order with steps of 5. For example, the first session is assigned 700 as route priority and then when you create the second session, the controller assigns it a route priority of 695 and a third session is assigned 690. Greater the number, greater is the route priority. With dynamically assigned priority, the priority is assigned in the order of reducing priority to the order of BGP sessions created. The BGP session created first gets the highest priority 700, the second session get the second highest priority 695 and so on. You can manually assign a route priority as well by assigning any number between 300 and 900.	(Integer Number)

You can perform the following actions for a gateway from the **BGP Sessions** page:

- Click the name of a BGP session to open the details page, which displays the detailed information about the BGP session. For more information, see [BGP Session Details View](#) on page 418.
- Create a BGP session by clicking **Create BGP Session**. For more information, see [Creating a BGP session](#) in the *Flow Virtual Networking Guide*.
- Update or delete an existing BGP session using the **Actions** dropdown menu. The **Actions** dropdown menu appears when one or more BGP sessions are selected. For more information, see [Updating a BGP session](#) or [Deleting a BGP session](#) in the *Flow Virtual Networking Guide*.
- Filter the gateway list based on various parameter values using the **Filters** pane. For more information, see [Filters Pane - BGP Sessions Page](#).

Filters Pane on the BGP Sessions Page

You can filter the information in the **BGP Sessions** page based on the following fields that are available in the **Filters** pane. For information on how to use the Filters pane, see the [Filters Pane](#) information.

Table 106: Filter Pane Field Description - BGP Sessions page

Field	Description	Values
Name	Filters based on the BGP session name. It returns a list of BGP sessions that satisfy the name condition/ string.	(BGP session name string)
Session Status	Filters based on the status of the eBGP session.	(Established or Down)

BGP Session Details View

The BGP Session details page provides detailed information about a BGP session.

The details page has the **Summary**, **Routes**, and **BGP Logs** tabs.

To access the details page of an individual BGP session:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Network & Security > Connectivity** from the **Navigation Bar**.

The **Gateways** page opens displaying the list of network gateways you have created and configured.

Note: For information on the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

3. Click the **BGP Sessions** tab.

The **BGP Sessions** page opens displaying the list of BGP sessions created for the clusters.

4. Click the name of a BGP session to open the details page of the session.

The **Summary** tab opens displaying the detailed information about the BGP session in widgets.

Summary Tab

The **Summary** tab provides detailed information about the BGP session in widgets.

The BGP session **Summary** tab has the following widgets:

Table 107: BGP Session Summary Tab Widgets

Parameter	Description	Values
Properties widget		
Session Status	Displays the overall status of the BGP session.	(Up or Down)
eBGP Status	Displays the eBGP status of the BGP session.	Established or Active
Route Priority	Displays an integer number that denotes the route priority. For more information about Route Priority, see BGP Sessions Summary View on page 416.	(Integer Number)
Local Gateway widget		

Parameter	Description	Values
Local Gateway	Displays the name of the local BGP gateway.	(Name)
eBGP ASN	Displays the Autonomous System Number (ASN) of the local BGP gateway used by the session. It would be an integer number in the 1-65534 range (per 32-bit ASN.1 standard).	(Number)
Note: Make sure that this ASN does not conflict with any of the other on-premises BGP ASNs.		
Remote Gateway widget		
Remote Gateway	Displays the name of the remote BGP gateway.	(Name)
eBGP ASN	Displays the ASN of the remote BGP gateway used by the session. It would be an integer number in the 1-65534 range (per 32-bit ASN.1 standard).	(Number)
Note: Make sure that this ASN does not conflict with any of the other on-premises BGP ASNs.		

You can perform the following actions from the **Summary** tab:

- View the detailed information of a BGP session. For the list of available parameters, see the *BGP Session Details - Summary Tab Fields* table above.
- Update an existing BGP session by clicking **Update**. For more information, see [Updating a BGP session](#) in the *Flow Virtual Networking Guide*.
- Delete an existing BGP session by clicking **Delete**. For more information, see [Deleting a BGP session](#) in the *Flow Virtual Networking Guide*.

Routes Tab

The **Routes** tab provides a list of the routes used by the BGP session with the corresponding **Next Hop** details. It has the following lists:

- **Advertised** (default): The **Routes** tab opens in the **Advertised** list. The **Advertised** list provides a list of the advertised routes with the corresponding **Next Hop** details.
- **Received**: This list provides list of the routes received from remote with the corresponding **Next Hop** details.

BGP Logs Tab

The **BGP Logs** tab provides detailed live logs for the BGP session. This information can be very useful in monitoring and debugging a BGP session.




Figure 139: BGP Session Details View - BGP Logs tab sample for a BGP session

Security Policies

Security policies are defined using Nutanix Flow that provides a policy-driven security framework to inspect traffic within the data center.

For information on how to create and apply security policies on Basic VLAN Subnets, see [Flow Network Security \(formerly Flow Microsegmentation\) Guide](#).

For information on how to create and apply security policies on (advanced) VLAN Subnets and Overlay Subnets, see [Flow Network Security Next-Gen Guide](#).

For information on how to view security policies in Prism Central, see [Security Policies Summary View](#) on page 540 or [Security Policy Details View](#) on page 544.

Security Dashboard

The Security Dashboard provides a dynamic summary of the security posture across all registered clusters.

The Security Dashboard provides a view of the most critical security parameters like cluster-based issue summary, STIG policy compliance, security hardening, and identified vulnerabilities. For more information, see [Security Dashboard](#) in the *Nutanix Security Guide*.

DATA PROTECTION AND RECOVERY ENTITIES

You can access dashboards for the following data protection types from the **Data Protection** entity of the **Infrastructure** application. For information about how to access the entity items available in **Data Protection** entity, see [Application-specific Navigation Bar](#) on page 61.

- [Protection Summary](#) on page 421
- [Protection Policies](#) on page 421
- [Recovery Plans](#) on page 421
- [VM Recovery Points](#) on page 421
- [VG Recovery Points](#) on page 421
- [Consistency Groups](#) on page 421

Protection Summary

Protection Summary provides a detailed information about the Nutanix Disaster Recovery entities in an AZ, and helps you to generate DR reports for the specified time. It also enables you to monitor the health of your DR deployments and the activities performed on the Nutanix Disaster Recovery entities. For more information about **Protection Summary**, see [Nutanix Disaster Recovery Guide](#).

Protection Policies

A Protection policies is a configurable policy that takes recovery points of the protected entities (guest VMs, volume groups, and consistency groups) in equal time intervals, and replicates those recovery points to the recovery AZs. For more information about protection policies, see [Nutanix Disaster Recovery Guide](#).

Recovery Plans

The Recovery plans are used in automated Disaster Recovery (DR) configurations to orchestrate the recovery of the protected entities to different Nutanix clusters at the same or different AZs. For more information about recovery plans, see [Recovery Plans View](#) information in *Nutanix Disaster Recovery Guide*.

VM Recovery Points

The protection policy associated with the VM creates the VM recovery points. For more information about VM recovery points, see [Nutanix Disaster Recovery Guide](#).

VG Recovery Points

VG Recovery Points lists the recovery points of all the protected volume groups (generated over time) on the local cluster. For more information about VG recovery points, see [Nutanix Disaster Recovery Guide](#).

Consistency Groups

A consistency group is a collection of the entities that are treated as a single group and backed up collectively in a recovery point. For information about Consistency Groups, see [Consistency Groups](#) information in *Nutanix Disaster Recovery Guide*.

HARDWARE ENTITIES

You can access dashboards for the following hardware components from the **Hardware** entity of the **Infrastructure** application. For information about how to access the entity items available in **Hardware** entity, see [Application-specific Navigation Bar](#) on page 61.

- Clusters (see [Clusters Summary View](#) on page 424)
- Hosts (see [Hosts Summary View](#) on page 445)
- Disks (see [Disks Summary View](#) on page 464)
- GPUs (see [GPUs Summary View](#) on page 470)

Cluster Management

In Prism Central, the clusters dashboard allows you to view summary information about all the registered clusters and access the detailed information about each cluster. For more information, see [Clusters Summary View](#) on page 424 and [Cluster Details View](#) on page 431.

The **Actions** menu appears when a cluster is selected. The following table provides the information about all the actions that you can perform to manage the clusters:

Note: The available actions appear in bold; other actions are disabled. For disabled options, you can see a tool tip explaining the reason. The available actions depend on the current state of the selected clusters.

Table 108: Cluster Actions

Action	Description	Applicable to Multiple Clusters
Launch Prism Element	Launch Prism element for that cluster in a separate tab or window (depending on your browser settings). Note: When you access a cluster from Prism Central, you log in through your Prism Central user account, not a cluster user account. As a result, the cluster user configuration options are different (more limited) than when logging directly into the cluster. The options that appear in the Prism Element main menu user dropdown list are REST API Explorer , About Nutanix , Support Portal , Help , Nutanix Next Community , and Sign Out .	No
Upgrade Software	Upgrade the AOS version on the cluster. For more information, see Prism Central Deployment on page 12.	Yes
Rack Configuration	Configure the rack awareness feature. This option appears only for clusters that satisfy the conditions for rack awareness. For instructions on configuring rack awareness, see Rack Fault Tolerance in the <i>Prism Element Web Console Guide</i> .	No
Manage Categories	Configure and manage cluster categories.	No

Action	Description	Applicable to Multiple Clusters
Enable Efficiency Measurement	Enable efficiency measurement for the cluster.	Yes
Disable Efficiency Measurement	Disable efficiency measurement for the cluster.	Yes
Enable Anomaly Detection	Enable Anomaly Detection for the cluster.	Yes
Disable Anomaly Detection	Disable Anomaly Detection for the cluster.	Yes
Run Playbook	Run or create a new playbook for the cluster entity type.	Yes
Manage & Backup Keys	Download the encryption keys for the cluster. If the cluster goes down, you can access your data with the backup key.	No
Enable Data-at-Rest Encryption	Enable Data-at-Rest Encryption for the cluster.	Yes
Manage KMS Types	Manage the Key Management Solution (KMS) for the cluster.	No
Enable Data-in-transit Encryption	Enable Data-in-transit Encryption for the cluster.	Yes
Disable Data-in-transit Encryption	Disable Data-in-transit Encryption for the cluster.	Yes
Manage Rebuild Capacity	Manage Rebuild Capacity for a cluster having a minimum of three nodes. For more information about rebuild capacity, see Rebuild Capacity Reservation in Prism Element Web Console Guide.	No
Manage Recycle Bin	Manage Recycle Bin for the cluster. You can choose to retain the deleted VMs and volume groups for a specific retention time (1 to 7 days). This feature is supported for clusters running AOS 6.6.1 and later.	No
Manage Cluster Fault Tolerance	Modify the fault tolerance of the cluster based on your cluster configuration. For more information, see Managing Cluster Fault Tolerance on page 439. Cluster fault tolerance determines the readiness of the cluster to recover from an unexpected hardware failure. For more information, see Cluster Fault Tolerance in the <i>Prism Element Web Console Guide</i>.	No
Maintenance Resiliency	Modify the rebuild preference setting of the cluster. By default, the cluster rebuild preference is set as <i>Smart</i>. To modify the cluster rebuild preference to <i>Immediate</i>, clear the <i>Smart</i> checkbox. For more information on cluster rebuild preference, see Cluster Rebuild Preference in the <i>Prism Element Web Console Guide</i>.	No
Expand Cluster	Expand a cluster. For information on expanding a cluster, see Expanding a Cluster through Prism Central on page 454.	No

Action	Description	Applicable to Multiple Clusters
Destroy Cluster	Destroy a cluster. For information on destroying a cluster, see Destroying a Cluster on page 437.	Yes
Collect Logs	Collect logs for troubleshooting purposes. For more information, see Collecting Logs on page 647.	Yes

Clusters Summary View

The **Summary** tab on the **Clusters** page provides a dashboard of all the registered clusters.

To access the summary view of all the clusters, perform the following steps:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and go to **Hardware** > **Clusters** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

The system displays the **List** tab by default with all the registered clusters in **Nutanix** environment.

3. Click the **Summary** tab. The system displays **Summary** page for all the clusters.

Note: To access the summary view of all the non-nutanix clusters managed by vCenter, select **Non-Nutanix** from the dropdown menu in **Clusters** page.

The **Summary**, **Alerts**, and **Events** tabs display the same information as for Nutanix-managed clusters. The **Metrics** tab displays CPU usage, memory usage, and VM operations. The **List** tab displays fields for name, host count, VM count, storage usage, and storage capacity. The **Profiles** tab allows you to view all the settings profiles created for the clusters.




Figure 140: Summary Page - All Clusters

The **Summary** page for all clusters displays information about all the registered clusters and allows you to access detailed information about each cluster.

The **Clusters** page includes the following tabs on the left: **Summary**, **List**, **Alerts**, **Events**, **Metrics**, and **Profiles** with a display area below the selected tab.

Note: This section describes the information and options that appear in the **Summary** page for all clusters. For instructions on how to view and organize that information in various ways, see [Prism Central GUI Organization](#) on page 49.

The **Summary** tab for all the clusters displays the following widgets:

- **Highlighted Entities:** Displays a list of clusters with the highest usage of the parameter you select from the dropdown menu on the right of the widget. The options are **CPU Usage**, **IO Latency**, **IOPS**, and **Memory Usage**. Click the **View all XX Clusters** link at the bottom to display the **List** tab (following section).
- **Storage Usage:** Displays a list of clusters along with the percentage usage (with respect to total usable capacity of the cluster) and data reduction. Total usage is calculated as the sum of used capacity and free reserved capacity. Data reduction ratio displays the ratio of how much the data size is reduced by enabling compression, deduplication, and erasure coding. Placing the cursor anywhere on the horizontal axis displays a breakdown view of the storage capacity usage.
- **Cluster Resiliency / Fault Tolerance:** Displays a list of clusters along with the fault domain and the current state of the fault tolerance. Fault tolerance of a cluster can have any one of the following states:
 - **OK:** This state indicates that the fault tolerance domain is highly resilient to safely handle a node or a disk (in single or two node clusters) failure.
 - **Warning:** This state indicates that the fault tolerance level is almost reaching to 0. Warning state is displayed if the cluster is not fault tolerant at the configured domain, but is fault tolerant at a lower domain. For example, if you have configured rack as the configured domain and the cluster can no longer handle any rack failures due to some reason but can still handle node (lower domain) failures, then fault tolerance state is displayed as Warning.
 - **Critical:** This state indicates that the fault tolerance level is 0, and the fault tolerance domain cannot handle a node or a disk (in single or two node clusters) failure.
 - **Computing:** This state indicates that the new fault tolerance level is being calculated. This state is displayed soon after a node or disk failure, before rebuild is initiated.
- **Alerts:** Displays a list of cluster-related alerts that occurred during the specified interval. Select either **Last 24 hours** (default), **Last 1 Hours**, or **Last week** from the dropdown menu. When an alert appears, you can click on the graph, which then displays a list of those alerts. Clicking on an alert displays the details page for that alert.
- **Storage Over-provisioning:** Displays a list of the clusters along with the storage over-provisioning ratio (calculated based on the provisioned storage and the available raw storage). The current value and state of storage over-provisioning ratio is displayed according to the configured limits.
- **Rebuild Capacity Reservation:** Displays a list of clusters with their total capacity, rebuild capacity, and resilient capacity. For more information about rebuild capacity, see [Rebuild Capacity Reservation](#) section in Prism Element Web Console Guide.
- **Recycle Bin Management:** Displays a list of clusters with their retention time in days. For more information on Recycle Bin, see [Recycle Bin](#) in the *Prism Element Web Console Guide*.
- **Anomalies:** Displays a graph of memory, I/O, CPU, networking, or disk anomalies that occurred during the specified interval. Select either **Last 24 hours** (default), **Last 1 Hours**, or **Last week** from the dropdown menu. When an anomaly appears, you can click on the graph, which then displays a list of those anomalies. Clicking on an anomaly displays the event page for that anomaly.

Lists Tab

The **List** tab displays the list of clusters.

To access the **Lists** tab:

1. Log in to Prism Central.

2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Hardware > Clusters** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

The system displays the **List** tab by default with clusters.

Note: You can use the **View by** and **Group by** options to create your own customized view and add the necessary columns to that view. For more information on how to **View by** and **Group by** options, see [Prism Central GUI Organization](#) on page 49.

The following table describes the fields that appear in the clusters **List** tab.

Note: The fields vary based on the **View by** and **Group by** options. A dash (-) is displayed in a field when a value is not available or applicable. For more information about **View by** and **Group by** options, see [Prism Central GUI Organization](#) on page 49.

Table 109: Clusters List Fields

Parameter	Description	Values
View by General Fields		
Name	Displays the cluster name. Clicking on the name displays the details page for that cluster. For more information, see Cluster Details View on page 431.	(cluster name)
AOS Version	Displays the version number of AOS running on the cluster.	(version number)
Upgrade Status	Displays the current upgrade status. There are various stages from scheduled to succeeded (or failed).	Pending, Downloading, Queued, PreUpgrade, Upgrading, Succeeded, Failed, Cancelled, Scheduled
Hypervisors	Displays the hypervisor type running in the cluster. In the case of a mixed cluster such as one running ESXi or Hyper-V that also includes NX-6035C nodes running AHV, both hypervisor types are listed.	AHV, ESX, Hyper-V
Host Count	Displays the number of hosts (nodes) in the cluster.	(number of nodes)
VM Count	Displays the total number of VMs in the cluster (in any state).	(number of VMs)
Cluster Runway	Displays the predicted runway (time period) before the cluster requires additional resources. For more information, see the <i>Capacity Tab</i> section in Cluster Details View on page 431.	(number of days)
Inefficient VMs	Displays the number of inefficient VMs in the cluster. For more information, see Behavioral Learning Tools in <i>Intelligent Operations Guide</i> .	(number)
View by Performance Fields		

Parameter	Description	Values
Name	Displays the cluster name.	(cluster name)
CPU Usage	Displays the percentage of CPU capacity in the cluster currently being used.	0 -100%
Memory Usage	Displays the percentage of memory capacity in the cluster currently being used.	0 -100%
IOPS	Displays total (both read and write) I/O operations per second (IOPS) for this cluster.	(number)
IO Bandwidth	Displays total I/O bandwidth used per second in this cluster.	xxx [MBps KBps]
IO Latency	Displays the average I/O latency in this cluster.	xxx [ms]
View by Capacity Fields		
Name	Displays the cluster name.	(cluster name)
Cluster Runway	Displays the predicted runway (time period) before the cluster requires additional resources. For more information, see the <i>Capacity Tab</i> section in Cluster Details View on page 431.	(number of days)
CPU Runway	Displays the predicted CPU runway (time period) before the cluster requires additional resources.	(number of days)
Memory Runway	Displays the predicted memory runway (time period) before the cluster requires additional resources.	(number of days)
Storage Runway	Displays the predicted storage memory runway (time period) before the cluster requires additional resources.	(number of days)
View by Encryption Fields		
Name	Displays the cluster name.	(cluster name)
Data-at-Rest Encryption	Displays the cluster encryption type.	Not Encrypted, Software, Hardware, Software and Hardware
KMS Type	Displays the key management system (KMS) type.	Not Set, Native (Local), Native (Remote), External
Data-in-Transit Encryption	Displays the cluster encryption type.	Not Encrypted, Software, Hardware, Software and Hardware
Host Count	Displays the number of hosts (nodes) in the cluster.	(number of nodes)

You can perform the following actions for the clusters in the **Lists** tab:

- Create a cluster. For information about how to create a cluster, see [Creating a Cluster](#) on page 435.

- Access the detailed information about an individual cluster. For more information, see [Cluster Details View](#) on page 431.
- Filter the clusters list based on available parameter values using **Filters** pane. For more information about **Filters** pane, see the following section *Filters Pane - Clusters Page*.
- Export the table that contains the list of clusters and their information to a file in a CSV format. For more information about **Export** option, see [Export](#) on page 54.
- Group the clusters based on pre-defined criteria. For information about how to group the clusters, see [Group by](#) on page 51.
- View clusters based on pre-defined criteria or create a custom view. For information about available views and how to create a custom view, see [View by](#) on page 51.
- Assign a label to the filtered criteria. For information about how to define a label, see [Label](#) on page 54.
- Perform the cluster-specific actions on a single or multiple clusters using the **Actions** dropdown menu. For more information, see [Cluster Management](#) on page 422.

Filters Pane - Clusters Page

The following table describes the fields available in the **Filters** pane:

Table 110: Filters Pane

Parameter	Description	Values
Labels	Filters on label name. Select one or more labels from the dropdown list. (If there are no labels currently, a message about how to create labels is displayed.)	(label names)
Name	Filters on the cluster name. Select a condition from the dropdown list (Contains , Doesn't contain , Starts with , Ends with , or Equal to) and enter a string in the field. It will return a list of clusters that satisfy the name condition/string.	(cluster name string)
AOS Version	Filters on AOS version. Select one or more versions to return a list of clusters running those version(s). The number of clusters currently running each version is displayed on the right of the line.	(Acropolis version numbers across clusters currently)
Hypervisors	Filters on the hypervisor type. Select one or more hypervisors to return a list of clusters running those hypervisor(s). The number of clusters currently running each hypervisor is displayed on the right of the line.	AHV, ESXi, Hyper-V, XenServer
Health	Filters on the cluster health state (good, warning, or critical). Select one or more states to return a list of clusters in that state(s). The number of clusters currently in each state is displayed on the right of the line.	Critical, Warning, Good

Parameter	Description	Values
CPU Usage	Filters on the amount of total CPU being used. Check the box for the desired range or enter a percentage range in the from <low> to <high> % field. It will return a list of clusters utilizing total CPU in that range (0-100%).	([xx] to [yy]% range)
Memory Usage	Filters on the amount of total memory being used. Check the box for the desired range or enter a percentage range in the from <low> to <high> % field. It will return a list of clusters utilizing total memory in that range (0-100%).	([xx] to [yy]% range)
IOPS	Filters on the total (both read and write) IOPS. Check the box for the desired range or enter a range in the from <low> to <high> iops field. It will return a list of clusters with total IOPS in that range.	([xx] to [yy] range)
IO Bandwidth	Filters on the total I/O bandwidth used. Check the box for the desired range or enter a range in the from <low> to <high> bps field. It will return a list of clusters with total I/O bandwidth usage in that range.	([xx] to [yy] range)
IO Latency	Filters on the average I/O latency. Check the box for the desired range or enter a range in the from <low> to <high> ms field. It will return a list of clusters with average I/O latency in that range.	([xx] to [yy] range)
Upgrade Status	Filters on the current upgrade status. There are various stages from scheduled to succeeded (or failed).	Pending, Downloading, Queued, PreUpgrade, Upgrading, Succeeded, Failed, Cancelled, Scheduled
Categories	Filters on the category type. Search for the category name. For example, ADGroup:\$Default	(category name string)
Data-at-rest Encryption Type	Filters on the encryption type. Check the box for the desired encryption type. It will return a list of clusters with the corresponding encryption type applied.	Not Encrypted, Software, Hardware, Dual
Data-at-rest Encryption Scope	Filters on the encryption scope. Check the box for the entities where encryption is applied on. It will return a list of clusters with the corresponding encryption scope applied.	Cluster, Entity
KMS Type	Filters on the key management system (KMS) type.	Not Set, Native (Local), Native (Remote), External
Data-in-Transit Encryption	Filters on the encryption type. Check the box for the desired encryption type. It will return a list of clusters with the corresponding encryption type applied.	Not Enabled, Enabled

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options that are available on the **Activity > Alerts** page, however it only displays the cluster-related alerts. For more information about alerts, see *Prism Central Alerts and Events Reference Guide*.

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options that are available on the **Activity > Events** page, however it only displays the cluster-related events. For more information about events, see *Prism Central Alerts and Events Reference Guide*.

Metrics Tab

The **Metrics** tab allows you to view usage metrics across the clusters. Click the **Metrics** tab to display dropdown menu of available metrics, and select the metric name to display the relevant performance information.

Note: The **Metrics** dropdown menu is hypervisor-specific, and might vary based on the hypervisors used in the cluster.

Table 111: Metrics Tab Fields

Metric	Description
CPU Usage	Displays a CPU usage table listing current values and total clusters (number). The current values are split into percentile intervals (for example, less than 25%, 25-50, 50-75, more than 75%). Clicking on a percentile interval displays the Summary tab filtered to just those clusters. Note: The same format also applies to the other metrics in this table with either percentile or quantity intervals.
Memory Usage	Displays a memory percentage usage table.
IOPS	Displays total, read, and write IOPS tables.
IO Latency	Displays total, read, and write I/O latency rate tables.
IO Bandwidth	Displays total, read, and write I/O bandwidth rate tables.

Profiles Tab

The **Profiles** tab allows you to view all the settings profiles created for the clusters. Click the **Profiles** tab to display the list of available profiles, and select a profile to display the relevant information about clusters, their compliance status, and Settings.

You can perform the following actions from the **Profiles** tab:

- Create a Settings Profile
- Updating a Settings Profile
- Deleting a Settings Profile
- Assigning clusters to a Setting Profile
- Disassociating Clusters from a Settings Profile

Cluster Details View

Summary Tab

The **Summary** tab of an individual cluster consists of a dashboard that provides the detailed information about the cluster.

To access the **Summary** tab of an individual cluster:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and go to **Hardware > Clusters** from the **Navigation Bar**.

For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61. The system displays the **List** tab by default with all the registered clusters in **Nutanix** environment. For information about how to access the list of clusters in **Non-Nutanix** environment (non-Nutanix clusters managed by external vCenter), see [Clusters Summary View](#) on page 424.

3. Click the target `<cluster_name>`.

Note: Replace `<cluster_name>` with the actual cluster name at your site.

The **Summary** tab of an individual cluster provides the following widgets:

- **Properties:** Displays summary information about the cluster. For information about the fields available in **Properties** widget, see [Cluster Properties Widget](#) on page 433.
- **Metrics:** Displays the cluster metrics such as CPU usage, memory usage, IOPS, IO latency, IO bandwidth, and power consumption Instant Watt.
- **Alert:** Displays a list of related alerts that occurred during the specified interval. You can select **Last week** (default), **Last 1 hour**, or **Last 24 hours** from the dropdown menu on the top right corner of the widget.
- **Anomalies:** Displays a graph of memory, I/O, CPU, networking, or disk anomalies that occurred during the specified interval. You can select **Last week** (default), **Last 1 hour**, or **Last 24 hours** from the dropdown menu on the top right corner of the widget. When an anomaly appears, you can click the graph to display a list of those anomalies. If you click an individual anomaly, the system displays the event page for that anomaly.
- **VM Efficiency:** Displays the number of VMs that are considered inefficient, broken down by category (over-provisioned, inactive, constrained, and bully). For more information, see [Behavioral Learning Tools](#) in *Intelligent Operations Guide*.
- **Tier-wise Usage:** Displays a pie chart divided into the percentage of storage space used by each disk tier in the cluster. Disk tiers can include SSD, HDD, and SSD-NVMe depending on the hardware model type.
- **Storage Over-provisioning:** Displays the storage over-provisioning ratio (calculated based on the provisioned storage and the available raw storage) in the cluster. You can configure a ratio limit to alert the administrator when the ratio changes excessively. If the ratio reaches 70% of the configured limit, the widget displays a warning alert in yellow color. The widget displays Critical alert in red color if the ratio exceeds 90%. Note that the time taken for the Storage Over-provisioning Ratio widget to reflect the changes made in the cluster varies according to the recent storage operations or activities performed.
- **Cluster Resiliency / Fault Tolerance Status:** Displays the current fault tolerance of the cluster. The information includes the type of cluster fault tolerance configured in the cluster, current failure domain (rack, block, disk, or node) and the level of fault tolerance (1 or 2). The widget indicates if any of the current fault domain changes from their configured state. If a data rebuild is triggered after a component failure, the widget also displays a rebuild progress indicator that enables you to track the estimated time until full resiliency has been restored to the cluster. For more information on the types of cluster fault tolerance see, [Cluster Fault Tolerance](#) in the *Prism Element Web Console Guide*.

- **Storage Summary:** Displays information about the physical storage space utilization (in GiB or TiB) and resilient capacity of the cluster. Placing the cursor anywhere on the horizontal axis displays a breakdown view of the storage capacity usage. You can also configure a threshold warning for the resilient capacity utilization in the cluster by clicking the gear icon at the top right corner. For more information, see [Configuring a Warning Threshold for Resilient Capacity](#).

The View Details link displays the resiliency status and storage information of all the individual nodes in the cluster.

The Storage Details page is divided into two sections: The right section displays a diagrammatic representation of the number of nodes present in the cluster along with the respective storage capacity used. You can view the warning threshold and resilient capacity of each of the nodes. The left section provides detailed storage information of the cluster as follows.

Table 112: Storage Information

Parameter	Description	Values
Failure Domain	Displays the entity (node, block, or rack) whose failure the cluster can tolerate while still running the guest VMs and responding to commands through the management console.	node, block, or rack
Total Capacity	Displays the total capacity of all the disks on all the hosts in the cluster.	xxx [GB TB]
Resilient Capacity	Displays the total resilient capacity of the cluster. Resilient capacity is the storage capacity available in the cluster after accounting for the storage space needed to rebuild and restore in case of any component failure.	xxx [GB TB]
Total Usage	Displays the sum of all the storage space used by the cluster. The total used capacity is calculated based on the following: • Used Capacity: The amount of used storage space in the cluster (by user data). • Snapshots: The total storage capacity in the cluster consumed by snapshots (sum of both local and remote). • Recycle Bin: The total storage capacity in the cluster consumed by the VMs deleted by the user. • Other: The total storage capacity occupied by VMs, VG disks, and images. • Free Reserved Capacity: The total capacity that can be used by selected storage containers.	xxx [GB TB]
Available Capacity	Displays the available storage capacity on all the disks on all the hosts in the cluster.	xxx [GB TB]

In addition, some actions also appear on the **Summary** tab. You can perform the cluster-specific actions on a single or multiple clusters using these actions. The first four actions are: **Launch Prism Element, Upgrade Software,**

Rack Configuration, and **Manage Categories**. The rest of the actions appear in the **More** dropdown menu. For more information, see [Cluster Management](#) on page 422.

Cluster Properties Widget

The following table describes the fields in the **Properties** widget. A dash (-) in a field indicates that there is not enough data to evaluate or a value is not assigned. The displayed fields vary by hypervisor.

Table 113: Cluster Properties Fields

Parameter	Description	Values
Health	Displays the cluster health state (good, warning, or critical).	Critical, Warning, Good
Storage Usage	Displays the amount of storage used in the cluster	xxx [GiB TiB]
Storage Capacity	Displays the total amount of storage capacity in this cluster.	xxx [GiB TiB]
Cluster Runway	Displays the predicted runway (time period) before the cluster requires additional resources (see the Capacity Tab section below).	(number of days)
VM Count	Displays the number of VMs in the cluster.	(number of VMs)
AOS Version	Displays the version number of AOS running on the cluster.	(version number)
Host Count	Displays the number of hosts (nodes) in the cluster.	(number of hosts)
Upgrade Status	Displays the status of the last (or current) upgrade attempt.	(status condition)
IP Address	Displays the virtual IP address for the cluster (if defined).	(IP address)
Hypervisors	Displays the hypervisor type running in the cluster. In case of a mixed cluster such as one running ESXi or Hyper-V that also includes NX-6035C nodes running AHV, both hypervisor types are listed.	AHV, ESX, or Hyper-V

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, except it is filtered to display just alerts for this cluster. For more information, see [Alerts Summary View \(Prism Central\)](#) in *Prism Central Alerts and Events Reference Guide*.

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options as the Events dashboard, except it is filtered to display just events for this cluster. For more information, see [Events Summary View \(Prism Central\)](#) in *Prism Central Alerts and Events Reference Guide*.

Metrics Tab

The **Metrics** tab displays usage metrics for the cluster. Click the **Metrics** tab and then the desired metric name (see following table) to display a graph for that metric below the tab. The graph is a rolling time interval performance or usage monitor. The baseline range appears as a blue band in the graph.

Note: The baseline range and identified anomalies are based on sophisticated machine-learning capabilities. For more information, see [Behavioral Learning Tools](#) in the *Intelligent Operations Guide*. The machine-learning algorithm uses 21 days of data to monitor and predict performance. A graph or baseline band might not appear if less than 21 days of data is available.

- Place the cursor anywhere on the horizontal axis to display the value at that time.
- Select the duration (time interval) from the dropdown list on the right (last 1 hour, last 24 hours, last week, last 21 days).
- [I/O-based metrics] Check the appropriate box or boxes to have the graph display total, read, or write usage (or any combination of the three).
- Click **Alert Settings** to configure an alert for this metric. For more information, see [Creating Custom Alert Policies](#) in *Prism Central Alerts and Events Reference Guide*.

The following table describes the available metrics. Some metrics are not available on all hypervisors.

Table 114: Metrics Tab Fields

Metric	Description
CPU Usage	Displays the percentage of CPU capacity currently being used by the cluster (0–100%).
Memory Usage	Displays the percentage of memory capacity currently being used by the cluster (0–100%).
IOPS	Displays separate graphs for total, write, and read I/O operations per second (IOPS) for the cluster.
I/O Latency	Displays separate graphs for total, write, and read average I/O latency (in milliseconds) for physical disk requests by the cluster.
I/O Bandwidth	Displays separate graphs for total, write, and read I/O bandwidth used per second (MBps or KBps) for physical disk requests by the cluster.
Power Usage	Displays the power consumption of the cluster in kWh.

Storage Usage Tab

The **Storage Usage** tab displays the following graphs:

- The **Cluster-wide Usage Summary** graph displays a rolling time interval monitor of total storage usage across the cluster that can vary from one to several hours depending on activity moving from right to left. Placing the cursor anywhere on the horizontal axis displays the value at that time. For more in-depth analysis, you can add the monitor to the analysis page by clicking the blue link in the upper right of the graph.
- The **Tier-wise Usage** graph displays a pie chart divided into the percentage of storage space used by each disk tier (SSD and DAS-SATA) across the cluster.

Hardware and Virtual Entities Tabs

These tabs display a list of hardware and virtual entity types that exist in this cluster. Clicking a hardware entry such as **Hosts** or a virtual entity such as **Containers** displays the information you would see on the **List** tab summary page for that the specified hardware or virtual entity except filtered to just those in this cluster. See [Hardware Entities](#) on page 422 and [Compute Entities](#) on page 99 for more information.

Capacity Tab

The **Capacity** tab displays current and historical usage information and provides resource planning tools. The capacity planning feature requires a Prism Pro license. For more information, see [Capacity Planning](#) in the *Intelligent Operations Guide*.

Security Dashboard Tab

The Security Dashboard provides dynamic summary of the security posture across all registered clusters. The Security Dashboard allows you to view the most critical security parameters like cluster-based issue summary, STIG policy compliance, security hardening, and identified vulnerabilities. For more information, see [Security Dashboard](#) in the *Nutanix Security Guide*.

Entity Relationship Widget

The entity relationship widget shows the relationship between related entities like clusters, hosts, and VMs instances. The widget allows quick access between the related entities. You can directly navigate to a target cluster, host, or VM instance through the respective dropdown menus.

Example: Viewing Cluster Instances

Click the **Cluster** dropdown menu to view the list of cluster instances. Alternatively, you search a cluster name.




Figure 141: Clusters

Note:

- The **Recent** label indicates the last accessed entity instances. The widget displays a maximum of three recently accessed entity instances.
- The filtered list of VMs display only the powered-on VM instances.
- If the VMs are not filtered on a host instance, all VMs on the selected cluster are displayed.

Creating a Cluster

Before you begin

Ensure that you meet the following prerequisites:

- All hosts are running AOS 6.7 or later versions.
- All hosts are reimaged with the same set of AOS and hypervisor versions. For information about reimaging the hosts using Foundation Central, see [Run Foundation Central](#).
- All hosts are reachable.
- All hosts are in the same subnet.

About this task

To create a cluster through Prism Central, perform the following steps:

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Hardware** > **Clusters**.
The system displays the **List** tab with all the registered clusters in your **Nutanix** environment.
3. Click **Create Cluster**.
4. Review the **Create Cluster Prerequisites**, and click **Next**.

5. Enter the following details in the **Create Cluster** wizard:

- a. In the **Select Host** section, enter the CVM IP of the hosts that you want to use for cluster creation, separated by a comma or a space and click **Discover Hosts**.
The list of discovered hosts, along with their details, is displayed at the bottom of the wizard.

- b. In the **Configure** section, enter the following details:

- **Cluster Name:** Name for the cluster.
- **Cluster external IP:** [Optional] External IP address for the cluster.
- **Cluster Fault Tolerance:** [Optional] Fault tolerance for the cluster. The default value is **1N/1D**. Consider the following when you select the fault tolerance for your cluster:
 - A one-node cluster supports 1N/1D mirroring inside single node.
 - A two-node cluster supports 1N/1D mirroring inside each node and 1N/1D across clusters.
 - Clusters with more than three nodes do not support any mirroring inside nodes.
 - You can choose 1N&1D if the cluster has a minimum of three nodes only. You can choose 2N/2D if the cluster has a minimum of five nodes only.

For more information on the types of cluster fault tolerance see, [Cluster Fault Tolerance](#) in the *Prism Element Web Console Guide*.

- **DNS Server(s) IP:** [Optional] IP address of DNS server, separated by a comma or a space if you are entering multiple IP addresses.
- **NTP Server(s) IP:** [Optional] IP address of NTP server, separated by a comma or a space if you are entering multiple IP addresses.
- **Container Name:** [Optional] Name of the container that you want to be associated with the cluster.
- **Cluster Fault Tolerance Domain:** [Optional] Select the fault tolerance domain for the cluster. The default value is **Disk** for a single node cluster, and **Node** for multiple-nodes cluster. You can choose **Disk**, **Block**, or **Rack** as the fault tolerance domain.
- **Enable Backplane Segmentation:** [Optional] Enable backplane segmentation. Backplane segmentation separates the CVM-to-CVM traffic from other kinds of traffic. This is applicable only if you have at least three hosts. If you enable backplane segmentation, you need to enter a few mandatory parameters such as subnet, netmask, and ports. For more information on backplane segmentation, see [Segmented and Unsegmented Networks](#).

- c. Click **Next**.

In the **Review** section, the pre-checks are run to validate that the hosts meet all the requirements mentioned under the Prerequisites. If the pre-checks are successful, the configuration details of the cluster are displayed for review.

- d. Click **Create Cluster**.

Cluster creation operation might take some time. Under **Recent Tasks**, you can verify that a cluster create task was initiated.

What to do next

You can register the newly created cluster with Prism Central. For more information, see [Registering a Cluster with Prism Central](#) on page 56.

Destroying a Cluster

Before you begin

Ensure the following prerequisites:

- The minimum supported versions for destroying a cluster through Prism Central are version 6.8 for Prism Element (with recommended AHV version) and version pc.2024.1 for Prism Central.
- You need to reclaim your license before destroying a cluster.

About this task

Caution: Destroying a cluster is irreversible. Upon successful completion of this operation, the cluster is unregistered, and all the nodes go back to the available node pool. All cluster settings, user VMs, and data policies are permanently deleted, and all Prism Central configurations for categories and policies applied to the cluster are lost. Nutanix recommends performing all the preparatory actions before destroying a cluster.

To destroy a cluster through Prism Central, perform the following steps:

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Hardware** > **Clusters** .
The system displays the **List** tab with all the registered clusters in your **Nutanix** environment.

3. Select the cluster to destroy and click **Actions** > **Destroy Cluster**.

Note: You cannot destroy a cluster that is hosting a Prism Central Instance.

4. Review the **Destroy Cluster Prerequisites** and click **Next**.

The pre-checks are run to validate that the cluster meets all the requirements mentioned under the Prerequisites. The result of the pre-checks and the cluster details are displayed for review.

Destroy Cluster

Precheck
Prechecks validate ongoing cluster operations. You cannot proceed with cluster destruction if precheck fails.
✓ Precheck successful! 100%

1 Selected Cluster

Name	Precheck Status	Additional Details
cluster_test	✓ Succeeded	Link

Destroy Cluster?
Are you sure you want to destroy the clusters for which precheck was successful? This action cannot be undone. To proceed, type DESTROY in the input field below.

DESTROY

Back Cancel Destroy Cluster

Figure 142: Destroy Cluster

5. Under **Destroy Cluster?**, type DESTROY to confirm the destroy cluster operation.
Note that this is a case-sensitive input, and you need to enter the word in upper-case.

6. Click **Destroy Cluster**.

Destroy Cluster operation might take some time. Under **Recent Tasks**, you can verify that a cluster destroy task has been initiated.

Managing Cluster Fault Tolerance

You can modify the fault tolerance of the cluster based on your cluster configuration.

About this task

- You can modify the fault tolerance of a cluster from 1N/1D to 2N/2D. Before you do this, ensure that the cluster has a minimum of five nodes.
- You cannot modify the fault tolerance of a cluster from 1N/1D to 1N&1D. The 1N&1D cluster fault tolerance must be configured when you create the cluster.
- You cannot modify the fault tolerance of a cluster with 1N&1D fault tolerance.
- You cannot reduce the cluster fault tolerance. For example, you cannot reduce the cluster fault tolerance from 2N/2D to 1N/1D. If you attempt to reduce cluster fault tolerance, Prism Central displays an error.
- Increasing the cluster fault tolerance might require at least 30 percent of your disk space.
- For more information on the types of cluster fault tolerance, see [Cluster Fault Tolerance](#) in the *Prism Element Web Console Guide*.

To modify the fault tolerance of a cluster through Prism Central, perform the following steps:

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#), and go to **Hardware** > **Clusters**.
The system displays the **List** tab with all the registered clusters in your **Nutanix** environment.
3. Select the checkbox associated with the cluster and click **Actions** > **Manage Cluster Fault tolerance**.
The system displays the **Manage Cluster Fault Tolerance** dialog box.
4. From the **Change Cluster Fault Tolerance** dropdown menu, select **2N/2D**.
The system increases the fault tolerance to 2N/2D and displays **Configured CFT : 2N/2D** in the **Cluster Resiliency / Fault Tolerance Status** widget in the **Summary** tab on the **Clusters** page of Prism Central.

Settings Profile for Clusters

Prism Central version pc.2024.3 onwards, you can create a cluster profile with system settings that you can apply to multiple clusters with a single click. When you apply a profile to a cluster, the settings in the profile are applied to the corresponding settings on the cluster.

The following system settings can be configured through the settings profile:

- NTP server
- Name server
- SNMP
- SMTP server
- Syslog server

- File system allowlists
- Pulse

For more information on cluster system settings, see [System Management](#) in the *Prism Element Web Console Guide*.

Viewing Cluster Profiles

Using the Prism Central web console, you can view the list of settings profile for clusters.

About this task

Summary tab on the Profiles page provides a dashboard of all the settings profiles created for the clusters.

To access the summary view of all the clusters, follow these steps:

Procedure

1. Log in to the Prism Central web console.
2. From the [Application Switcher](#), select the **Infrastructure** application, and navigate to **Hardware > Clusters**. The system displays the **List** tab with all the registered clusters in your Nutanix environment.
3. Select the **Profiles** tab.
The system displays a list of all the profiles along with the number of assigned clusters, drifting clusters (clusters where local settings override the profile settings), number of system settings in the profile, and the date and time when the profile was last modified.
4. Click the name of a profile to view the details.
The system displays the profile details under three tabs: **Summary**, **Compliance**, and **Settings**. The **Summary** tab displays basic details about the Profile. The **Compliance** tab displays details such as compliance status, number of drifting cluster settings, actions to be taken for the drifting settings. The **Settings** tab displays all the settings included in the profile, listed under respective categories.

Creating a Cluster Settings Profile

Using the Prism Central web console, you can create a settings profile for clusters.

Before you begin

Ensure that all hosts are running AOS 7.0 or later versions.

About this task

To create a settings profile for clusters, follow these steps:

Procedure

1. Log in to the Prism Central web console.
2. From the [Application Switcher](#), select the **Infrastructure** application, and navigate to **Hardware > Clusters**. The system displays the **List** tab with all the registered clusters in your Nutanix environment.
3. Click the **Profiles** tab.
4. Select **Create Profile**.

5. In the **Create Settings Profile** window, enter the following details under the Setup Profile section:

- **Profile Name:** Enter a name for the settings profile. Alphabets, numbers, dots, hyphens, and underscores are allowed.
- **Profile Description:** Enter a description for the settings profile.
- **Configuration Type:** Do one of the following:
 - **Configure new settings values:** To configure new settings for the profile, follow these steps:
 1. Click **Add Settings** to add the settings and configure values for each setting added to the profile.
 2. In the **Add Settings** window, click the blue plus (+) icon to select the settings or a group of settings to configure as part of this profile from **Available Settings** on the left. The selected settings or group of settings move to the right side under **Settings in Profile**.




Figure 143: Create Profile - Configure New Settings

3. Click **Save**. Under the **Configure Settings** section, you can see all the selected settings listed along with their setting type, configuration status, and applicable actions.
- **Import Settings Values from Existing Cluster:** To import settings, follow these steps:
 1. Select a cluster from which the settings are to be imported.
 2. Click **Next**.

Under the **Configure Settings** section, all the settings in the selected cluster are replicated to the new profile.
 3. The configuration for SNMP and SMTP Server requires authentication, you must enter the authentication information to complete the configuration.
6. [Optional] Click **Configure** to configure a setting, and **Remove** to remove a setting from the profile.
Pulse is configured by default, you can either edit the setting or remove it.
7. [Optional] Select **Allow Override** to allow a cluster admin to change the settings for an individual cluster through Prism Element web console.
In this case, the changes made by the cluster admin are not considered as drifted settings.
8. Click **Create Profile**.
The system displays the newly created profile under the **Profiles** tab.

Updating a Settings Profile

Using the Prism Central web console, you can update a settings profile for clusters.

About this task

To update a settings profile for clusters, follow these steps:

Procedure

1. Log in to the Prism Central web console.
2. From the [Application Switcher](#), select the **Infrastructure** application, and navigate to **Hardware > Clusters**.
The system displays the **List** tab with all the registered clusters in your Nutanix environment.
3. Select the **Profiles** tab.
The system displays a list of all the profiles.
4. From the **Actions** drop-down menu, select the settings profile, then select **Update**.
5. In the **Update Settings Profile** window, modify one or more of the following:
 - Profile name and description
 - Setting configurations
 - List of clusters assigned to the profile
6. Click **Save**.

Deleting a Settings Profile

Using the Prism Central web console, you can delete a settings profile for clusters.

Before you begin

You must disassociate all the clusters from a profile before you delete the profile.

About this task

To delete a settings profile for clusters, follow these steps:

Procedure

1. Log in to the Prism Central web console.
2. From the [Application Switcher](#), select the **Infrastructure** application, and navigate to **Hardware > Clusters**. The system displays the **List** tab with all the registered clusters in your Nutanix environment.
3. Select the **Profiles** tab. The system displays a list of all the profiles.
4. From the **Actions** drop-down menu, select the settings profile, then select **Delete**. The **Delete** action is enabled only for the profiles that do not have any associated clusters.

Assigning Clusters to a Settings Profile

Using the Prism Central web console, you can assign a settings profile to one or more clusters. A cluster can be assigned only one profile at a time.

About this task

To assign a settings profile to one or more clusters, follow these steps:

Procedure

1. Log in to the Prism Central web console.
2. From the [Application Switcher](#), select the **Infrastructure** application, and navigate to **Hardware > Clusters**. The system displays the **List** tab with all the registered clusters in your Nutanix environment.
3. Select the **Profiles** tab. The system displays a list of all the profiles.
4. From the **Actions** drop-down menu, select the settings profile, then select **Assign Clusters**.
5. In the **Assign Settings Profile** page, click **Select Clusters**.
6. In the **Select Clusters** page, select the clusters where the profiles setting is to be applied.
7. Click **Save**.
A pre-check script verifies if the version of all the selected clusters are compatible with the settings included in the settings profile. After a successful verification, the clusters are assigned to the profile.

Disassociating Clusters from a Settings Profile

Before you can assign a profile to a cluster that already has another profile assigned to it, you must first disassociate the cluster from the earlier profile, then you can assign the new profile to the cluster.

About this task

To disassociate a cluster from a settings profile, follow these steps:

Procedure

1. Log in to the Prism Central web console.
2. From the [Application Switcher](#), select the **Infrastructure** application, and navigate to **Hardware > Clusters**. The system displays the **List** tab with all the registered clusters in your Nutanix environment.

3. Select the **Profiles** tab.
The system displays a list of all the profiles.
4. Click a profile.
5. In the **Compliance** tab, select a cluster to disassociate and click **Disassociate Clusters**.
6. When a prompt appears, click **Disassociate Clusters** again.

After the clusters are disassociated, the Compliance Status in Prism Central for disassociated clusters does not appear. The settings for the attached clusters remain unchanged, and you can modify the individual settings directly from the Prism Element.

Reapplying Profile to a Cluster

If a settings profile is assigned to a cluster and the same settings are later modified on cluster level, the modified settings are displayed as drifted, and the cluster is displayed as non-compliant. Using the Prism Central web console, you can overwrite these drifted settings by reapplying the profile.

About this task

To reapply a settings profile on a cluster, follow these steps:

Procedure

1. Log in to the Prism Central web console.
2. From the [Application Switcher](#), select the **Infrastructure** application, and navigate to **Hardware > Clusters**.
The system displays the **List** tab with all the registered clusters in your Nutanix environment.
3. Select the **Profiles** tab.
The system displays a list of all the profiles.
4. Select a settings profile.
5. From the list of associated clusters, select the clusters to reapply the profile and click **Reapply**.
After you reapply a profile to a cluster, the drifted settings are overwritten, and the cluster becomes compliant with the settings profile.

Host Management

In Prism Central, the hosts dashboard allows you to view summary information about hosts across registered clusters and access detailed information about each host. For more information, see [Hosts Summary View](#) on page 445 and [Host Details View](#) on page 449.

You can perform the following actions to manage hosts in Prism Central:

- Assign a category to the host. For more information, see [Assigning a Category to an Entity](#) in the *Prism Central Admin Center Guide*.
- Run playbook to perform operations on host. For more information, see [Task Automation - Playbooks](#) in *Intelligent Operations Guide*.
- Rename an AHV host. For more information, see [Renaming an AHV Host](#) on page 454.
- Expand a cluster. For more information, see [Expanding a Cluster through Prism Central](#) on page 454.
- Remove a node from the cluster. For more information, see [Removing a Node through Prism Central](#) on page 461.

- Put the host in maintenance mode. For more information, see [Putting a Host into Maintenance Mode Using Prism Central](#) on page 462.
- Remove the host from maintenance mode. For more information, see [Exiting a Host from Maintenance Mode Using Prism Central](#) on page 463.

Hosts Summary View

The **Summary** tab on the **Hosts** page provides a dashboard of the hosts across all the registered clusters.

To access the summary view of all the hosts, perform the following steps:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Hardware > Hosts** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

The system displays the **List** tab by default with all the hosts across registered clusters in **Nutanix** environment. For information about how to access the list of hosts in **Non-Nutanix** environment (non-nutanix hosts managed by external vCenter), see [Host Details View](#) on page 449.

3. Click the **Summary** tab. The system displays **Summary** page for all the hosts.

Note: To access the summary view of all the hosts managed by vCenter, select **Non-Nutanix** environment from the dropdown menu in **Hosts** page.

The **Summary**, **Alerts**, and **Events** tabs display the same information as for Nutanix-managed clusters. The **Metrics** tab displays a subset of the full list (5 of the 7 metrics). The **List** tab displays fields for name, host IP, VM count, memory capacity, and cluster name.

Using playbooks, you can manage the hosts in **Non-Nutanix** environment. For more information about playbooks, see [Task Automation - Playbooks](#) in *Intelligent Operations Guide*.

Clicking the **Summary** tab displays the following three widgets:

- **Suggested:** Displays a list of the hosts with the highest usage of the parameter you select from the dropdown menu on the right of the widget. The options are **CPU Usage**, **Memory Usage**, **IO Latency**, and **IOPS**. Click the **View all XX Hosts** link at the bottom to display the **List** tab (following section).
- **Alert:** Displays a list of host-related alerts that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the dropdown menu.
- **Anomalies:** Displays a graph of memory, I/O, CPU, networking, or disk anomalies that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the dropdown menu. When an anomaly appears, you can click on the graph, which then displays a list of those anomalies. Clicking on an anomaly displays the event page for that anomaly.

List Tab

Clicking the **List** tab, which appears by default when you first open the page, displays a list of the hosts across the registered clusters. The following table describes the fields that appear in the hosts list. The fields vary based on the **Focus** menu selection, which is either **General** or **Performance**. A dash (-) is displayed in a field when a value is not available or applicable.

Table 115: Hosts List Fields

Parameter	Description	Values
General Focus Fields		
Name	Displays the name of the host. Clicking on the name displays the details page for that host. For more information, see Host Details View on page 449.	(host name)
Host IP	Displays the IP address assigned to the hypervisor running on the host.	(IP address)
CVM IP	Displays the IP address assigned to the Controller VM.	(IP address)
Hypervisor	Displays the hypervisor type running on the host.	AHV, ESX, or Hyper-V
Memory Capacity	Displays the memory capacity of the host.	xxx [MB GB]
Cluster	Displays the name of the cluster in which the host resides.	(cluster name)
Performance Focus Fields		
Name	Displays the name of the host.	(host name)
CPU Usage	Displays the percentage of CPU capacity currently being used by this host.	0 - 100%
Memory Usage	Displays the percentage of memory capacity currently being used by this host.	0 -100%
IOPS	Displays I/O operations per second (IOPS) for this host.	[0 - unlimited]
Disk IO Bandwidth	Displays I/O bandwidth used per second for this host.	xxx [MBps KBps]
IO Latency	Displays the average I/O latency (in milliseconds) for this host.	xxx [ms]
Cluster	Displays the name of the cluster in which the host resides.	(cluster name)

You can filter the hosts list based on a variety of parameter values. The following table describes the filter options available when you open the Hosts view **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters. Some parameter filters require additional context such as a constraint string or a range.

Table 116: Filters Pane - Hosts Page

Parameter	Description	Values
Name	Filters on the host name. Select a condition from the dropdown menu (Contains , Doesn't contain , Starts with , Ends with , or Equal to) and enter a string in the field. It will return a list of hosts that satisfy the name condition/string.	(host name string)

Parameter	Description	Values
Cluster	Filters on the cluster name. Enter a string in the field. It will return a list of hosts that reside in the clusters which satisfy the name string.	(cluster name string)
Health	Filters on the host health state (good, warning, or critical). Select one or more states to return a list of hosts in that state(s). The number of hosts currently in each state is displayed on the right of the line.	Critical, Warning, Good
Categories	Filters on category names. Enter a category name in the field and then check the box. As you type a dropdown menu appear to help you select the correct category. A new field appears where you can add more categories to the filter. The number of hosts tagged to each selected category is displayed on the right of the line.	(category name)
Hypervisor	Filters on the hypervisor type. Select one or more hypervisors to return a list of clusters running those hypervisor(s). The number of clusters currently running each hypervisor is displayed on the right of the line.	AHV, ESXi, Hyper-V
Memory Capacity	Filters on the host memory capacity. Check the box for the desired range or enter an amount range in the <i>from <low> to <high> GiB</i> field. It will return a list of hosts with memory capacity in that range.	([xx] to [yy] GiB range)
CPU Usage	Filters on the amount of CPU being used. Check the box for the desired range or enter a percentage range in the <i>from <low> to <high> %</i> field. It will return a list of hosts utilizing CPU in that range (0-100%).	([xx] to [yy]% range)
Memory Usage	Filters on the amount of total memory being used. Check the box for the desired range or enter a percentage range in the <i>from <low> to <high> %</i> field. It will return a list of clusters utilizing total memory in that range (0-100%).	([xx] to [yy]% range)
GPUs	Filters for GPU configuration information such as model name. Select a condition from the dropdown menu (Contains , Doesn't contain , Starts with , Ends with , or Equal to) and enter a string in the field. As you type a dropdown menu appears to help you select the correct configuration information. It will return a list of hosts that satisfy the GPU condition/string.	(configuration info)
IOPS	Filters on the IOPS. Check the box for the desired range or enter a range in the <i>from <low> to <high> iops</i> field. It will return a list of hosts with IOPS in that range.	([xx] to [yy] range)

Parameter	Description	Values
IO Bandwidth	Filters on the I/O bandwidth used. Check the box for the desired range or enter a range in the <i>from <low> to <high> bps</i> field. It will return a list of hosts with I/O bandwidth usage in that range.	([xx] to [yy] range)
IO Latency	Filters on the average I/O latency. Check the box for the desired range or enter a range in the <i>from <low> to <high> ms</i> field. It will return a list of hosts with average I/O latency in that range.	([xx] to [yy] range)

You can perform the following actions for the hosts in the **List** tab:

- Access the detailed information about an individual host. For more information, see [Host Details View](#) on page 449.
- Filter the host list based on available parameter values using **Filters** pane. For more information about **Filters** pane, see [Filters Pane - Hosts Page](#).
- Export the table that contains the list of hosts and their information to a file in a CSV format. For more information about **Export** option, see [Export](#) on page 54.
- Group the hosts based on pre-defined criteria. For information about how to group the clusters, see [Group by](#) on page 51.
- View hosts based on pre-defined criteria or create a custom view. For information about available views and how to create a custom view, see [View by](#) on page 51.
- Perform the host-specific actions on a single or multiple hosts using the **Actions** dropdown menu. For more information, see [Host Management](#) on page 444 .

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, except it is filtered to display just host-related alerts across the registered clusters. For more information, see [Alerts Summary View \(Prism Central\)](#) in *Prism Central Alerts and Events Reference Guide*.

Audits Tab

The **Audits** tab displays a table of audits. This tab provides the same features and options as the Audits dashboard, except it is filtered to display just host-related events across the registered clusters. For more information, see [Audits Summary View](#) on page 493.

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options as the Events dashboard, except it is filtered to display just host-related events across the registered clusters. For more information, see [Events Summary View \(Prism Central\)](#) in *Prism Central Alerts and Events Reference Guide*.

Metrics Tab

The **Metrics** tab allows you to view usage metrics across the hosts. Clicking the **Metrics** tab displays a list of available metrics; click the metric name to display the relevant information to the right. The following table describes the available metrics. (Some metrics are not available on all hypervisors.)

Table 117: Metrics Tab Fields

Metric	Description
CPU Usage	Displays a CPU usage table listing current values and total hosts (number). The current values are split into percentile intervals (for example, less than 25%, 25-50, 50-75, more than 75%). Clicking on a percentile interval displays the Summary tab filtered to just those hosts.
Note: The same format also applies to the other metrics in this table with either percentile or quantity intervals.	
Memory Swap	Displays memory swap-out and swap-in rate tables.
Memory Usage	Displays a memory percentage usage table.
IOPS	Displays total, read, and write IOPS tables.
IO Latency	Displays total, read, and write I/O latency rate tables.
IO Bandwidth	Displays total, read, and write I/O bandwidth rate tables.
Network Packets Dropped	Displays the number of network packets that have been received and then dropped.

Host Details View

Summary Tab

The **Summary** tab of an individual host consists of a dashboard that provides the detailed information about the host.

To access the **Summary** tab of an individual host:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Hardware > Hosts** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

By default, the system displays the **List** tab with all the hosts.

3. Click the target `<host_name>` to view the **Summary** tab of an individual host.

Note: Replace `<host_name>` with the actual host name at your site.

The **Summary** tab of an individual host provides the following widgets and actions:

- **Properties:** Displays summary information about the host. For more information, see the [Host Properties](#) on page 450 table.
- **Alert:** Displays a list of related alerts that occurred during the specified interval. Select either **Last week** (default), **Last 1 hour**, or **Last 24 hours** from the dropdown menu on the top right corner of the widget.
- **Metrics:** Displays the metrics information: **CPU Usage**, **Memory Usage**, **IOPS**, **IO Latency**, **IO Bandwidth**, and **Power Consumption Instant Watt**.
- **Anomalies:** Displays a graph of memory, I/O, CPU, networking, disk, or other anomalies that occurred during the specified interval. Select either **Last week** (default), **Last 1 hour**, or **Last 24 hours** from the dropdown menu on the top right corner of the widget. When an anomaly appears, you can click the graph to display a list of those anomalies. If you click an individual anomaly, the system displays the event page for that anomaly.

- **Rename:** Renames an AHV host. For more information, see [Renaming an AHV Host](#) on page 454.
- **Manage Categories:** Manages categories for the host. For more information, see [Category Management](#) in the *Prism Central Admin Center Guide*.
- **Run Playbook:** Runs a playbook. For more information, see [Running a Playbook \(Manual Trigger\)](#) in *Intelligent Operations Guide*.

The host name and the following set of tabs appear on the left: **Summary, Alerts, Events, Metrics, Usage, Hardware, and Virtual Entities**. Click a tab to display that information on the right. Click the **Back to Hosts** link to return to the summary view.

Host Properties

The following table describes the fields in the **Properties** widget. A dash (-) in a field indicates there is not enough data to evaluate or a value is not assigned. The displayed fields vary by hypervisor.

Table 118: Host Properties Fields

Parameter	Description	Values
Memory Capacity	Displays the total memory capacity for this host.	xxx [MB GB]
Disk Capacity	Displays the total amount of disk capacity on this host.	xxx [GB TB]
Cluster	Displays the name of the cluster in which the host resides. Clicking the name displays the details page for that cluster . For more information, see Cluster Details View on page 431.	(cluster name)
Host IP	Displays the host IP address.	(IP address)
Hypervisor	Displays the hypervisor name.	(hypervisor name)
VM Count	Displays the number of VMs running on this host.	(number)
Block Model	Displays the block model number.	(model series number)
Serial Number	Displays the block serial number.	(block serial number)
CPU Capacity	Displays the total CPU capacity for this host.	xxx [GHz]
CVM IP	Displays the IP address assigned to the Controller VM.	(IP address)
IPMI Address	Displays the IP address of the Intelligent Platform Management Interface (IPMI) port. An IPMI port is used for the hypervisor host console.	(IP address)
Node Serial	Displays the node serial number. The node serial is a unique number passed through from the manufacturer. (The form can vary because it is determined by each manufacturer.)	(manufacturer serial number)
Oplog Disk %	Displays the percentage of the operations log (oplog) capacity currently being used. The oplog resides on the metadata disk.	[0 - 100%]

Parameter	Description	Values
Oplog Disk Size	Displays the current size of the operations log. (The Oplog maintains a record of write requests in the cluster.) A portion of the metadata disk is reserved for the Oplog, and you can change the size through the nCLI.	xxx [GB]
Monitor Enabled	Displays whether the host is high availability (HA) protected. A Yes value means HA is active for this host. A No value means VMs on this host are not protected (will not be restarted on another host) if the host fails. Normally, this value should always be Yes . A No value is likely a sign of a problem situation that should be investigated.	[Yes No]
Disks	Displays the number of disks in each storage tier in the host. Tier types vary depending on the Nutanix model type.	DAS-SATA: (number), SSD-SATA: (number), SSD-PCIe: (number)
GPUs	Displays the number and type of GPUs in the host. For example, if the host contains four Tesla M10 GPUs, this field displays "Tesla M10 (4)".	(GPU type and number)
Datastore(s)	Displays the names of any datastores.	(names)

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, except it is filtered to display just alerts for this host. For more information, see [Alerts Summary View \(Prism Central\)](#) in *Prism Central Alerts and Events Reference Guide*.

Audits Tab

The **Audits** tab displays a table of audits. This tab provides the same features and options as the Audits dashboard, except it is filtered to display just host-related events across the registered clusters. For more information, see [Audits Summary View](#) on page 493.

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options as the Events dashboard, except it is filtered to display just events for this host. For more information, see [Events Summary View \(Prism Central\)](#) in *Prism Central Alerts and Events Reference Guide*.

Metrics Tab

The **Metrics** tab allows you to view usage metrics for the host. Click the **Metrics** tab and then the desired metric name (see following table) to display a graph for that metric on the right. The graph is a rolling time interval performance or usage monitor. The baseline range appears as a blue band in the graph.

Note: The baseline range and identified anomalies are based on sophisticated machine-learning capabilities. For more information, see [Behavioral Learning Tools](#) in *Intelligent Operations Guide*. The machine-learning algorithm uses 21 days of data to monitor and predict performance. A graph or baseline band may not appear if less than 21 days of data is available.

- Place the cursor anywhere on the horizontal axis to display the value at that time.
- Select the duration (time interval) from the dropdown list on the right (last 1 hour, last 24 hours, last week, last 21 days).

- [I/O-based metrics] Check the appropriate box(es) to have the graph display total, read, or write usage (or any combination of the three).
- Click **Alert Settings** to configure an alert for this metric. For more information, see [Creating Custom Alert Policies](#) in *Prism Central Alerts and Events Reference Guide*.




Figure 144: Host Metrics Tab

The following table describes the available metrics. Some metrics are not available on all hypervisors.

Table 119: Metrics Tab Fields

Metric	Description
CPU Usage	Displays the percentage of CPU capacity currently being used by the host (0 - 100%).
Memory Usage	Displays the percentage of memory capacity currently being used by the host (0 - 100%).
Controller IOPS	Displays graph for total I/O operations per second (IOPS) for the host.
Controller Read IOPS	Displays graph for read I/O operations per second (IOPS) for the host.
Controller Write IOPS	Displays graph for write I/O operations per second (IOPS) for the host.
Controller AVG I/O Latency	Displays graph for average I/O latency (in milliseconds) for physical disk requests by the host.
Controller AVG Read I/O Latency	Displays graph for average read I/O latency (in milliseconds) for physical disk requests by the host.
Controller AVG Write I/O Latency	Displays graph for average write I/O latency (in milliseconds) for physical disk requests by the host.

Metric	Description
Controller I/O Bandwidth	Displays graph for I/O bandwidth used per second (MBps or KBps) for physical disk requests by the host.
Controller I/O Read Bandwidth	Displays graph for read I/O bandwidth used per second (MBps or KBps) for physical disk requests by the host.
Controller I/O Write Bandwidth	Displays graph for write I/O bandwidth used per second (MBps or KBps) for physical disk requests by the host.
Power Usage	Displays the power consumption of the host in kWh.

Storage Usage Tab

The **Storage Usage** tab displays the following graphs:

- The **Usage Summary** graph displays a rolling time interval monitor of host storage usage that can vary from one to several hours depending on activity moving from right to left. Placing the cursor anywhere on the horizontal axis displays the value at that time. For more in depth analysis, you can add the monitor to the analysis page by clicking the blue link in the upper right of the graph.
- The **Tier-wise Usage** graph displays a pie chart divided into the percentage of host storage space used by each disk tier (SSD and DAS-SATA).




Figure 145: Host Usage Tab

Hardware and Virtual Entities Tabs

Clicking these tabs displays a list of hardware and virtual entity types that exist on this host. Clicking a hardware entry such as **Disks** or a virtual entity such as **VMs** displays the information you would see on the **List** tab summary page for that the specified hardware or virtual entity except filtered to just those on this host. For more information, see [Hardware Entities](#) on page 422 and [Compute Entities](#) on page 99 topics.

Entity Relationship Widget

The entity relationship widget shows the relationship between related entities like clusters, hosts, and VMs instances. The widget provides quick access between the related entities. You can directly navigate to a target cluster, host, or VM instance through the respective dropdown menus.

Example: Viewing Hosts on a Cluster

Click the **Host** dropdown menu to view the list of host instances on the selected cluster. Alternatively, you can search the host instance name belonging to the target cluster.

Note:

- The **Recent** label indicates the last accessed entity instances. The widget displays a maximum of three recently accessed entity instances.

- The filtered list of VMs display only the powered-on VM instances.
- If the VMs are not filtered on a host instance, all VMs on the selected cluster are displayed.

Renaming an AHV Host

You can rename an AHV host from Prism Central.

About this task

The minimum supported versions to rename an AHV host from Prism Central are AOS 6.6 for the registered Prism Element (with recommended AHV version) and 2022.9 for Prism Central. For information on the recommended AHV version based on the AOS release and hardware model, see [Compatibility and Interoperability Matrix](#).

To rename an AHV host from Prism Central, perform the following steps:

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Hardware** > **Hosts** .
The system displays the **List** tab by default with all the hosts across registered clusters.
3. Select the target host checkbox, and choose **Rename** from the **Actions** dropdown menu.
The system displays the **Rename Host** window.
4. Enter a new name for the host in the **New Host Name** field, and click **Save**.

You must adhere to the following rules for the host name:

- Allowed characters are: uppercase letters (A-Z), lowercase letters (a-z), decimal digits (0-9), dots (.), and hyphens (-).
- The host name must start and end with a number or letter.
- A minimum of one character and a maximum of 63 characters are allowed.
- Successive dots are not allowed. Each dot-separated string must follow the first two rules.

Note: Once the AHV host rename task is complete, it might take a few minutes for the changes to propagate from the host to Prism Central.

Expanding a Cluster through Prism Central

A cluster is a collection of nodes. You can expand a cluster by adding new nodes to the cluster after connecting them to the network on the same subnet as the cluster. The cluster expansion process compares the AOS version on the existing and new nodes and performs any upgrades necessary for all the nodes to run the same AOS version.

Before you begin

Ensure that the following prerequisites are met before you add a node to the cluster:

- Review the relevant sections in [Prerequisites for adding a node](#) before attempting to add a node to the cluster. The process for adding a node varies depending on several factors. This section covers specific considerations based on your AOS, hypervisor, encryption, and hardware configuration.
- Check the Health of the cluster. If any health checks are failing, resolve them before adding any nodes. As a final check, run NCC to ensure that the cluster is healthy.

- Allow any current expand cluster operations to complete.
- All nodes are in the correct metadata state by checking the Hardware dashboard. If any nodes show Metadata store disabled on the node or Node is removed from metadata store, enable the metadata store by clicking **Enable Metadata Store**.
- Follow the SSDs' requirements for Hybrid HCI Node and All-Flash HCI Node specified in [HCI Node Field Requirements](#) in the *Acropolis Advanced Administration Guide*.

About this task

You can add new nodes to a cluster using Prism Central or Prism Element web console. For information about expanding a cluster through Prism Element web console, see [Expanding a Cluster](#) in the *Prism Element Web Console Guide*.

The minimum supported versions to add new hosts using Prism Central web console is Prism Central version 2023.3. From Prism Central version pc.2024.3 onwards, you can expand a cluster without registering the same with Prism Central.

To add one or more nodes to an existing cluster through Prism Central, do the following:

Procedure

1. Log in to Prism Central.
2. From the [Application Switcher Function](#) on page 40, select the **Infrastructure** application, and navigate through **Navigation Bar** to **Hardware** > **Clusters**.
For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.
The system displays the **List** tab with all the registered clusters in your **Nutanix** environment.
3. Select the cluster that you want to expand, and click **Expand Cluster**.
The system searches for hosts in the entire network and all the discovered hosts are listed in the **Expand Cluster** page.
4. (Optional) Click **Discover Hosts Manually** to discover any hosts that are not listed after automatic host discovery, and enter the hostname or CVM IP.

Note: ESXi or AHV compute-only nodes and hosts with deactivated IPv6 support are not discovered automatically. You must initiate manual host discovery for such hosts.

5. In the **Select Host** page, select the hosts that you want to add to the cluster, and click **Next**.
6. In the **Choose Host Type** page, select **Host Type** from the dropdown menu for each of the hosts you have selected, and then click **Next**.
You can select any one of the following host types:
 - **HCI:** Standard host in the cluster.
 - **Storage:** Host optimized for storage usage. This type of host allows seamless expansion of storage capacity of your cluster.
 - **Compute:** Host optimized for computing purposes (CPU and memory).
7. In the **Configure Host** page, enter or update the host details such as hostname, CVM IPv4, CVM IPv6, hypervisor IPv4, hypervisor IPv6, IPMI IPv4, IPMI IPv6. Click **Next**.
Host name is required for the Hyper-V clusters, while it is optional for AHV and ESXi clusters.

8. (Optional) In the **Networking** page, select the **Active Uplink** and **Backup Uplink** for each host and click **Next**. To configure the switches, expand the row containing each host.
9. In the **Software Check** page, review the AOS and hypervisor versions and their compatibility status.

The new hosts must run the same AOS and hypervisor versions as the cluster. If the host's AOS version is lower than the cluster's AOS version, an upgrade is performed.

If there is a version mismatch between the new host and ESXi or Hyper-V cluster, you are prompted to upload the ISO in the `~/software_downloads/hypervisor_installer` directory in any one CVM of the cluster using `wget` command. If the `~/software_downloads/hypervisor_installer` directory does not exist on the CVM, you must create the directory and then upload the ISO file. After you upload the ISO file, enter the ISO file name in the text box under **hypervisor Compatibility** and click **Validate**. The system checks for the file and displays a message.
10. Click **Next**.
11. Review the model and serial information displayed for all the selected hosts, and click **Expand Cluster**.
12. (Optional) You can click **Run Prechecks** before clicking **Expand Cluster** to know any issues that you must fix before proceeding with the Expand Cluster operation.

Under Recent tasks list, you can verify that the **Expand Cluster** operation has started. The operation might take some time if it involves the re-imaging or a software upgrade at the node.

Prerequisites for expanding a Cluster through Prism Central

The process for expanding a cluster varies depending on the AOS version, hypervisor host type, data-at-rest encryption status, and certain hardware configuration factors.

AOS Considerations

The following apply to any cluster:

- Ensure that the total number of nodes per cluster does not exceed the Cluster Maximums defined in [Maximum System Values](#) in *Prism Element Web Console Guide*. Note that the maximum number of nodes per cluster differ per the hypervisor type and in case it is a pure hypervisor cluster (cluster with only one type of hypervisor) or mixed hypervisor cluster (cluster with more than one type of hypervisor).
- The expand cluster process does not support compute-only node preparation.
- The expand cluster discovery step (finding nodes to add) requires that IPv6 multicast packets are allowed through the physical switch. Therefore, the expand cluster process does not work when IPv6 is disabled in your network. If IPv6 is not enabled currently, do one of the following:
 - Enable IPv6 in your network and retry the expand cluster operation.
 - If enabling IPv6 is not an option, the expand cluster procedure allows you to enter the IP addresses of nodes to add manually and run discovery using IPv4. This requires that you have the IP addresses before starting the expand cluster operation.
- If the Controller VM memory on a new node is less than the current nodes in the cluster, the expand cluster process increases the memory on the new node to the same base value as the current nodes. The new Controller VM is upgraded to a maximum of 32 GB.

The Controller VM is upgraded to a maximum of 28 GB for ESXi nodes with 64 GB or less of total physical memory. With total physical memory greater than 64 GB, the existing Controller VM memory is increased by 4 GB.
- A new node is reimaged automatically before being added under certain conditions. The following table describes those conditions.

Table 120: Node Imaging Criteria

Configuration	Description
Same AOS and hypervisor versions	The node is added to the cluster without reimaging it.
Same hypervisor version but different AOS version	The node is automatically reimaged before it is added. However, if the AOS version on the node is higher than the version on the cluster, you can upgrade the cluster to the higher version. If you do not upgrade the base cluster to match the node AOS version, the node is reimaged automatically to match the lower AOS version of the cluster. For more information about how to upgrade your cluster, see the Life Cycle Manager Guide .
Same AOS version but different hypervisor version	The node is automatically reimaged before it is added.

Note: If you are expanding a cluster (now) on which the network is segmented only by traffic type (management and backplane), see [Network Segmentation During Cluster Expansion](#) in the *Security Guide*.

AHV Considerations

The following apply to clusters running AHV:

- If the Controller VMs in the cluster reside in a VLAN configured network, the discovery process still finds any factory-prepared nodes regardless of their current VLAN status (configured or not configured).
- Network configuration has the following restrictions and requirements:
 - You cannot migrate management from br0 to other bridges.
 - You can only have the Controller VM management interface (eth0) and hypervisor management interface deployed on the br0 bridge.

ESXi Considerations

The following apply to clusters running ESXi:

- Before adding a host running ESXi 7.0U2 and later versions, with Trusted Platform Module (TPM) 2.0 enabled, to a cluster, Nutanix recommends that you backup the recovery key created when encrypting the host with TPM. For information on how to generate and backup the recovery key, see [KB 81661](#) in the *VMware documentation*. Ensure that you use this recovery key to restore the host configuration encrypted by TPM 2.0 if it fails to start after adding the host to your cluster. For information on how to restore an encrypted host, see [KB 81446](#) in the *VMware documentation*. If you don't have the recovery key, and if the host fails to start, contact Nutanix Support.
- If the ESXi root user password was changed from the default, the expand cluster operation might fail. In this case, reset the ESXi root user password to the default, and then retry the expand cluster procedure. For default cluster credentials, see [KB 1661](#).
- While expanding a Nutanix cluster running NSX enabled ESXi hosts, add the newly imaged node to the Nutanix cluster where the host and CVM management network are configured with a standard vSwitch, and then add the node in NSX manager. Otherwise, the cluster expansion operation fails with the following error. For information on expanding a cluster, see [Expanding a Cluster](#) in *Prism Element Web Console Guide*.

Failed to get VLAN tag of node <MAC Address of the node>

- If network segmentation is not enabled, expand cluster supports mixed (ESXi + storage-only) node clusters. However, when network segmentation is enabled, you cannot use expand cluster for a mixed cluster.

- Network configuration has the following restrictions and requirements:
 - You cannot configure the network when either a target node or the base cluster has LACP enabled. Prepare LACP nodes offline using a Foundation VM.
 - You cannot migrate management from vSwitch0 to another standard vSwitch.
 - Management interfaces must either be on a VSS or a DVS; a mixed setup is not supported.
 - If on VSS, the Controller VM management interface (eth0) and the hypervisor management interface must be deployed on vSwitch0 and connected to port group VM Network and Management Network, respectively.
 - If on DVS, all Controller VM management interfaces and all host management interfaces must be connected to the same distributed virtual switch and same port group. (However, the Controller VM interfaces can be connected to a different DVS port group than the host management interfaces.)
 - Segmented network interfaces (backplane, volume, DR) must be on same vSwitch type (VSS or DVS) as the management.
 - If network segmentation is enabled on the base cluster and the backplane is deployed on a separate vSwitch than management, you can create vSwitches (VSS or DVS) and prepare the required Controller VM interfaces. (Manual switch configuration is required for expand now but is integrated into the expand later work flow.)
 - If the base cluster is on DVS and you are doing an expand later, you can migrate the target nodes from the default vSwitch0 to that DVS.
- If the Controller VMs in the cluster reside in a VLAN configured network, you must first configure the new nodes in the same VLAN before attempting to add them. Otherwise, the discovery process does not find these nodes. For more information about VLAN configuration instructions, see [Discovering Nodes in a VLAN-Segmented Network](#) in the *Field Installation Guide*.
- To expand a cluster (now) configured with DVS for Controller VM external communication, ensure that you do the following:
 - Expand DVS with the new node.
 - Make sure both the host and the CVM are configured with DVS.
 - Make sure that host to CVM and CVM to CVM communications are working.
 - Follow the cluster expansion procedure.
- After adding the new nodes, note the following:
 - The common Nutanix datastores are mounted on the new nodes by default after cluster expansion.
 - The target storage containers must be set to mount on the new hosts. You can check the mount status from the **Storage** dashboard. For more information, see [Storage Table View](#) in *Prism Element Web Console Guide*. Click the **Storage Container** tab, select the target storage container, click the **Update** button, and verify that **Mount on all ESXi Hosts** (or the new hosts are checked in **Mount/Unmount on the following ESXi Hosts**) is selected in the **NFS DATASTORE** field.
 - If an added node has an older processor class than the existing nodes in the cluster, cluster downtime is required to enable EVC (enhanced vMotion compatibility) with the lower feature set as the baseline. For an indication of the processor class of a node, see the **Block Serial** field in the **Hardware** dashboard. For more information, see [Hardware Diagram View](#) or [Hardware Table View](#) in *Prism Element Web Console Guide*. For more information on enabling EVC, see [vSphere EVC Settings](#) in the *vSphere Administration Guide for Acropolis*.

Caution: If you mix processor classes without enabling EVC, vMotion/live migration of VMs is not supported between processor classes. If you add the host with the newer processor class to vCenter Server before enabling

EVC, cluster downtime is required to enable EVC later because all VMs (including the Controller VM) must be shut down.

- Add the new nodes to the appropriate vCenter Server cluster. If an added node has a newer processor class (for example, Haswell) than the existing nodes in the cluster (Ivy Bridge or Sandy Bridge), enable EVC with the lower feature set as the baseline before adding the node to vCenter.
- If you are adding multiple nodes to an existing EVC-enabled vCenter cluster, which requires powering off the Controller VM for each node to complete the addition, add just one node at a time and wait for data resiliency to return to *OK* before adding the next node to vCenter.

Caution: Adding multiple nodes to vCenter simultaneously can cause a cluster outage when all the Controller VMs are powered off at the same time.

- If you are adding a node to a cluster where HA is enabled with APD and VMCP is enabled, you must enable APD and APD timeout on the new host.
- If you are adding new nodes to vCenter EVC configured cluster, ensure the following requirements.
 - Enabling EVC requires ESXi version 6.0 or above.
 - All ESXi nodes should be on same version and build.

Hyper-V Considerations

The following apply to clusters running Hyper-V:

- Network preparation is not supported. To bypass this step and prepare the nodes otherwise, click the **Skip Networking** button when you get to the **Networking** tab. In addition, imaging to AHV or ESXi is not possible if the node or base cluster is Hyper-V.
- If the Controller VMs in the cluster reside in a VLAN configured network, you must first configure the new nodes in the same VLAN before attempting to add them. Otherwise, the discovery process does not find these nodes. For more information about VLAN configuration instructions, see [Discovering Nodes in a VLAN-Segmented Network](#) in the *Field Installation Guide*.
- After adding the new nodes, if you manage your Hyper-V cluster by using Microsoft System Center VM Manager (SCVMM), do not use the Prism Element web console or Microsoft Failover Cluster Manager to add the new node to the failover cluster. Instead, use SCVMM to add the node to the Hyper-V cluster and then perform the following steps in the SCVMM user interface.
 1. Open the SCVMM user interface.
 2. Refresh the cluster in SCVMM. The new node is displayed under the failover cluster in the **Pending** state.
 3. Right-click the node and select **Add to host cluster**.
 4. Choose a run-as account that has the local administrator permissions on the new node.
 5. Click **OK**. The SCVMM agent is installed on the node and file shares are registered to the new node.
 6. Update the networking and other settings of the node to match your standard configuration.

For Microsoft Windows Server 2012 R2 deployments, after adding the node to the cluster in SCVMM, ensure that node disks are not added as clustered resources. Open the Failover Cluster Manager and click **Storage > Disks** to check.

Nutanix Clusters Considerations

The following apply to clusters hosted on a cloud platform:

- When expanding a cluster on AWS, the nodes are not added to the Cassandra ring (wait in a queue) until there are enough nodes to extend the ring. In addition, new nodes are added to the Cassandra ring only when they do not break domain awareness. (Other services remain unaffected.)

Hardware Considerations

Note the following when it applies to the nodes you are adding:

- Ensure that you use the minimum version of Foundation required by your hardware platform. To determine whether Foundation needs an upgrade for a hardware platform, see the respective system specifications guide. If the nodes you want to include in the cluster are of different models, determine which of their minimum Foundation versions is the most recent version, and then upgrade Foundation on all the nodes to that version.
- When adding partially populated all-SSD nodes to an existing cluster, the minimum number of partially populated all-SSD nodes added must be equal to the maximum RF in the cluster.
- If you are adding a node with a different processor class to the cluster, ensure that there are no running VMs on the node and the host has the following configuration:
 - ESXi: Verify that EVC is enabled on the cluster.
 - Hyper-V: If you want to move the VMs between the nodes, ensure that you have selected the **Migrate to a physical computer with a different processor version** option for each VM by browsing to **Settings > Processor > Compatibility** in the **Action** pane of the Hyper-V Manager.

Note: Do not shut down more than one Controller VM at the same time.

- If you expand a cluster by adding a node with older generation hardware to a cluster that was initially created with later generation hardware, power cycle (do not reboot) any guest VMs before migrating them to the added older generation node or before upgrading the cluster.

Guest VMs are migrated during hypervisor and firmware upgrades (but not AOS upgrades).

For example, if you are adding a node with G4 Haswell CPUs to a cluster that also has newer G5 nodes with Broadwell CPUs, you must power cycle guest VMs hosted on the G5 nodes before you can migrate the VMs to the node with G4 CPUs. Power cycling the guest VMs enables them to discover a CPU set compatible with older G4 processors.

In rare cases, certain CPU features might be deprecated in the new generation of CPUs. For example, Intel introduced MPX in Skylake class of CPUs and deprecated it with Ice Lake. In such cases, introduction of Ice Lake (newer) CPUs to an all-Skylake cluster can cause problems with existing VMs that are running with MPX. Such VMs must be power cycled.

Power cycle guest VMs from the Prism Element web console VM dashboard. Do not perform a Guest Reboot; a VM power cycle is required in this case.

- If you physically add a node to a block (for example, a single node shipped from Nutanix is placed into an empty slot in an existing chassis), log on to the Controller VM for that node, and update the following parameters in the `/etc/nutanix/factory_config.json` file:
 - `rackable_unit_serial`: Set it to the same value as the other Controller VMs in the same block.
 - `node_position`: Set it to the physical location of the node in the block (A, B, C, D).

After changing the configuration file, restart Genesis with the `genesis restart` command.

Limitations for Expanding a Cluster through Prism Central

Consider the following limitations before you expand a cluster through Prism central.

- Encryption and Rack Awareness functionalities are not yet supported for expanding a cluster through Prism Central.
- Expand cluster functionality does not have an option to upload the hypervisor image through Prism Central web console. If there is a version mismatch between the new host and ESXi or Hyper-V cluster, you are prompted to

upload the ISO in the ~/software_downloads/hypervisor_installer directory in any one CVM of the cluster using wget command, and enter the file details in the Prism Central web console.

Removing a Node through Prism Central

You can remove a node from a cluster using Prism Central.

Before you begin

Review [Prerequisites for Removing a Node through Prism Central](#) on page 461 before attempting to remove a node from a cluster.

About this task

From Prism Central version pc.2024.3 onwards, you can remove a node without registering the hosting cluster with Prism Central. The minimum supported version to remove a host using Prism Central web console is pc.2023.3.

To remove a host from a cluster using Prism Central web console, perform the following steps:

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Hardware** > **Hosts** .
The system displays the **List** tab by default with all the hosts across registered clusters.
3. Select the target host checkbox, and choose **Remove Host** from the **Actions** dropdown menu.
The system displays the **Remove Host** window. As a result of successful execution of remove host operation, the host gets detached from the cluster and the VMs on the selected host are migrated to other hosts. Some VMs might be automatically powered off due to lack of resources.
4. In the **Remove host** window, click **Run Precheck** to run the prechecks for remove host operation. The system checks several points including the space availability for VMs of the node to be removed, fault tolerance, upgrade status on the node. If any of the prechecks fail, fix the issue and then click **Remove** to remove the host.
You can verify the progress of the remove host operation under recent tasks list.

Caution: Do not shut down the CVM or put the CVM into maintenance mode while the node removal is in progress.

Prerequisites for Removing a Node through Prism Central

Consider the following set of prerequisites before you remove a node from a cluster:

- For information about reclaiming or rebalancing your cluster licenses after node removal, see the [License Manager Guide](#).
- Removing a node (host) takes some time because data on that node must be migrated to other nodes before it can be removed from the cluster. You can monitor progress through the dashboard messages. Removing a node implicitly removes all the disks in that node.
- (Hyper-V only) Initiating a removal of a node running Hyper-V fails if the node is running as a part of a Hyper-V failover cluster and the following message appears.

Node *node id* is a part of a Hyper-V failover cluster *failover cluster name*. Please drain all the roles, remove the node from the failover cluster and then mark the node for removal.

If this message is displayed in either nCLI or in web interface, as a cluster administrator, you must use the management tools provided by Microsoft such as *Failover Cluster Manager* to drain all the highly-available

roles off the node. Then remove the node from the failover cluster followed by removing the node from the AOS cluster.

- (ESXi only) Before removing the node, ensure that there are no guest VMs running on the respective ESXi host. Note that CVM should be running on the node to initiate the node removal.
- (ESXi only) Temporarily disable DRS on the cluster before the node removal. The DRS must be re-enabled after the node is removed. Update the HA configuration to exclude the node you want to remove.
- (ESXi only) As a cluster administrator, you must use the management tools provided by VMware to first migrate all the guest VMs off the node/host, then remove the node from the AOS cluster of ESXi hosts. After that disconnect and remove the node (host) from the vCenter server.

Caution: Ensure that you migrate the guest VMs before removing a host or node. Verify that the target cluster has enough available compute capacity before actually migrating the VMs. Removing a node or host without first migrating the guest VMs may result in loss of service.

1. Migrate the guest VMs that need to be migrated.
2. Click **Remove Host** from Prism.
3. After successful removal from the AOS cluster, put the host in maintenance mode.

Caution: When you put the host in maintenance mode, the maintenance mode process powers down or migrates all the VMs that are running on the host.

For more information, see [Node Maintenance](#) in the *vSphere Administration Guide for Acropolis*.

4. Remove the host from the vCenter server.

Limitations for Removing a Node through Prism Central

Consider the following limitations before you remove a node from a cluster.

- You can remove only one node at a time.
- For a cluster with fault tolerance 1 (FT1), the minimum number of nodes required in a cluster is three. Therefore, you can remove a node only if the cluster has a minimum of four nodes.
- For a cluster with fault tolerance 2 (FT2), the minimum number of nodes required in a cluster is five. Therefore, you can remove a node only if the cluster has a minimum of six nodes.

Putting a Host into Maintenance Mode Using Prism Central

This section describes how to put the host into maintenance mode using Prism Central.

Before you begin

Ensure that the following prerequisites are met before you put a host in maintenance mode:

- All the hosts are running, and no host in the cluster is in maintenance mode or undergoing existing maintenance mode operation.
- The cluster comprises at least two hosts.
- The selected host uses the AHV or ESXi hypervisor only.
- The **Data Resiliency** status of the cluster is **OK**. For more information, see [Clusters Summary View](#) on page 424.

About this task

You can put only one host of the cluster in maintenance mode from Prism Central at a time.

Note: If you put the ESXi host in maintenance mode using vCenter, Prism Central does not show the host maintenance mode status.

As the host enters the maintenance mode, the following high-level tasks are performed internally:

- The AHV host initiates the maintenance mode process.
- The HA VMs are live migrated.
- The system prompts the user to confirm the shut down action for pinned VMs and for VMs with replication factor 1.
- The AHV host completes the maintenance mode process.

Note: At this stage, the AHV host is not shut down.

- The CVM enters the maintenance mode.
- The CVM shuts down.

Procedure

To put the node into maintenance mode, follow these steps:

1. Log on to Prism Central.
2. From the [Application Switcher Function](#), select the **Infrastructure** application, and navigate through the **Navigation Bar** to **Hardware > Hosts**.
For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#).
The system displays the **List** tab with all the hosts across registered clusters.
3. Select the target host checkbox and from the **Actions** dropdown menu, choose **Enter Maintenance Mode**. The system displays the **Host Maintenance** window and performs the checks to identify if there are any non-migratable VMs on the host.

Note: For the ESXi cluster, the system prompts you to enter the vCenter credentials before it performs any check to identify the non-migratable VMs in the cluster. After you enter the vCenter credentials, the system prompts you to perform the following actions based on the VM condition:

- Non-migratable VMs: The system prompts you to shut down the non-migratable VMs. Select the **Power off VMs that cannot migrate (x unmigratable VMs)** checkbox.
- No non-migratable VMs: The system prompts you to confirm the enter maintenance mode action.

Note: Click **View these x VMs** to display the list of non-migratable VMs. x indicates the number of non-migratable VMs

For more information on the conditions that impact VM live migration, see VM [Live Migration Restrictions](#) in the *AHV Administration Guide*.

4. Click **Enter Maintenance Mode**.
The system puts the host in maintenance mode. You can check the status of the **Enter Maintenance Mode** operation from the **Tasks View** page. For more information, see [Tasks View](#) on page 496.

Exiting a Host from Maintenance Mode Using Prism Central

This section describes how to exit the host from maintenance mode using Prism Central.

About this task

You can exit maintenance mode from Prism Central only if you have put the host in maintenance mode from Prism Central or Prism Element.

Note:

- If you put the host in maintenance mode using aCLI (for AHV host) or using vCenter (for ESXi host), you cannot exit the maintenance mode from Prism Central.
- If you put the ESXi host in maintenance mode using vCenter, Prism Central does not show the host maintenance mode status.

As the host exits the maintenance mode, the following high-level tasks are performed internally:

- The CVM is powered on.
- The CVM is taken out of maintenance mode.
- The host is taken out of maintenance mode.

Note: The AHV host is shut down while [Putting a Host into Maintenance Mode Using Prism Central](#) on page 462 and you must start the AHV host. For information on how to start the AHV host, see [Starting a Node in a Cluster \(AHV\)](#) in *AHV Administration Guide*.

After the host exits the maintenance mode, the VMs with replication factor 1 continue to be powered on and the VMs migrate to restore the host locality.

Procedure

To exit the node from maintenance mode, follow these steps:

1. Log on to Prism Central.
2. From the [Application Switcher Function](#), select the **Infrastructure** application, and navigate through the **Navigation Bar** to **Hardware > Hosts**.
For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#).
The system displays the **List** tab with all the hosts across registered clusters.
3. Select the target host checkbox and from the **Actions** dropdown menu, choose **Exit Maintenance Mode**.
The system displays the **Host Maintenance** window and prompts you to confirm the exit maintenance mode action.
4. Click **Exit Maintenance Mode**.
The system exits the host from maintenance mode. You can check the status of the **Exit Maintenance Mode** operation from the **Tasks** page. For more information, see [Tasks View](#) on page 496.

Disks Summary View

Summary View of All Disks

The **Summary** tab on the **Disks** page provides a dashboard of the disks across all the registered clusters.

To access the **Summary** view of all the disks:

1. Log in to Prism Central.

2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Hardware > Disks** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.
3. Click the **Summary** tab. The system displays **Summary** page for all the disks.

Clicking the **Summary** tab displays the following three widgets:

- **Highlighted Entities:** Displays a list of the disks with the highest usage of the parameter you select from the dropdown menu on the right of the widget. The options are **IO Bandwidth**, **IOPS**, **IO Latency**, and **Disk Usage**. Click the **View all XX Disks** link at the bottom to display the **List** tab (following section).
- **Alert:** Displays a list of disk-related alerts that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the dropdown menu.
- **Anomalies:** Displays a graph of memory, I/O, CPU, networking, or disk anomalies that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the dropdown menu. When an anomaly appears, you can click on the graph, which then displays a list of those anomalies. Clicking on an anomaly displays the event page for that anomaly.

Detailed View of an Individual Disk

For information about how to access the detailed view of an individual host, see [Disk Details View](#) on page 468.

List Tab

Clicking the **List** tab, which appears by default when you first open the page, displays a list of the disks across the registered clusters. The following table describes the fields that appear in the disks list. The fields vary based on the **Focus** menu selection, which is either **General** or **Performance**. A dash (-) is displayed in a field when a value is not available or applicable.

Table 121: Disks List Fields

Parameter	Description	Values
General Focus Fields		
Serial Number	Displays the disk serial number.	(serial number)
Host	Displays the name of the host in which this disk resides.	(host name)
Tier	Displays the disk type (tier name). Nutanix models can contain disk tiers for PCIe solid state disks (SSD-PCIe), SATA solid state disks (SSD-SATA), and direct attach SATA hard disk drives (DAS-SATA) depending on the model type.	[SSD-PCIe SSD-SATA DAS-SATA]
Mode	Displays the operating state of the disk.	online, offline
Disk Usage	Displays the percentage of disk space used and total capacity of this disk.	[0 - 100%] of xxx [GB TB]
Cluster	Displays the name of the cluster in which the host resides.	(cluster name)
Performance Focus Fields		
Serial Number	Displays the disk serial number.	(serial number)

Parameter	Description	Values
Disk Usage	Displays the percentage of disk space used and total capacity of this disk.	[0 - 100%] of xxx [GB TB]
Disk Capacity	Displays the total physical space on the drive.	xxx [GB TB]
IOPS	Displays I/O operations per second (IOPS) for this disk.	[0 - unlimited]
IO Bandwidth	Displays I/O bandwidth used per second for this disk.	xxx [MBps KBps]
IO Latency	Displays the average I/O latency (in milliseconds) for this disk.	xxx [ms]
Cluster	Displays the name of the cluster in which the disk resides.	(cluster name)

You can filter the disks list based on a variety of parameter values. The following table describes the filter options available when you open the **Filter** pane. To apply a filter, select a parameter and check the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters. Some parameter filters require additional context such as a constraint string or a range.

Table 122: Filter Pane Fields

Parameter	Description	Values
Serial Number	Filters on the disk serial number. Select a condition from the dropdown list (Contains , Doesn't contain , Starts with , Ends with , or Equal to) and enter a string in the field. It will return a list of disks that satisfy the serial number condition/string.	(serial number string)
Host	Filters on the host name. Enter a string in the field. It will return a list of disks in the hosts that satisfy the host name condition/string.	(host name string)
Cluster	Filters on the cluster name. Enter a string in the field. It will return a list of disks in the clusters that satisfy the cluster name condition/string.	(cluster name string)
Mode	Filters on whether the disk is online or offline. Check the box for one or both of these modes. The number of disks currently in each state is displayed on the right of the line.	Online, Offline
Tier	Filters on whether the disk is in the solid state (SSD-SATA) or hard disk (DAS-SATA) tier. Check the box for one or both of these modes. The number of disks currently in each tier is displayed on the right of the line.	DAS-SATA, SSD-SATA
Health	Filters on the disk health state (good, warning, or critical). Select one or more states to return a list of disks in that state(s). The number of disks currently in each state is displayed on the right of the line.	Critical, Warning, Good

Parameter	Description	Values
Disk Usage	Filters on the used capacity. Enter a percentage range in the " from <low> to <high> % " field. It will return a list of disks with used capacity in that range (0-100%).	([xx] to [yy]% range)
Disk Capacity	Filters on the total capacity. Enter an amount range in the " from <low> to <high> GiB " field. It will return a list of disks with total capacity in that range.	([xx] to [yy] GiB range)
IOPS	Filters on the IOPS. Enter a range in the " from <low> to <high> iops " field. It will return a list of disks with IOPS in that range.	([xx] to [yy] range)
IO Bandwidth	Filters on the I/O bandwidth used. Enter a range in the " from <low> to <high> bps " field. It will return a list of disks with I/O bandwidth usage in that range.	([xx] to [yy] range)
IO Latency	Filters on the average I/O latency. Enter a range in the " from <low> to <high> ms " field. It will return a list of disks with average I/O latency in that range.	([xx] to [yy] range)

You can group the disks list in the following ways:

- The **Color** dropdown menu allows you to color code the disk entries by tier type, mode, or health state. (You can only choose one.) A legend appears at the bottom to indicate what each color means in that grouping.
- The **Group** dropdown menu allows you to group the disk entries by host, tier type, mode, cluster, or health state. (You can only choose one.)
- [Tiles and Circles views only] The **Sort** dropdown menu allows you to group the disk entries by the information parameters (fields), which vary depending on whether you selected the **General** or **Performance** focus. (You can only choose one parameter.)

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, except it is filtered to display just disk-related alerts across the registered clusters. For more information, see [Alerts Summary View \(Prism Central\)](#) in *Prism Central Alerts and Events Reference Guide*.

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options as the Events dashboard, except it is filtered to display just disk-related events across the registered clusters. For more information, see [Events Summary View \(Prism Central\)](#) in *Prism Central Alerts and Events Reference Guide*.

Metrics Tab

The **Metrics** tab allows you to view usage metrics across the hosts. Clicking the **Metrics** tab displays a list of available metrics; click the metric name to display the relevant information to the right. The following table describes the available metrics.

Table 123: Metrics Tab Fields

Metric	Description
IOPS	Displays total, read, and write IOPS tables listing current values and total disks (number). The current values are split into intervals (for example, less than 20, 20-40, 40-60, more than 60). Clicking on an interval displays the Summary tab filtered to just those disk.
Note: The same format also applies to the other metrics in this table.	
IO Latency	Displays total, read, and write I/O latency rate tables.
IO Bandwidth	Displays total, read, and write I/O bandwidth rate tables.

Disk Details View

Summary Tab

The **Summary** tab of an individual disk consists of a dashboard that provides the detailed information about the disk.

To access the **Summary** tab of an individual disk:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Hardware > Disks** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

By default, the system displays the **List** tab with all the disks.

3. Click the target `<disk_name>` to view the **Summary** tab of an individual disk.

Note: Replace `<disk_name>` with the actual host name at your site.

Summary Tab

The **Summary** tab of an individual disk provides the following widgets:

- A **Properties** widget that displays summary information about the disk. For more information, see the [Table 124: Disk Properties Fields](#) on page 468 table.
- An **Alert** widget that displays a list of related alerts that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the dropdown menu.
- An **Anomalies** widget that displays a graph of memory, I/O, CPU, networking, or disk anomalies that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the dropdown menu. When an anomaly appears, you can click on the graph, which then displays a list of those anomalies. Clicking on an anomaly displays the event page for that anomaly.

The following table describes the fields in the **Properties** widget. A dash (-) in a field indicates there is not enough data to evaluate or a value is not assigned. The displayed fields vary by hypervisor.

Table 124: Disk Properties Fields

Parameter	Description	Values
Disk Usage	Displays the amount of used space on the drive.	xxx [GB TB]

Parameter	Description	Values
Cluster	Displays the name of the cluster in which the disk resides.	(cluster name)
Host	Displays the name of the host in which the disk resides.	(host name)
Host IP	Displays the IP address of the host.	(IP address)
Tier	Displays the disk type (tier name). Nutanix models can contain disk tiers for PCIe solid state disks (SSD-PCIe), SATA solid state disks (SSD-SATA), and direct attach SATA hard disk drives (DAS-SATA) depending on the model type.	[SSD-PCIe SSD-SATA DAS-SATA]
Mode	Displays whether the disk is currently online or offline.	[online offline]
Disk Capacity	Displays the total physical space on the drive.	xxx [GB TB]
Status	Displays the operating status of the disk. Possible states include the following: • Normal. Disk is operating normally. • Data migration initiated. Data is being migrated to other disks. • Marked for removal, data migration is in progress. Data is being migrated in preparation to remove disk. • Detachable. Disk is not being used and can be removed.	Normal; Data migration initiated; Marked for removal, data migration is in progress; Detachable
Self Encryption Drive	Displays whether this is a self-encrypted drive.	Not Present, Present

Alerts Tab

The **Alerts** tab displays a table of alerts. This tab provides the same features and options as the Alerts dashboard, except it is filtered to display just alerts for this disk. For more information, see [Alerts Summary View \(Prism Central\)](#) in *Prism Central Alerts and Events Reference Guide*.

Events Tab

The **Events** tab displays a table of events. This tab provides the same features and options as the Events dashboard, except it is filtered to display just events for this disk. For more information, see [Events Summary View \(Prism Central\)](#) in *Prism Central Alerts and Events Reference Guide*.

Metrics Tab

The **Metrics** tab allows you to view usage metrics for the disk. Click the **Metrics** tab and then the desired metric name (**IOPS**, **IO latency**, and **IO Bandwidth**) to display a graph for that metric on the right. The graph is a rolling time interval performance or usage monitor. The baseline range (based on the machine-learning algorithm) appears as a blue band in the graph.

Note: The machine-learning algorithm uses 21 days of data to monitor and predict performance. A graph or baseline band may not appear if less than 21 days of data is available.

- Check the appropriate box(es) to have the graph display total, read, or write usage (or any combination of the three).
- Place the cursor anywhere on the horizontal axis to display the value at that time.
- Select the duration (time interval) from the dropdown list on the right (last 1 hour, last 24 hours, last week, last 21 days).
- Click **Alert Settings** to configure an alert for this metric. For more information, see [Creating Custom Alert Policies](#) in *Prism Central Alerts and Events Reference Guide*.

Storage Usage Tab

The **Usage** tab displays the following graph:

- The **Usage Summary** graph displays a rolling time interval monitor of disk storage usage that can vary from one to several hours depending on activity moving from right to left. Placing the cursor anywhere on the horizontal axis displays the value at that time. For more in depth analysis, you can add the monitor to the analysis page by clicking the blue link in the upper right of the graph.

GPUs Summary View

Summary View of All GPUs

The **Summary** tab on the **GPUs** page provides a dashboard of the GPUs across all the registered clusters.

To access the **Summary** view of all the GPUs:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Hardware > GPUs** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.
3. Click the **Summary** tab.

The system displays **Summary** page for all the GPUs.

The following table describes the fields that appear in the GPUs list. A dash (-) is displayed in a field when a value is not available or applicable.

Table 125: GPU List Fields

Parameter	Description	Values
Type	Displays the GPU model type.	Tesla M10, Tesla M60, Tesla M60 compute
Cluster	Displays the name of the cluster in which the GPU resides.	(cluster name)
Mode	Displays the mode in which the GPU is operating.	None, vGPU, passthrough
Allocation	Indicates the number of VMs allocated to the GPU.	"No VM allocated", "x of y VMs allocated"

You can filter the GPUs list based on several parameter values. The following table describes the filter options available when you open the GPUs view **Filter** pane. To apply a filter, select a parameter and check the box of the

desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters. Some parameter filters require additional context such as a constraint string or a range.

Table 126: Filter Pane Fields

Parameter	Description	Values
Type	Filters on the GPU model type. Select a condition from the dropdown menu (Contains , Doesn't contain , Starts with , Ends with , or Equal to) and enter a string in the field. It returns a list of GPUs that satisfy the type condition/string.	Tesla M10, Tesla M60, Tesla M60 compute
Host	Filters on the host name. Enter a string in the field. It returns a list of GPUs in the selected hosts.	(host name string)
Cluster	Filters on the cluster name. Enter a string in the field. It returns a list of GPUs in the selected clusters.	(cluster name string)
Mode	Filters on the GPU operation mode. Check the box for one or more of these modes. The number of GPUs currently in each mode is displayed on the right of the line.	None, vGPU, Passthrough

You can group the GPUs list in the following ways:

- The **Group** dropdown menu allows you to group the GPU entries by cluster, host, or mode. (you can only choose one.)
- The **Sort** dropdown menu allows you to group the GPU entries by type, cluster, mode, or allocation. (you can only choose one parameter.)

Clicking the **Summary** tab displays the following three widgets:

- **Highlighted Entities:** Displays a list of the disks with the highest usage of the parameter you select from the dropdown menu on the right of the widget. The options are **IO Bandwidth**, **IOPS**, **IO Latency**, and **Disk Usage**. Click the **View all XX Disks** link at the bottom to display the **List** tab (following section).
- **Alert:** Displays a list of disk-related alerts that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the dropdown menu.
- **Anomalies:** Displays a graph of memory, I/O, CPU, networking, or disk anomalies that occurred during the specified interval. Select either **Last 24 hours** (default) or **Last week** from the dropdown menu. When an anomaly appears, you can click on the graph, which then displays a list of those anomalies. Clicking on an anomaly displays the event page for that anomaly.

Detailed View of an Individual GPU

For information about how to access the detailed view of an individual host, see [GPU Details View](#) on page 471.

GPU Details View

Summary Tab

The **Summary** tab of a GPU consists of a dashboard that provides the detailed information about the GPU.

To access the **Summary** tab of an individual GPU:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Hardware > GPUs** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.
By default, the system displays the **List** tab with GPUs across all the registered clusters.
3. Click the target `<GPU_name>` to view the **Summary** tab of an individual GPU.

Note: Replace `<GPU_name>` with the actual GPU name at your site.

Clicking the **Summary** tab displays the following:

- A section on the left that displays summary information about the GPU. For more information, see the [Table 127: GPU Summary Fields](#) on page 472 table.
- A section of the right that displays GPU performance metrics. The graphs are rolling time interval performance monitors that can vary from one to several hours depending on activity moving from right to left. Placing the cursor anywhere on the horizontal axis displays the value at that time. This section includes the following graphs:
 - **GPU Usage:** Displays the percentage of GPU capacity being used.
 - **GPU Framebuffer Usage:** Displays the percentage of GPU framebuffer (RAM) capacity being used.
- **Action** on the upper right (which appears regardless of which tab is selected):
 - Click the [Help icon](#) to open a help page in a separate tab or window.
 - Click the X icon to close the details page.

The following table describes the GPU summary information fields. A dash (-) is displayed in a field when a value is not available or applicable.

Table 127: GPU Summary Fields

Parameter	Description	Values
GPU Type	Displays the GPU type for this entry.	Tesla M10, Tesla M60, Tesla M60 compute
Cluster Name	Displays the name of the cluster in which the GPU resides.	(cluster name)
Host	Displays the name of the host in which the GPU resides.	(host name)
Mode	Displays the GPU operational mode. If it is vGPU, an additional field appears that displays the vGPU profile used.	None, vGPU, Passthrough
Allocation	Displays the number of VMs allocated to this GPU. This field does not appear if no VMs are allocated.	"x of y VMs allocated"
Framebuffer (RAM) Per GPU	Displays the framebuffer (RAM) size per GPU.	xx GiB
ID	Displays the GPU ID number.	(ID number)

VMs Tab

Clicking the **VMs** tab displays a table of VMs allocated (attached) to the GPU. The table includes the following fields:

- **Name:** Displays the VM name. Click the name to display the details page for that VM. For more information, see [Viewing VM Details](#) on page 113.
- **GPU Usage:** Displays the percentage of GPU capacity used by this VM.
- **GPU Framebuffer Usage:** Displays the percentage of GPU framebuffer (RAM) capacity used by this VM.

NICs and NIC Profile Management

You can manage the NICs attached to the hosts managed by the Prism Central, and the NIC Profiles for SR-IOV and Network Offload capabilities.

In Prism Central, the **NICs** page includes the **List** tab that displays the summary information about all the NICs available across the registered clusters, and provides access to the detailed information about each NIC. The **NICs** page includes the **NIC Profiles** tab that displays the summary information about all the NIC profiles available across the registered clusters, and provides access to the detailed information about each NIC profile.

You can perform the following actions to manage NICs and NIC profiles in Prism Central:

1. View the list of NICs in the registered clusters. For information, see [Viewing the NICs and NIC Profiles](#) on page 473.
2. View the details of a selected NIC. For information, see [Viewing the NIC Port Details](#) on page 478.
3. View the list of NIC profiles in the registered clusters. For information, see [Viewing the NICs and NIC Profiles](#) on page 473.
4. View the details of a selected NIC profile. For information, see [Viewing the NIC Profile Details](#) on page 481.
5. Create a NIC profile. For information, see [Creating a NIC Profile](#) on page 485.
6. Update a NIC Profile, including managing the assignment of NIC Ports to NIC Profiles. For information, see [Updating a NIC Profile](#) on page 486.
7. Delete a NIC Profile. For information, see [Deleting a NIC Profile](#) on page 489.

For information on attaching a NIC with SR-IOV NIC profile to a VM, see [Attaching an SR-IOV NIC to a VM](#) on page 490.

For information on attaching a NIC to a Network Offload NIC profile to a VM, see [Attaching a Network Offload NIC Profile to a VM](#) on page 490.

Viewing the NICs and NIC Profiles

The NICs page displays the summary information on NICs on the **List** tab and the summary information on NIC profiles on the **NIC Profiles** tab.

About this task

The **List** tab of the NIC page provides a summary view of the NICs available across all the registered clusters.

The **NIC Profile** tab provides a summary view of the NIC profiles configured for NICs across all the registered clusters.

To access the summary view of all the NICs, follow these steps:

Procedure

1. Log in to Prism Central.

- From [Application Switcher Function](#) on page 40, select the **Infrastructure** application and go to **Hardware > NICs** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

The **NICs** page appears, opening the **List** tab by default.

The **NICs** page displays the list of NIC profiles in the **NIC Profile** tab.

NIC List Tab Attributes

The **List** tab displays a list of the NICs across the registered clusters.

The fields in the NICs list vary based on the **Focus** menu selection, which is **General** by default. You can create a custom view with specific or all fields included. The following table describes all the available fields. A dash (-) is displayed in a field when a value is not available or applicable.

Table 128: NICs List Tab Attributes

Attribute	Description	Values
Name	Displays the name of the NIC. For more information on the NIC, see Viewing the NIC Port Details .	(NIC name)
MAC Address	Displays the MAC address of the NIC.	(MAC address)
NIC Family	Displays the name of the family of NICs that the NIC belongs to. It denotes the family of the NICs provided by the NIC OEMs. The NIC Family information is derived for processing in the <code>vendor_id:device_id</code> format from the name and model that you select in the NIC Family field when you create a NIC profile. For more information, see Creating a NIC Profile .	NIC Family and OEM information
Max / Configured Speed (Gbps)	Displays the maximum speed supported by the NIC provider OEM and the configured speed which is the speed negotiated with the peer.	xxx/xxx (Gbps)
Category	(N/A for current version)	-
Capability	Displays the capability (physical functions) of the NIC.	SR-IOV or Network Offload
Operational Status	Displays the operation status of the NIC.	Up or Down
Host	Displays the name of the host that the NIC belongs to.	(host name)
Cluster	Displays the name of the cluster that the NIC belongs to.	(cluster name)

Attribute	Description	Values
NIC Profile	Displays the NIC profile assigned to the NIC. The supported NIC profiles are SR-IOV and Network Offload .	SR-IOV or Network Offload

Actions

You can perform the following actions on a NIC in the **List** tab:

- Assign a NIC Profile: Assign an NC Profile to a NIC by selecting the NIC in the list and clicking **Actions** > **Assign NIC Profile**. This action is unavailable if the selected NIC does not have SR-IOV or Network Offload capabilities.

Important: You can assign a NIC Profile to a NIC only if the NIC supports the capability (SR-IOV or Network Offload as the case may be) configured in the NIC Profile to be assigned.

- Remove a NIC Profile: Remove a NIC profile already assigned to a NIC by selecting the NIC in the list and clicking **Actions** > **Delete NIC Profile**. This action is unavailable if a NIC profile was not previously assigned to the selected NIC.

You can perform the following actions in the **List** tab:

- Access the detailed information about an individual host. For more information, see NIC Port Detail View.
- Filter the host list based on available parameter values using the **Filters** pane. For more information on the Filters pane, see *Filters*.
- Export the table that contains the list of hosts and their information to a file in a CSV format. For more information on the **Export** option, see the *Export* section in [Prism Central GUI Organization](#) on page 49.
- View hosts based on pre-defined criteria or create a custom view. For information about available views and how to create a custom view, see *View by* section in [Prism Central GUI Organization](#) on page 49.

Filters

You can filter the NICs list based on a variety of attribute values. The following table describes the filter options available when you open the NICs view Filter pane. To apply a filter, select an attribute and check the box of the required value (or multiple values) to use as a filter. You can apply filters across multiple attributes. Some attribute filters require additional context such as a constraint string or a range.

Table 129: NICs Filter Pane Attributes

Attribute	Description	Value
Name	Filters on the NIC name. Select a condition from the dropdown menu (Contains , Doesn't contain , Starts with , Ends with , Equal to , or Not equal to) and enter a string in the field. It returns a list of NICs that satisfy the name condition/string.	(NIC name string)

Attribute	Description	Value
Host	Filters on the host name. Enter a string in the field. It returns a list of NICs that belong to the hosts that satisfy the name string.	(host name string)
Cluster	Filters on the cluster name. Enter a string in the field. It returns a list of NICs that belong to the clusters which satisfy the name string.	(cluster name string)
NIC Profile	Filters on the NIC Profile. Enter a string in the field. It returns a list of NICs that are assigned the selected NIC Profile. Note: NIC Profiles are available for selection only if any of the listed NICs are assigned any NIC Profile.	(NIC profile selection)
MAC Address	Filters on the MAC address of the NIC. Select a condition from the dropdown menu (Contains, Doesn't contain, Starts with, Ends with, Equal to, or Not equal to) and enter a part of MAC addresses of the NICs in the field. It returns a list of NICs that satisfy the MAC address condition.	(MAC address or part thereof)
Maximum Speed (GBPS)	Filters on the Maximum speed. Select the appropriate pre-defined speed range or enter the starting and ending values in the custom speed range by selecting the custom speed range. It returns the list of NICs that have maximum speeds in the selected range. The unit of the speed value is Gbps.	(Range selection or enter custom range)
Configured Speed (GBPS)	Filters on the configured or negotiated speed. Select the appropriate pre-defined speed range or enter the starting and ending values in the custom speed range by selecting the custom speed range. It returns the list of NICs that have the configured or negotiated speeds in the selected range. The unit of the speed value is Gbps.	(Range selection or enter custom range)

NIC Profile Tab Attributes

The **NIC Profile** tab displays a list of NIC profiles available across all the registered clusters.

The fields in the NIC Profile list vary based on the **Focus** menu selection, which is **General** by default. You can create a custom view with specific or all fields included.

You can perform the following actions on a NIC Profile in the **NIC Profile** tab:

- **Create NIC Profile:** Click the Create NIC Profile button to create a NC Profile. For more information, see [Creating a NIC Profile](#).
- **Actions:** Perform the following actions on a NIC Profile using the **Actions** dropdown menu:
- **Update NIC Profile:** For more information, see [Updating a NIC Profile](#)
- **Manage NIC Port(s):** For more information, see [Managing NIC Ports for a NIC Profile](#).
- **Delete NIC Profile:** For more information, see [Deleting a NIC Profile](#).

The following table describes all the available fields in the **NIC Profile** tab.

A dash (-) is displayed in a field when a value is not available or applicable.

Table 130: NIC Profile Tab Attributes

Attribute	Description	Value
Name	Displays the name of the NIC Profile. Clicking the name displays the details page for that NIC Profile. For more information, see <i>NIC Profile Details View</i> .	(NIC Profile name)
NIC Family	Displays the name of the family of NICs that the NIC belongs to. It denotes the family of the NICs provided by the NIC OEMs. The NIC Family information is derived for processing in the <code>vendor_id:device_id</code> format from the name and model that you select in the NIC Family field when you create a NIC profile. For more information, see <i>Creating a NIC Profile</i> .	NIC Family and OEM information.
NIC Ports	Displays the operational status of the NIC ports assigned to the NIC profile. Multiple NIC ports might be assigned to one NIC Profile. The status separately displays the number of NICs that are not operating properly.	(Status string)
Capabilities	Displays the capabilities (physical functions) of the NIC	SRIOV or DP_OFFLOAD

Attribute	Description	Value
Attached to vNICs	Displays the number of vNICs attached to the NIC Profile.	(Integer number)
Virtual Functions	Displays the number of Virtual Functions available for the NIC Profile. The number is displayed as the number of VFs available for use out of the number of VFs provided by the NIC hardware. For example, 7 of 8 Available meaning seven VFs are available for use by VMs out of the eight VFs that are provided by that NIC hardware.	(Integer number) of (integer number) Available

Viewing the NIC Port Details

You can view the NIC port details on the details page of the selected NIC.

About this task

To access the details of the selected NIC, follow these steps:

Procedure

1. Log in to Prism Central.
2. From [Application Switcher Function](#) on page 40, select the **Infrastructure** application and go to **Hardware > NICs** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

The **NIC** page appears, opening the **List** tab by default.

3. Click the name of the NIC to view the details page of the selected NIC.
The details page of the selected NIC provides the details of the NIC in the **Summary** tab, and the alerts for the NIC in the **Alerts** tab.

NIC Details Summary Tab Attributes

The Summary tab of the selected NIC summarizes the details of the NIC port.

Widgets

The **Summary** tab of the selected NIC provides the following widgets:

- **Properties** widget: Displays summary information about the NIC. For more information, see the NIC Properties table.
- **Attached VM NICs** widget: Displays a list of VM NICs attached to the NIC port.

Widget Attributes

The following table describes the fields in the **Properties** widget. A dash (-) in a field indicates there is not enough data to evaluate or a value is not assigned.

Table 131: Properties Widget Attributes

Attribute	Description	Value
Name	Displays the name of the NIC. Clicking the name displays the details page for that NIC.	(NIC name)
MAC Address	Displays the MAC address of the NIC.	(MAC address)
NIC Family	Displays the name of the family of NICs that the NIC belongs to. It denotes the family of the NICs provided by the NIC OEMs. For example, Mellanox Technologies or Intel Corporation. The NIC Family information is derived for processing in the <code>vendor_id:device_id</code> format from the name and model that you select in the NIC Family field when you create a NIC profile. For more information, see Creating a NIC Profile.	NIC Family and OEM information.
Configured Speed	Displays the configured speed which is the speed negotiated with the peer.	xxx (Gbps)
Max Speed	Displays the maximum speed supported by the NIC provider OEM.	xxx (Gbps)
Driver Version	Displays the driver type (for example, virtIO) and the version (for example, 1.0.0).	(Driver Version)
Firmware Version	Displays the firmware on the NIC and its version.	Firmware Version)
NIC Profile	Displays the NIC Profile assigned to the NIC.	
Nutanix Validated Capabilities	Displays the capabilities validated by Nutanix for the NIC.	SRIOV and Network Offload
Enabled Capabilities	Displays the capability (Physical Functions) enabled on the NIC. You can enable only one capability on the NIC, either SRIOV or Network Offload, using the appropriate NIC Profile.	SRIOV and Network Offload

Attribute	Description	Value
Number Of Virtual Functions	Displays the number of Virtual Functions available for the NIC Profile. The number is displayed as the number of VFs available for use out of the number of VFs provided by the NIC hardware. For example, 7 of 8 Available meaning seven VFs are available for use by VMs out of the eight VFs that are provided by that NIC hardware.	(Integer number) of (integer number) Available
Status	Displays the operational status of the NIC.	Up or Down
DHCP	Displays the status of DHCP on the NIC.	Enabled or Disabled
MTU	Displays the MTU applied on the NIC.	(Configured MTU value)
Discovery Protocol	Displays the details of the link layer protocol discovered in the cluster that the NIC belongs to. For example, Link Level Discovery Protocol (LLDP).	(Discovery protocol name)
Ring Size	Displays the RX and TX ring buffer size configured on the host that the NIC belongs to. For information on changing ring buffer sizes on an AHV host, see Change Interface Ring Buffer Sizes .	Rx: xxxx bytes, Tx: xxxx bytes

The following table describes the fields in the **Attached VM NICs** widget. A dash (-) in a field indicates there is not enough data to evaluate or a value is not assigned.

Table 132: Attached VM NICs Widget Attributes

Attribute	Description	Value
VM Name	Displays the name of the VM that the NIC is attached to.	(VM-name)
MAC Address	Displays the MAC Address of the NIC.	(MAC-address)
Subnet	Displays the name of the subnet that the VM NIC is configured on.	(Subnet-name)

Actions

The **Summary** tab of the selected NIC provides the following actions:

- **Assign NIC Profile**

- **Remove NIC Profile**

NIC Details Alerts Tab Attributes

The **Alerts** tab displays a table of alerts.

The **Alerts** tab provides the same features and options as the **Alerts** dashboard, except that it is filtered to display just alerts for this disk. For more information, see [Alerts Summary View \(Prism Central\)](#) in *Prism Central Alerts and Events Reference Guide*.

Alerts Tab Attributes

The **Alerts** tab displays a table of alerts.

The **Alerts** tab provides the same features and options as the **Alerts** dashboard, except that it is filtered to display just alerts for this disk. For more information, see [Alerts Summary View \(Prism Central\)](#) in *Prism Central Alerts and Events Reference Guide*.

Viewing the NIC Profile Details

You can view the NIC profile details on the NIC profile details page.

About this task

To access the details page of the selected NIC profile:

Procedure

1. Log in to Prism Central.
2. From [Application Switcher Function](#) on page 40, select the **Infrastructure** application and go to **Hardware > Clusters** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

The **NICs** page appears, opening the **List** tab by default.

3. Select the **NIC Profile** tab.
4. Click the name of the NIC profile to view the details page of the selected NIC profile. The details page opens in the **Summary** tab by default.

NIC Profile Details Page Attributes

The NIC profile details page displays multiple tabs of information.

The **NIC Profile** page includes the following tabs:

- The **NIC Profile** page includes the following tabs:
 - **Summary**—displayed by default when you open the **NIC Profile** details page. This tab provides information about the NIC profile summarized into widgets.
 - **NIC Ports**—displays the information on the NIC ports that the NIC profile is assigned to. You can assign a single NIC profile to multiple NICs or NIC ports.
 - **Attached VM NICs**—displays the information on the VM NICs attached to the NIC profile.
 - **Alerts**—displays the details of the alerts raised for the NIC port. For information on the **Alerts** tab, see [Alerts Summary View \(Prism Central\)](#) in the *Prism Central Alerts and Events Reference Guide*.

Summary Tab Attributes

The NIC profile Summary tab summarizes the details of the NIC profile in widgets.

The **Summary** tab of the selected NIC profile page provides the **Properties** widget. This widget summarizes the information about the NIC profile.

Table 133: Properties Widget Attributes

Attribute	Description	Value
Name	Displays the name of the NIC profile.	(NIC profile name)
Description	(Optional) Displays the description provided when the NIC profile was created.	(Description string)
NIC Family	Displays the name of the family of NICs that the NIC belongs to. It denotes the family of the NICs provided by the NIC OEMs. For example, Mellanox Technologies or Intel Corporation. The NIC Family information is derived for processing in the <code>vendor_id:device_id</code> format from the name and model that you select in the NIC Family attribute when you create a NIC profile. For more information, see NIC Family on page 483.	NIC Family and OEM information.
Capabilities	Displays the capabilities (physical functions) of the NIC.	SRIOV or DP_OFFLOAD
NIC Ports	Displays the operational status of the NIC ports assigned to the NIC Profile. Multiple NIC ports might be assigned to one NIC Profile. The status separately displays the number of NICs that are not operating properly.	(Status string)
NIC Port Hosts	Displays the details of the AHV hosts that host the NICs assigned to the NIC Profile.	
Virtual Functions	Displays the number of Virtual Functions available for the NIC Profile. The number is displayed as the number of VFs available for use out of the number of VFs provided by the NIC hardware. For example, 7 of 8 Available , meaning seven VFs are available for use by VMs out of the eight VFs that are provided by that NIC hardware.	(Integer number) of (integer number) Available

Attribute	Description	Value
Attached to vNICs	Displays the number of vNICs attached to the NIC Profile.	(Integer number)
Created by	Displays the username of the creator of the NIC profile.	(username string)

Actions

The **Summary** tab of a selected NIC profile page provides the following actions:

- **Update Profile:** Update the NIC Profile by clicking this button.
The attributes in the **Update NIC Profile** pane are identical to the attributes in the **Create NIC Profile** pane. For more information, see Creating a NIC Profile.
- **Delete:** You can delete a NIC Profile by clicking this button. For more information, see Deleting a NIC Profile.

NIC Family

The NIC Family attribute in the **Summary** tab of the NIC profile details page provides information about the family of the NICs provided by the NIC OEMs.

When you select the NIC family in the **NIC Family** field, the selection is translated to information on the NIC hardware in the `vendor_id:device_id` format to apply the firmware, OFED and other information of the NIC hardware. The selected NIC hardware and its `vendor_id:device_id` information is available to the host.

The following sample table provides the NIC vendor and device information for a set of NIC hardware devices:

Table 134: NIC Family Details

NIC Family	Vendor ID	Device ID
Mellanox MT2894 Family [ConnectX-6 Lx]	15b3	101f
Mellanox MT2892 Family [ConnectX-6 Dx]	15b3	101d
Mellanox MT2910 Family [ConnectX-7]	15b3	1021

For example, when you select Mellanox MT2894 Family [ConnectX-6 Lx] in the **NIC Family** field, the **NIC Profile** creation process derives the `vendor_id:device_id` information for this selection as `15b3:101f`.

NIC Ports Tab Attributes

The **NIC Ports** tab displays a table of all the NIC ports that the NIC profile is assigned to.

Table 135: NIC Ports Tab Attributes

Attribute	Description	Value
NIC Port	Displays the name of the NIC port attached to the NIC Profile.	(NIC-port)

Attribute	Description	Value
Host	Displays the name of the host that the NIC is attached to.	(host-name)
Cluster	Displays the name of the cluster that the host belongs to.	(cluster-name)
MAC Address	Displays the MAC address of the NIC.	(MAC address)
Max Speed / Configured Speed (Gbps)	Displays the maximum speed and the configured speed in Gbps in the Max Speed / Configured Speed format.	xxx / xxx
Driver Version	Displays the version of the NIC driver in the driver-name:version format.	(driver-name:version)
Firmware Version	Displays the version of the firmware of the NIC hardware.	(Firmware version)
Compliance Status	Displays the status of the NIC, whether complaint or not compliant.	(Compliant)
Operational Status	Displays the operational status of the NIC port, whether Up or Down .	(Up or Down)

Actions

The **NIC Ports** tab provides the following actions:

- **Manage NIC Port(s)**: Add NIC ports to the NIC profile by clicking this button.
- **Remove NIC Ports**: Remove a previously added NIC port from the NIC profile by clicking this button.

Attached VM NICs Tab

The **Attached VM NICs** tab displays a table of the VM NICs attached to the NIC profile.

Table 136: Attached VM NICs Tab Attributes

Attribute	Description	Value
VM Name	Displays the name of the VM that the attached NIC belongs to.	(VM Name)
MAC Address	Displays the MAC address of the VM NIC attached to the NIC profile.	(MAC Address)
NIC Port	Displays the name of the VM NIC port attached to the NIC profile. For example, eth6.	(NIC name)
Host	Displays the name of the AHV host that hosts the VM.	(host name)

Attribute	Description	Value
Cluster	Displays the name of the cluster that the host and VM belong to.	(cluster name)
VM Power Status	Displays the status of the power to the VM, whether the power is turned On or Off.	On or Off

Creating a NIC Profile

You can create a NIC profile on the **NIC Profile** tab of the **NICs** page.

About this task

To create a NIC profile, follow these steps.

Procedure

1. Log in to Prism Central.
2. From [Application Switcher Function](#) on page 40, select the **Infrastructure** application and go to **Hardware** > **NICs** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

The **NICs** page appears, opening the **List** tab by default.

3. Select the **NIC Profile** tab.
4. Click **Create NIC Profile**.

On the **Create NIC Profile** pane, provide the following information:

- a. **General**: Provide appropriate values in the following fields:

- **Name**: Provide a name to the NIC Profile.
- **Description (Optional)**: Optionally, provide a description of the NIC profile.
- **NIC Family**: Select the NIC family to which the NIC hardware associated with this NIC profile belongs.

- b. **NIC Capability**: Select the appropriate option to partition the physical NIC into virtualized interfaces (virtual functions or VFs) allowing multiple VMs to share the NIC.

- **SR-IOV**: Select this option to create VFs for the SR-IOV passthrough capability.
- **Network Offload**: Select this option to create VFs for processing the packets for the network or datapath offload capability.

5. Perform one of the following actions:

- » Click **Create NIC Profile** to create the NIC profile without assigning it to any NIC ports.

The **Create NIC Profile** pane closes and returns to the **NIC Profile** tab.

You can assign NIC ports to the NIC profile after the NIC profile is created. For steps to assign the NIC ports to NIC profile after creating the NIC profile, see [Adding NIC Ports for a NIC Profile](#) on page 487.

- » Click **Create and Assign to NIC Ports** to go to the **Assign NIC Ports** pane. Go to the next step.

6. On the **Assign NIC Ports** pane, click **Add NICs**.

7. On the **Add NIC Ports** dialog box, follow these steps:

a. In the **Select Host and Cluster** tab, select the required host or hosts or clusters as follows:

- If the required hosts are not displayed in **Viewing (x) hosts**, type the initial letters of the name of the required host or hosts or clusters in the **Filter** field.

Select the required host or hosts in **Viewing (x) hosts**.

- If the required host or hosts or clusters are already displayed in **Viewing (x) hosts**, select the host or hosts or clusters.

b. In the **Select NIC Ports** tab, select the NIC ports as follows:

- » If the required NICs are not displayed in **Viewing (x) NIC Ports**, type the initial letters of the name of the required host or hosts or clusters in the **Filter** field.

You can modify the filter to search for the required NIC ports. Click **Modify Filters** to use **Name** or **MAC Address** as the filter and click **Apply** to search. A successful search displays the required NIC ports in **Viewing (x) NIC Ports**.

Select the required port or ports in **Viewing (x) NIC Ports**.

- » If the required NIC ports are already displayed in **Viewing (x) hosts**, select the NIC ports.

Important: If you select **SR-IOV** as the capability for the NIC profile, ensure that the selected NIC port is not configured on any virtual switch, and is configured on an unmanaged subnet. For information on SR-IOV and its limitations in a Nutanix environment, see [SR-IOV and Network Offload](#).

If you select **Network Offload** as the capability for the NIC profile, ensure that the selected NIC port is configured on a virtual switch. For information on Network Offload and its considerations in a Nutanix environment, see [SR-IOV and Network Offload](#).

c. Click **Add** to add the selected NIC ports.

8. On the **Assign NIC Ports** page, click **Save**

The NIC ports you added are displayed on the **NIC Profile** tab of the **NICs** page and the **NIC Ports** tab of the selected NIC profile page.

Updating a NIC Profile

You can update only the **Name** and **Description** of the NIC Profile.

About this task

The fields in the **Update NIC Profile** pane are identical to the fields in the **Create NIC Profile** pane. No fields, other than **Name** and **Description**, are available for update.

Procedure

1. Log in to Prism Central.
2. From [Application Switcher Function](#) on page 40, select the **Infrastructure** application and go to **Hardware** > **NICs** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

The **NICs** page appears, opening the **List** tab by default.

3. Select the **NIC Profile** tab.

4. Open the **Update NIC Profile** pane in one of the following methods:
 - » Select the NIC profile that you want to add NIC ports to and select **Actions > Update NIC Profile**.
 - » Click the name of the NIC profile that you want to add NIC ports to.On the selected NIC profile's details page, click the **Update Profile** button.
5. Update the necessary values in the respective fields.
6. Click **Save** to update the NIC profile.

Adding NIC Ports for a NIC Profile

You can add NIC ports to a NIC Profile after you create the NIC profile.

About this task

When you add NIC ports to a NIC profile, the NIC ports are assigned to the NIC profile to enable virtual functions (VFs) for the physical functions (PFs) of the NIC hardware that provides the NIC ports.

To add NIC ports to a NIC profile, follow these steps:

Procedure

1. Log in to Prism Central.
2. From [Application Switcher Function](#) on page 40, select the **Infrastructure** application and go to **Hardware > NICs** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

The **NICs** page appears, opening the **List** tab by default.
3. Select the **NIC Profile** tab.
4. Open the **Manage NIC Port(s)** pane in one of the following methods:
 - » Select the NIC profile that you want to add NIC ports to and select **Manage NIC Port(s)** in the **Actions** dropdown list.
 - » Click the name of the NIC profile that you want to add NIC ports to and, on the selected NIC profile details page, click the **NIC Ports > Manage NIC Port(s)** button.
5. On the **Manage NIC Port(s)** pane, click the **Add NICs** button.

The **Add NIC Ports** dialog box appears.
6. Follow the steps, starting with step 7 on page 486 of the [Creating a NIC Profile](#) on page 485 procedure.

Removing NIC Ports from a NIC Profile

You can remove NIC ports from a NIC profile after you create the NIC profile.

About this task

The steps to remove NIC ports assigned to a NIC Profile is the same as the steps provided in [Adding NIC Ports for a NIC Profile](#) on page 487 except step 7 on page 486.

To remove a NIC port from a NIC profile, follow these steps:

Procedure

1. Follow the steps provided in [Adding NIC Ports for a NIC Profile](#) on page 487 until step 7 on page 486.

2. In the **Select NIC Ports** tab, clear the checkboxes of the selected NIC ports as follows:

- » If the required NICs are not displayed in **Viewing (x) NIC Ports**, type the initial letters of the name of the required host or hosts or clusters in the **Filter** field.

You can modify the filter to search for the required NIC ports. Click **Modify Filters** to use **Name** or **MAC Address** as the filter and click **Apply** to search. A successful search displays the required NIC ports in **Viewing (x) NIC Ports**.

Clear the checkboxes for NIC ports in **Viewing (x) NIC Ports**.

- » If the required NIC ports are already displayed in **Viewing (x) hosts**, clear the checkboxes of the NIC ports.

Clear the checkboxes for NIC ports in **Viewing (x) NIC Ports**.

Important: If you select **SR-IOV** as the capability for the NIC profile, ensure that the selected NIC port is not configured on any virtual switch, and is configured on an unmanaged subnet. For information on SR-IOV and its limitations in a Nutanix environment, see [SR-IOV and Network Offload](#).

If you select **Network Offload** as the capability for the NIC profile, ensure that the selected NIC port is configured on a virtual switch. For information on Network Offload and its considerations in a Nutanix environment, see [SR-IOV and Network Offload](#).

3. Click **Add** to add the selected NIC ports.

4. On the **Assign NIC Ports** page, click **Save**

What to do next

Check the **NIC Profile** tab to verify that the NIC ports are removed.

Updating a NIC Profile for NIC Replacement

You must follow this procedure when you replace a physical NIC.

About this task

Before you replace a physical NIC on a host, follow these steps:

Procedure

1. Remove the NIC ports of the physical NIC from the NIC profile.

For more information, see [Removing NIC Ports from a NIC Profile](#) on page 487.

2. Replace the physical NIC.

For more information, see the [NIC Replacement documentation](#).

3. Add the NIC ports on the new physical NIC to the NIC profile.

For more information, see [Adding NIC Ports for a NIC Profile](#) on page 487.

4. Verify that the NIC ports are displayed on the **NIC profile** tab.

What to do next

If the NIC ports on the new physical NIC are not displayed, see [KB-3261](#).

Updating a Ubuntu VM

You must update a Ubuntu Guest VM to use its NIC that is assigned to an SR-IOV or Network Offload NIC profile.

About this task

When you assign an SR-IOV or Network Offload NIC profile to a NIC that is associated with a guest VM with the Ubuntu operating system, the VM operating system displays the status of the associated virtual function (VF) interface as **Down** by default.

For more information on Virtual Functions, see [SR-IOV and Network Offload](#).

To bring the interface up, follow these steps:

Procedure

1. Log on to the terminal of the Ubuntu guest VM.
2. Check the name of the interface that is displayed as down.
3. Bring the interface up.

```
$ ifconfig <interface-name> up
```

4. Create a text file with the interface name, such as interface-name.network, in the /etc/systemd/network/ folder with the following content:

```
$ cat /etc/systemd/network/<file-name>.network
[Match]
Driver=mlx5_core
[Network]
DHCP=ipv4
LinkLocalAddressing=ipv6
```

What to do next

Verify if the VM operating system displays the status of the associated virtual function (VF) interface as **Up**

Deleting a NIC Profile

You can delete a NIC profile after meeting a condition.

Before you begin

Remove the NIC ports from the NIC profile.

About this task

To delete a NIC profile, follow these steps:

Procedure

1. Log in to Prism Central.
2. From [Application Switcher Function](#) on page 40, select the **Infrastructure** application and go to **Hardware > NICs** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.
The **NICs** page appears, opening the **List** tab by default.
3. Select the **NIC Profile** tab.

4. Open the **Delete** *nic-profile-name* pane in one of the following methods:
 - » Select the NIC profile that you want to add NIC ports to and select **Actions** > **Delete NIC Profile**.
 - » Click the name of the NIC profile that you want to add NIC ports to.
On the selected NIC profile's details page, click the **Delete** button.
5. Click **Delete** in the **Delete** *nic-profile-name* dialog box.

Attaching an SR-IOV NIC to a VM

You must attach a NIC with SR-IOV NIC profile to a VM to use the SR-IOV capability on the physical NIC.

About this task

To attach a NIC with SR-IOV NIC profile to a VM, follow these steps:

Procedure

1. Log in to Prism Central.
2. From [Application Switcher Function](#) on page 40, select the **Infrastructure** application and go to **Hardware** > **VMs** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.
The **VMs** page appears.
3. Select the **Update** > **Network** using any of the following methods:
 - » Select the checkbox of the target VM in the **List** tab, and choose **Update** > **Network** from the **Actions** dropdown menu.
 - » Right-click on the target VM in the **List** tab, and select **Update** > **Network** from the dropdown menu.
 - » Select the required action in the **Summary** page of an individual VM. For information about how to access the **Summary** page of an individual VM, see [Viewing VM Details](#) on page 113.The **NICs** tab of the VM appears.
4. Click **More Network Options** > **Add SR-IOV**.
The **Create SR-IOV NIC** dialog box appears.
5. Select the required SR-IOV NIC profile in the **Select NIC Profile & Family** pane
6. (Optional) Add a VLAN ID in the **Associate to VLAN ID (Optional)** field to associate it with the NIC.
7. Click **Save**.
The SR-IOV NIC appears in the **SR-IOV NIC(s)** pane of the **NICs** tab.

Attaching a Network Offload NIC Profile to a VM

You must attach a Network Offload NIC profile to NIC attached to a VM to offload the VM traffic flows.

About this task

To attach a NIC Profile to a VM, follow these steps:

Procedure

1. Log in to Prism Central.

2. From [Application Switcher Function](#) on page 40, select the **Infrastructure** application and go to **Hardware > VMs** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

The **VMs** page appears.

3. Select the **Update > Network** using any of the following methods:
 - » Select the checkbox of the target VM in the **List** tab, and choose **Update > Network** from the **Actions** dropdown menu.
 - » Right-click on the target VM in the List tab, and select **Update > Network** from the dropdown menu.
 - » Select the required action in the **Summary** page of an individual VM. For information about how to access the **Summary** page of an individual VM, see [Viewing VM Details](#) on page 113.

The **NICs** tab of the VM appears.

4. Click **Attach to Subnet**.

The **Create Normal NIC** dialog box appears.

5. In the **Network Offload** section, select **Attach NIC Profile**.

The **Attach to physical NICs** dialog box appears. The dialog box provides a table of available NIC profiles that you can attach.

Note: You can attach only one NIC profile to a VM.

- a. Select the checkbox for the NIC identified for attachment.
 - b. Select the checkbox for the NIC identified for attachment.
6. Click **Create**.

ACTIVITY ENTITIES – ALERT AND EVENT MONITORING

In Prism Central, alerts and events monitoring provide centralized monitoring for all the registered clusters.

You can use these functionalities to observe and manage system notifications, faults, and performance-related statistics from a single console. You can monitor alerts and events under the **Activity** menu. Some entities, such as SQL Server instances, virtual switches, or vCenter, have their own **Alerts** and **Events** tabs that display a filtered list of notifications relevant only to that specific entity.

For more information about Alerts and Events in Prism Central, see [Prism Central Alerts and Events Reference Guide](#).

ACTIVITY ENTITIES – TASKS AND AUDITS

You can access dashboards for the following activity monitors from the **Activity** entity of the **Infrastructure** application. For information about how to access the entity items available in **Activity** entity, see [Application-specific Navigation Bar](#) on page 61.

- Audits (see [Audits Summary View](#) on page 493)
- Tasks (see [Tasks View](#) on page 496)

Audits Summary View

The **Audits** summary view provides a dashboard of the actions performed across the registered clusters.

Note:

- The retention period for audit entries is four weeks by default.
- Audit information appears only for those registered clusters running AOS 5.10 or later.
- Audit logs with default values are generated when updates to VMs are initiated, either by Prism Central **Self Service** users or by using Nutanix v3 API calls for the first time.
- For all Prism Central related events, see [Prism Central Logs](#) in *Prism Central Alerts and Events Reference Guide*. Additionally, you can refer to the [Audit Log Events](#) in *Prism Central Alerts and Events Reference Guide* for the complete list of captured audit events.

To access the **Audits** summary view, perform the following steps:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Activity** > **Audits** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

The system displays the **Audits** dashboard.

The following table describes the fields that appear in the audits list. A dash (-) is displayed in a field when a value is not available or applicable.

Table 137: Recovery Plans Fields

Parameter	Description	Values
Action Description	Describes the action taken such as "deleted VM <i>vm-name</i> " or "added disk <i>disk-name</i> "	(action description)
User Name	Displays the name of the user who requested the action.	(user name)
Entity Affected	Displays the entity name. Click the name to go to the details page for that entity.	(entity name)
Entity Type	Displays the entity type such as VM or host.	(entity type)

Parameter	Description	Values
Operation Type	Displays the type of operation that took place. The possible operation types depend on the entity type and can include create, update, delete, and power state change.	(operation type)
Request Time	Displays the time the user requested the action.	(time and date)
Cluster	Displays the name of the cluster in which the action took place. Click the cluster name to display the details page for that cluster.	(cluster name)

To filter the list, click **Modify Filters** (upper right). This displays a pane for selecting filter values. The following table describes the filter options available. You can apply multiple filters together. For example, if you want to filter by a particular VM name, you can filter by *entity type* = *VM* and then use the *name contains* filter.

Table 138: Filter Pane Fields

Field	Description	Values
User IP	Enter a user IP address and then click Add to filter for actions requested by that user. You can add multiple user IP addresses.	(IP address)
Cluster	Enter a cluster name in the field to filter for actions in the cluster.	(cluster name)
Entity Type	Check the boxes of one or more entities to filter for actions on those entity types.	VM, Storage Container, Catalog Item, Image, Cluster, Host, Disk, GPU, Security Policy, NGT Policy, Project, Role, User, Category, Availability Zone, Protection Policy, Recovery Plan, Recoverable Entity, Report
Operation Type	Check the boxes of one or more operations to filter on those operations.	Create, Update, Delete, Power State Change
Request Time	Check an interval box to filter for actions that were requested during that time period. For the custom interval option (from xxx to xxx), click in each field and select a date from the pop-up calendar.	Last 1 hour, Last 24 hours, Last week, From xxx to xxx
Cluster	Enter a cluster name in the field to filter for actions in the cluster.	(cluster name)
Note: Entering <i>Prism Central</i> in the cluster filter criteria does not populate any result.		
User IP	Enter a user IP address and then click the Add to filter for actions requested by that user. You can add multiple user IP addresses.	(IP address)

Audit Details View

The details page of an individual action provides the detailed information about the action.

To access the details page of an action:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Activity > Audits** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.
3. Click the target `<action_name>` to view the details page of an individual action.

Note: Replace `<action_name>` with the actual action description at your site.

The details page includes the following:

- The action description (upper left). You can switch from one action details page to another by selecting from the dropdown list.
- A section on the left that displays summary information about the action (see following table).
- A section on the right that displays a table of information specific to that action. The table shows the attributes that were changed during the action and the current value that is set for the attribute. The attributes vary depending on the specifics of the action (see following examples).

The following table describes the action summary information fields. A dash (-) is displayed in a field when a value is not available or applicable.

Table 139: Action Summary Fields

Parameter	Description	Values
Action Description	Displays the action description.	(description)
User Name	Displays the name of the user who requested the action.	(user name)
Target Entity	Displays the name of the entity that was the action target. Click the entity name to display the details for that entity.	(entity name)
Affected Entities	Displays the names of the entities that were affected by the action. Click an entity name to displays the details page for that entity.	(one or more entity names)
Operation Type	Displays the type of operation that took place. The possible operation types depend on the entity type and can include create, update, delete, and power state change.	(operation type)
Request Time	Displays the time the user requested the action.	(time and date)
User IP	Displays the IP address of the user.	(IP address)
Cluster	Displays the name of the cluster in which the action took place. Click the cluster name to display the details page for that cluster.	(cluster name)
Status	Displays the status of the action.	Succeeded, Failed

Tasks View

The **Tasks** page provides a dashboard that displays information about all tasks across the registered clusters. You can view a task if you have the view permission for any one of the clusters affected by the task.

Note: Prism Central version pc.2024.3 introduces an enhanced **Tasks** dashboard. You can access this enhanced dashboard only if all the connected Prism Element clusters are upgraded to version 7.0.

To access the **Tasks** dashboard:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Activity > Tasks** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.
3. The system displays the **Tasks** dashboard.

An icon appears in the main menu when there are current tasks (running or completed within the last 48 hours). The icon appears blue when a task runs normally, yellow when it generates a warning, or red when it fails. Clicking the icon displays a dropdown menu of current tasks.

You can see the tasks dashboard by doing one of the following:

- Clicking the **View All Tasks** link at the bottom of the current tasks dropdown menu.
- Clicking the **View All Tasks** link in the Tasks widget in the main dashboard.
- Selecting **Activity > Tasks** from the **Navigation Bar**.

The following table describes the fields on the tasks page. Each task appears on the list for a minimum of one hour after completion. The duration for which a task stays on the list varies based on several factors. However, tasks may be removed from the list sooner when new tasks are added, especially during periods of high activity. In some instances, a task may remain visible for longer than two weeks if it is the last task associated with a specific component.

Table 140: Tasks List Fields

Parameter	Description	Values
Task	Specifies which type of operation the task is performing. The number of child tasks are also displayed in this column, if applicable.	(name of operation such as garbage collection of logs)
Entity Affected	Lists the entity name or entity UUID. If this is a live link, click it to either see more details or go to the entity details page. If multiple entities are affected, only upto five entities are displayed under this column.	(name or description)
Cluster	Specifies the cluster name(s) on which the task is running. A dash (-) appears when the task applies to Prism Central rather than one of the registered clusters.	(cluster names)

Parameter	Description	Values
Status	Indicates the status of the task: • Canceled: The task has been canceled by a user. • Canceling: The task cancellation has been initiated by a user, but cancellation is yet to be completed. • Failed: The task has failed. • Queued: The task has been queued by the system. • Running: The task is running. The progress percentage is displayed for the running tasks. • Succeeded: The task has succeeded. • Suspended: The task has been suspended by the system.	Cancelled, Canceling, Failed, Queued, Running, Succeeded, Suspended
Initiator	Indicates the user who initiated the task.	(user name)
Start Time	Displays when the task began.	[date],[time]
Duration	Displays how long the task has been running or took to complete.	xx [seconds minutes hours days]

When you select a task, a side panel appears at the right side, displaying the details of the task, such as:

- Overall status of the task
- Names of all the entities affected (a maximum of 300 entities can be displayed in the side pane)
- Cluster
- Initiator (user who initiated the task)
- Duration of the task
- Details of the child tasks
- Brief description of the sub-steps included in the selected task (if applicable)
- Errors or warnings related to the task (if applicable)

You can do the following in the Tasks dashboard:

- Filter the list by entering a criteria in the filter field at the top of the page.
- Save your favorite filters.
- Filter the list by clicking **Modify Filters**. This displays a filter pane on the right of the screen. For more information, see [Filters Pane - Tasks Page](#).
- Click the X icon to close the filter pane.

- Export the entire task list or a filtered list in CSV format by clicking **Export**.

Filters Pane - Tasks Page

The following table describes the fields available in the **Filters** pane:

Table 141: Filter Pane Fields

Parameter	Description	Values
TASK	Filters based on task name. Select a condition from the dropdown menu (Contains, Equal to, Not Equal to, Doesn't contain, Starts with, Ends with), and enter a string in the field.	(task name)
ENTITIES AFFECTED	Filters based on the entities affected. Select a condition from the dropdown menu (Contains, Equal to, Not Equal to, Doesn't contain, Starts with, Ends with), and enter a string in the field.	(entities name string)
INITIATOR	Filters based on the initiator of the task. Select a condition from the dropdown menu (Contains, Equal to, Not Equal to, Doesn't contain, Starts with, Ends with), and enter a string in the field.	(initiator string)
CLUSTER	Filters based on the cluster name. Select a cluster name from the dropdown list to filter the result.	(cluster name)
STATUS	Filters based on the status of the task. Select one or more status from the list to filter the result.	(Canceled, Canceling, Failed, Queued, Running, Suspended, or Succeeded)
TIME RANGE	Filters based on the time range. Select date and time for Start and End fields.	(date and time)

OPERATIONS ENTITIES

Starting from Prism Central version 2023.3, NCM Intelligent Operations (formerly AIOps) now has a new user interface. The following Operations Entities have been moved from Prism Central Infrastructure app to the new user interface created for NCM Intelligent Operations:

- Analysis
- App Discovery
- Monitoring Configurations
- Operations Policies
- Planning
- Playbooks
- Reports
- Settings and Configurations

The documentation for the above-mentioned entities has been moved to the newly created *Intelligent Operations Guide*. For more information, see [Intelligent Operations Guide](#).

ADMINISTRATION ENTITIES

You can manage the following entity items from the **Administration** entity of the **Infrastructure** application:

- Users, Roles, and Projects. For more information on how to manage users, roles, and projects in Prism Central, see [Prism Central Admin Center Guide](#).
- LCM. For information on LCM, see [LCM Documentation](#).
- Availability Zones. For information on Availability Zones in Prism Central, see [Availability Zones](#) on page 500.

Note: Starting with Prism Central version pc.7.3, you can create and manage categories from Admin Center. For information on how to manage categories from Prism Central Admin Center, see [Category Management](#) in the *Prism Central Admin Center Guide*.

For information on how to access the entities items available in **Administration** entity, see [Application-specific Navigation Bar](#) on page 61.

Availability Zones

Nutanix Disaster Recovery works with sets of physically isolated locations called availability zones (AZs). For information about Availability Zones in Prism Central, see [Nutanix Disaster Recovery Guide](#) .

POLICIES IN INFRASTRUCTURE

This section describes the policies framework in Prism Central.

Policies are used to define the systematic configuration guidelines for the following requirements:

- Controlled management of **VMs, Images, Storage, and Network & Security** entities. For more information, see the following sections:
 - [VM Policy Management](#) on page 501 for VM-related policies.
 - [Image Policy Management](#) on page 527 for image-related policies.
 - [Storage Policy Management](#) on page 545 for storage-related policies.
 - [Security Policy Management](#) on page 540 for policies related to network security framework setup.
- System services customization. For more information, see [Operations Policy Management](#) in *Intelligent Operations Guide*.
- Disaster Recovery setup. For information about how to create the protection policies in Disaster Recovery setup, see [Nutanix Disaster Recovery Guide](#).

VM Policy Management

In Prism Central, you can create and manage policies for VMs to control the placement and behaviour of VMs in a cluster.

You can create and manage the following types of policies for the VMs:

- VM-Host Affinity Policies (see [VM-Host Affinity Policies Defined in Prism Central](#) on page 501)
- VM-VM Anti-Affinity Policies (see [VM-VM Anti-Affinity Policies Defined in Prism Central](#) on page 512)
- NGT Policies (see [NGT Policies](#) on page 522)

VM-Host Affinity Policies Defined in Prism Central

In Prism Central, you can define the category-based VM-Host affinity policies, where a set of VMs can be affinity (correlated) to run only on a particular set of hosts. Category-based affinity policy enables you to easily manage affinities for a large number of VMs. In case of any changes to the affinity hosts, you only need to update the category of the host, and it updates the affinity policy for all the affected VMs.

This policy checks and hard enforces where a VM can be hosted when you start or migrate the VM. If there are no resources available on any of the affinity hosts, the VM is not started.

Note:

If you create a VM-Host affinity policy for a VM that is configured for asynchronous replication, you must create similar categories and corresponding policies on the remote site as well. If you define similar categories and policies on the remote site, the system applies the affinity policies when the VMs are migrated to the remote site.

For information on how the category-based VM-Host affinity policies are handled during disaster recovery, see [Affinity Policies Handling - PC Based DR Solution with Physically Isolated Locations](#).

VM-Host Affinity Policies Summary View

The VM-Host affinity policies summary view enables you to access a list of all the user-defined affinity policies across registered clusters.

To access the summary view of all VM-Host affinity policies:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Compute > VMs** from the **Navigation Bar**. For information on the **Navigation Bar**, see [Application-specific Navigation Bar](#).

The system displays the **List** tab by default.

3. Select **VM-Host Affinity Policies** in the **Policies** tab dropdown menu.

The system displays a summary view of VM-Host affinity policies across all registered clusters.




Figure 146: Summary View - All Affinity Policies

Table 142: VM-Host Affinity Policies Page Field Description

Parameter	Description	Values
Name	Displays the VM-Host affinity policy name.	(name)
VMs	Displays the count of VMs associated with the selected VM-Host affinity policy.	(number of VMs)
Hosts	Displays the count of hosts associated with the selected VM-Host affinity policy.	(number of hosts)
VM Compliance Status	Displays the compliance status of the VMs associated with this policy. If the policy is being applied and the compliance status is not yet known, the status is displayed as Pending. If a VM is part of multiple VM-Host affinity policies, the oldest policy is applied on the VM. For rest of the policies, the VM is displayed as non-compliant.	(number of VMs Compliant/Non Compliant/Pending)
Modified By	Displays the name of the user who modified the selected VM-Host affinity policy last time.	(user)
Last Modified	Displays the date and time when the selected VM-Host affinity policy is modified last time.	(date & time)

You can perform the following actions for the affinity policies in the **VM-Host Affinity Policies** summary view:

- Access the detailed information about an individual VM-Host affinity policy. For more information, see [VM-Host Affinity Policies Details View](#) on page 503.
- Create a VM-Host affinity policy. For more information, see [Creating a VM-Host Affinity Policy](#) on page 509.

- Use the **Actions** dropdown menu to update, delete, or re-enforce VM-Host affinity policy. For more information, see [Managing VM-Host Affinity Policies](#) on page 510.

VM-Host Affinity Policies Details View

The VM-Host affinity policy details view allows you access the detailed information about an individual VM-Host affinity policy. It includes four tabs: **Summary**, **Categories**, **Entities** and **Audit**.

To access the details view of an individual VM-Host affinity policy, perform the following steps:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Compute > VMs** from the **Navigation Bar**. For information on the **Navigation Bar**, see [Application-specific Navigation Bar](#).

The system displays the **List** tab by default.

3. Select **VM-Host Affinity Policies** in the **Policies** tab.

The system displays a summary view of VM-Host affinity policies across all registered clusters.

4. Click the target `<Affinity_Policy_Name>` to view the **Summary** tab of an individual VM-Host affinity policy.

Note: Replace `<VM-Host_Affinity_Policy_Name>` with the actual VM-Host affinity policy name at your site.




Figure 147: VM-Host Affinity Policies Details view

The **Summary** tab of an individual VM-Host affinity policy provides the following widgets:

- **Overview** - Displays an overview information about the VM-Host affinity policy.
- **Associations**- Displays the associated VM and Host entities and their categories.
- **Compliance Status** - Displays the compliance status of the VMs associated with this policy. If the policy is being applied and the compliance status is not yet known, the status is displayed as Pending. If a VM is part of multiple VM-Host affinity policies, the oldest policy is applied on the VM. For rest of the policies, the VM is displayed as non-compliant.

For information on the fields available in the widgets, see [VM-Host Affinity Policy Widgets Field Details](#) on page 503.

`<Action>` available above the widgets. Click the appropriate `<Action>` to run that administrative action on the VM-Host affinity policy. For more information on how to perform any `<Action>`, see [Managing VM-Host Affinity Policies](#) on page 510.

VM-Host Affinity Policy Widgets Field Details

The following table describes the fields in the **Overview**, **Associations**, and **Compliance Status** widgets.

Table 143: VM-Host Affinity Policy Widgets Field Description

Field	Description	Values
Overview widget		
Description	Displays the VM-Host affinity policy description specified while creating the VM-Host affinity policy.	(description)
Last Modified By	Displays the name of the user who modified the VM-Host affinity policy last time.	(user)
Last Modified	Displays the date and time when the policy is modified last time.	(date & time)
Created By	Displays the name of the user who created the VM-Host affinity policy	(user)
Create time	Displays the date and time when the VM-Host affinity policy is created.	(date & time)
Associations widget		
VM Categories	Displays the count of VM categories mapped to the selected VM-Host affinity policy.	(number of VM categories)
VMs	Displays the count of VMs associated with the selected VM-Host affinity policy.	(number of VMs)
Host Categories	Displays the count of host categories mapped to the selected VM-Host affinity policy.	(number of host categories)
Hosts	Displays the count of hosts associated with the selected VM-Host affinity policy.	(number of hosts)
Compliance Status widget		
Total VMs	Displays the total number of VMs in the selected VM-Host affinity policy.	(total number of VMs)
Compliance	Displays the compliance status in percentage.	(percentage of complaint VMs)
Compliant, In Progress, or Non compliant	Displays the number of VMs that are compliant and non-compliant with the VM-Host affinity policy. If the compliance check is in progress, the system displays the compliance status as <i>In progress</i> .	Compliant, In Progress, or Non compliant

Categories Tab

The **Categories** tab displays the VMs and hosts categories information.




Figure 148: VM-Host Affinity Policies Categories (VMs)




Figure 149: VM-Host Affinity Policies Categories (Hosts)

The following table describes the information displayed in the **Categories** tab

Table 144: Categories Tab Field Description

Field	Description	Values
VMs tab		
Name	Displays the VM category name and its value.	(VM Category Name: Value)
VMs	Displays the count of VMs associated with the selected VM-Host affinity policy.	(number of VMs)
Status	Displays the compliance status.	(number of VMs Compliant Non-compliant Pending)
Host tab		
Name	Displays the host category name and its value.	(host category Name: Value)
Hosts	Displays the count of hosts associated with the selected VM-Host affinity policy.	(number of hosts)

You can export the table that contains the list of VM categories and host categories and their information to a file in a CSV format. For more information about **Export** option, see [Export](#) on page 54.

Entities Tab

The **Entities** tab displays the VMs and hosts entity-related information.




Figure 150: VM-Host Affinity Policies Entities (VMs)




Figure 151: Affinity Policies Entities (Hosts)

The following table describes the information displayed in the **Entities** tab

Table 145: Entities Tab Field Description

Field	Description	Values
VMs tab		
Name	Displays the VM name associated with the selected VM-Host affinity policy.	(VM name)
Host	Displays the host name associated with the selected VM-Host affinity policy.	(host name)
Cluster	Displays the name of the cluster (on which the host resides) associated with the selected VM-Host affinity policy.	(cluster name)
Associated via Categories	Displays the VM category name and its value.	(VM category name: Value)
Compliance Status	Displays the VM compliance status with the selected VM-Host affinity policy.	Compliant Non-compliant Pending
Host tab		
Host	Displays the name of the host associated with the selected VM-Host affinity policy.	(host name)

Field	Description	Values
Cluster	Displays the name of the cluster (on which the host resides) associated with the selected VM-Host affinity policy.	(cluster name)
Associated via Categories	Displays the host Category name and its value.	(host category name: Value)

You can export the table that contains the list of VM entities and host entities related information to a file in a CSV format. For more information about **Export** option, see [Export](#) on page 54.

Audits Tab

The **Audits** tab displays the user action-related information for the VM-Host affinity policies.




Figure 152: Affinity Policies Audits

The following table describes the information displayed in the **Audits** tab

Table 146: Audits Tab -Field Description

Field	Description	Values
Action Description	Displays the user-action for VM-Host affinity policies.	(action description)
User Name	Displays the name of the user who performed the action.	(host name)
Operation Type	Displays the type of operation performed by the user for VM-Host affinity policies. For example, Create	(cluster name)
Request Time	Displays the date and time of the user action.	(date and time)

You can export the table that contains the list of user actions related information to a file in a CSV format. For more information about **Export** option, see [Export](#) on page 54.

Limitations of VM-Host Affinity Policies

VM-Host affinity policies created in Prism Central have the following limitations:

- The minimum supported versions for VM-Host affinity policies are version 6.1 for Prism Element and version 2022.1 for Prism Central.

- Host category attach or detach takes around five minutes to get reflected in the applicable VM-Host affinity policies.

When you assign a category to a host and map the host category to the VM-Host affinity policy, you can observe that the host count gets updated immediately in the [Entities](#) tab. However, the system takes approximately 5 minutes to update the host count in [VM-Host Affinity Policies Summary View](#) on page 501.

The delay in host count update is due to the usage of different APIs to derive the host count in [Entities](#) tab and [VM-Host Affinity Policies Summary View](#) on page 501.

For information about how to create a category, see [Creating a Category](#).

For information on how to assign a category to host, see [Associating Hosts with Categories](#) on page 509

For information on how to create the VM-Host affinity policy and map the host category to the VM-Host affinity policy, see [Creating a VM-Host Affinity Policy](#) on page 509.

VM-Host Affinity Policy Configuration Workflow

About this task

To set up VM-Host affinity policy, perform the following steps:

Procedure

1. Create categories for the following entities:

- a. VMs

- b. Hosts

For information on how to create a category, see [Creating a Category](#).

2. Apply the VM categories to the VMs and host categories to the hosts.

For information on how to associate a category to the VMs, see [Associating VMs with Categories](#). For information on associating categories with hosts, see [Associating hosts with Categories](#).

3. Create the VM-Host affinity policy. For more information, see [Creating an Affinity Policy](#).

Associating VMs with Categories

About this task

To associate categories to VMs, perform the following steps:

Procedure

1. Log in to Prism Central.

2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Compute > VMs** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#).

The system displays the **List** tab by default.

3. Select the target VMs checkboxes that you want to associate with a category, and choose **Manage Categories** from the **Actions** dropdown menu.

The system displays the **Manage VM Categories** window.

4. Type the name of the category in **Set Categories** field.

The system displays the list of matching categories based on the typed entry.

5. Use the [Add icon](#) and [Remove icon](#) to add and remove the required categories.

Note: If the VM category you selected is already part of any affinity policy, the system displays the associated affinity policy details.

6. Click **Save**.

Associating Hosts with Categories

About this task

To associate categories with hosts, perform the following steps:

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Hardware > Hosts** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#).
The system displays the **List** tab by default.
3. Select the target host checkboxes that you want to associate with a category, and choose **Manage Categories** from the **Actions** dropdown menu.
The system displays the **Manage Host Categories** window.
4. Type the name of the category in **Set Categories** field.
The system displays the list of matching categories based on the typed entry.
5. Use the [Add icon](#) and [Remove icon](#) to add and remove the required categories.
6. Click **Save**.

Creating a VM-Host Affinity Policy

About this task

This section describes how to create a VM-Host affinity policy in Prism Central.

Before you begin

Ensure that the following prerequisites are met before you create the VM-Host affinity policy:

- VM and Host categories are created. For information on how to create a category, see [Creating a Category](#).
- VMs are associated with VM categories, and hosts with host categories. For more information, see [Associating VMs with Categories](#) on page 508 and [Associating Hosts with Categories](#) on page 509.

Note:

- The system also allows you to associate VMs with VM category and hosts with host category after creation of VM-Host affinity policy.
- If you have configured any legacy affinity policy (non-category-based affinity policy) associated with the VMs, you must first remove those legacy affinity policies to allow the creation of category-based VM-Host affinity policies associated with the same VMs.

Procedure

To create the VM-Host affinity policy, perform the following steps:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Compute > VMs** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#).
The system displays the **List** tab by default.
3. Select **VM-Host Affinity Policies** in the **Policies** tab dropdown menu, and click **Create VM-Host Affinity Policy**.
The system displays the **Create VM-Host Affinity Policy** window.
4. Specify the following information in the **Create VM-Host Affinity Policy** window:
 - **Name:** Enter the VM-Host affinity policy name.
 - (optional) **Description:** Enter the description for the VM-Host affinity policy.
 - **VM Categories:** Type the name of the VM category. The system displays the list of matching categories based on the typed entry. Use the [Add icon](#) and [Remove icon](#) to add and remove the required categories.
 - **Host Categories:** Type the name of the host category. The system displays the list of matching categories based on the typed entry. Use the [Add icon](#) and [Remove icon](#) to add and remove the required categories.
5. Click **Create**.

Managing VM-Host Affinity Policies

This section describes how to manage the existing VM-Host affinity policies in Prism Central.

About this task

You can perform the following actions to manage the existing VM-Host affinity policies in Prism Central:

- Update a VM-Host affinity policy.
- Delete a VM-Host affinity policy.
- Re-enforce a VM-Host affinity policy.

Procedure

To manage the VM-Host affinity policies, perform the following steps:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Compute > VMs** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#).
The system displays the **List** tab by default.
3. Select **VM-Host Affinity Policies** in the **Policies** tab dropdown menu.
The system displays the affinity policies created across registered clusters.
4. Select the target VM-Host affinity policy checkbox, and choose **Update** from the **Actions** dropdown menu.

5. Update the fields in the **Update VM-Host Affinity Policy** window as per your requirement, and click **Save**. For information on the fields available in **Update VM-Host Affinity Policy** window, see [Creating a VM-Host Affinity Policy](#) on page 509.

To delete the VM-Host affinity policy, select the target VM-Host affinity policy checkbox and choose **Delete** from the **Actions** dropdown menu. The system prompts you to confirm the delete action. Click **Delete** to confirm the delete affinity policy action.

To re-enforce the VM-Host affinity policy, select the target VM-Host affinity policy checkbox and choose **Re-enforce** from the **Actions** dropdown menu. The system prompts you to confirm the action. Click **Re-enforce** to apply the updated VM-Host affinity policy.

Role Based Access Control for VM-Host Affinity Policies

This section describes the Role Based Access Control (RBAC) requirements and limitations for the VM-Host affinity policies. For information on RBAC, see [Controlling User Access \(RBAC\)](#) in *Security Guide*.

Role Based Access Control Requirements for VM-Host Affinity Policies

Before you allocate the [VM-Host Affinity Policy Related Permissions](#) on page 511, ensure that the following basic permissions are allocated to the admin user or a custom user to access the VM-Host affinity policies framework in Prism Central:

- View VM
- View category

VM-Host Affinity Policy Related Permissions

In addition to the basic permissions to view VM and categories as specified in [Role Based Access Control Requirements for VM-Host Affinity Policies](#) on page 511, you can also assign the following permissions to the admin or custom user to manage the VM-Host affinity policies in Prism Central:

Table 147: VM-Host Affinity Policy Related Permissions

Permission	Description
Create_VM_Host_Affinity_Policy	Admin or custom user can create the VM-Host affinity Policies.
View_VM_Host_Affinity_Policy	Admin or custom user can view the VM-Host affinity policies.
Update_VM_Host_Affinity_Policy	Admin or custom user can update the existing VM-Host affinity policies.
Reenforce_VM_Host_Affinity_Policy	Admin or custom user can re-enforce the existing VM-Host affinity policies.
Delete_VM_Host_Affinity_Policy	Admin or custom user can delete the VM-Host affinity policies.

These permissions also facilitate the admin or custom user to manage the legacy affinity policies (affinity policies that are defined in the Prism Element) from Prism Central.

Role-Based Access Control Limitations for VM-Host Affinity Policies

The following table describes the limitations that apply to RBAC for the VM-Host affinity policies:

Table 148: RBAC Limitations for VM-Host Affinity Policies

Permission Assigned to the Admin or Custom user	Admin or Custom User Action	System Behavior	Limitation
Create_VM_Host_Affinity_Policy Update_VM_Host_Affinity_Policy	Admin or custom user creates or updates the VM-Host affinity policy.	The system applies the VM-Host affinity policy to all the VMs associated with the mapped VM categories.	The system does not check whether the admin or the custom user has permission to access the VMs or hosts mapped in the VM-Host affinity policy.
Update VM Categories	Admin or custom user updates a VM category (for example, attaches a VM to a VM category).	The system applies the VM-Host affinity policy to which the VM category is assigned.	The system does not check whether the admin or the custom user has permission to create or update the VM-Host affinity policy.
Update Host Categories	Admin or custom user updates a Host category (for example, attaches a host to a Host category).	The system applies the affinity policy to which the host category is assigned.	The system does not check whether the admin or the custom user has permission to create or update the VM-Host affinity policy.

VM-VM Anti-Affinity Policies Defined in Prism Central

This topic describes the VM-VM anti-affinity policies feature available in Prism Central.

In Prism Central, you can define category-based VM-VM anti-affinity policies that run a set of VMs on different hosts, preventing single host failure from impacting all the VMs. The category-based anti-affinity policy allows you easily to manage anti-affinities for a large number of VMs.

The VM-VM anti-affinity policy is soft enforced. The system does not block any operations, such as VM creation, VM updation, upgrades, host maintenance mode, or manual live migration of the VM, even if such operations result in a policy violation.

Disaster Recovery Considerations

If you create a VM-VM anti-affinity policy for a VM that is configured for replication, and you want those policies to be honored on the remote site, then you must create similar categories and corresponding policies on the remote site as well. When you define similar categories and policies on the remote site, the system applies the VM-VM anti-affinity policies to the VMs that you migrate to the remote site. For more information, see [Management of VM-VM Anti-Affinity Policies for PC-Based Disaster Recovery Solution](#) on page 521.

VM-VM Anti-Affinity Policies Summary View

Use VM-VM anti-affinity policies summary view to access a list of all user-defined VM-VM anti-affinity policies.

To access the summary view of all the VM-VM anti-affinity policies, follow these steps:

1. Log in to Prism Central.
2. From the [Application Switcher Function](#), select **Infrastructure** application and from the navigation bar, select **Compute > VMs**.

The system displays the **List** tab by default.

3. From the **Policies** tab dropdown menu, select **VM-VM Anti-affinity Policies**.

The system displays a summary view of VM-VM anti-affinity policies.

Table 149: VM-VM Anti-affinity Policies Page Field Description

Parameter	Description	Values
Name	Displays the VM-VM anti-affinity policy name.	(name)
Categories	Displays the count of categories associated with the selected VM-VM anti-affinity policy.	(number of categories)
VMs	Displays the count of VMs associated with the selected VM-VM anti-affinity policy.	(number of VMs)
VM Compliance Status	Displays the compliance status of the VMs associated with the VM-VM anti-affinity policy. If the policy is being applied and the compliance status is not yet known, the system displays the compliance status as Pending .	(number of VMs Compliant/Non-Compliant/Pending)
Modified By	Displays the name of the user who last modified the selected VM-VM anti-affinity policy.	(user)
Last Modified	Displays the date and time when the selected VM-VM anti-affinity policy is last modified.	(date & time)

You can perform the following actions for the VM-VM anti-affinity policies in the **VM-VM Anti-affinity Policies** summary view:

- Access the detailed information about an individual VM-VM anti-affinity policy. For more information, see [VM-VM Anti-Affinity Policies Details View](#) on page 513.
- Create a VM-VM anti-affinity policy. For more information, see [Creating a VM-VM Anti-Affinity Policy](#) on page 517.
- Update a VM-VM anti-affinity policy. For more information, see [Updating a VM-VM Anti-Affinity Policy](#) on page 518.
- Delete a VM-VM anti-affinity policy. For more information, see [Deleting a VM-VM Anti-Affinity Policy](#) on page 519.
- Remove a legacy VM group. For more information, see [Removing a Legacy VM Group](#) on page 519.

VM-VM Anti-Affinity Policies Details View

Use VM-VM anti-affinity policy details view to access detailed information about an individual VM-VM anti-affinity policy.

To access the details view of an individual VM-VM anti-affinity policy, follow these steps:

1. Log in to Prism Central.
2. From the [Application Switcher Function](#), select **Infrastructure** application and from the navigation bar, select **Compute > VMs**.

The system displays the **List** tab by default.

3. From the **Policies** tab dropdown menu, select **VM-VM Anti-affinity Policies**.

The system displays a summary view of VM-VM anti-affinity policies across all registered clusters.

- Click the target <VM-VM Anti-Affinity_Policy_Name> to view the **Summary** tab of an individual VM-VM Anti-affinity policy.

Note: Replace <VM-VM Anti-Affinity_Policy_Name> with the actual anti-affinity policy name at your site.

The **Summary** tab of an individual VM-VM anti-affinity policy provides the following widgets:

- Overview:** Displays an overview information about the VM-VM anti-affinity policy.
- Compliance Status:** Displays the compliance status of the VMs associated with this policy. If the policy is being applied and the compliance status is not yet known, the status is displayed as Pending. If a VM is part of multiple VM-VM anti-affinity policies, the oldest policy is applied on the VM. For rest of the policies, the VM is displayed as non-compliant.

For more information, see [VM-VM Anti-Affinity Policy Widgets Field Details](#) on page 514.

VM-VM Anti-Affinity Policy Widgets Field Details

The following table describes the fields in the **Overview** and **Compliance Status** widgets.

Table 150: VM-VM Anti Affinity Policy Widgets Field Description

Field	Description	Values
Overview widget		
Description	Displays the VM-VM anti-affinity policy description specified while creating the anti-affinity policy.	(description)
Last Modified By	Displays the name of the user who modified the VM-VM anti-affinity policy last time.	(user)
Last Modified	Displays the date and time when the VM-VM anti-affinity policy is modified last time.	(date & time)
Created By	Displays the name of the user who created the VM-VM anti-affinity policy.	(user)
Create Time	Displays the date and time when the VM-VM anti-affinity policy is created.	(date & time)
Categories	Displays the count of VM categories mapped to the selected VM-VM anti-affinity policy.	(number of VM categories)
Compliance Status widget		
Total VMs (Unique)	Displays the total number of VMs in the selected VM-VM anti-affinity policy.	(total number of VMs)
Compliance	Displays the compliance status in percentage .	(percentage of complaint VMs)
Compliant, In Progress, or Non compliant	Displays the number of VMs that are compliant and non-compliant with the VM-VM anti-affinity policy. If the compliance check is in progress, the system displays the compliance status as <i>In progress</i> .	Compliant, In Progress, or Non compliant

Clusters Tab

The **Clusters** tab displays the associated clusters information.

The following table describes the information displayed in the **Clusters** tab:

Table 151: Clusters Tab Field Description

Field	Description	Values
Name	Displays the cluster name.	(cluster name)
VMs	Displays the count of VMs associated with the selected VM-VM anti-affinity policy.	(number of VMs)
VM Compliance Status	Displays the number of VMs that are compliant and non-compliant with the VM-VM anti-affinity policy. If the compliance check is in progress, the system displays the compliance status as <i>In progress</i> .	(number of VMs Compliant Non-compliant Pending)

You can export the table that contains the list of VMs and their compliance status information to a file in a CSV format. For more information about **Export** option, see [Export](#) on page 54.

Categories Tab

The **Categories** tab displays the associated VM categories information.

The following table describes the information displayed in the **Categories** tab:

Table 152: Categories Tab Field Description

Field	Description	Values
Name	Displays the VM category name and its value.	(VM Category Name: Value)
VMs	Displays the count of VMs associated with the selected VM-VM anti-affinity policy.	(number of VMs)
VM Compliance Status	Displays the number of VMs that are compliant and non-compliant with the VM-VM anti-affinity policy. If the compliance check is in progress, the system displays the compliance status as <i>In progress</i> .	(number of VMs Compliant Non-compliant Pending)

You can export the table that contains the list of associated VM categories and their information to a file in a CSV format. For more information about **Export** option, see [Export](#) on page 54.

VMs Tab

The **VMs** tab displays the VMs information.

The following table describes the information displayed in the **VMs** tab:

Table 153: VMs Tab Field Description

Field	Description	Values
Name	Displays the VM name associated with the selected VM-VM anti-affinity policy.	(VM name)

Field	Description	Values
Host	Displays the host name associated with the selected VM-VM anti-affinity policy.	(host name)
Cluster	Displays the name of the cluster (on which the host resides) associated with the selected VM-VM anti-affinity policy.	(cluster name)
Category	Displays the VM category name and its value.	(VM category name: Value)
VM Compliance Status	Displays the compliance status of the VM. If the policy is being applied and the compliance status is not yet known, the system displays the compliance status as Pending VMs that do not adhere to the anti-affinity placements are displayed as non-compliant. The following are the reasons for displaying a VM as non-compliant: - Cluster Not Supported: When the VM is running on clusters earlier to 6.9 version. - Conflicting Legacy Anti-Affinity Policy: When the VM is already part of the PE based legacy policy. - Conflicting Anti-Affinity Policy: When is VM is part of multiple VM-VM anti-affinity policies, the policy having the oldest creation time is applied on the VM. - Not Enough Hosts: The number of powered on VMs is greater than the number of schedulable hosts. - Not Enough Resources: The scheduler is unable to place the VMs in a manner that conforms with the policy.	(Compliant/Pending/Cluster Not Supported/Conflicting Legacy Anti-Affinity Policy/Conflicting Anti-Affinity Policy/Not Enough Hosts/Not Enough Resources)

You can export the table that contains the list of VM entities related information to a file in a CSV format. For more information about **Export** option, see [Export](#) on page 54.

Audits Tab

The **Audits** tab displays the user action-related information for the VM-VM anti-affinity policies.

The following table describes the information displayed in the **Audits** tab:

Table 154: Audits Tab Field Description

Field	Description	Values
Action Description	Displays the user-action for the VM-VM anti-affinity policies.	(action description)
User Name	Displays the name of the user who performed the action.	(host name)
Operation Type	Displays the type of operation performed by the user for the VM-VM anti-affinity policies.	(cluster name)

Field	Description	Values
Request Time	Displays the date and time of the user action.	(date and time)

You can export the table that contains the list of user actions related information to a file in a CSV format. For more information about **Export** option, see [Export](#) on page 54.

Limitations of VM-VM Anti-Affinity Policies

This topic outlines the limitations of the VM-VM anti-affinity policies.

The VM-VM anti-affinity policies created in Prism Central have the following limitations:

- When you attach or detach a VM category, Prism Central takes approximately seven minutes to reflect the change in the applicable VM-VM anti-affinity policies.
- When you assign a category to a VM and map the VM category to the VM-VM anti-affinity policy, you can observe that the VM count updates immediately in the [VMs Tab](#) on page 515. However, Prism Central takes approximately seven minutes to update the VM count in the [VM-VM Anti-Affinity Policies Summary View](#) on page 512.
- You can associate a maximum of 20 VM categories to a single VM-VM anti-affinity policy.
- When the VM is part of multiple VM-VM anti-affinity policies, the policy with the oldest creation time is applied on the VM.
- When the VM is part of a legacy VM-VM anti-affinity policy and PC based VM-VM anti-affinity policy, the legacy VM-VM anti-affinity policy takes precedence over PC based VM-VM anti-affinity policy.

VM-VM Anti-Affinity Policy Configuration Workflow

This topic describes the workflow to create a VM-VM anti-affinity policy.

About this task

To configure a VM-VM anti-affinity policy, follow these steps:

Procedure

1. Create a category.
For more information, see [Creating a Category](#).
2. Apply the category to the VMs.
For more information on how to associate a category to the VMs, see [Associating VMs with Categories](#).
3. Create the VM-VM anti-affinity policy.
For more information, see [Creating a VM-VM Anti-Affinity Policy](#) on page 517.

Creating a VM-VM Anti-Affinity Policy

This task describes how to create a VM-VM anti-affinity policy from Prism Central.

Before you begin

Ensure that you meet following prerequisites before you create the VM-VM anti-affinity policy:

- You are on AOS 7.0 or later and pc.2024.3 or later.
- A category is created. For information on how to create a category, see [Creating a Category](#).

- VMs are associated with VM categories. For more information, see [Associating VMs with Categories](#) on page 508.

Note: Prism Central also allows you to associate VMs with VM category after creation of VM-VM anti-affinity policy.

- If you configured any legacy VM-VM anti-affinity policy (non-category-based anti-affinity policy) associated with the VMs, you must first remove those legacy anti-affinity policies to allow the creation of category-based anti-affinity policies associated with the same VMs. For more information, see [Removing a Legacy VM Group](#) on page 519.

About this task

To create a VM-VM anti-affinity policy, follow these steps:

Procedure

1. Log in to Prism Central.
2. From the [Application Switcher Function](#), select **Infrastructure** application and from the navigation bar, select **Compute > VMs**.
The system displays the **List** tab by default.
3. From the **Policies** tab dropdown menu, select **VM-VM Anti-affinity Policies**.
4. Click **Create VM-VM Anti-affinity Policy**.
5. Enter the following information in the **Create VM-VM Anti-affinity Policy** window:
 - **Name:** Enter the VM-VM anti-affinity policy name.
 - (Optional) **Description:** Enter the description for the VM-VM anti-affinity policy.
 - **VM Categories:** Type the name of the VM category. The system displays the list of matching categories based on the typed entry. Use the [Add icon](#) and [Remove icon](#) to add and remove the required categories.
6. Click **Create**.

Updating a VM-VM Anti-Affinity Policy

This task describes how to update a VM-VM anti-affinity policy from Prism Central.

About this task

To update a VM-VM anti-affinity policy, follow these steps:

Procedure

1. Log in to Prism Central.
2. From the [Application Switcher Function](#), select **Infrastructure** application and from the navigation bar, select **Compute > VMs**.
The system displays the **List** tab by default.
3. From the **Policies** tab dropdown menu, select **VM-VM Anti-affinity Policies**.
The system displays the VM-VM anti-affinity policies created across registered clusters.
4. Select the target VM-VM anti-affinity policy checkbox, and from the **Actions** dropdown menu, choose **Update**.

5. Update the fields in the **Update VM-VM Anti-affinity Policy** window as needed and click **Update**.
For information on the fields available in the **Update VM-VM Anti-affinity Policy** window, see [Creating a VM-VM Anti-Affinity Policy](#) on page 517.

Deleting a VM-VM Anti-Affinity Policy

This task describes how to delete a VM-VM anti-affinity policy from Prism Central.

About this task

To delete a VM-VM anti-affinity policy, follow these steps:

Procedure

1. Log in to Prism Central.
2. From the [Application Switcher Function](#), select **Infrastructure** application and from the navigation bar, select **Compute > VMs**.
The system displays the **List** tab by default.
3. From the **Policies** tab dropdown menu, select **VM-VM Anti-affinity Policies**.
The system displays the VM-VM anti-affinity policies created across registered clusters.
4. Select the target VM-VM anti-affinity policy checkbox and from the **Actions** dropdown menu, choose **Delete**.
The system prompts you to confirm the delete action.
5. Click **Delete**.

Removing a Legacy VM Group

This task describes how to remove a legacy VM Group from Prism Central.

About this task

The VM-VM anti-affinity policies configured from aCLI are called legacy anti-affinity policies and the VM groups mapped to these policies are called legacy VM groups in Prism Central. If you previously created any legacy VM groups that are associated with VM-VM anti-affinity policies, you must first remove those legacy VM groups to use the new Prism Central based VM-VM anti-affinity policies for those VMs.

To remove the legacy VM Group from Prism Central, follow these steps:

Procedure

1. Log in to Prism Central.
2. From the [Application Switcher Function](#), select **Infrastructure** application and from the navigation bar, select **Compute > VMs**.
The system displays the **List** tab by default.
3. From the **Policies** tab dropdown menu, select **VM-VM Anti-affinity Policies**.
The system displays a summary view of VM-VM anti-affinity policies across all registered clusters.
4. Click **You have Legacy VM Group defined through the CLI** banner.
The system displays **Legacy VM-VM Anti-affinity Policies** window.
5. Select the target **VM Groups** checkbox, and select **Delete selected group**.
6. Click **Delete**.

Role Based Access Control for the VM-VM Anti-Affinity Policies

This topic describes the Role Based Access Control (RBAC) requirements and limitations for VM-VM anti-affinity policies.

Using RBAC, you can provide the necessary permissions for users to access the VM-VM anti-affinity policies. For information on RBAC, see [Controlling User Access \(RBAC\)](#) in *Security Guide*.

RBAC Requirements for the VM-VM Anti-Affinity Policies

Before you allocate the [VM-VM Anti-Affinity Policy Related Permissions](#) on page 520, ensure that the view VM permission is allocated to the admin user or a custom user to access the VM-VM anti-affinity policies in Prism Central.

VM-VM Anti-Affinity Policy Related Permissions

In addition to the view VM permission, you can also assign the following permissions to the admin or custom user to manage the VM-VM anti-affinity policies in Prism Central:

Table 155: VM-VM Anti-Affinity Policy Related Permissions

Permission	Description
Create_VM_Anti_Affinity_Policy	Admin or custom user can create anti-affinity policies.
View_VM_Anti_Affinity_Policy	Admin or custom user can view anti-affinity policies.
Update_VM_Anti_Affinity_Policy	Admin or custom user can update the existing anti-affinity policies.
View_VM_Anti_Affinity_Policy_VM_Compliance_State	Admin or custom user can view the compliance state of anti-affinity policies.
Delete_VM_Anti_Affinity_Policy	Admin or custom user can delete anti-affinity policies.
View_Legacy_VM_Anti_Affinity_Policy	Admin or custom user can view legacy anti-affinity policies.
Delete_Legacy_VM_Anti_Affinity_Policy	Admin or custom user can delete legacy anti-affinity policies.

RBAC Limitations for the VM-VM Anti-Affinity Policies

The following table describes the limitations that apply to RBAC for the VM-VM anti-affinity policies:

Table 156: RBAC Limitations for VM-VM Anti-Affinity Policies

Permission Assigned to the Admin or Custom User	Admin or Custom User Action	System Behavior	Limitation
Create_VM_Anti_Affinity_Policy Update_VM_Anti_Affinity_Policy	Admin or custom user creates or updates the VM-VM anti-affinity policy.	The system applies the VM-VM anti-affinity policy to all VMs associated with the mapped VM categories.	The system does not check whether the admin or the custom user has permission to access the VMs mapped in the VM-VM anti-affinity policy.

Permission Assigned to the Admin or Custom User	Admin or Custom User Action	System Behavior	Limitation
Update_VM_Categories	Admin or custom user updates a VM category (for example, attaches a VM to a VM category).	The system applies the VM-VM anti-affinity policy to which the VM category is assigned.	The system does not check whether the admin or the custom user has permission to create or update the VM-VM anti-affinity policy.

Management of VM-VM Anti-Affinity Policies for PC-Based Disaster Recovery Solution

This section describes how to manage VM-VM anti-affinity policies in a Prism Central-based Disaster Recovery (DR) setup.

Managing Failover Operation

This section describes how to set up VM-VM anti-affinity policies at the recovery site for a failover operation.

About this task

After the failover is completed, the VM categories assigned to the VM on the primary site (primary AZ) are synchronized to the recovery site (recovery AZ). The synchronized VM categories are also reflected in the Prism Central instance of the recovery site.

To set up the VM-VM anti-affinity policies at the recovery site, perform the following actions before a disaster occurs:

Procedure

1. Create the VM categories at the recovery site with the same name as defined at the primary site.
For more information, see [Creating a Category](#) in *Prism Central Infrastructure Guide*.
2. Reconfigure the VM-VM anti-affinity policies at the recovery site using the VM categories created at the recovery site.
For more information, see [Creating a VM-VM Anti-Affinity Policy](#) in *Prism Central Infrastructure Guide*.

Important: If the legacy VM-VM anti-affinity policies are defined using Prism element, the policies are not replicated during failover, and you need to manually redefine the anti-affinity policies at the recovery site after failover.

Overview of Failback Operation

This section describes how failback operation functions for VM-VM anti-affinity policies at the primary site in a Prism Central-based Disaster Recovery setup.

After you perform failback, the VM categories are synchronized back to the primary site, and the anti-affinity policies created on the primary site (primary AZ) are enabled for the VMs.

Managing Synchronous Replication

This section describes how to set up VM-VM anti-affinity policies at the recovery site for a synchronous replication operation.

About this task

You can configure the VM-VM anti-affinity policy at the recovery site for the VMs or VM categories that are protected with synchronous replication schedule in the primary site using the following workflow:

Procedure

1. Create the VM categories at the recovery site with the same name as defined at the primary site.
For more information, see [Creating a Category](#) in *Prism Central Infrastructure Guide*.
2. Create the VM-VM anti-affinity policies at the recovery site using the VM categories created at the recovery site.
For more information, see [Creating a VM-VM Anti-Affinity Policy](#) in *Prism Central Infrastructure Guide*.

Note: For more information about how to configure the protection policy with a synchronous replication schedule, see [Creating a Protection Policy](#) in *Nutanix Disaster Recovery Guide*.

NGT Policies

NGT policies are used to define the VM reboot action after you install or upgrade the NGT.

If a restart is required post installation or upgrade, you can define a policy on when the restart should happen. You can use these policies when you have different set of VMs which you would like to be restarted at different times, for example if they are in different time zones or if their down times are expected to be different. Policy workflow are defined on the categories. You have to attach the category to the VM and then you can create a policy for the VM.

Important: The NGT policies are applied only when you select **Skip restart** option during installation or upgradation of NGT. However, if you select a different restart option during installation or upgradation of NGT, the NGT policy is not applied. For more information about how to install and upgrade NGT, see [NGT Management in Prism Central](#) on page 192.

NGT Policies Summary View

The NGT policies summary view enables you to access a list of all the user-defined Nutanix Guest tools (NGT) policies (VM reboot policies) across registered clusters.

To access the NGT policies summary view, perform the following steps:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Compute > VMs** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

The system displays the **List** tab by default.

3. Select **NGT Policies** in the **Policies** tab dropdown menu.

The system displays a view of NGT policies across all registered clusters.




Figure 153: NGT Policies

Table 157: NGT Policies Page - Field Description

Field	Description	Values
Name	Displays the NGT policy name.	(name)
Policy Type	Displays the type of NGT policy.	REBOOT

You can perform the following actions for the NGT policies in the **NGT Policies** summary view:

- Create an NGT policy. For more information, see [Creating a NGT Policy](#) on page 523.
- Filter the NGT policies list based on the NGT policy name using **Filters** pane. For more information about **Filters** pane, see [Filters Pane](#) on page 50.
- Use the **Actions** dropdown menu to update and delete an NGT policy. For more information, see [Managing NGT Policies](#) on page 525.

Creating a NGT Policy

This section describes how to create a NGT Policy (VM reboot policy).

Procedure

To create a NGT policy, perform the following steps:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Compute > VMs** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.
The system displays the **List** tab by default.

3. Select **NGT Policies** in the **Policies** tab dropdown menu, and click **New VM Reboot Policy**. The system displays the **New VM Reboot Policy** window.

4. Specify the following information in the **Create Affinity Policy** window:

- **Policy Name:** Enter the NGT policy (VM reboot policy) name.
- (optional) **Policy Description:** Enter the description for the NGT policy.
- **Guest Restart Schedule:** Select one of the following guest VM restart schedules:
 - **Restart as soon as the install/upgrade is completed:** Select this option if you want to restart the VMs of the selected category right after the install process is completed.
 - **Restart at specific day and time after the upgrade is completed:** Select this option and choose the date and time on which you want the restart to happen.

Note: The VMs restart depends on the Prism Element timezone settings.

- **Add Category:** Click **Add Category** and select the category that you want to apply to the NGT policy. Type the name of the VM category. The system displays the list of matching categories based on the typed entry. Use the [Add icon](#) and [Remove icon](#) to add and remove the required categories.

For information about how to create a VM category, see [Creating a Category](#) in the *Prism Central Admin Center Guide*.

New VM Reboot Policy ? X

Policy Name

VM reboot policy -Site A

Description

VM reboot policy for Site A

Guest Restart Schedule

☒ Restart as soon as the install/upgrade is completed

☐ Restart at specific day and time after the upgrade is completed

Associated Categories Update

Name	Value
AppTier	Default

Cancel Save

Figure 154: Create VM Reboot Policy

5. Click **Save**.

Note: Any VM reboot policy you create using this procedure is overridden if you choose to restart the VM during installation or upgradation of NGT by using a different option compared to the option configured in the policy for

VM restart. If you choose not to restart the VM during installation or upgradation of NGT, the system continues to apply the existing reboot policy for the VM. For example, if you select the **Restart as soon as the install/upgrade is completed** option as part of the NGT policy, and if you select the **Restart at specific date and time after the install is completed** option during installation or upgradation of NGT, the **Restart at specific date and time after the install is completed** setting take precedence and the VM is restarted at the scheduled date and time. However, if you select the **Skip restart** option during installation or upgradation of NGT, the VM reboot policy takes precedence and the VM restart occurs immediately after an installation or upgrade.

The policy is created and saved in the NGT Policies pane.

Managing NGT Policies

This section describes how to manage the existing NGT policies (VM reboot policies) in Prism Central.

About this task

You can perform the following actions to manage the existing NGT policies in Prism Central:

- Update an NGT policy.
- Delete an NGT policy.

Procedure

To modify a NGT policy, perform the following steps:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Compute > VMs** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.
The system displays the **List** tab by default.
3. Select the target NGT policy checkbox, and choose **Update** from the **Actions** dropdown menu.
The system displays the NGT policies created across registered clusters.
4. Update the fields in the **Update VM Reboot Policy** window as per your requirement, and click **Save**. For field details, see [Creating a VM-Host Affinity Policy](#) on page 509.

Update VM Reboot Policy

Policy Name
VM reboot policy - Site A

Description
VM reboot policy for Site A

Guest Restart Schedule
☒ Restart as soon as the install/upgrade is completed
☐ Restart at specific day and time after the upgrade is completed

Associated Categories [Update](#)

Name	Value
ApplTier	Default

Figure 155: Update VM Reboot Policy

To delete an NGT policy, select the target NGT policy checkbox and choose **Delete** from the **Actions** dropdown menu. The system prompts you to confirm the delete action. Click **Delete** to confirm the delete NGT policy action.

VM Template Policy Management

In Prism Central, you can create and manage policies for VM templates.

You can create a policy to manage the placement of the VM templates on the clusters.

Template Placement Policies

You can create and manage policies for VM templates to control the placement of the templates on the clusters.

When you create a VM template, it is located with the source VM in the same cluster. To deploy a VM to any other cluster through the template, the system first replicates the template to the target cluster and then deploys the VM. A template placement policy automates this task of template replication. Using a template placement policy, you can fetch the active version of a template in advance on your preferred clusters for faster VM deployment.

Creating a Template Placement Policy

Create a template placement policy to map the templates to the target clusters using categories.

Before you begin

Ensure that you have assigned the required categories to the template. Assign the same categories to the cluster where you want to deploy the VMs. For information on assigning a category, see [Assigning a Category to an Entity](#) in the *Prism Central Admin Center Guide*.

About this task

To create a template placement policy, follow these steps:

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and from the **Navigation Bar**, navigate to **Compute > Templates**.
For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.
The system displays the **List** tab by default.
3. Select the **Policies** tab, and then from the dropdown menu, select **Placement Policies**.
The system displays the placement policies across all the registered clusters.
4. Click **Create Placement Policy**.
5. In the **Create Template Placement Policy** dialog box, enter the following information:
 - a. In the **Policy Name** field, enter the name of the template placement policy.
 - b. In the **Description** field, enter the description for the template placement policy.
 - c. In the **Template Category** field, select the template categories to specify the associated templates that needs to be replicated to the target cluster.
Your selected categories are listed along with the number of associated templates. If multiple templates are associated with the selected categories, you can remove the templates that are not required for this policy.
 - d. In the **Cluster Category** field, select the cluster category to specify the associated clusters where the template needs to be replicated.
Your selected categories are listed along with the number of associated clusters. If multiple clusters are associated with the selected categories, you can remove the clusters that are not required for this policy.
6. Click **Save**.

Managing a Template Placement Policy

You can update or delete a template placement policy from Prism Central.

About this task

To manage a template placement policy, follow these steps:

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and from the **Navigation Bar**, navigate to **Compute > Templates**.
For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.
The system displays the **List** tab by default.
3. Select the **Policies** tab, and then from the dropdown menu, select **Placement Policies**.
The system displays the template placement policies on the Prism Central instance.
4. Do one of the following:
 - To update a template placement policy:
 - a. Select the template placement policy, and choose **Update** from the **Actions** dropdown menu.
 - b. Update the information as per your requirement in the **Update Template Placement Policy** window, and click **Save**.
For more information on the fields, see [Creating a Template Placement Policy](#) on page 526.
 - To delete a template placement policy:
 - a. Select the template placement policy, and choose **Delete** from the **Actions** dropdown menu.
 - b. Click **Delete**.

Image Policy Management

In Prism Central, you can create and manage the following types of policies for the images:

- **Image Placement Policies:** Prism Central enables you to configure policies that govern which clusters receive the images that you upload. These policies, called image placement policies, map images to target clusters using categories associated with both those entities. For more information, see [Image Placement Policies](#) on page 527.
- **Bandwidth Throttling Policies:** The bandwidth throttling policy allows you to limit the bandwidth consumed during image creation using the URL option in specific clusters. For more information, see [Bandwidth Throttling Policies](#) on page 535.

Image Placement Policies

Prism Central enables you to upload VM images from your workstation or from a remote server. It also enables you to manually specify a subset of registered clusters as targets for those images. However, image placement decisions are often driven by compliance requirements, policies, and regulations. For example, you might have to confine a VM image to a Nutanix cluster in a specific location. Or, a cluster in one region might require a Windows VM image with a specific set of applications while a cluster in another region might require a Linux VM with its own set of applications.

Prism Central enables you to configure policies that govern which clusters receive the images that you upload. These policies, called image placement policies, map images to target clusters using categories associated with both those entities. For example, an image placement policy states that the Linux images associated with the `OS:Linux` and `distribution:Centos` categories must be uploaded to the clusters associated with the `Location:SFO` category.

Image placement policies also specify how strictly the policy must be enforced. Soft enforcement allows clusters not identified by the policy to use the images, when required, by checking them out. Hard enforcement disallows such usage. By using both soft and hard enforcement on the same category of clusters, you can address more specific use cases. For more information, see [Sample Scenarios and Configurations](#) on page 534.

Image placement policies maintain a state of compliance across the clusters registered with Prism Central. They detect and correct violations introduced by changes in category assignments and updates to their configuration (including a switch from soft enforcement to hard enforcement). For more information, see [Sample Scenarios and Configurations](#) on page 534.

Image Placement Policies Summary View

The image placement policies summary view allows you to access a list of all the user-defined image placement policies across registered clusters.

To access the summary view of all image placement policies:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Compute** > **Images** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

The system displays the **List** tab by default.

3. Select **Placement Policies** in the **Policies** tab dropdown menu.

The system displays a summary view of image placement policies across all registered clusters.




Figure 156: Summary View - All Image Placement Policies

Table 158: Image Placement Policies Page - Field Description

Field	Description	Values
Name	Displays the image placement policy name.	(name)
Description	Displays the purpose of the image placement policy.	(text string)
Status	Displays the status of the image placement policy.	Active Suspended

You can perform the following actions in the **Image Placement Policies** summary view:

- Access the detailed information about an individual image placement policy. For more information, see [Image Placement Policy Details View](#) on page 529.
- Create an image placement policy. For more information, see [Creating an Image Placement Policy](#) on page 531.
- Use the **Actions** dropdown menu to update, delete, suspend enforcement (if status is active) or resume enforcement (if status is suspended). For more information, see [Managing Image Placement Policies](#) on page 532.
- Filter the image placement policies list using **Filters** pane. For more information about **Filters** pane, see [Filters Pane](#) on page 50.

For more information about fields, see [Table 158: Image Placement Policies Page - Field Description](#) on page 528.

- Export the table that contains the list of image placement policies and their information to a file in a CSV format. For more information about **Export** option, see [Export](#) on page 54.
- View image placement policies based on pre-defined criteria or create a custom view. For information about available views and how to create a custom view, see [View by](#) on page 51.

Image Placement Policy Details View

To access the details view of an individual image placement policy, perform the following steps:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Compute > Images** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

The system displays the **List** tab by default.

3. Select **Placement Policies** in the **Policies** tab.

The system displays a summary view of image placement policies across all registered clusters.

4. Click the target `<Image_Placement_Policy_Name>` to access the details view individual image placement policy.

Note: Replace `<Image_Placement_Policy_Name>` with the actual image placement policy name at your site.




Figure 157: Image Placement Policy Details view

The details view of an individual image placement policy provides the following information:

- Policy name that appears in the upper left. You can switch from one policy to another by selecting the policy name from the dropdown menu.
- Policy enforcement status appears below the name and indicates whether the enforcement is hard or soft.
- Two columns that specify the target images (on the left) and target clusters (on the right). Click **Expand all** (or **Collapse all**) to display (or hide) the list of images or clusters.

<Action> available in the top right corner. Click the appropriate <Action> to run that administrative action on the image placement policy. For more information about how to perform any <Action>, see [Managing Image Placement Policies](#) on page 532.

Configuration Workflow

About this task

To set up an image placement policy, perform the following steps:

Procedure

1. Create categories for the following entities:

- a. Clusters (for example, create categories based on cluster location or region)
- b. Images (for example, create categories based on operating system and size)

For information about how to create a category, see [Creating a Category](#) in the *Prism Central Admin Center Guide*.

2. Apply the cluster categories to clusters and image categories to images.

For information about how to associate categories with a Nutanix cluster, see [Associating Clusters with Categories](#) on page 530.

For information about how to associate categories with an image, see [Associating an Image with Categories](#).

3. Configure the image placement policy. For more information, see [Configuring an Image Placement Policy](#).

Associating Clusters with Categories

About this task

To associate categories with a registered Nutanix cluster, perform the following steps:

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Hardware > Clusters** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.
The system displays the **List** tab by default.
3. Select the target cluster checkboxes that you want to associate with a category, and choose **Manage Categories** from the **Actions** dropdown menu.
The system displays the **Manage Cluster Categories** window.
4. Type the name of the category in **Set Categories** field.
The system displays the list of matching categories based on the typed entry.
5. Use the [Add icon](#) and [Remove icon](#) to add and remove the required categories.
6. Click **Save**.

Associating Images with Categories

About this task

To associate categories with an image, perform the following steps:

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Compute > Images** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.
The system displays the **List** tab by default.
3. Select the target image checkbox(es) that you want to associate with a category, and choose **Manage Categories** from the **Actions** dropdown menu.
The system displays the **Manage Image Categories** window.
4. Type the name of the category in **Set Categories** field.
The system displays the list of matching categories based on the typed entry.
5. Use the [Add icon](#) and [Remove icon](#) to add and remove the required categories.
6. Click **Save**.

Creating an Image Placement Policy

About this task

This section describes how to create an image placement policy in Prism Central.

Before you begin

Ensure that the following prerequisites are met before you create an image placement policy:

- Cluster and Image categories are created. For information about how to create a category, see [Creating a Category](#) in the *Prism Central Admin Center Guide*.
- Images are associated with the Image categories, and clusters with the cluster categories. For more information, see [Associating Images with Categories](#) on page 530 and [Associating Clusters with Categories](#) on page 530.

Note: The system also allows you to associate images with image category and clusters with cluster category after creation of image placement policy.

- Review existing image placement policies to avoid causing conflicts. For information about issues caused by conflicting policies, see [How Prism Central Handles Conflicting Policies](#) on page 533

Procedure

To create an affinity policy, perform the following steps:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Compute > Images** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.
The system displays the **List** tab by default.
3. Select **Placement Policies** in the **Policies** tab dropdown menu, and click **Create Placement Policy**.
The system displays the **Create Image Placement Policy** window.

4. Specify the following information in the **Create Image Placement Policy** window:

- **Policy Name:** Enter the image placement policy name.
- (optional) **Description:** Enter the description for the image placement policy.
- **Assign Images With All Of The Following Categories:** Type the name of the image category. The system displays the list of matching categories based on the typed entry. Use the [Add icon](#) and [Remove icon](#) to add and remove the required categories.
- **To Clusters With All Of The Following Categories:** Type the name of the cluster category. The system displays the list of matching categories based on the typed entry. Use the [Add icon](#) and [Remove icon](#) to add and remove the required categories.
- **Policy Enforcement:** Select one of the following mechanisms for policy enforcement:
 - **Soft** - The images are uploaded to the set of clusters identified by the image placement policy, and no restriction is placed on using those images on clusters outside the identified set. For example, if clusters A, B, and C are registered with Prism Central, and an image placement policy with a soft enforcement policy uploads an image to clusters A and B, you can use the image to create a VM on cluster C. The checkout to cluster C is not blocked.
 - **Hard** - The images are uploaded to the set of clusters identified by the image placement policy, however you cannot use the image to create a VM on any cluster outside the identified set. For example, if clusters A, B, and C are registered with Prism Central, and an image placement policy with a hard enforcement policy uploads an image to clusters A and B, the image is available only on clusters A and B. Cluster C cannot check out the image.

Create Image Placement Policy

Policy Name
Image-Placement-Policy-Site A

Description
Image-Placement-Policy for Site A

Assign Images With All Of The Following Categories
Image Cat SiteA:Site A
Search for a category

To Clusters With All Of The Following Categories
Cluster Cat SiteA:Site A
Search for a category

Policy Enforcement
Soft enforcement will allow you to manually place images on other clusters if needed. Hard enforcement will ensure that these images are only placed on the selected clusters.
Enforcement
Soft

Cancel Save

Figure 158: Create Image Placement Policy

5. Click **Save**.

Managing Image Placement Policies

This section describes how to manage the image placement policies in Prism Central.

About this task

You can perform the following actions to manage the image placement policies in Prism Central:

- Update an image placement policy.
- Delete an image placement policy.
- Suspend Enforcement
- Resume Enforcement

Updates to an image placement policy can result in a policy violation. Prism Central attempts to correct the violation by running a series of actions. For example, if you change the cluster category such that the policy now identifies a different set of clusters as targets, Prism Central first copies the images to the clusters in the new category, and then removes the images from the clusters in the previously mentioned category. If you change the policy enforcement setting from Soft to Hard, and if any clusters that are not identified as targets had previously checked out the images, Prism Central corrects the violation by deleting the images from those clusters. An update to a policy might also result in a conflict with one or more of the other policies that apply to the same category of images.

For more information about how Prism Central handles conflicting policies, see [How Prism Central Handles Conflicting Policies](#) on page 533.

Procedure

To update an image placement policy, perform the following steps:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Compute > Images** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.
The system displays the **List** tab by default.
3. Select **Placement Policies** in the **Policies** tab dropdown menu.
The system displays the placement policies across all the registered clusters.
4. Select the target image placement policy checkbox, and choose **Update** from the **Actions** dropdown menu.
5. Update the fields in the **Update Image Placement Policy** window as per your requirement, and click **Save**.
For field details, see [Creating an Image Placement Policy](#) on page 531.

To delete the image placement policy, select the target image placement policy checkbox and choose **Delete** from the **Actions** dropdown menu. The system prompts you to confirm the delete action. Click **Delete** to confirm the delete action.

To suspend enforcement of the image placement policy (if Active), select the target image placement policy checkbox and choose **Suspend Enforcement** from the **Actions** dropdown menu. The system prompts you to confirm the **Suspend Enforcement** action. Click **Confirm** to suspend the policy enforcement.

To resume enforcement of the image placement policy (if Suspended), select the target image placement policy checkbox and choose **Resume Enforcement** from the **Actions** dropdown menu. The system prompts you to confirm the **Resume Enforcement** action. Click **Confirm** to resume the policy enforcement.

How Prism Central Handles Conflicting Policies

Policies are said to be in conflict if they identify the same category of images but their resulting cluster placements are different.

For example, the following policies are conflicting policies if they apply to the same category of images:

- Policy P1, which identifies clusters C1 and C2 as target clusters, and the policy enforcement setting is Soft.

- Policy P2, which identifies clusters C3 and C4 as target clusters, and the policy enforcement setting is Hard.

Prism Central handles such conflicts in the following manner:

- Prism Central ignores conflicting policies and does not perform any action on the images. If policy P1 was configured first, Prism Central stops enforcing P1 as soon as P2 is configured, and it does not enforce P2 either. When there is a policy conflict, the conflict is noted in the **Policies** tab of the details view for that image.
- Prism Central continues to enforce other policies that are not in conflict. For example, if policies P3 and P4 apply to the same categories of images as P1 and P2, Prism Central continues to enforce P3 and P4 as long as they are not in conflict with P1 and P2.

Sample Scenarios and Configurations

The following table describes how to configure image management policies to achieve a desired result. The examples use a Prism Central instance to which three clusters A, B, and C are registered (possibly among other clusters).

Table 159: Sample Scenarios and Configurations

Desired Result	Image Placement Policy Configuration
The images must be available on any one of the clusters (A, B, or C) and can be checked out to the other two clusters and to any other clusters that might be registered with Prism Central now or in the future.	Create an image placement policy that identifies any one of the clusters (A, B, or C) as the target cluster (by the use of categories) and set the policy enforcement to Soft.
The images must be available on all of the clusters A, B, and C and can be checked out to any other clusters that are registered with Prism Central now or in the future.	Create an image placement policy that identifies clusters A, B, and C as the target clusters (by the use of categories) and set the policy enforcement to Soft.
The images must be available on all of the clusters A, B, and C and cannot be checked out to any other clusters that are registered with Prism Central now or in the future.	Do the following: • Create an image placement policy that identifies clusters A, B, and C as the target clusters (by the use of categories) and set the policy enforcement to Soft. (This policy ensures that the images are available on all of the three clusters) • Create a second image placement policy that identifies clusters A, B, and C as the target clusters (by the use of categories) and set the policy enforcement to Hard. (This policy makes sure that the images are available only on these three clusters)
The images must be available only on cluster A and cannot be checked out to clusters B and C now, and the images also cannot be checked out to any other clusters that are registered with Prism Central in the future.	Create an image placement policy that identifies cluster A as the target cluster (by the use of categories) and set the policy enforcement to Hard.

Desired Result	Image Placement Policy Configuration
The images are available on cluster A and must now be available on clusters B and C. Cluster A can use the images if required. The policy enforcement is currently Soft.	Do the following: • Dissociate the categories specified in the image placement policy from cluster A and associate them with clusters B and C. • Do not update the image placement policy (no changes are required). Prism Central takes the following corrective actions: 1. Prism Central copies the images to clusters B and C. If clusters B and C had checked out the images before the reassignment of categories to the clusters, the images are not copied. 2. Additionally, because the Soft policy enforcement setting is retained, Prism Central does not remove the image from cluster A.
The images are available on cluster A and must now be available only on clusters B and C. Cluster A must not be permitted to use the images. The policy enforcement is currently Soft.	Do the following: • Dissociate the categories specified in the image placement policies from cluster A and associate them with clusters B and C. • Change the policy enforcement setting in the image placement policy to Hard. Prism Central takes the following corrective actions: 1. Prism Central copies the images to clusters B and C. If clusters B and C had checked out the images before the reassignment of categories to the clusters, the images are not copied. 2. To honor the Hard policy enforcement setting, Prism Central removes the image from cluster A.

Bandwidth Throttling Policies

Prism Central uses the bandwidth throttling policy feature to manage the usage of bandwidth when you create a new image or transfer an image from one cluster to another cluster using Prism Central. The bandwidth throttling policy allows you to limit the bandwidth consumed during image creation using the URL option in specific clusters. For information about how to add a image using the URL option, see [Adding Images from a Remote Server](#) on page 269.

Without bandwidth throttling policy, an image creation from the remote server consumes as much bandwidth as available. High bandwidth consumption for creating a new image limits the bandwidth availability for the other cluster operations.

Bandwidth Throttling Policies Summary View

The bandwidth throttling policies summary view allows you to access a list of all the user-defined bandwidth-throttling policies across registered clusters.

Summary View of all Bandwidth Throttling Policies

To access the bandwidth throttling summary view, perform the following steps:

1. Log in to Prism Central.

2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Compute > Images** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

The system displays the **List** tab by default.

3. Select **Bandwidth Throttling Policies** in the **Policies** tab dropdown menu.

The system displays a view of bandwidth throttling policies across all registered clusters.




Figure 159: Bandwidth Throttling Policies

Table 160: Bandwidth Throttling Policies Page - Field Description

Field	Description	Values
Policy Name	Displays the policy name.	(name)
Description	Describes the purpose of the policy.	(description)
Bandwidth Throttle Limit	Describes the bandwidth throttle limit configured for the policy in MBps	xx MBps
Last Updated On	Describes the date and time when the bandwidth throttling policy was created or updated.	(date and time)

You can perform the following actions for the bandwidth throttling policies in the **Bandwidth Throttling Policies** summary view:

- Access the detailed information about an individual bandwidth throttling policy. For more information, see [Bandwidth Throttling Policy Details View](#) on page 537.
- Create a bandwidth throttling policy. For more information, see [Creating a Bandwidth Throttling Policy](#) on page 538.
- Filter the bandwidth throttling policies list using **Filters** pane. For more information about **Filters** pane, see [Filters Pane](#) on page 50

For more information about fields, see [Table 160: Bandwidth Throttling Policies Page - Field Description](#) on page 536.

- Export the table that contains the list of bandwidth throttling policies and their information to a file in a CSV format. For more information about **Export** option, see [Export](#) on page 54.
- View bandwidth throttling policies based on pre-defined criteria or create a custom view. For information about available views and how to create a custom view, see [View by](#) on page 51.
- Use the **Actions** dropdown menu to update and delete a bandwidth throttling policy. For more information, see [Managing Bandwidth Throttling Policies](#) on page 540.

Bandwidth Throttling Policy Details View

To access the details view of an individual bandwidth throttling policy, perform the following steps:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Compute > Images** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

The system displays the **List** tab by default.

3. Select **Bandwidth Throttling Policies** in the **Policies** tab.

The system displays a summary view of bandwidth throttling policies across all registered clusters.

4. Click the target `<Bandwidth_Throttling_Policy_Name>` to access the details view individual bandwidth throttling policy.

Note: Replace `<Bandwidth_Throttling_Policy_Name>` with the actual bandwidth throttling policy name at your site.




Figure 160: Bandwidth Throttling Policy Details view

The details view of an individual bandwidth throttling policy provides the following information:

- Policy name that appears in the upper left.
- **Enrolled Cluster** tab that displays the following cluster information:
 - **Cluster Name** - Displays the name of the cluster that is configured with bandwidth throttling policy.
 - **Effective Bandwidth Limit** - The selected cluster uses the **Effective Bandwidth Limit** configured to add a new image next time. In a scenario where multiple bandwidth throttling policies are applied to the same cluster, the minimum of all the applicable bandwidth throttling policies is enforced on the cluster.

For example, you have a cluster named **PE-123**, and you have created multiple bandwidth throttling policies with different **Bandwidth Limit**, then during the cluster association with the bandwidth throttling policy (step 6), the cluster is associated with the bandwidth throttling policy which has the least **Bandwidth Limit** configured on it.
 - **Hypervisor** - Displays the Hypervisor name.
 - **Health** - Displays the cluster health.
 - **Host Count** - Displays the number of hosts in the cluster

`<Action>` available above the detailed view column. Click the appropriate `<Action>` to run that administrative action on the bandwidth throttling policy. For more information about how to perform any `<Action>`, see [Managing Bandwidth Throttling Policies](#) on page 540 .

Requirements for Bandwidth Throttling Policy

Ensure that the following prerequisites are met to configure the bandwidth throttling policy feature in Prism Central:

- Prism Central release 2021.9 or later and AOS version 6.0 or later are deployed at your site.
- The clusters are registered with Prism Central.
- A cluster category is created and associated with the cluster.

For information about how to create a cluster category, see [Creating a Category](#) in the *Prism Central Admin Center Guide*.

For information about how to associate a category to a cluster, see [Assigning a Category to an Entity](#) in the *Prism Central Admin Center Guide*.

- AHV Hypervisor is deployed at your site.

Note: The bandwidth throttling policy is not supported with ESXi and Hyper-V hypervisors.

Limitations with Bandwidth Throttling Policy

The following limitations apply to the bandwidth throttling policies:

- Prism Central enforces bandwidth throttling policy only for new images created by using the URL option on clusters registered with Prism Central. For information about how to add a image using the URL option, see [Adding Images from a Remote Server](#) on page 269.
- You cannot enforce a bandwidth throttling policy if you create an image directly from a Prism Element.
- The least applicable bandwidth throttling policy on the source and destination cluster is enforced when you transfer an image from one cluster to another cluster.

Creating a Bandwidth Throttling Policy

About this task

This section describes how to create a bandwidth throttling policy in Prism Central.

Perform the following procedure to create a new bandwidth throttling policy.

Before you begin

Ensure that the following prerequisites are met before you create a bandwidth throttling policy in Prism Central:

- Requirements that are described in [Requirements for Bandwidth Throttling Policy](#) on page 537 are followed.
- Limitations that are described in [Limitations with Bandwidth Throttling Policy](#) on page 538 are observed.

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Compute > Images** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.
The system displays the **List** tab by default.

3. Select **Bandwidth Throttling Policies** in the **Policies** tab dropdown menu, and click **Create Policy**. The system displays the **Create Image Bandwidth Throttling Policy** window.

Create Image Bandwidth Throttling Policy

Policy Name

Bandwidth Throttling Policy -Site A

Description

Bandwidth Throttling Policy for Site A

Cluster Categories

Cluster-Cat-SiteA:Sit... x

Bandwidth Limit

1000 KBps

Save

Figure 161: Create Bandwidth Throttling Policy

4. Specify the following information in the **Create Image Bandwidth Throttling Policy** window:

- **Policy Name:** Enter the image placement policy name.
- (optional) **Description:** Enter the description for the image placement policy.
- **Cluster Categories:** Type the name of the cluster category. The system displays the list of matching categories based on the typed entry. Use the [Add icon](#) and [Remove icon](#) to add and remove the required categories.

For more information on categories, see [Category Management](#) in the *Prism Central Admin Center Guide*.

- **Bandwidth Limit:** Enter the bandwidth throttling value in KBps. The system considers this value when you add the image using the URL option and while enforcing image placement policy to transfer image transfer from one cluster to another.

Note: The bandwidth limit task is serialized on each cluster to strictly enforce the bandwidth limit on each cluster. For example, when you create a new image, only one image create task runs on the cluster at a time. The second image creation task configured with another bandwidth limit can start only after completion of the first image creation task.

5. Click **Save**.

Managing Bandwidth Throttling Policies

This section describes how to manage the bandwidth throttling policies in Prism Central.

About this task

You can perform the following actions to manage the bandwidth throttling policies in Prism Central:

- Update an bandwidth throttling policy
- Delete an bandwidth throttling policy.

Procedure

To manage the bandwidth throttling policies, perform the following steps:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Compute > Images** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.
The system displays the **List** tab by default.
3. Select **Bandwidth Throttling Policies** in the **Policies** tab dropdown menu.
The system displays the bandwidth throttling policies across all the registered clusters.
4. Select the target bandwidth throttling policy checkbox, and choose **Update** from the **Actions** dropdown menu.
5. Update the information as per your requirement in the **Update Image Bandwidth Throttling Policy** window, and click **Save**. For more information about fields, see [Creating a Bandwidth Throttling Policy](#) on page 538.
To delete the bandwidth throttling policy, select the target bandwidth throttling policy checkbox and choose **Delete** from the **Actions** dropdown menu. The system prompts you to confirm the delete action. Click **Delete** to confirm the delete action.

Security Policy Management

Security policies inspect traffic that originates and terminates within a data center and help eliminate the need for additional firewalls within the data center.

Security policies are defined using Nutanix Flow which provides a policy-driven security framework. You can enable Nutanix Flow using the Flow-specific settings from the **Prism Central Settings** page. For more information, see [Prism Central Settings \(Infrastructure\)](#) on page 44.

For information on how to create and apply security policies on Basic VLAN Subnets, see [Flow Network Security \(formerly Flow Microsegmentation\) Guide](#).

For information on how to create and apply security policies on (advanced) VLAN Subnets and Overlay Subnets, see [Flow Network Security Next-Gen Guide](#).

For information about how to view the security policies across all clusters in prism Central, see [Security Policies Summary View](#) on page 540.

For information about how to view the detailed information about an individual security policy, see [Security Policy Details View](#) on page 544.

Security Policies Summary View

The Security policies summary view allows you to access a list of all the user-defined security policies across registered clusters.

To access the security summary view, perform the following steps:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Network & Security > Security Policies** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

The system displays the **Policies** tab by default with consolidated view of security policies across all registered clusters.

3. Select **Security Policies** in the **Policies** tab dropdown menu.

The system displays a view of security policies across all registered clusters.




Figure 162: Security Policies

Table 161: Security Policies Page - Field Description

Field	Description	Values
Name	Displays the security policy name. The policy is one of three types: application, quarantine, or isolation.	(name), Application, Quarantine, Isolation
Purpose	Describes the purpose of the security policy.	(text string)
Policy	Displays a high-level view of the policy configuration.	(boxed text)
Status	Displays the current status of the security policy as either Enforced or in Monitoring mode.	Enforced, Monitoring
Last Modified	Displays the time-line for the security policy. Last modified time-line or the creation time-line if the policy has never been modified.	Time-line. For example, 4 days ago, 6 minutes ago

You can perform the following actions for the security policies in the **Security Policies** summary view:

- Access the detailed information about an individual security policy. For more information, see [Security Policy Details View](#) on page 544.
- Create, import, or export a security policy. For more information, see [Flow Microsegmentation Guide](#).
- Filter the security policies list using **Filters** pane. For more information about **Filters** pane, see [Filters Pane - Security Policies Page](#) on page 543.
- Use the **Actions** dropdown menu to update, delete, enforce (if initial selection is monitoring), or monitor (if initial selection is enforce) a security policy. For more information, see [Flow Microsegmentation Guide](#).

Services Tab

Service is a group of protocol-port combination. You can use any of the default services or create a custom service.

The **Services** tab displays the list of all system and user-defined (custom) services.




Figure 163: Security Services Tab

The following table describes the fields that appear in the **Services** tab of **Security** page:

Table 162: Security Services Tab - Field Description

Field	Description	Values
Name	Displays the service name.	(service name)
Description	Displays the description of the service.	(description, string)
Services	Displays the protocol and port combination of the service	(Protocol-Port combination) For example, TCP 1027
Policies	Displays the security policies in which the services are attached.	(policy name)

You can perform the following actions for the services in the **Services** tab:

- Filter the services list based on the available field values using **Filters** pane. For more information about **Filters** pane, see [Filters Pane - Security Policies Page](#) on page 543.
- Create a service group. For more information about how to create a service group, see [Flow Microsegmentation Guide](#).
- Use the **Actions** dropdown menu to update, delete, or clone a service. For more information, see [Flow Microsegmentation Guide](#).

Note: The **Actions** dropdown menu is grayed out unless you select a service or if you select multiple services.

Addresses Tab

Address Group is a way to group one or many IP address or range. It can be used when you define security policies.

The **Addresses** tab displays the list of all user-defined addresses.



Figure 164: Security Addresses Tab

The following table describes the fields that appear in the **Addresses** tab of **Security** page:

Table 163: Security Addresses Tab - Field Description

Field	Description	Values
Name	Displays the address group name.	(address group name)
Description	Displays the description of the address group.	(description, string)
Subnets	Displays the subnet information	(subnet IP)
Policies	Displays the security policies in which the addresses are used.	(policy name)

You can perform the following actions for the addresses in the **Addresses** tab:

- Filter the services list based on the available field values using **Filters** pane. For more information about **Filters** pane, see [Filters Pane - Security Policies Page](#) on page 543.
- Create Address. For more information about how to create a address, see [Flow Microsegmentation Guide](#).
- Use the **Actions** dropdown menu to update or delete an address. For more information, see [Flow Microsegmentation Guide](#).

Note: The **Actions** dropdown menu is grayed out unless you select a address or if you select multiple addresses.

Filters Pane - Security Policies Page

The following table describes the fields available in the **Filters** pane:

Table 164: Filter Pane Fields - Security Policies

Field	Description	Values
Filter pane fields in Policies tab.		
Name	Filters based on the security policy name. Select a condition from the dropdown menu and enter a string in the field. The system returns a list of security policies that satisfy the security policy name condition/string.	(security policy name string)
Note: In this field, the condition menu options are Contains, Doesn't contain, Starts with, Ends with, and Equal to.		
Type	Filters on the security policy type. Select the checkbox for one or more of the policy types (application, quarantine, isolation, and VDI).	Application, Quarantine, Isolation, VDI
Status	Filters on the security policy status. Select the checkbox for Enforced, or Monitoring, or both.	Enforced, Monitoring
Filter pane fields in Services tab.		
Name	Filters based on the service name. Select a condition from the dropdown menu and enter a string in the field. The system returns a list of services that satisfy the service name condition/string.	(service name)
Note: In this field, the condition menu options are Contains, Doesn't contain, Starts with, Ends with, and Equal to.		
Description	Filters based on the description of the service.	(description, string)
Protocol	Filters based on the service protocol; TCP, UDP, and ICMP.	TCP UDP ICMP
Filter pane fields in Addresses tab.		
Name	Filters based on the address name. Select a condition from the dropdown menu and enter a string in the field. The system returns a list of addresses that satisfy the address name condition/string.	(address name)
Note: In this field, the condition menu options are Contains, Doesn't contain, Starts with, Ends with, and Equal to.		
Description	Filters based on the description of the address.	(description, string)

Security Policy Details View

To access the details view of an individual security policy, perform the following steps:

1. Log in to Prism Central.

2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Network & Security > Security Policies** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

The system displays the **Policies** tab by default with consolidated view of security policies across all registered clusters. .

3. Select **Security Policies** in the **Policies** tab dropdown menu.

The system displays a view of security policies across all registered clusters.

4. Click the target `<Security_Policy_Name>` to access the details view individual security policy.

Note: Replace `<Security_Policy_Name>` with the actual security policy name at your site.




Figure 165: Security Policy Details view




Figure 166: Security Policy Details View: Enforced Policy Example

You can perform the following actions in the **Security Policies** details view:

- Use toggle **Show Blocked Traffic** to view or hide the blocked traffic data. This toggle is available when policy status is **Enforced**.
- Use toggle **Show Discovered Traffic** to view or hide the discovered traffic data. This toggle is available when policy status is **Monitoring**.
- Update, delete, enforce (if initial selection is monitoring), or monitor (if initial selection is enforce) a security policy. For more information, see [Flow Microsegmentation Guide](#).
- Click the [Help icon](#) to open a help page in a separate tab or window.

The available actions appear in bold; other actions are grayed out. For grayed out options, the system displays a tool tip with the reason.

Storage Policy Management

Storage policies in Prism Central let you manage the storage attributes, such as replication factor, encryption, compression, and QoS, of entities like VMs and volume groups (VGs). You can configure entity-centric storage policies to manage the storage attributes of the entities as required.

A single storage policy can manage the attributes of several entities that are associated with various categories. When you install or upgrade the Prism Central instance to Prism Central pc.2024.1 or later versions, Prism Central configures a predefined or default storage policy.

For more information on the default storage policy, see [Default Storage Policy](#) on page 548.

A storage policy uses categories to associate with entities such as VMs and VGs. Using a storage policy, you can manage parameters of the entities such as Replication Factor, encryption, type of data compression (such as inline, no compression, or post-process compression), and IOPS or throughput throttling to be applied to the entities. When you apply storage policies to the entities using categories, the attributes are configured for the entities, and the entity data is transformed to comply with the storage policy. This mechanism is described in [Storage Policy Compliance](#) on page 557.

Important Considerations

- Prism Central supports storage policies only on AHV.
- Prism Central supports Data-at-rest encryption at the entity level with storage policies. In other words, when Prism Central applies a storage policy on an entity such as a VM or VG, it applies the policy even if data-at-rest encryption is not enabled on the cluster. Storage policies support entity-centric data-at-rest encryption only if the encryption on the underlying clusters is enabled during deployment.

Note: After you enable encryption using storage policies, you cannot turn it off. However, you can delete the storage policy on which the encryption is enabled. In such cases, any new data written into the entities is not encrypted. The data that was written prior to the deletion of the storage policy remains encrypted.

- Prism Central supports storage policies for VMs and VGs.
- Prism Central applies storage policies to Categories. You need to tag the VMs and VGs in Categories. You might also attach VGs to VMs. When you apply different storage policies through different categories to a VG and VM tagged to each other, the storage policies are applied independently. The storage policy applied to the tagged VG does not affect the tagged VM and vice versa.
- The storage policy engine takes 30 minutes or more to apply a newly created storage policy to the associated entities like VMs. The time taken depends on the number of entities associated with the storage policy.
- Storage policies use categories to associate with entities like VMs and VGs. You cannot associate an entity such as a VM or a VG directly with the storage policy. Create a category for the entities, and then assign the category with the storage policy.

When you associate a category with a storage policy, Prism Central associates all the entities that are associated with that category with the storage policy.

For information on how to create a category and assign a category to a VM or VG, see [Category Management](#) in the *Prism Central Admin Center Guide*.

- A cloned entity like a cloned VM or VG inherits the category and the attached storage policy of the original entity from which it is cloned. Thus, the storage policy engine applies the storage policy that is associated with the inherited category to the cloned entities.
- A cloned entity like a cloned VM does not inherit the compliance state of the storage policy applied to the entity from which it is cloned. The cluster runs a full compliance cycle to establish compliance with the storage policy of the cloned entity. In this case, the system displays the **Compliance State** of the cloned entity as **In Progress**, and not inherit the **Compliant** state.

For example, if VM1 has a storage policy applied to it and is in a **Compliant** state and VM2 is cloned from VM1, then VM2 does not automatically inherit the **Compliant** state of VM1.

- Shared data between the encrypted VMs and unencrypted VMs is encrypted.
- Shared data between a VM with replication factor 2 and another VM with Replication Factor 3 configurations complies with replication factor 3. In other words, in such a case, the shared data has three copies.

- For disaster recovery (DR), data-at-rest encryption enabled through entity-level storage policies at the source site, encrypts data at the destination site also for asynchronous replication only.
- For DR, when you use a snapshot or recovery point of a VM on which you have applied a storage policy, the snapshot restores the category of the VM and, therefore, the storage policy as well.

Note: The recovery point or snapshot of a VG does not restore the category, and therefore the storage policy, of the VG.

- After a DR (failover), you must create and apply storage policies and categories to the entities. The newly recovered data will not be encrypted until a policy is applied.

The policies and categories must be available after the failover.

Requirements

Ensure that you meet the following requirements to complete the built-in pre-checks and successfully create a storage policy:

- A unique name is assigned to every storage policy. Prism Central displays an error if the storage policy name matches any existing storage policy entry.
- A *defined* (non-default) value is provided for at least one storage policy attribute or property. For example, for **Compression** the default value is **Inherit from Cluster**, so you might instead select **On** with **Inline** as the value. You can select or provide a defined value for any one or more properties of the storage policy.

For the list of properties that you can configure during creation or updating of a storage policy, see [Creating a Storage Policy](#) on page 560.

- The throttled IOPS value is set greater than 99 (to enable rate-limiting or throttling) or equal to -1 (to clear the throttling) in the storage policies.
- The replication factor that you set is supported by the fault tolerance of the cluster that hosts the VMs or VGs.

For example, if the fault tolerance of the cluster is 1N/1D, then select the replication factor to **2** in the policy.

Note: The storage policy engine performs the pre-checks based on these requirements in the sequence listed below. For example, if you repeat an existing storage policy name for a new storage policy, the policy engine fails the storage policy creation, generates a failure message stating, Duplicate name , and stops other checks. The engine triggers the storage policy creation only after all the pre-checks are validated.

Limitations

- Prism Central does not support storage policies for entities running on ESXi and Hyper-V.
- Prism Central does not support storage policies for Nutanix Disaster Recovery with on-premises recovery AZ. For more information, see [Nutanix Disaster Recovery Guide](#).
- Nutanix supports storage policies for Nutanix Disaster Recovery with single Prism Central deployments. However, the storage properties defined for the container are applied to the snapshots replicated to the replication cluster. In other words, the snapshot data replicated to the remote site is transformed based on the settings defined for the remote site containers. The storage policies are applied only after the VM failover.
- You can associate multiple categories with one storage policy. However, you can associate one category with only one storage policy. After associating a category with one storage policy, you cannot associate the same category with another storage policy. When you associate the same category with another storage policy, Prism Central adds the category to the **Associated Categories** list but displays an alert (an exclamation mark) and turns off the **Save** button until you remove the category. The following message pops up when you hover on or click the alert (exclamation mark):

This category has an existing storage policy (<name of storage policy>)

associated with it.

- You can create a maximum of 100 storage policies irrespective of the type of Prism Central deployed (whether x-large, large, small, or x-small Prism Central with scale-out or single instance).

For information on Prism Central scalability and the type of Prism Central deployments, see the *Prism Central Scalability* topic in the release notes for the [Prism Central](#) version to be installed.

- You cannot create or apply a storage policy for a snapshot. Prism Central applies a storage policy to a live vDisk and the entire chain of snapshots of that vDisk.
- Prism Central does not encrypt the snapshots of the vDisk that are in the storage if the live vDisk is deleted before AOS applies the storage policy to the vDisk.
- Role-based access control (RBAC) can be configured for **Storage Policies**. Prism Central supports only Super Admin, Prism Admin, or Prism Viewer roles for Storage Policies. Prism Central does not support granular RBAC for storage policies. Do not attempt to configure permissions using APIs. Use **Super Admin** role permissions for **Storage Policies** operations.

You can create custom roles to assign specific permissions such as create, update, delete, or view permissions.

- If two entities that share data have different replication factors, the higher Replication Factor is applied to the shared data.
- If two entities share data and one or both entities have a storage policy with encryption enabled, the shared data is encrypted.
- If two entities share data and only one of the entities has a storage policy with compression enabled, the shared data is not compressed.
- Suppose you associate an entity with multiple storage policies. In that case, the policy engine sorts the storage policy list using a method based on the policy name and applies the highest-ranking policy from the sorted list to the entity. The highest-ranking storage policy is the one with the lexicographically smallest name in the list. If you add another storage policy later, and this newly added policy becomes the highest-ranking policy in the newly sorted list, the engine continues to apply the previously applied storage policy. To apply the newly added storage policy to the entity, remove the association of the entity with the previously applied storage policy.
- You cannot migrate a vDisk to other storage container if any storage policy is configured for the guest VM associated with the vDisk. For more information, see [Live vDisk Migration Across Storage Containers](#) in the *AHV Administration Guide*.

Default Storage Policy

Storage Policies in Prism Central lets you manage the storage attributes like Replication factor, encryption, compression, and QoS of entities like Virtual Machines (VMs) and Volume Groups (VGs). You can configure custom entity-centric storage policies to manage the storage attributes of the entities as required.

When you install or upgrade the Prism Central instance to pc.2024.1 or later versions, Prism Central configures a predefined or default storage policy. The default storage policy only defines the Replication Factor (RF) and Compression parameters. It allows the values for the other parameters to be taken from the cluster storage settings.

The default storage policy requires a minimum Prism Central version of pc.2024.1. AOS 6.8 and later versions fully support the default storage policy. Older versions of AOS do not fully support the default storage policy, even if they support storage policies broadly.

The **Storage Policies** dashboard in Prism Central provides an introductory dialog that describes the **Default Storage Policy**.

What's New



Introducing 'Default Storage Policy'

Quickly configure your VMs and Volume Groups by choosing 'Default Storage Policy' - a set of most commonly used storage configurations; during entity creation.

Policy Configuration

▪ Encryption	Inherit from Cluster	▪ Compression	Inline
▪ Replication Factor	2	▪ Storage QOS - IOPS	-

How does it work?

When creating a new VM or Volume Group, enable 'Default Storage Policy' under Management. This applies only when no other user-created storage policy has been chosen via another category.

Version Support: VMs- Clusters with AOS 6.1 or above and Volume Groups- Clusters with AOS 6.7 or above.

[View Release Notes](#)

Ok, Got It

Figure 167: Introducing 'Default Storage Policy'

You can apply the default storage policy to any entity, such as VMs and VGs, by enabling the **Enable 'Default-Storage' policy** toggle switch when you create or update a VM or VG. The default storage policy is associated with the **Storage: \$Default** category. When you turn on the **Enable 'Default-Storage' policy** toggle switch, the default storage policy is applied, like any other storage policy, to the VM or VG that is associated with the **Storage: \$Default** category.

The standard configuration of the default storage policy is as follows:

- Replication Factor (RF) is set to **2**.
- Encryption is set to **Inherit from Cluster**.
- Compression is set to **Inline**.
- No values are set for QoS.

Limitations for Default Storage Policy

- If the name of an existing storage policy conflicts with the **Default Storage** policy during its creation, Prism Central creates the **Default Storage** policy but does not issue any alert. Rename the existing policy. The **Default Storage** policy cannot be updated.
- When you create or update an entity like a VM or a VG,
 - The default storage policy (**Enable 'Default-Storage' policy** toggle switch) is not enabled for the entity by default.
 - If you assign a category associated with a non-default storage policy to an entity and also enable the **Enable 'Default-Storage' policy** toggle switch, the non-default storage policy gets precedence over the default storage policy. The default storage policy is not applied to the entity in such a case.

- **Storage: \$Default Category**

The **Storage: \$Default** category is used to assign the default storage policy to an entity such as a VM or a VG. Nutanix recommends that you do not associate any other policy with the **Storage: \$Default** category or make any other changes to this category.

Do not assign any other category that is already associated with another storage policy to the entity, such as a VM or VG, if you want to enable the **Enable 'Default-Storage' policy** switch for that entity. If you enable the **Enable 'Default-Storage' policy** switch for an entity that is already associated with another category, the **Storage: \$Default** category is enabled, but the storage policy associated with the other category overrides the default storage policy.

- **Edit the Default Storage Policy**

You cannot edit the default storage policy. You can only apply it to an entity by associating the entity with the **Storage: \$Default** category, enabling the **Enable 'Default-Storage' policy** switch for that entity, and ensuring that no other categories with storage policies are applied to that entity, thus overriding the default storage policy.

Storage Policies Summary View

The Storage policies summary view allows you to access a list of all the user-defined storage policies across registered clusters.

Note: The **Storage Policies** page is available if you are logged on as a view-only user. However, you cannot create a storage policy or modify an existing one.

To access the storage policies summary view, perform the following steps:

1. Log in to Prism Central.

2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Storage > Storage Policies** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

The system displays the **List** tab by default with a summary view of storage policies across all registered clusters. When you open the **Storage Policies** page for the first time, Prism Central displays the **Introducing 'Default Storage Policy'** page. Select **OK, Got it** to go to the **List** tab of the **Storage Policies** summary view.

For more information on the default storage policy, see [Default Storage Policy](#) on page 548.

Table 165: Storage Policies Page - Field Description

Field	Description	Values
Name	Displays the name of the Storage Policy. You can click the name to view the details page of the selected storage policy. The default storage policy (see Default Storage Policy on page 548) is named Default Storage and marked SYSTEM .	(Name of storage policy)
Replication Factor	Displays the Replication Factor selected.	Inherit from Container or 2 or 3
Encryption	Displays the encryption mode selected.	Enabled or Inherit from Cluster
Compression	Displays the compression mode selected.	Inline or Post Process (if On) Off Inherit from Cluster
Throttled Throughput (IOPS)	Displays the throttled throughput value in terms of IOPS.	(Integer number)
Throttled Throughput (MB/s)	Displays the throttled throughput value in MBps.	(Integer number)
Total Categories	Displays the total number of categories that are associated with the storage policy.	(Integer number)

Field	Description	Values
Realized Entities	Indicates the number of entities to which the storage policy is applied.	(Integer number)
Note: A single storage policy can be applied to multiple entities. Multiple storage policies cannot be applied to a single entity.		

You can perform the following actions for the storage policies in the **Storage Policies** summary view:

- Access the detailed information about an individual storage policy. For more information, see [Storage Policy Details View](#) on page 553.
- Create a storage policy. For more information, see [Creating a Storage Policy](#) on page 560.
- Update, clone, or delete a storage policy. For more information, see [Managing Storage Policies](#) on page 564.

Note: You cannot update or delete the **Default Storage** policy. You can only clone the **Default Storage** policy. Cloning the default storage policy does not create another default policy. It creates another non-default storage policy that you can modify to suit the needs of your deployment.

- Filter the storage policies list using **Filters** pane. For more information on **Filters** pane, see [Filters Pane](#) on page 50.

For field details, see [Table 165: Storage Policies Page - Field Description](#) on page 551.

- Export the table that contains the list of storage policies and their information to a file in a CSV format. For more information on **Export** option, see [Export](#) on page 54.
- Use the **Actions** dropdown menu to update, clone, or delete a storage policy. For more information, see [Managing Storage Policies](#) on page 564.

Associations Tab

The **Associations** tab displays the compliance state of the storage policies across the registered clusters.

The following table describes the fields that appear in the **Associations** tab of **Storage** page:

Note: A dash (-) is displayed in a field when a value is not available or applicable.

Table 166: Storage Associations Tab - Field Description

Parameter	Description	Values
Name	Indicates the name of the Storage Policy. You can click the name to view the details page of the selected storage policy.	(Name of storage policy)
Associated Categories	Displays the number of categories associated with the storage policy.	(Integer number)
Unrealized Entities		

Parameter	Description	Values
VMs	Displays the number of VMs that the system considers for compliance check and to apply the storage policy.	(Integer number)
VGs	Displays the number of VGs that the system considers for compliance check and to apply the storage policy.	(Integer number)
Realized Entities		
VMs	Indicates the number of VMs to which the storage policy is applied.	(Integer number)
VGs	Indicates the number of VGs to which the storage policy is applied.	(Integer number)
Compliance for Realized Entities	Provides a link to a dialog window that displays the state of compliance of the VMs and VGs covered by the policy.	 (Integer) In Progress  (Integer) Non Compliant  (Integer) Compliant




Figure 168: Storage Policy Compliance Pop-up Window

You can filter the storage policies list based on the storage policy name using **Filters** pane in the **Associations** tab. For more information, see [Filters Pane](#) on page 50.

For more information about Storage Policies, see [Storage Policy Management](#) on page 545.

- For information about creating or updating storage policies, see [Creating a Storage Policy](#) on page 560.
- For information about deleting storage policies, see [Deleting a Storage Policy](#) on page 566.

Storage Policy Details View

To access the details page for a storage policy, go to the storage policy **List** tab (see [Storage Policy Management](#) on page 545) and click the name of the storage policy you want to view the details of. You can also access the

details page by clicking the name of the storage policy wherever that name appears, such as in a dashboard widget or search result.




Figure 169: Storage Policies details view

The storage policy details page includes the following tabs:

- **Summary**—displayed by default when you open the storage policy details page. This tab provides information about the storage policy summarized into widgets.
- **Entities**—displays the list of entities including VMs and VGs associated with the storage policy.

The **Summary** tab of an individual storage policy provides the following widgets:

- **Configuration** - Displays the configuration attributes of the storage policy.
- **Associations**- Displays the count of associated realized and unrealized VMs, VGs, and their categories.
- **Compliance Status** - Displays the compliance status of the entities associated with this policy. If the policy is being applied and the compliance status is not yet known, the status is displayed as **In Progress**. For information about the fields available in the widgets, see [Storage Policy Widgets- Field Details](#) on page 554.

<Action> available above the widgets. Click the appropriate <Action> , such as **Update**, **Clone**, or **Delete** to run that administrative action on the storage policy. For more information about how to perform any <Action>, see [Managing Storage Policies](#) on page 564.

Storage Policy Widgets- Field Details

The following table describes the fields in the **Configuration**, **Associations**, and **Compliance Status** widgets.

Table 167: Storage Policy Widgets- Field Description

Field	Description	Values
Configuration widget		
Replication Factor	Indicates the Replication Factor selected.	Inherit from Container or 2 or 3
Encryption	Indicates the encryption mode selected.	Enabled or Inherit from Cluster

Field	Description	Values
Compression	Indicates the compression mode selected.	Inline or Post Process (if On) Off Inherit from Cluster
QoS Metric (Throttled IOPS)	Indicates the throttled throughput value in terms of IOPS.	(Integer number)
QoS Metric (Throttled Throughput)	Indicates the throttled throughput value in MBps.	(Integer number)
Associations widget		
Categories	Displays the count of VM and VG categories mapped to the selected storage policy.	(number of VM and VG categories)
Realized Entities	Indicates the number of entities that the storage policy is applied to. The storage policy acts on these entities and defines their storage properties.	(Integer number)
Non-realized Entities	Indicates the number of entities that the storage policy is not applied to. The storage policy is associated with these entities but does not act on them. When a storage policy is not applied to an entity like a VM or a VG, then the entity is listed as Non-realized Entity .	(Integer number)
Compliance Status widget		
Realized Entities	Indicates the number of entities that the storage policy is applied to.	(Integer number)
	Note: A single storage policy can be applied to multiple entities. Multiple storage policies cannot be applied to a single entity.	
Compliance State	Indicates the number of entities in each of the following compliance states: - In Progress: Indicates the number of entities that the policy has already been applied to and data transformation of entity's data (compliance) is in progress. - Compliant: Indicates the number of entities that are in compliance. - Non Compliant: Indicates the number of entities that are not in compliance	Compliant, In Progress, or Non compliant

Entities Tab

The **Entities** tab enables you to view the realized and non-realized VMs for the selected storage policy:

- Select **Realized** from the **Entities** tab dropdown menu to view the entities that the storage policy is applied to.

- Select **Unrealized** from the **Entities** tab dropdown menu to view the entities that the storage policy is not applied to.




Figure 170: Entities Tab

The following table describes the fields that appear in the **Entities** tab of storage policies details view:

Table 168: Entities Tab - Field Details

Field	Description	Values
Name	Indicates the name of the entity like VMs and VGs.	(entity name such as the name of the VM or VG)
Compliance Status	Indicates the compliance status of the entity as In Progress , Compliant or Non Compliant .	In Progress Compliant Non Compliant
Associated Via	Indicates the categories to which the entity and the storage policy are concurrently associated. The storage policy is associated with the entity because the storage policy is associated with this category that is associated with the entity. For example, StoragePolicy1 is associated with CategoryA which is in turn associated with VM1. The category displayed in Associated via is CategoryA. The other categories associated with VM1 but not associated with StoragePolicy1 are not displayed here.	(category name)
Cluster	Indicates the name of the cluster on which the entity resides.	(cluster name)
IP Addresses	Indicates the IP addresses configured for the entity.	(IP address)

You can perform the following actions for the storage policies in the **Entities** tab:

- Filter the storage policies list using **Filters** pane. For more information, see [Filters Pane](#) on page 50.
For field details, see [Table 168: Entities Tab - Field Details](#) on page 556.
- Export the table that contains the list of storage policies and their information to a file in a CSV format. For more information about this option, see [Export](#) on page 54.

Storage Policy Compliance

When you create a policy like a storage policy and apply it to entities like VMs and VGs, there is no way to know whether the application configured or transformed the entities to the parameters set in the policy. You need a mechanism to track the application of the policy and the changes made to the entities in alignment with the parameters set in the policies. Compliance provides that mechanism to track the application of policies and the transformation of the entities and provide feedback on the status at an entity level.

The compliance process begins when you associate a storage policy to a category and, therefore, to entities that the category is assigned to. The association applies the storage policies to the entities that the category is assigned to. For the several reasons, a storage policy is not applied to an entity.

After application, the parameters of the entities such as replication factor, encryption, and compression are transformed according to the configuration of the storage policy. Such transformation (compliance) does not occur for some entities that the category may be assigned to, any of the following, but not restricted to, reasons:

- Incompatible hypervisor: The node is running a hypervisor or version that does not support storage policies. Only AHV supports storage policies.
- Incompatible AOS version: The cluster is running an AOS version that does not support storage policies.
- Multiple storage policies applied: When you apply multiple storage policies to an entity, only one is applied. Other storage policies report that entity as non compliant.

Compliance, thus, tracks the process from application to transformation and reports the status in the storage policy **Summary** tab. AOS applies the storage policy parameters in a specific order that **Compliance** also uses to report status. This order of implementation and tracking of the storage policy parameters is as follows:

- Replication Factor: Replication Factor indicates data availability which is the most factor for an entity.
- Encryption: Encryption indicates data security which is the second most important factor for an entity.
- Compression: There may be cases where compression may not be possible for an entity. So in such a case, if Replication Factor and Encryption parameters are applied and entity data is transformed accordingly, then the compliance engine reports the state for the entity as **Compliant**.
- Throughput: This parameter has lower priority in implementation and compliance reporting. The implementation is similar to the Compression parameter.

Compliance reports the states separately for VMs and VGs, even if the VMs or VGs are tagged to each other. For example, if VM1 is tagged to VG1. VM1 is associated with Category1 and VG1 is associated with Category2. Storage policies SP1 and SP2 are associated to Category1 and Category2 respectively. **Compliance** reports the states of VM1 and VG1 separately. It reports the states separately even if VM1 and VG1 are associated to the same Category1, and hence storage policy SP1.

Compliance reports the following states:

Realized Entities

The compliance engine reports this state when a storage policy is applied to the entity. Application of storage policy means that the parameters set in the storage policy are identified for the entity that it is applied to. So when the system reads the storage policy and identifies the Replication Factor, Encryption, Compression and throttled throughput QOS metric parameters for the entity, the compliance engine reports the entity as **Realized Entity**.

Non-realized Entities

There may be several reasons why a storage policy could not be applied to (realized on) an entity. For example, storage policies cannot be applied to ESXi VMs (non-Nutanix entity).

In Progress

Compliance engine reports this state when the one or more parameters set in the storage policy are applied to an entity and the entity data is being transformed to comply with those parameters.

Compliant

The compliance engine reports this state when the transformation of the entity data is complete and all the high priority parameters like Replication Factor and Encryption have been set for the entity in accordance with the settings in the storage policy.

Non Compliant

The compliance engine reports this state when the transformation of the entity data could not be completed and one or more high priority parameters like Replication Factor and Encryption could not be set for the entity in accordance with the settings in the storage policy. **Non Compliant** also indicates unknown state of transformation.

The compliance states are displayed on the storage policy details page. For more information, see [Storage Policy Details View](#) on page 553.

Storage Policy Based Replication Factor

Replication Factor is the number of replication copies of any data made on the cluster. For example if you have a Replication Factor of 2 for a container, then the data in that container is written twice on nodes in different fault tolerance domains. This is dependent on how you have configured the fault tolerance domains.

You can use a storage policy to apply Replication Factor to entities like VMs and VGs. When you apply a storage policy to an entity, the Replication Factor defined in the storage policy is applied to the entity. Storage policy compliance checks whether the entity data is compliant with the Replication Factor setting in the storage policy. You can use a storage policy to apply Replication Factor with Deduplication enabled.

Note: Prism Central supports Policy-based replication factor configuration for entities on storage containers that have Deduplication enabled.

Limitations

Configuring Replication Factor for entities using storage policies is not possible in some cases.

- Storage Policies do not support setting Replication Factor of 1. You can only use a storage policy to set Replication Factor (RF) 2 or 3.
- Prism Central does not support Policy-based replication factor configuration for entities on storage containers that have Erasure Coding.
- A storage policy Replication Factor setting is not applied when the RF setting in the storage policy is higher than the Replication Factor setting in the cluster.

For example, if the RF of a cluster is 2, but a storage policy with RF 3 is applied to an entity in that cluster, the entity is reported as **Realized** but **Non-compliant**.

- Shared data between a VM with Replication Factor 2 and another VM with Replication Factor 3 configurations complies with Replication Factor 3. In other words, the shared data has 3 copies.
- A cloned entity like a cloned VM or VG inherits the category, and the storage policy attached thereto, of the original entity that it is cloned from. Thus, the storage policy engine applies the storage policy that is associated with the inherited category to the cloned entities.

Storage Policy-Based Encryption

Storage policies allow you to make data-at-rest encryption choices at the entity level.

Deployments can continue to have data-at-rest encryption capabilities scoped for the entire cluster. Storage policies provide the additional option to control the encryption scope decisions at the entity (VM or VG) level.

Note: Prism Central does not support policy-based encryption for storage containers with erasure coding or deduplication enabled.

The key manager system (KMS) choice is configured as a cluster-wide setting, using either the Nutanix built-in native key manager or a third-party key manager. As part of the key manager configuration, the scope of encryption-at-rest is defined either as cluster-wide or specific entity scoping. Entity scoping means that the encryption-at-rest scope is defined at the container or the data policy level.

Important:

- When you select the cluster, you can manage the KMS type by clicking **Actions > Manage KMS Type**.
- From Prism Central, if you try to configure storage policy-based encryption for any single VM or Category entity without enabling **Entity Encryption** on that cluster, the entity is marked as **Non Compliant**, and Prism Central does not transform the entity data to be encrypted.

Log on to Prism Central and perform the following steps:

1. From [Application Switcher Function](#) on page 40, select the **Infrastructure** application and navigate to **Hardware > Clusters**
2. Select the target cluster, and from the **Actions** dropdown menu, choose **Enable Data-at-Rest Encryption**.
3. In the **Data-at-Rest Encryption** window, specify the following information:

- **Select Encryption Type:** Select **Entity Encryption** and click **Save Encryption Type**.
- **Select Key Management Server (KMS):** Select **Native KMS (Local)** and click **Enable Encryption**.

The system prompts you to confirm the action.

4. In the **Type SET to Confirm** window, enter **SET** and click **Set Encryption**

When a storage policy is used, AOS data for the VM or VG is encrypted on the disk for all the vDisks attached to the VM or VG. The main features include:

- The data-at-rest encryption, enabled through entity-level storage policies at the source site, encrypts data at the destination site also for asynchronous replication only.
- After encryption is enabled in a policy, it cannot be disabled. If the VM or VG moves out of the policy or the policy is deleted, then the new writes to the VM or VG are unencrypted and existing data remains encrypted.
- The system automatically generates a unique key for each storage container.
- Multi-tenant requirements that have a unique key per tenant can be fulfilled by ensuring that the set of VMs belonging to each tenant is mapped to its own container. Nutanix recommends having only 50 tenants or less per cluster.
- Encryption can be enabled for new and existing VMs or VGs.
- Storage policies supports volume groups. For any VG attached to a VM that is associated with a storage policy enabling encryption, the data stored in the VG is encrypted only if the storage policy associated with the VG also enables encryption. In other words, data encryption in a VG attached to a VM is independent of the storage policy associated with the VM. It is not derived from the storage policy applied to the VM that the VG is attached to.
- Unencrypted VM must remain unencrypted during the cloning process.
- Shared data between the encrypted entities and unencrypted entities is encrypted.

For information about containers, see the [Storage Management](#) section in the *Prism Element Web Console Guide*.

For information about configuring key manager, see the [Security Guide](#).

Creating a Storage Policy

You can create or update a storage policy in Prism Central only.

Before you begin

Ensure that you have created the necessary categories and assigned the entities such as VMs and VGs to the categories.

For more information about how to create a category, see [Creating a Category](#) in the *Prism Central Admin Center Guide*.

For more information about how to assign a category to the entity, see [Assigning a Category to an Entity](#) in the *Prism Central Admin Center Guide*.

About this task

Prism Central provides a two-step framework to create a storage policy that involves the following steps:

- **Configuration** step - Enables you to configure the data replication (Replication Factor), data security (Encryption), data reduction (Compression) and QoS (IOPS or Throughput throttling) aspects of the storage policy.
- **Association** step - Enables you to associate the storage policy with categories that the entities (like VM or VG) are associated with.

Procedure

To create a storage policy, perform the following steps:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and go to **Storage > Storage Policies** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

The system displays the **List** tab by default with summary view of storage policies across all registered clusters.

3. Click **Create Storage Policy**.

The system displays the **Create Storage Policy** window with the **Configuration** step.

Figure 171: Create Storage Policy - Configuration Step

Configure the following fields in the **Configuration** step. Select a defined (non-default) value for a least one of the properties.

Table 169: Configuration Step - Field Description

Field	Description and Values
Name	Enter a unique name for the storage policy. Note: Do not repeat a name previously used for another storage policy. The built-in pre-check fails if you enter a non-unique name.
Data Redundancy	

Field	Description and Values
Replication Factor	Select one of the following values from the dropdown list. • Inherit from Container (default)—the category or entity inherits the replication factor (2 or 3) configuration from the container. This is the default value for Replication Factor. • 2—to enable replication factor 2 (replication factor 2—Original plus 1 copy) through the storage policy. • 3—to enable Replication Factor 3 (replication factor 3—Original plus 2 copies) through the storage policy. Note: The replication factor selected must be supported by the fault tolerance of the cluster. For example, if the fault tolerance of the cluster is 1N/1D, then selecting replication factor 3 in the policy leads to non-compliance.
Data Security	
Encryption	Select one of the following values from the dropdown list.
Value	• Enabled—to enable encryption through the storage policy. Note: After you enable encryption using storage policies, you cannot disable it. • Inherit from Cluster (default)—the category or entity inherits the encryption configuration from the cluster. This is the default value for Encryption.
Data Reduction	
Compression	Select one of the following values from the dropdown list.
Value	• On—to enable compression. If you select On then you must select the type of compression that you want to apply. You can select Inline compression (default) or Post Process compression. If you select Post Process, the data is compressed with some delay (up to 3600 seconds) that is set by the system default. • Off—to disable compression. • Inherit from Cluster (default)—the category or entity inherits the compression configuration from the cluster. This is the default value for compression.
IO Performance (Also see the <i>Relationship between Block Size, IOPS and Throughput</i> table for information about how Throttle IOPS or Throughput is calculated based on block size)	
Quality Metric	Select one of the following values from the dropdown list. • Throughput • IOPS

Field	Description and Values
Throttled Value	Provide a numeric value.
Note: For IOPS the value must be greater than 99 or equal to -1 (minus 1).	

Throttle IOPS or Throughput Relationship: Set the Throttle IOPS or Throughput based on the block size. Nutanix uses a fixed block size of 32k. For example, with a block size of 32K, if you set a Throttle IOPS value of 1000 IOPS, the Throttle Throughput would be 31.25 MBps. See the table.

Table 170: Relationship between Block Size, IOPS and Throughput

Block Size	IOPS	Throughput (MBps)
8k	1000	7.81
16K	1000	15.63
32K	1000	31.25
64K	500	31.25
128K	250	31.25

4. Click Next.

If you did not select or enter an appropriate value and the pre-checks failed, the relevant error is displayed below the field that failed the pre-check. Select or enter an appropriate value, and click **Next** again.

The system displays the **Association** step.

5. In the **Association** step, select the category or categories you want to associate with the storage policy in the **Add Category** field.

Type the first 2-3 characters of the category name you want to associate with in the **Add Category** field. The system displays a dropdown list of the matching categories based on the typed entry.

Use the [Add icon](#) and [Remove icon](#) to add and remove the required categories in the **Associated Categories** list.

Note: You can add multiple categories for one storage policy. When you associate a category with one storage policy, you cannot associate the same category with another storage policy.




Figure 172: Create Storage Policy - Association Step

6. Click **Save** to save the storage policy.

Managing Storage Policies

This section describes how to update, clone, and delete the storage policies in Prism Central

About this task

You can perform the following actions to manage the Storage policies in Prism Central:

- Update a Storage policy.

Note: You cannot update the **Default Storage** policy.

- Clone a Storage policy.

Note: You can clone the **Default Storage** policy.

- Delete a Storage policy.

Note: You cannot delete the **Default Storage** policy.

When you delete a storage policy, the entities associated with the deleted storage policy return to their default state, or the configurations from any other storage policy associated with the category are applied to the entities. Nutanix recommends that you remove all the categories and entities associated with the storage policy.

When a storage policy is deleted, it is marked for removal in the database. Prism Central stops displaying that storage policy. The process stops applying the storage attributes of the policy and instead applies the storage attributes configured for the container. However, the VM continues to display the storage policy as if it still applies until the next system scan occurs and the policy is removed from the VM information. A Kanon service scan occurs every 1800 seconds and refreshes the information.

- Manage associations.

Note: You cannot update the category association of the **Default Storage** policy. The **Default Storage** policy is associated with the **Storage:\$Default** category.

Procedure

To update a storage policy, perform the following steps:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Storage > Storage Policies** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61. The system displays the **List** tab by default.
3. Select the target storage policy check box, and choose **Update** from the **Actions** dropdown menu.
4. Update the information as per your requirement in the **Update Storage Policy** window, and click **Save**. For field details, see [Creating a Bandwidth Throttling Policy](#) on page 538.

To delete the storage policy, select the target storage policy check box and choose **Delete** from the **Actions** dropdown menu. The command prompt prompts you to confirm the deletion action. Click **Delete** to confirm the delete action.

Note: You can also click **Preview changes** to view the categories and entities that are affected by the deletion of the storage policy.

To clone the storage policy, select the target storage policy check box and choose **Clone** from the **Actions** dropdown menu. The system displays the **Create copy of <Storage Policy-Name>** window with the same values populated in it as specified for the selected storage policy.

Note: The **<Storage Policy-Name>** indicates the actual name of the storage policy you selected to clone.

During cloning, the categories to which the original storage policy is attached are removed. The clone storage policy does not inherit the categories from the original storage policy.

Enter the field details as per your requirement, and click **Save**. For field details, see [Creating a Storage Policy](#) on page 560.

To update the category associations only in an existing storage policy, navigate to the **Entities** tab in [Storage Policy Details View](#) on page 553 and click **Manage Associations**. Type the first 2-3 characters of the

category name in the **Add Category** field and use the [Add icon](#) and [Remove icon](#) to add and remove the required categories in the **Associated Categories** list.

Note: You can associate multiple categories with one storage policy. When you associate a category with only one storage policy. You cannot associate one category with multiple storage policies

The **Storage: \$Default** category is associated with the **Default Storage** policy. Do not associate the **Storage: \$Default** category with any other storage policy.

Deleting a Storage Policy

You can delete a storage policy. When you do so, the entities associated with the deleted storage policy return to their default state or the configurations from any other storage policy that is associated with the category are applied to the entities.

Before you begin

Nutanix recommends that you remove all the categories and entities associated with the storage policy.

About this task

When a storage policy is deleted, it is marked for removal in the database. Prism Central stops displaying that storage policy. The process stops applying the storage attributes of the policy and instead applies the storage attributes configured for the container. However, the VM continues to display the storage policy as if it still applies, until the next system scan occurs and removes the policy from VM information.

Remove associations of the storage policy to categories in the **Categories** or **Entities** tabs of the storage policy details page.

Delete a storage policy from the **Storage Policies** dashboard or the **Summary** tab of the storage policy details page.

Do the following to delete a storage policy:

Procedure

1. Click the name of the storage policy. On the storage policy details page, click **Delete**.

The **Delete** <storage-policy-name> dialog box is displayed as follows:

Delete [redacted]

This will remove the enforcement of the storage policy '[redacted]' on 0 associated entities across [redacted] Category. Any policy overridden by [redacted] might get enforced on the entities.

Are your sure you want to delete the storage policy '[redacted]'? This action cannot be undone.

☐ I understand the impact of this action.

[Preview Changes](#)

Cancel

Delete

Figure 173: Delete Storage Policy

2. (Optional) Click **Preview Changes** to view the **Changes Preview** dialog box.

The **Changes Preview** dialog box lists the categories and entities that are affected by the deletion of the storage policy.

Tip: If you have already removed all the associations of the storage policy, then the list of categories and entities is empty.




Figure 174: Change Preview

After you preview the impact, click **Back** to go back to the **Delete <storage-policy-name>** window.

3. Select the **I understand the impact of this action.** checkbox.

The **Delete** is enabled only after you select the **I understand the impact of this action.** checkbox.

4. Click **Delete**.

PRISM SELF SERVICE ADMINISTRATION

This section describes how to configure and administer Prism Self Service.

The Prism Self Service feature allows you to create projects where consumers of IT infrastructure within an enterprise—individual users or teams such as development, test, and DevOps—can provision and manage VMs in a self-service manner, without having to engage IT in day-to-day operations.

For information about user roles and configuration workflow for Prism Self Service, see [Prism Self Service Setup](#) on page 569.

Prism Self Service Setup

Prism Self Service represents a special view within Prism Central. While Prism Central enables infrastructure management across clusters, Prism Self Service allows end users to consume that infrastructure in a self-service manner. Prism Self Service uses the resources provided by a single AHV cluster.

Important: Prism Self Service is only supported with AHV hypervisor, and not supported with other hypervisors.

Note: A SelfServiceContainer storage container is created on the target cluster and used by Prism Self Service for storage and other feature operations. To ensure proper operation of these features, do not delete this storage container. For more information about storage containers, see [Storage Container Management](#) on page 289

Prism Self Service Roles

There are three roles to consider when configuring Prism Self Service:

1. *Prism Central administrator.* The Prism Central administrator adds an Active Directory or OpenLDAP directory service that includes the pool of self-service users and (optionally) creates one or more project manager. Prism Central administrators also create VMs, images, and network configurations that may be consumed by self-service users.
2. *project manager.* The project manager performs the following tasks:
 - Creates a project for each team that needs self service and adds users and groups to the projects.
 - Configures roles for project members. A project member can access only the entities or perform only the actions defined in the role assigned to that project member.
 - Publishes VM templates and images to the catalog.
 - Monitors resource usage by various projects and its VMs and members, and then adjusts resource quotas as necessary.

It is optional to configure a separate project manager because a Prism Central administrator can perform any of these tasks. However, if you would like to authorize users to administer end-user VM or application provisioning, you can use this role to give them access to virtual infrastructure without giving them access to physical infrastructure.

Caution: Project managers have full access to all VMs running on the Nutanix cluster, including infrastructure VMs not tied to a project. project managers can assign infrastructure VMs to project members, add them to the catalog, and delete them even if they do not have administrative access to Prism Central. Consider these privileges when appointing project managers, and make sure to communicate to project managers the need to exercise caution when working with infrastructure VMs.

After a Prism Central administrator has designated a user as a project manager, the Prism Central administrator cannot limit the user's privileges. Therefore, if you plan to delegate the responsibilities of a project manager to an Active Directory or OpenLDAP directory group, be sure that you want to delegate the responsibility to all the users in the group. If the user group is large or includes users that must not have project manager privileges, Nutanix recommends that you create a separate Active

Directory or OpenLDAP directory group for the users to whom you want to delegate project manager responsibilities.

3. *Project user*. These are the users assigned to a project by a self-service administrator. They can perform any action that the project manager grants them. The permissions are determined by the roles assigned to the users and groups in the project. When project users log in, they see a custom self-service GUI interface that shows only what the role permissions allow. Project users create and manage only what they need.

Prism Self Service Workflow

To configure Prism Self Service, perform the following tasks:

Table 171: Prism Self Service Workflow

Sequence	Tasks	Description
1	Specify an Active Directory or OpenLDAP directory service for Prism Self Service.	For more information, see Configuring Centralized Authentication (Active Directory/OpenLDAP) information in <i>Security Guide</i> .
2	Add one or more project managers.	For more information, see Managing Prism Self Service Admins information in <i>Prism Central Admin Center Guide</i> .
3	Create a project for each team that needs self service and add users and groups to the projects.	For more information, see Project Overview information in <i>Prism Central Admin Center Guide</i> .
4	Configure roles for project members. You can also allow project members to create their own VMs	For more information, see Controlling User Access (RBAC) in <i>Security Guide</i> . Note: Role Based Access Control (RBAC) is an independent feature, so you can configure RBAC without configuring self service or creating projects. However, projects allow you to enforce RBAC in a more granular way.
5	Add VM templates and images to the catalog.	For more information, see Catalog Management on page 283.
6	Create VMs as needed and assign them to project members as appropriate.	For information about how to create a VM, see Creating a VM (Self Service) on page 571 and Creating a VM from Catalog Items (Self Service) on page 571. For information about how to assign a VM to a project member, see Assigning a VM to a Project Member on page 578.
7	Monitor resource usage for the projects and adjust resource quotas as needed.	For more information, see the following information in <i>Prism Central Admin Center Guide</i> : • Project Details View - To monitor resource usage for a project. • Quota Policy Overview - To define the resource quotas.

Creating a VM (Self Service)

About this task

When you are logged in to Prism Central as a project manager or a project member, you can create a VM if you have the required permissions to create a VM.

Procedure

In Prism Central Self Service framework, you can create a VM in the following two ways:

- Using the workflow that an admin user follows. For more information, see [Creating a VM through Prism Central \(AHV\)](#) on page 149.
- Using a source file (image or VM template) that is stored in the Prism Central catalog. For more information, see [Creating a VM from Catalog Items \(Self Service\)](#) on page 571.

For information about Prism Self Service, see [Prism Self Service Administration](#) on page 569.

Creating a VM from Catalog Items (Self Service)

This section describes how to create a VM from catalog items.

About this task

When you are logged on as a project manager or a project member with permission to create a VM, you can create a VM based on a source file stored in the Prism Central catalog. For information about how to add a source file (image or VM template) to a catalog, see [Adding a Catalog Item](#) on page 285.

For information about Prism Self Service, see [Prism Self Service Administration](#) on page 569.

Note: If you are not logged in as a project manager or a project member and you want to create a VM without the source file stored in the Prism Central catalog, you can follow the procedure as described in [Creating a VM through Prism Central \(AHV\)](#) on page 149.

Procedure

To create a VM from catalog items, perform the following steps:

1. Log in to Prism Central as a project manager or a project member.

2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Compute > VMs** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.
The system displays the **List** tab by default with all the VMs across registered AHV clusters.




Figure 175: VM Summary View - Self Service Administrator login

3. Click **Create VM from Catalog Item**.

Note: The **Create VM from Catalog Item** option only appears if you are logged in as project manager or a project member with permissions to create a VM.

The system displays the **Create VM from Catalog Item** page.

4. In **Catalog Type**, select one of the following source for the VM:


- » **VM Template** - Used to create a VM from a VM template in the catalog.

A VM template includes all the configuration information required to create a VM. You can use a template to quickly deploy a VM without specifying all the configuration information.

- » **Image** - Used to create a VM from a mounted disk image.

Disk images can be CD-ROM images such as installer ISO images or images of hard drives that contain pre-installed applications and data. Disk images enable you to share data with other VMs, however you need to specify configuration information for the VM that you create using the disk image.

5. In **Catalog Item** field, perform either of the following actions based on the source you select in the **Catalog Type** field:
 - If **VM Template** is selected as the source in the **Catalog Type** field, select the target VM template in the **Catalog Item** field. The available VM templates from the catalog are listed. If the target VM image does not appear in the list, you can search for it by name in the search field.



Create VM from Catalog Items [X]

Catalog Type
VM Template

Catalog Item
kkkk

CPU	Memory	Storage	Last Updated
4 vCPU	3 GiB	0 GiB	02/08/22, 07:30:45 AM

[Cancel] [Begin]

Figure 176: Create VM from Catalog Items

- If **Image** is selected as the source in the **Catalog Type** field, select one or more of the disk images. The available disk images from the catalog are listed. If the target disk image does not appear in the list, you can search for it by name in the search field.

Create VM from Catalog Items

Catalog Type

Image

Catalog Item

virtio-win-0.1-74.iso × virtio-win-0.1-74.iso ×

☒ virtio-win-0.1-74.iso

☒ virtio-win-0.1-74.iso

CDROM	59 MiB	02/04/22, 12:11:17 AM

virtio-win-0.1-74.iso

Type	Size	Last Updated
CDROM	59 MiB	02/04/22, 12:11:16 AM

Cancel

Begin

Figure 177: Create VM Window (Configuration)

- Click **Begin**.

The **Create VM** page appears.

7. In the **Configuration** step, specify the following information:
 - a. **Name:** Enter a name for the VM.
 - b. **Project:** Select the project associated with this VM from the dropdown menu.
 - c. **Cluster:** Select the target cluster from the dropdown menu on which you intend to create the VM.
 - d. **Number of VMs:** Enter the number of VMs you intend to create. The VM names are suffixed with sequential numbers (1 to 5).
 - e. **vCPU(s):** Enter the number of virtual CPUs to allocate to this VM.
 - f. **Number of Cores per vCPU:** Enter the number of cores assigned to each virtual CPU.
 - g. **Memory:** Enter the amount of memory (in GB) to allocate to this VM.

Create VM

1 Configuration 2 Resources 3 Management 4 Review

Name
test-catalog-item-vm

Description
(Optional)

Project
project-1

Cluster
auto_cluster_nested_61f2e0f082e14f222b771a4f

Number of VMs
1

Cancel Next

Figure 178: Create VM Window (Configuration)

8. In the **Resources** step, specify the following information.

- **Disks:** Displays the disks and CD-ROMs attached to this VM.
 - If **VM Template** is selected as the source in **Catalog Type** field, the system displays a fixed list of devices. Select the device from which you want to boot the VM.
 - If **Image** is selected as the source in **Catalog Type** field, there are **+ New Disk** and **+ New CDRM** links above the list.
 - Click **+ New Disk** to add a disk. The system displays a new line at the bottom of the list. Specify the disk size and then click **Save** at the end of line to add the disk. The name and type field values are entered automatically.
 - Click **+ New CDRM** to add a CD-ROM. The system displays a new line at the bottom of the list. Click **Save** at the end of line to add the CD-ROM. All the field values are entered automatically.

Repeat this action until you have added all the desired disks and CD-ROMs. When the list is complete, select the device from which you want to boot the VM.

9. In the **Resources** step, perform the following actions to configure the network interface for the VM:

- a. Click **Attach to Subnet**. The system displays the **Attach to Subnet** window.

VLAN ID	IPAM	Virtual Switch
0	Not Managed	br0

Figure 179: Attach to Subnet Window

- b. Select the target subnet from the **Subnet Name** dropdown menu.
The list includes all defined networks. For more information, see [Network Configuration for VM Interfaces](#) information in *Prism Element Web Console Guide*.
- c. **VLAN ID:** This is a read-only field that displays the VLAN ID.
- d. **IPAM:** This is a read-only field that informs you if the subnet is IPAM managed or not.
- e. **Virtual Switch:** This is a read-only field that displays the name of the virtual switch associated with the subnet.
- f. Select the state for the network that you want it to operate in after VM creation in **Network Connection State** field. The options are Connected or Disconnected.
- g. Click **Save** to create a network interface for the VM. The system returns to the **Create VM** window.
- h.
Repeat this step to create additional network interfaces for the VM.

10. In the **Resources** step, perform the following action to specify boot configuration for the VM:
Select one of the following firmware to boot the VM in **Boot Configuration** field:
 - » **Legacy BIOS Mode:** Select legacy BIOS to boot the VM with legacy BIOS firmware.
 - » **UEFI BIOS Mode:** Select UEFI to boot the VM with UEFI firmware. UEFI firmware supports larger hard drives, faster boot time, and provides more security features. For more information about UEFI firmware, see the [UEFI Support for VM](#) section in the *AHV Administration Guide*.
11. (Optional) In the **Management** step, perform the following action to assign a VM category.
Enter the category name or select one from the **Categories** dropdown menu. The policies associated with the category value are assigned to the VM.

This field acts like a search field; it provides a list of matching categories as you enter a string. Select the desired category value when you see it in the list. You can repeat this step for as many categories as desired.
12. In the **Management** step, perform the following actions for guest customization:

The screenshot shows the 'Create VM' window in the Management step. The progress bar indicates the current step is 'Management'. The 'Categories' section has a search bar with the text 'AnalyticsExclusions:AnomalyDetection'. Below this is a 'Guest Customization' section with two dropdown menus: 'Script Type' set to 'No Customization' and 'Configuration Method' set to 'Custom Script'. At the bottom are 'Back', 'Cancel', and 'Next' buttons.

Figure 180: Create VM Window (Management)

- Select Cloud-init (for Linux VMs) or Sysprep (for Windows VMs) in the **Guest Customization** dropdown menu.
- Select the VM script customization type from the **Script Type** dropdown menu. The options are *No Customization*, *Sysprep(Windows)*, or *Cloud-init (Linux)*.
- Select the VM script customization method in the **Configuration Method** dropdown menu. This field is activated when the **Script Type** is either *Sysprep(Windows)* or *Cloud-init (Linux)*. The options are *Custom Script* or *Guided Script*.

The system displays the fields required to configure Cloud-init and Sysprep, such as options to specify a configuration script and the option to upload a script.

13. To specify a user data file (Linux VMs) or answer file (Windows VMs) the script for unattended provisioning, perform one of the following actions:

- » If the file is available on your local computer, click **Upload Script**, and select and upload the file.
- » Create or paste the contents of the file in the text box below the **Upload Script**.

Note: The script type supports the following file formats.

- Sysprep: XML
- Cloud-init (Linux): YAML, JSON, or Shell.

14. In the **Review** step, Click **Create VM** to create the VM and close the **Create VM** window. The system displays the new VM in the [VMs Summary View](#) on page 100.

Managing a VM (Self Service)

About this task

After you log in to Prism Central as a project manager or a project member you can perform various actions on the VM to manage the VM configuration.

For information about all the actions, see [Managing a VM through Prism Central \(AHV\)](#) on page 161.

Note:

- The system allows you to perform only those actions for which you have permission.
- The available actions appear in bold; other actions are grayed out. The available actions depend on the current state of the VM and your permissions.

Procedure

You can perform these actions using any of the following methods:

- Select the target VM in the **List** tab of the **VMs** page, and choose the required action from the **Actions** dropdown menu. For more information, see [VMs Summary View](#) on page 100.
- Right-click on the target VM in the **List** tab of the **VMs** page, and select the required action from the **Actions** dropdown menu.
- Go to the details page of the target individual VM, and perform the required action. For more information, see [Viewing VM Details](#) on page 113.

Assigning a VM to a Project Member

This section describes how to assign a VM to a project member.

About this task


A VM is assigned to a member within the scope of a particular project. You can assign a VM to only one project member.

Procedure

To assign a VM to a project member, perform the following steps:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Compute > VMs** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.
The system displays the **List** tab by default with all the VMs across registered clusters in **Nutanix** environment. For information about how to access the list of non-nutanix VMs managed by an external vCenter, see [VMs Summary View](#) on page 100.
3. Select the **Manage Ownership** action using any of the following methods:
 - Select the target VM checkbox in the **List** tab, and choose the **Manage Ownership** from the from the **Actions** dropdown menu.
 - Right-click on the target VM in the **List** tab, and select **Manage Ownership** from the displayed dropdown menu.
 - Select the **Manage Ownership** action from the **More** dropdown menu in [Viewing VM Details](#) on page 113.The system displays the **Manage VM Ownership** window.

4. Enter the following information in the **Manage VM Ownership** window, and click **Save**.
- **Project:** Select the target project from the dropdown menu.
 - **User:** Enter a user name. A list of matches appears as you enter a string; select the user name from the list when it appears.



The screenshot shows a window titled "Manage VM Ownership" with a question mark icon and a close (X) icon in the top right corner. Below the title bar, there are two input fields. The first is labeled "Project" and is a dropdown menu currently showing "Project-SiteA". The second is labeled "User" and is a text input field containing the text "sspgroup1". At the bottom right of the window, there are two buttons: "Cancel" and "Save".

Figure 181: Assign VM to a Project Member

The selected VM is assigned to the project member.

ADDITIONAL OPERATIONS

After installing Prism Central, you can perform the following additional operations:

Note: For information about how to install Prism Central, see [Prism Central Deployment](#) on page 12.

- Log in and out through any supported browser. For more information, see [Logging Into Prism Central](#) on page 32 and [Logging Out of Prism Central](#) on page 632.
- Monitor the status of this Prism Central instance. For more information, see [Managing Prism Central](#) on page 581.
- Check version information. For more information, see [Finding the Prism Central Version](#) on page 616.
- Modify login page and automatic log off settings. For more information, see [Configuring Prism Central UI Settings](#) in *Prism Central Admin Center Guide*.
- Register clusters with this Prism Central instance. For more information, see [Registering a Cluster with Prism Central](#) on page 56.
- Expand Prism Central from a single VM to a multi-VM instance. For more information, see [Expanding \(Scale Out\) Prism Central](#) on page 619.
- Shutdown or Start Up a Prism Central VM. For more information, see [Shutting Down or Starting Up Prism Central VM](#) on page 625.
- Protect Prism Central with up to three AHV or ESXi clusters registered to it. For more information, see [Backing up Prism Central](#).
- Recover Prism Central from Prism Element running on the registered AHV or ESXi cluster. For more information, see [Restoring Prism Central \(One-Click Recovery\)](#).

Managing Prism Central

Prism Central consists of one or more VMs treated as a single instance that monitors and manages registered clusters. You can view information about the Prism Central VMs by clicking the gear icon in the main menu and then selecting **Prism Central Management** from the **Settings** menu (see [Prism Central Settings \(Infrastructure\)](#) on page 44). This displays the **Manage Prism Central** page.




Figure 182: Manage Prism Central Page (1-VM instance)

The **Prism Central Management** page contains the following widgets:

- A **Prism Central Summary** widget on the upper left that displays fields for
 - Cluster name, which is **Unnamed** by default.
 - Virtual IP. If you have configured a IP address as the **Virtual IP** for Prism Central, then this field is displayed with the configured IP address.
 - Domain name. A dash appears if a domain name is not specified.
 - Version. This is the Prism Central version number.
 - ID number. Click the copy link to copy the ID number so you can paste it elsewhere.
- A **Prism Central Capacity** widget on the lower left that displays fields for
 - Number of VMs monitored currently by this Prism Central instance.
 - Number of additional VMs (remaining capacity) this Prism Central instance can monitor.
 - Number of clusters registered to this Prism Central instance; click the number to display the clusters page. For more information, see [Clusters Summary View](#) on page 424.
 - Number of VMs that comprise this Prism Central instance; click the number to display the VMs page. For more information, see [VMs Summary View](#) on page 100.
- A **Prism Central VMs** widget in the middle that displays network address information (network name, subnet mask, and default gateway address) at the top and information below about each VM in the Prism Central instance (VM name and name of storage container in which it is located, IP address, number of vCPUs, and memory size). In addition, single VM instances include a **Scale Out PC** option to scale out this Prism Central instance . For more information, see [Expanding \(Scale Out\) Prism Central](#) on page 619.
- An **Alerts** widget on the right that displays a list of Prism Central-related alerts broken into Critical, Warning, and Info sections. For more information about alerts, see [Prism Central Alerts and Events Reference Guide](#).

When the Prism Central instance consists of three VMs, the display changes slightly. The **Scale Out PC** option disappears, and a new **Add PCVM** option appears (top right). Prism Central instances are limited to a maximum of three VMs, but if you lose one of those VMs for any reason, you can add a replacement by clicking **Add PCVM**. The system displays the **Add PC VMs** page (similar to the **Scale Out PC** page) from which you can add the new VM.

Note: Verify that the prerequisites in [Expanding \(Scale Out\) Prism Central](#) on page 619 are satisfied before adding a Prism Central VM.

To add or change the Prism Central domain name or virtual IP, click the **Edit** link in the **Prism Central Summary** widget. The **Cluster Details** window appears. Do the following:

1. In the **Cluster Name** field, enter a name for the cluster (default "Unnamed"). After naming the cluster, the field no longer appears when subsequently opening the window. Naming the cluster is optional.
2. In the **FQDN** field, enter the fully qualified domain name (FQDN) for the Prism Central cluster. This requires an administrator to configure the domain name in the DNS server to resolve to all the external IPs of the Prism Central VMs. To provide load balancing, you must configure a round-robin DNS record that includes the Prism Central VM (PCVM) IP addresses.
3. In the **Virtual IP** field, enter an IP address that will be used as a virtual IP for the cluster. This is relevant if you have a multi-VM Prism Central.

Note: A virtual IP (VIP) provides resiliency but does not provide load balancing, while an FQDN handles both load balancing and resiliency. It is recommended not to use FQDN and VIP simultaneously, as it will cause a conflict.

4. When the fields are correct, click **Update** to save the changes and close the window.

Cluster Details ×

Virtual IP / FQDN is used to access the PC VM Cluster.

Cluster Name

Unnamed

FQDN

Virtual IP

10.51.132.4

Cancel Update

Figure 183: Cluster Details

Microservices Infrastructure

Microservices Infrastructure provides a common framework and services to deploy the container-based services for Prism Central components.

Microservices Infrastructure (also referred to as Controller Microservices Platform (CMSP)) deploys services like Identity and Access Management (IAM), Load Balancing (LB), and Virtual Private Networking (VPN). Such services are packaged in containers as microservices. Microservices Infrastructure is a shared Kubernetes cluster that containerizes multiple Prism Central services and the resources associated with the Prism Central services.

Microservices Infrastructure is enabled by default in the following cases:

- When you install Prism Central version pc.2023.3 or later on a compatible AOS version.
- When you upgrade to Prism Central version pc.2022.9 or later on a compatible AOS version.
- When you install Prism Central pc.2023.3 or later, or upgrade Prism Central to Prism Central version pc.2022.9 or later on a compatible AOS version in a dark site. The Prism Central installation and upgrade bundles contain the

Microservices Infrastructure package necessary to install or upgrade Microservices Infrastructure and to enable it by default.

A dark site is a site that does not have any access to the Internet.

Note: If Microservices Infrastructure is not automatically enabled during a successful Prism Central installation or upgrade, contact Nutanix Support.

- For information on compatible AOS versions, see the [Compatibility and Interoperability Matrix](#).
- For information on upgrading Prism Central to a version in which Microservices Infrastructure by default, see [Prism Central Deployment](#) on page 12.
- For information on installing Prism Central in a dark site, see [Installing Prism Central Through Prism Element web console](#) on page 17.
- For information on upgrading Prism Central in a dark site, see [Upgrading Prism Central](#).

Important Considerations

The following important considerations apply to Microservices Infrastructure in Prism Central deployments.

- You cannot disable Microservices Infrastructure.
- Prism Central upgrade to a version in which Microservices Infrastructure is enabled by default takes up to 40 minutes longer than a normal Prism Central upgrade.
- If you enabled Objects in the Prism Central managing the cluster, ensure that the Objects VMs are healthy and running when Prism Central upgrade deploys Microservices Infrastructure. If the Objects VMs are not healthy and running, the task for enabling Microservices Infrastructure might fail.

To start the Objects VMs, see [Starting the Objects VMs](#) in the *Objects User Guide*.

Microservices Infrastructure Prerequisites and Considerations

Fulfill these requirements to ensure that Microservices Infrastructure is successfully deployed in a Prism Central instance.

The Prism Central and cluster deployment must fulfill some prerequisites for Microservices Infrastructure to be enabled by default.

Make sure you meet the prerequisites listed before you enable the Microservices Infrastructure.

Note: Microservices Infrastructure uses several ports and protocols for its operations. For information about the ports that are used by Microservices Infrastructure and need to be open in the firewalls accordingly, see [Ports and Protocols for Microservices Infrastructure](#) and [Ports and Protocols for Prism Central](#). For more information on the ports, protocols, IP addresses, and subnets that Microservices Infrastructure uses and needs to be managed in the firewalls, see the [Firewall Access to URLs](#) section in this topic.

Microservices Infrastructure Prerequisites

- **Only if you have Prism Central Backup and Restore enabled**

If you have deployed Prism Central Backup and Restore in Prism Central, then follow the steps provided in [Managing Prism Central Backup and Restore](#) on page 590

- **Minimum AOS Versions**

- Nutanix supports fresh installation or deployment of Prism Central clusters with Microservices Infrastructure deployed by default on AOS versions. With a minimum Prism Central version of pc.2023.3, fresh installation or deployment is supported on AOS versions 6.5.3.x and later, and 6.6 and later.
- To upgrade a Prism Central instance to *minimum supporting Prism Central version*, the hosting Prism Element cluster must run *minimum supporting AOS version* 6.5.x or 6.6. If the hosting Prism Element cluster is not

running a *minimum supporting AOS version*, then upgrade the AOS version to a *minimum supporting AOS version*.

This requirement of *minimum supporting AOS version* is not applicable if the Prism Central instance is registered to the hosting Prism Element cluster.

- When upgrading a Prism Central instance to *minimum supporting Prism Central version*:
 - If the Prism Central instance is not registered to the hosting Prism Element cluster, then register the Prism Central instance to the hosting Prism Element cluster and then run the upgrade workflow.
 - If the Prism Central instance cannot be registered to the hosting Prism Element cluster, enable Microservices Infrastructure to establish trust between the Prism Central instance and the hosting Prism Element cluster.
 - If the Prism Central instance is not registered to the hosting Prism Element cluster or Microservices Infrastructure is not deployed (with trust setup) on the Prism Central instance, then the Prism Central upgrade fails.




Figure 184: Process to Upgrade or Install Minimum Supporting Prism Central

• Hypervisor

Clusters running AHV or ESXi only support Microservices Infrastructure. For ESXi clusters, enter your vCenter credentials (user name and password) and a network for deployment.

- **Prism Element Cluster and Prism Central IP Addresses**

Ensure that you have configured virtual IP address and iSCSI data services IP address on the Prism Element that hosts Prism Central. For more information, see [Modifying Cluster Details](#) information in the *Prism Element Web Console Guide*).

- You can only change these IP addresses of Prism Central or Prism Element cluster after Microservices Infrastructure is enabled. You cannot remove any of these IP addresses. If you remove the current IP address, you must add another IP address before you save the configuration.
- If a Prism Central instance (on which Microservices Infrastructure is already deployed or to be deployed during upgrade) manages multiple Prism Element clusters, configure the iSCSI Data Service IP address on the Prism Element cluster hosting the Prism Central instance. Ensure that the iSCSI Data Service IP address is configured on the hosting Prism Element cluster before you install the or upgrade to the Prism Central that deploys Microservices Infrastructure by default.

- **Cold Migration and CLI-based Prism Central Backup and Restore Support**

Cold Migration and CLI-based Prism Central BR of Prism Central VMs are not supported.

- **Name and NTP Servers**

Ensure that you have configured at a minimum, one Name server and one NTP server in both the host Prism Element cluster and Prism Central.

Warning: Do not remove the Name and NTP server configurations when Microservices Infrastructure is enabled.

- **Prism Central Memory**

[KB 8932](#) lists the Prism Central VM memory requirements. If your Prism Central VM resources do not meet the recommended amount, manually increase the memory.

- **Prism Central Registration and Trust Configuration**

Ensure that you have registered the Prism Central instance to the Prism Element cluster. For more information, see [Registering a Cluster with Prism Central](#) on page 56.

If you upgrade the Prism Central version even though it is not registered to the Prism Element cluster, the upgrade process configures trust between the Prism Element cluster and the Prism Central instance and deploys Microservices Infrastructure. There is no impact to Prism Central. Ensure that Prism Element cluster is running AOS versions compatible with Prism Central version pc.2022.9 or later.

Note: Trust is automatically configured between the Prism Element cluster and the Prism Central instance in the following scenarios:

- When you must not register the Prism Central instance with the Prism Element.
- Or
- When you must unregister the Prism Element cluster from the Prism Central instance after Microservices Infrastructure is deployed.

- **Default Private VXLAN**

If the Prism Central that is required to be upgraded is registered to the Prism Element cluster, then the upgrade process deploys Microservices Infrastructure by default using private VXLAN network. If the private VXLAN is used, the ports 9440, 3205, and 3260 are required to be opened from Prism Central to all the Controller VM IP addresses and Controller VM virtual IP addresses in the clusters managed by the Prism Central, and the iSCSI Data Services IP (DSIP) addresses configured for the clusters managed by the Prism Central.

- **Firewall Access to URLs**

Microservices infrastructure requires access to ports and URLs listed in the [Ports and Protocols](#).

Allowlist the listed URLs in the firewalls to ensure that Prism Central and Prism Element clusters have access to these URLs.

Microservices Infrastructure requires the following IP address ranges or subnets open (in the firewall) to access the Prism Element Controller VM IP addresses, Controller VM virtual IP addresses, and the Prism Central VIP addresses of port 0, for ICMP traffic. Microservices Infrastructure also requires these IP address ranges or subnets to access the Prism Element Controller VM IP addresses and Controller VM virtual IP addresses, the Prism Element iSCSI Data Services IP (DSIP), and the Prism Central virtual IP addresses on ports 9440, 3205, and 3260. Therefore, ports 0, 3205, 3260 and 9440 must be open in the firewall for these IP address ranges or subnets to reach all the Prism Element Controller VM IP addresses and Controller VM virtual IP addresses, Prism Element iSCSI Data Services IP (DSIP) address, and the Prism Central virtual IP addresses as required.

The firewall must also allow any specific IP address range or subnets that are or might be configured as the private networks for Microservices infrastructure.

Note: Do not use the IP addresses in the subnets specified in the [Table 172: Default Subnets for Microservices Infrastructure](#) on page 589 table for operational purposes such as DNS or CVM networks. This prevents any negative impact and maintains seamless operations in the networking.

For more information, see [Ports and Protocols](#) to see the access requirements for all the Nutanix software.

Table 172: Default Subnets for Microservices Infrastructure

Subnet	Purpose
10.100.0.0/16	Reserved for Kubernetes pod network
10.200.32.0/24	Reserved for Kubernetes services network
192.168.5.0/24	Reserved for Kubernetes infrastructure network

Note: Before installing or upgrading to a Prism Central version that deploys Microservices Infrastructure by default, check if these subnets are already in use in your environment. If these subnets are already in use, see [KB-15379](#) and contact Nutanix Support to configure the alternative subnet details for successful Microservices Infrastructure deployment or upgrade.

- **Prism Central Service Domain Name.**

Microservices Infrastructure uses a shared Kubernetes cluster that containerizes multiple Prism Central services and the services resources associated with the services. During the Microservices Infrastructure deployment process, a Domain Name Service (DNS) is deployed within Prism Central. Microservices Infrastructure uses this DNS to discover and manage the available Prism Central microservices. The Prism Central Domain is set up for discovery of the microservices and communication of the microservices with each other. The Prism Central Domain facilitates the discovery and communication using the DNS service deployed for the Microservices Infrastructure within Prism Central. This Domain Name and DNS is internal to the Prism Central instance and is not used outside the instance.

Prism Central Domain Name configuration is no longer manual. The default value for Prism Central Domain Name is `prism-central.cluster.local`. No conflict occurs in a setup with multiple Prism Centrals interacting with each other. If a conflict occurs in your setup and Microservices Infrastructure fails to deploy, contact Nutanix Support.

- **Flow Microsegmentation container**

When Flow Microsegmentation is enabled, a Kafka container called `flow_data` is automatically created on the cluster that hosts Prism Central. The `flow_data` container stores data that is required and essential for Flow visualization. Do not delete this container.

- **Network Latency**

If the Prism Element and Prism Central instances are on different subnets, ensure to have less than 5 ms of round-trip time latency between the two instances.

Limitations

The following limitations apply after Microservices Infrastructure is enabled:

- **Hypervisor Support**

Nutanix supports Microservices Infrastructure on an on-premises Prism Central deployment hosted on an AOS cluster running AHV or ESXi.

Nutanix supports Microservices Infrastructure on on-premises Prism Central VM deployed in a non-Nutanix ESXi environment.

Clusters running other hypervisors are not supported.

- **Nutanix Move**

Moving a Prism Central enabled with Microservices Infrastructure from ESXi host to AHV host using Nutanix Move is not supported.

- **Scale-out and Single PCVM deployments**

Scale-out three-node Prism Central VMs deployment: Small, large, or x-large Prism Central VMs are supported. Reverting the deployment to a single Prism Central VM deployment is not supported.

Single Prism Central VM deployment: Small, large, or x-large Prism Central VM are supported. You can expand a single PCVM deployment to a scale-out three-node Prism Central deployment on a minimum Prism Central version of pc.2022.1 and minimum AOS version of 6.1. Expanding a single node Prism Central with Microservices Infrastructure is supported only on a minimum Prism Central version of pc.2022.1 and minimum AOS version of 6.1.

For information on Prism Central scalability and the type of Prism Central deployments, see the *Prism Central Scalability* topic in the release notes for the [Prism Central](#) version to be installed.

- **IP Address Restrictions**

Ensure that the IP addresses in subnet that you configure for Microservices Infrastructure do not conflict with the IP addresses in the management subnet.

Note: Do not use the IP addresses in the subnet 10.200.32.0/24 for operational purposes such as DNS or CVM network.

Managing Prism Central Backup and Restore

If you have deployed Prism Central Backup and Restore in Prism Central, you must disable Prism Central BR before the upgrade and reenable it after the upgrade.

About this task

To manage Prism Central BR during Microservices Infrastructure upgrade, follow these steps:

Procedure

1. Disable Prism Central BR in Prism Central. For more information, see [Disabling Prism Central Backup and Restore](#) on page 611.

Note: When you disable Prism Central BR, you cannot enable it again after the Prism Central upgrade until you upgrade at least one Prism Element in the Prism Central cluster (Availability Zone) to *minimum supporting AOS version*.

2. Upgrade Prism Central to *minimum supporting Prism Central version* (pc.2022.9 or later version that deploys Microservices Infrastructure by default).
3. Upgrade at least one Prism Element in the Prism Central cluster (Availability Zone) to *minimum supporting AOS version*.
4. Enable Prism Central Backup and Restore.

Microservices Infrastructure Cluster Upgrade

When you enable Microservices Infrastructure, it creates a Kubernetes cluster with a version that is deployed by the Microservices Infrastructure controller.

When you upgrade Prism Central version, it upgrades the Microservices Infrastructure controller version. The cluster is not upgraded. The controller upgrade process upgrades the base services running on the cluster instead of upgrading the cluster.

Starting with Prism Central pc.2024.3, the Prism Central upgrade process upgrades the Microservices Infrastructure controller version, the Microservices Infrastructure Kubernetes cluster version and the base services running on the cluster.

Note: The Prism Central upgrade fails if the Microservices Infrastructure cluster version upgrade fails. In other words, the Microservices Infrastructure cluster version upgrade failure leads to Prism Central upgrade failure. The Prism Central upgrade task failure details reflect the failure of the Microservices Infrastructure cluster upgrade failure. The error message is prefixed with CMSP cluster upgrade failed:.

Displaying the Cluster Version

About this task

Perform the following steps to ascertain the Microservices Infrastructure cluster version.

Procedure

1. SSH to the Prism Central VM.
2. Run the following command:

```
nutanix@pcvm$ mspctl cls version
```

The following is a sample output of this command:

COMPONENT	VERSION
Cluster	1.x.x

Prism Central Backup, Restore, and Migration

The Prism Central Backup and Restore (PCBR) solution protects Prism Central instances, including scale-out configurations, against unforeseen events such as natural disasters, network outages, or power failures.

You can also use PCBR to migrate Prism Central instances from one cluster to another.

To protect Prism Central instances and ensure resiliency of your instances and service configurations across on-premises environments and S3-compatible object-based storage such as AWS S3 and Nutanix Objects, PCBR provides the following automated backup types:

Continuous Backup

Enables real-time backup of Prism Central instances and various service configurations within it to up to three AHV or ESXi clusters registered to the same Prism Central instance, ensuring minimal data loss and swift recovery after disruptions.

Point-in-Time Backup

Enables scheduled backups of Prism Central instances and various service configurations within it to one S3 compatible object-based storage on AWS or Nutanix Objects, allowing for recovery to specific points in time and protection against ransomware attacks and datacenter failures.

The recovery point objective (RPO) and recovery time objective (RTO) for recovery from a continuous backup are 30 minutes and 90 minutes, respectively. You can recover the Prism Central instance only from the latest backup.

While the RTO for recovery from a point-in-time backup is 90 minutes, the RPO can be set in values ranging from one hour to 24 hours. The first restore point is created immediately, then restore points are created every RPO cycle. You can recover the Prism Central instance from any available restore point. Prism Central supports up to 30 days of restore points for restoring the Prism Central instance.

With PCBR, you can backup a Prism Central instance to at most three AHV or ESXi clusters and one S3 compatible object-based storage such as AWS S3 or Nutanix Objects bucket. Clusters running Hyper-V are not supported.

You can back up a Prism Central instance and various service configurations within it in the following combinations to up to four locations:

- Continuous backup: Up to three AHV or ESXi clusters
- Point-in-time backup: One Object-based storage. You can use an AWS S3 or a Nutanix Objects bucket to back up multiple Prism Central instances.
- Hybrid (continuous and point-in-time): A combination of up to three AHV or ESXi clusters and one S3-compatible object-based storage, which can be either AWS S3 or a Nutanix Objects bucket

One-click recovery restores the important service configurations along with the entire Prism Central instance, minimizing downtime and operational impact. For more information, see *Supported Services for Backup and Recovery* and *Unsupported Services for Backup and Recovery* in [Implementation Considerations and Limitations](#) on page 593.

Prism Central Backup and Restore Requirements

Before protecting or backing up a Prism Central instance to Nutanix clusters or an S3-compatible object-based storage like AWS S3 or Nutanix Objects, ensure that you meet the following requirements.

- For continuous backup:
 - The Prism Central instance that you plan to back up must run on a minimum version of pc.2021.7 or later.
 - The Nutanix clusters where you plan to back up the Prism Central instance must meet these requirements:
 - The Nutanix clusters are registered to the Prism Central instance you plan to back up.
 - The Nutanix clusters are running a minimum AOS 6.0 or later version.
 - The Nutanix clusters are configured with a Data Services IP.
 - At least one Nutanix cluster where you plan to back up the Prism Central instance must be running AOS 6.5.3.1 or later version.

Prism Central instances can be restored from a continuous backup only on clusters running AOS 6.5.3.1 or later version. For more information, see [Restoring Prism Central](#) on page 604.

- For point-in-time backup to AWS S3:
 - The Prism Central instance that you plan to back up must run on a minimum version of pc.2024.1 or later.
 - The Nutanix clusters where you plan to back up the Prism Central instance must be running a minimum AOS 6.0 or later version.
 - At least one cluster where you plan to back up the Prism Central instance must be running AOS 6.8 or later version.

Prism Central instances can be restored from a AWS S3 point-in-time backup only on clusters running AOS 6.8 or later version. For more information on supported AOS versions, see [Compatibility and Interoperability Matrix](#).

- While configuring an AWS S3 bucket, perform the operations noted in [Configuring Amazon S3 Object Lifecycle and Policies](#) on page 598.
 - For point-in-time backup to Nutanix Objects:
 - The Prism Central instance that you plan to back up must run on a minimum version of pc.7.3 or later.
 - The Nutanix clusters where you plan to back up the Prism Central instance must be running a minimum AOS 7.3 or later version.
 - At least one cluster where you plan to back up the Prism Central instance must be running AOS 7.3 or later version.
- Prism Central instances can be restored from a Nutanix Objects point-in-time backup only on clusters running AOS 7.3 or later. For more information on the supported AOS versions, see [Compatibility and Interoperability Matrix](#).
- While configuring a Nutanix Objects bucket, perform the operations noted in [Configuring Nutanix Objects Lifecycle and Policies](#) on page 602.
 - The Prism Central instance that you plan to back up must have an NTP configuration to synchronize time between the Prism Central instance and the registered clusters. For more information, see [Configuring NTP Servers \(Prism Central\)](#) in the *Prism Central Admin Center Guide*.
 - Note the Files version. In the case of a disaster, you might need it later to restore Files to the pre-disaster version.

Implementation Considerations and Limitations

Consider the following recommendations and limitations for implementing Prism Central Backup and Restore (PCBR):

- Protecting Prism Central instances using protection domain-based disaster recovery, protection policy-based workflows, third-party backup software, or any unsupported methods renders those instances unrecoverable.
- When Prism Central upgrade is initiated on an instance that is configured for PCBR, the system pauses the backup of Prism Central configuration data. The backup resumes after the upgrade process is complete.
- You can use the PCBR capability to migrate Prism Central instances to other Nutanix clusters.

For more information, see [Migrating a Prism Central Instance](#) on page 610. While the PCBR workflow ensures the successful migration of important service configurations to the remote cluster, the services listed in the *Unsupported Services for PCBR* section of this document are not migrated.

Note: Migration Prism Central instances using vMotion is not supported.

- For point-in-time backups to Nutanix Objects, Nutanix recommends securely storing the object store certificate including complete certificate chain because those certificates are required to restore the Prism Central instance.

By default, the system deploys self-signed certificates on Nutanix Objects. To replace these self-signed certificates, see [Managing FQDN and SSL Certificate](#) in the *Objects User Guide*.

- For PCBR, you can use Nutanix Objects public IP address or FQDN. If you use an IP address, ensure that the Nutanix Objects public IP address is present in the alternative names of the certificate.
- For FQDN, Nutanix recommends configuring FQDN to IP mapping in the DNS server. If you cannot use the DNS server, then use the IP address of Nutanix Objects object store while configuring Nutanix Objects as a backup target and recovering the Prism Central instance.
- For IP addresses while configuring backups, use the same IP address only while recovering. If FQDN is used while configuring backups, then use the same FQDN only during recovery.

Supported Services for PCBR

The following services can be backed up and the data associated with them can be recovered:

- Nutanix Disaster Recovery
- Networking (such as AHV-based VLAN and virtual switch configurations)
- Flow Virtual Networking (virtual private clouds (VPCs) and virtual networks using ANC network controller and network gateway).

To ensure that traffic is not disrupted when the IP addresses of the Prism Central VMs change during PCBR, update the IP addresses of the network gateways deployed for Flow Virtual Networking. For more information, see [KB-19776](#).

- Legacy Flow Network Security
- Intelligent Operations
- VM management
- Cluster management
- Categories
- Licensing
- Reporting templates
- Files Storage (only configuration)
- Nutanix Objects (only configuration)
- LCM
- Domain Manager (Projects and Marketplace)

Unsupported Services for PCBR

Although the following services continue to run, their data is not backed up and, therefore, is not recoverable:

- Foundation Central

Important: Contact Nutanix Support for more information.

- Nutanix Kubernetes Platform
- Nutanix Kubernetes Engine

- Cisco Device Connector
- Flow Network Security Next-Gen
- Nutanix Objects data
- Files storage data
- Encryption keys for PCKMS.

Note: Nutanix recommends backing up the encryption keys before performing PCBR. After the recovery, you can restore the encryption keys manually through nCLI on the newly restored Prism Central instance.

- NCM Self-Service

Important: To manually back up NCM Self-Service-related entities and restore those entities on the new Prism Central instance, see [Backup and Restore in Self-Service](#) in the *Self-Service Administration and Operations Guide*.

- Catalog
- Power Monitor
- OVAs
- VM images and templates

For Nutanix Disaster Recovery jobs (RPJs) in progress, see this [Note](#). Metrics older than 90 days cannot be recovered.

Backing up Prism Central

You can protect or back up a Prism Central instance to its registered Nutanix clusters, S3-compatible object-based storage like AWS S3 and Nutanix Objects, or Nutanix clusters and an S3-compatible object-based storage together.

- To enable real-time backup of your Prism Central instances, and various service configurations within it, configure continuous backup.

For more information, see [Configuring Continuous Backup](#) on page 595.

- To enable scheduled backups of your Prism Central instances and various service configurations within it, configure point-in-time backup to S3 object-based storage on AWS or Nutanix Objects, allowing for recovery to specific points in time.
 - For point-in-time backup to AWS S3, see [Configuring Point-in-Time Backup to AWS S3](#) on page 596
 - For point-in-time backup to Nutanix Objects, see [Configuring Point-in-Time Backup to Nutanix Objects](#) on page 600

Configuring Continuous Backup

You can protect or back up a Prism Central instance to its registered Nutanix clusters, S3-compatible object-based storage like AWS S3 and Nutanix Objects, or Nutanix clusters and an S3-compatible object-based storage together.

Before you begin

See [Prism Central Backup and Restore Requirements](#) on page 592 and [Implementation Considerations and Limitations](#) on page 593 before you start.

About this task

Backup to registered Nutanix clusters is called continuous backup, because it happens continuously at an RPO of 30 minutes. To protect or back up a Prism Central instance to Nutanix clusters, follow these steps:

Procedure

1. Log in to the Prism Central web console.
2. From [Application Switcher Function](#), select the **Infrastructure** application and click the [Settings icon](#).
3. From the **Settings** menu, navigate to **General > Prism Central Backups**.

Tip: The **What's New** page appears when you go to the **Prism Central Backups** page for the first time.

4. Navigate to **Prism Central Backups > Continuous Backup** and click **Protect Now**.

The **Protect Now** option is available only when one or more clusters are registered to the Prism Central instance that you plan to back up.

The **Protect Prism Central** window appears showing the services data that will be backed up and those that will not be backed up. For the detailed list of services, see *Supported Services for PCBR* and *Unsupported Services for PCBR* in [Implementation Considerations and Limitations](#) on page 593.

5. Click **Continue**.
6. Select at most three clusters to back up the Prism Central instance.

Note: You can configure a continuous backup to up to three AHV or ESXi clusters.

7. (Optional) To back up the Prism Central instance to more Nutanix clusters, click **+ Add Backup** in the **Prism Central Backups** page.

The number of entities backing up the Prism Central instance must be at most four: three on-prem clusters and one S3-compatible object-based storage. You can use a cluster only once to back up the same Prism Central instance.

8. (Optional) To remove a Nutanix cluster or a bucket from backing up the Prism Central instance, click **Remove** in front of the cluster, type Remove (not case sensitive), and click **Remove**.
9. Click **Proceed**.

The system synchronizes Prism Central configuration data to the selected clusters. While the configuration data is synchronizing, the **Status** of the Prism Central instances is shown Sync in Progress in the **Prism Central Backups** dashboard. The **Status** turns to Synced when the configuration data is synchronized. The first backup creation on the selected clusters takes at least 30 minutes. After the first backup, the system synchronizes Prism Central configuration data with the selected clusters every 30 minutes.

Configuring Point-in-Time Backup to AWS S3

You can protect or back up a Prism Central instance to its registered Nutanix clusters, S3-compatible object-based storage like AWS S3 and Nutanix Objects, or Nutanix clusters and an S3-compatible object-based storage together.

Before you begin

- See [Prism Central Backup and Restore Requirements](#) on page 592 and [Implementation Considerations and Limitations](#) on page 593 before you start.
- Perform the operations noted in [Configuring Amazon S3 Object Lifecycle and Policies](#) on page 598.

- Ensure that the firewalls allow port 443 for communication and replication between the Prism Central instance, Prism Element instance, and the AWS S3 bucket.
 - To protect or back up a Prism Central instance to an AWS S3 bucket, allow port 443 between the Prism Central instance to `bucket-name.S3.amazonaws.com` and `bucket-name.S3.region-name.amazonaws.com`.
 - To restore a Prism Central instance from an AWS S3 bucket, allow port 443 between the Prism Element instance to `bucket-name.S3.amazonaws.com` and `bucket-name.S3.region-name.amazonaws.com`
- Replace `bucket-name` with the name of your AWS S3 bucket.
- Replace `region-name` with the region name of your AWS S3 bucket.

Tip: If you do not know the bucket name or region name, allow port 443 between the Prism Central instance or Prism Element instance and `*.S3.amazonaws.com` and `*.S3.*.amazonaws.com`. For more information on the port requirements, see [Ports and Protocols](#).

About this task

Backups taken at specified intervals ranging from one hour to 24 hours are called point-in-time backups. These backups are supported on Nutanix Objects and AWS S3. To protect or back up a Prism Central instance to AWS S3, follow these steps:

Procedure

1. Log in to the Prism Central web console.
2. From [Application Switcher Function](#), select the **Infrastructure** application, and click the [Settings icon](#).
3. From the **Settings** menu, navigate to **General > Prism Central Backups**.

Tip: The **What's New** page appears when you go to the **Prism Central Backups** page for the first time.

4. Navigate to **Prism Central Backups > Point-in-Time Backup** and click **Protect Now**.

The system supports one point-in-time backup target at a time. If a point-in-time backup is configured already, the **Protect Now** does not appear.

The **Protect Prism Central** window appears, showing the services data that will be backed up and those that will not be backed up. For the detailed list of services, see *Supported Services for PCBR* and *Unsupported Services for PCBR* in [Implementation Considerations and Limitations](#) on page 593.

5. Click **Continue** and specify the following details in the **Protect Prism Central** window:

- a. **Endpoint Type:** Select AWS.
- b. **Region Name:** Enter an AWS region name.
For more information, see *Amazon S3 documentation*.
- c. **Bucket Name:** Enter a bucket name.
For more information, see *Amazon S3 documentation*.
- d. (Optional for NC2 environments) **Access Key:** Enter your access key.
- e. (Optional for NC2 environments) **Secret Access Key:** Enter your secret access key.
- f. **Restore Point Objects (RPO):** Enter the time frequency in hours to create restore points.

RPO is the frequency at which the system creates restore points. The first restore point is created immediately, then restore points are created according to the specified RPO, which can be 1, 2, 4, 6, 8, 12, or 24 hours. Up to 30 days of restore points are supported for restoring the Prism Central instance.

Tip: You can update the RPO, secret keys, and certificates for backup and restore to S3-compatible object storage. Use **Edit RPO** to change the restore point frequency, and use **Edit Authentication** to update credentials and certificates. For more information, see [Modifying Recovery Point Objective \(RPO\) for a Point-in-Time Backup](#) on page 603 and [Modifying Authentication for a Point-in-Time Backup](#) on page 603 respectively.

6. Click **Proceed**.

The system synchronizes Prism Central configuration data to the AWS S3 bucket. While the configuration data is synchronizing, the **Status** of the Prism Central instances is shown Sync in Progress in the **Prism Central Backups** dashboard. The **Status** turns to Synced when the configuration data is synchronized. The process of taking the first backup starts immediately when you submit the AWS S3 bucket details and takes at least 15 minutes. After the first backup, the system synchronizes Prism Central configuration data with the AWS S3 bucket and creates restore points every RPO cycle.

Configuring Amazon S3 Object Lifecycle and Policies

Create an AWS S3 bucket in a suitable region to store the backup data, then configure the following policies on the bucket.

Procedure

1. Enable bucket versioning in AWS.

A version of the file is created each time the system writes the seed data file on the new key path to the bucket. The previous version becomes non-current, and the currently written file becomes current. For more information on the current and non-current items and how to enable bucket versioning, see *Amazon S3 documentation*.

2. Enable the object lock feature for objects to enable the write-once-read-multiple (WORM) configuration and to create point-in-time snapshots of the Prism Central configuration.

Configure the object lock feature when you create a bucket. To ensure that the backup data put into the bucket is retained for some time and deleted after that time, the system retains the data for 31 days. Enable the default retention mode, configure the retention mode as Governance, and set the default retention period to 31 days.

For more information on how to enable object lock, see *Amazon S3 documentation*.

3. Configure an object lifecycle rule to auto-delete older backup data.

To enable the automatic deletion of objects after a certain time as they become stale, configure an object lifecycle rule from the Amazon S3 console that applies to all objects in the bucket. For more information on how to configure a lifecycle rule, see *Amazon S3 documentation*.

Note: Configure a maximum time of 31 days for deleting noncurrent objects. Nutanix recommends setting the expiration of current objects versions, permanent deletion of noncurrent objects versions, deletion of expired object delete markers or incomplete multipart uploads, and the deletion of incomplete multipart uploads to a day.

The rule you created is listed as follows. Ensure that the scope is seen as Entire bucket.




Figure 185: Lifecycle Configuration

4. Enable access to S3 buckets from the Amazon S3 console.

By default, public access is blocked for buckets in the AWS console. To enable access, log on as the root user of the AWS account and navigate to the IAM service. For more information, see *AWS S3 documentation*.

5. Create a user from the Amazon S3 console.

Do not provide console access to this user because an access key is to be generated for this user to enable programmatic access from outside of AWS.

6. Create a policy for the user with an identifiable name and paste the following JSON in the text box:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "s3:PutObject",
        "s3:GetObject",
        "s3:GetObjectVersion",
        "s3>DeleteObject",
        "s3>DeleteObjectVersion"
      ],
      "Resource": "arn:aws:s3:::Bucket_name/*"
    },
    {
      "Effect": "Allow",
      "Action": [
        "s3:ListBucket",
        "s3:GetBucketLocation",
        "s3:GetLifecycleConfiguration",
        "s3:GetBucketObjectLockConfiguration"
      ],
      "Resource": "arn:aws:s3:::Bucket_name"
    }
  ]
}
```



```
}
```

Replace *Bucket_name* with the name of your Amazon S3 bucket.

7. Find the policy that you created in the Step 6 and attach it to the created user.
8. Create an access key with a descriptive name.
9. From the AWS console, select **Application running outside AWS**.
10. Save the access key.

Nutanix recommends downloading the CSV file. The same key must be entered on the Prism Central web console while configuring this S3 bucket as the target endpoint for storing your Prism Central configuration backups.

For more information, see *AWS S3 documentation*.

Configuring Point-in-Time Backup to Nutanix Objects

You can protect or back up a Prism Central instance to its registered Nutanix clusters, S3-compatible object-based storage like AWS S3 and Nutanix Objects, or Nutanix clusters and an S3-compatible object-based storage together.

Before you begin

- See [Prism Central Backup and Restore Requirements](#) on page 592 and [Implementation Considerations and Limitations](#) on page 593 before you start.

- Create an object store.

For more information, see [Creating or Deploying an Object Store on Prism Central](#) in the *Objects User Guide*.

- Perform the operations noted in [Configuring Nutanix Objects Lifecycle and Policies](#) on page 602.
- Enable access to the Nutanix Objects bucket.

For more information, see [Enabling Nutanix Objects Bucket Access for a User](#) on page 602.

- Ensure that the firewalls allow port 443 for communication and replication between the Prism Central instance, Prism Element instance, and the Nutanix Objects bucket.
 - To protect or back up a Prism Central instance to a NutanixObjects bucket, allow port 443 between the Prism Central instance to *bucket-name.ObjectStore_FQDN* and *ObjectStore_IPaddress:443*.
 - To restore a Prism Central instance from the Nutanix Objects bucket, allow port 443 between the Prism Element instance to *bucket-name.ObjectStore_FQDN* and *ObjectStore_IPaddress:443*.

Replace *ObjectStore_FQDN* with the FQDN you configure on Nutanix Objects.

Replace *ObjectStore_IPaddress* with the public IP address of the object store in Nutanix Objects.

Tip: If you do not know the bucket name or region name, allow port 443 between the Prism Central instance or Prism Element instance and **.ObjectStore_FQDN* and *ObjectStore_IPaddress*. For more information on the port requirements, see [Ports and Protocols](#).

About this task

Backups taken at specified intervals ranging from one hour to 24 hours are called point-in-time backups. These backup are supported on Nutanix Objects and AWS S3.

To protect or back up a Prism Central instance to Nutanix Objects, follow these steps:

Procedure

1. Log in to the Prism Central web console.
2. From [Application Switcher Function](#), select the **Infrastructure** application, and click the [Settings icon](#).
3. From the **Settings** menu, navigate to **General > Prism Central Backups**.

Tip: The **What's New** page appears when you go to the **Prism Central Backups** page for the first time.

4. Navigate to **Prism Central Backups > Point-in-Time Backup** and click **Protect Now**.

The system supports one point-in-time backup target at a time. If a point-in-time backup is configured already, the **Protect Now** does not appear.

The **Protect Prism Central** window appears, showing the services data that will be backed up and those that will not be backed up. For the detailed list of services, see *Supported Services for PCBR* and *Unsupported Services for PCBR* in [Implementation Considerations and Limitations](#) on page 593.

5. Click **Continue** and specify the following details in the **Protect Prism Central** window:

- a. **Endpoint Type:** Select Nutanix Objects.
- b. **IP Address or Host Name (https):** Enter the public IP address or FQDN of the object store in Nutanix Objects.
- c. **Bucket Name:** Enter a Nutanix Objects bucket name.

Note: Ensure that the Nutanix Objects bucket names do not contain a dot (.). For more information, see [Bucket Naming Conventions](#) in the *Nutanix Objects User Guide*.

- d. (Optional) **Upload Certificate:** Upload the certificate you downloaded from the Nutanix Objects instance.

Note: If you do not upload a certificate, the system deploys the default Nutanix Objects with self-signed certificates. To replace these self-signed certificates, see [Managing FQDN and SSL Certificate](#) in the *Objects User Guide*.

- e. **Access Key:** Enter your access key.
- f. **Secret Access Key:** Enter your secret access key.
- g. **Restore Point Objects (RPO):** Enter the time frequency in hours to create restore points.

RPO is the frequency at which the system creates restore points. The first restore point is created immediately, then restore points are created according to the specified RPO, which can be 1, 2, 4, 6, 8, 12, or 24 hours. Up to 30 days of restore points are supported for restoring the Prism Central instance.

Tip: You can update the RPO, secret keys, and certificates for backup and restore to S3-compatible object storage. Use **Edit RPO** to change the restore point frequency, and use **Edit Authentication** to update credentials and certificates. For more information, see [Modifying Recovery Point Objective \(RPO\) for a Point-in-Time Backup](#) on page 603 and [Modifying Authentication for a Point-in-Time Backup](#) on page 603 respectively.

6. Click **Proceed**.

The system synchronizes Prism Central configuration data to the Nutanix Objects. While the configuration data is synchronizing, the **Status** of the Prism Central instances is shown Sync in Progress in the **Prism Central Backups** dashboard. The **Status** turns to Synced when the configuration data is synchronized. The process of taking the first backup starts immediately when you submit the Nutanix Objects bucket details and takes at least 15 minutes. After the first backup, the system synchronizes Prism Central configuration data with Nutanix Objects bucket and creates restore points every RPO cycle.

Configuring Nutanix Objects Lifecycle and Policies

Create a Nutanix Objects bucket to store the backup data, then configure the following policies on the bucket.

About this task

Procedure

1. Enable bucket versioning in Nutanix Objects.

A version of the file is created each time the system writes the seed data file on the new key path to the bucket. The previous version becomes non-current, and the currently written file becomes current. For more information on the current and non-current items and how to enable bucket versioning, see [Object Versioning in Objects](#) in the *Objects User Guide*.

2. Enable the object lock feature for objects to enable the write-once-read-multiple (WORM) configuration and to create point-in-time snapshots of the Prism Central configuration.

Configure the object lock feature when you create a bucket. The system retains the backup data for 31 days to ensure that it is retained for a certain time and then deleted after that time. For more information on how to enable object lock, see [Applying WORM Policy to Buckets from Prism Central](#) in the *Objects User Guide*.

Note: Ensure that the object store has enough storage.

3. Configure an object lifecycle rule to auto-delete older backup data.

To enable the automatic deletion of objects after a certain time as they become stale, configure an object lifecycle rule from the Nutanix Objects instance that applies to all objects in the bucket. For more information on how to configure a lifecycle rule, see [Lifecycle Policies in Objects](#) in the *Objects User Guide*.

Note: Nutanix recommends setting up the following for all objects:

- **Expire current objects version after:** 31 days
- **Expire previous objects version after:** 1 day
- Select **Delete markers automatically** if they are the only object version.
- **Multi-part uploads:** 1 day

You cannot combine delete markers and multi-part uploads in a rule with an object version policy. For such policies, create discrete rules.

Enabling Nutanix Objects Bucket Access for a User

By default, the user who has created a bucket has access to all operations in the Nutanix Objects bucket. However, the user performing Prism Central Backup and Restore (PCBR) does not require permissions for all the operations. Therefore, you must enable access to the Nutanix Objects bucket with selective permissions.

Procedure

1. Create a user for the object store.

Securely store the access keys for future use. The access keys are required to restore the Prism Central instance.

2. Click the bucket you plan to use to back up the Prism Central instance and go to the **User Access** tab.

3. Click **Edit User Access**, select the user, and choose **Custom Permission**.

For more information, see [Editing User Access in Nutanix Objects](#) in the *Nutanix Objects User Guide*.

4. Choose only the following operational permissions:

- DeleteObject
- DeleteObjectVersion
- GetObject
- GetObjectVersionTagging
- ListBucket
- PutObject
- GetBucketLocation
- GetLifecycleConfiguration
- GetBucketObjectLockConfiguration
- GetBucketVersioning

For more information, see [Sharing a Bucket in Objects](#) in the *Nutanix Objects User Guide*.

Modifying Recovery Point Objective (RPO) for a Point-in-Time Backup

You can modify the RPO to change how often Prism Central creates restore points for point-in-time backups.

About this task

To modify the RPO, follow these steps:

Procedure

1. Log in to the Prism Central web console.
2. From [Application Switcher Function](#), select the **Infrastructure** application, and click the [Settings icon](#).
3. From the **Settings** menu, navigate to **General > Prism Central Backups > Point-in-Time Backup**.
4. In the list of configured backups, click **Edit RPO** for the backup that you plan to modify.
5. From the dropdown menu in the **Edit RPO** window, select a restore point frequency that you plan to set for the backup.

RPO is the frequency at which the system creates restore points. The first restore point is created immediately, then restore points are created according to the specified RPO, which can be 1, 2, 4, 6, 8, 12, or 24 hours. Up to 30 days of restore points are supported for restoring the Prism Central instance.
6. Click **Save**.

Modifying Authentication for a Point-in-Time Backup

You can modify authentication details for point-in-time backups to S3-compatible object storage.

About this task

To modify the authentication, follow these steps:

Procedure

1. Log in to the Prism Central web console.
2. From [Application Switcher Function](#), select the **Infrastructure** application, and click the [Settings icon](#).
3. From the **Settings** menu, navigate to **General > Prism Central Backups > Point-in-Time Backup**.
4. In the list of configured backups, click **Edit Authentication** for the backup that you plan to modify.
5. In **Edit Authentication** window, specify the following details:
 - a. (Optional for NC2 environments) **Access Key**: Enter your access key.
 - b. (Optional for NC2 environments) **Secret Access Key**: Enter your secret access key.
6. Click **Save**.

Restoring Prism Central

If you protected or backed up a Prism Central instance, you can use an AHV or ESXi cluster to restore the Prism Central instance.

- For continuous backups, you can recover the Prism Central instance from a registered cluster only. To recover a Prism Central instance from a continuous backup, see [Recovering Prism Central from Continuous Backup](#) on page 604.
- For point-in-time backups, you can recover the Prism Central instance on a Nutanix cluster, whether or not it is registered with any Prism Central instance that fulfills the requirements.
 - For recovering from a point-in-time backup available on AWS S3, see [Recovering Prism Central from Point-in-Time Backup on AWS S3](#) on page 605.
 - For recovering from a point-in-time backup available on Nutanix Objects, see [Recovering Prism Central from Point-in-Time Backup on AWS S3](#) on page 605.

Recovering Prism Central from Continuous Backup

You can use only a registered AHV or ESXi cluster to restore the Prism Central instance from a continuous backup. You can recover the Prism Central instance only from the latest backup. The recovery point objective (RPO) and recovery time objective (RTO) for recovery from a continuous backup are 30 minutes and 90 minutes, respectively.

Before you begin

See [Prism Central Backup and Restore Requirements](#) on page 592 and [Implementation Considerations and Limitations](#) on page 593 before you start.

About this task

To restore a Prism Central instance from continuous backup, follow these steps:

Procedure

1. Log in to any Prism Element web console registered to the Prism Central instance that you plan to restore.

The Prism Element dashboard shows the Prism Central widget, which contains Prism Central information (IP address and connection status). If the Prism Element instance is a fresh instance that you created and is not yet registered to a Prism Central instance, it does not appear. Instead, you see the **Register or Deploy Prism Central**.

2. Click the Settings icon and from the **Settings** menu, navigate to **Data Resiliency > Restore Prism Central**.
For more information, see [Settings Menu](#) in the *Prism Element Web Console Guide*.
The **Restore Prism Central** page appears. This page provides you with options to restore the Prism Central instance from a Prism Element cluster or S3-compatible object storage.
3. Select **Restore Prism Central from Prism Element** and click **Restore Now**.

Note: If you configured only continuous backup, the **Restore Now** option is available only when Prism Central is in a disconnected state; the Prism Central instance is shown as **Disconnected** in the Prism Element web console.

The **Restore Prism Central** window appears, showing the services data that will be backed up and those that will not be backed up. For the detailed list of services, see *Supported Services for PCBR* and *Unsupported Services for PCBR* in [Implementation Considerations and Limitations](#) on page 593.

4. Click **Continue** and specify the following information according to the source from which you plan to recover.
5. Select the cluster where you plan to restore the Prism Central instance.
6. Verify the version of the Prism Central instance that you plan to use on the selected cluster.
7. Select the network where you plan to restore and install Prism Central instances.
The **Subnet Mask**, **Gateway**, and **DNS Address(s)** fields show the relevant information associated with the selected network.
8. Enter the name and IP address of the Prism Central instance that you plan to restore and click **Save**.
9. Review the summary and click **Recover**.
The Prism Central instance is restored in 60 to 90 minutes, depending on the configuration data of the hosts. The restoration involves the deployment of Prism Central and the restoration of its configuration data from the backup, which can take between 60 and 120 minutes, depending on the size of the data. The restored Prism Central instance takes an additional 30 to 40 minutes to show all the guest VMs, disks, and metrics. Wait to perform any actions on the restored Prism Central instance until all the recovery tasks are complete on the cluster. You can see the restoration status and the related processes in the **Tasks** window.

What to do next

See [Post Restoration Considerations](#) on page 609.

Recovering Prism Central from Point-in-Time Backup on AWS S3

You can use a registered or unregistered AHV or ESXi cluster to restore the Prism Central instance from a point-in-time backup on AWS S3. While the recovery time objective (RTO) for recovery from a point-in-time backup is 90 minutes, the recovery point objective (RPO) can be set in values ranging from one hour to 24 hours. Prism Central supports up to 30 days of restore points for restoring the Prism Central instance.

Before you begin

- See [Prism Central Backup and Restore Requirements](#) on page 592 and [Implementation Considerations and Limitations](#) on page 593 before you start.
- Ensure that the old Prism Central is disconnected and powered off before initiating the restore operation.
The restoration creates a new Prism Central instance that replaces the older instance.
- Ensure to initiate a restore from only one Prism Element.

About this task

You can recover the Prism Central instance from any available restore point. To restore a Prism Central instance from a point-in-time backup on AWS S3, follow these steps:

Procedure

1. Log in to any Prism Element web console registered to the Prism Central instance that you plan to restore.
The Prism Element dashboard shows the Prism Central widget, which contains Prism Central information (IP address and connection status). If the Prism Element instance is a fresh instance that you created and is not yet registered to a Prism Central instance, it does not appear. Instead, you see the **Register or Deploy Prism Central**.
2. Click the [Settings icon](#) and from the **Settings** menu, navigate to **Data Resiliency > Restore Prism Central**.
For more information, see [Settings Menu](#) in the *Prism Element Web Console Guide*.
The **Restore Prism Central** page appears. This page provides you with options to restore the Prism Central instance from a Prism Element cluster or S3-compatible object storage.
3. Select **Restore any Prism Central from S3-compatible object storage** and click **Restore Now**.
The **Restore Prism Central** window appears, showing the services data that will be backed up and those that will not be backed up. For the detailed list of services, see *Supported Services for PCBR* and *Unsupported Services for PCBR* in [Implementation Considerations and Limitations](#) on page 593.
4. Click **Continue**.
5. In the **Restore Prism Central** page under **Connect**, specify the following information according to the source from where you plan to recover:
 - **Endpoint Type**: From the drop-down menu, select AWS.
 - **Region Name**: Enter an AWS region name that hosts the Prism Central backup that you plan to restore. For more information, see *AWS documentation*.
 - **Bucket Name**: Enter a bucket name that hosts the Prism Central backup that you plan to restore. For more information, see *AWS documentation*.
 - (Optional for NC2 environments) **Access Key**: Enter your access key for the bucket that hosts the Prism Central backup that you plan to restore.
 - (Optional for NC2 environments) **Secret Access Key**: Enter your secret access key for the bucket that hosts the Prism Central backup that you plan to restore.
 - Click **Next**.
6. From the **Source** tab, select a point-in-time backup to restore the Prism Central instance and click **Next**.
The **Source** tab lists the available Prism Central backups in the configured S3 Bucket.
7. From the **Restore Point** tab, select a point-in-time backup to restore the Prism Central instance and click **Next**.
The **Restore Point** tab shows the restore points for the selected point-in-time..
8. From the **Installation** tab, verify the version of the Prism Central instance to be restored and click **Next**.

9. From the **Configuration** tab, specify the following:

- **Network:** From the dropdown menu, select the network where you plan to restore and install Prism Central instances.
- **Subnet Mask:** Enter the subnet mask address in the selected network.
- **Gateway IP:** Enter the gateway IP address in the selected network.
- **DNS Address(es):** Enter the DNS addresses in FQDN, IPv4, or IPv6 format.

For multiple addresses, separate each address by a comma (,).

The **NTP Address(es)**, **Container**, and **Virtual IP** fields show the relevant information associated with the selected network.

The system fetches the **VM name** and **IP address** of the Prism Central instance that you plan to restore.

10. Click **Next**.

11. From the **Summary** tab, review the backup and Prism Central information and click **Restore**.

The Prism Central instance is restored in 60 to 90 minutes, depending on the configuration data of the hosts.

The restoration involves the deployment of Prism Central and the restoration of its configuration data from the backup, which can take anywhere between 60 and 120 minutes, depending on the size of the data. The restored Prism Central instance takes an additional 30 to 40 minutes to show all the guest VMs, disks, and metrics. Wait to perform any actions on the restored Prism Central instance until all the recovery tasks are complete on the cluster. You can see the restoration status and the related processes in the **Tasks** window.

What to do next

See [Post Restoration Considerations](#) on page 609.

Recovering Prism Central from Point-in-Time Backup on Nutanix Objects

You can use a registered or unregistered AHV or ESXi cluster to restore the Prism Central instance from a point-in-time backup on Nutanix Objects. While the recovery time objective (RTO) for recovery from a point-in-time backup is 90 minutes, the recovery point objective (RPO) can be set in values ranging from one hour to 24 hours. Prism Central supports up to 30 days of restore points for restoring the Prism Central instance.

Before you begin

- See [Prism Central Backup and Restore Requirements](#) on page 592 and [Implementation Considerations and Limitations](#) on page 593 before you start.
- Ensure that the old Prism Central is disconnected and powered off before initiating the restore operation.
The restoration creates a new Prism Central instance that replaces the older instance.
- Ensure to initiate a restore from only one Prism Element.

About this task

You can recover the Prism Central instance from any available restore point. To restore a Prism Central instance from a point-in-time backup on Nutanix Objects, follow these steps:

Procedure

1. Log in to any Prism Element web console registered to the Prism Central instance that you plan to restore.
The Prism Element dashboard shows the Prism Central widget, which contains Prism Central information (IP address and connection status). If the Prism Element instance is a fresh instance that you created and is not yet registered to a Prism Central instance, it does not appear. Instead, you see the **Register or Deploy Prism Central**.
2. Click the [Settings icon](#) and from the **Settings** menu, navigate to **Data Resiliency > Restore Prism Central**.
For more information, see [Settings Menu](#) in the *Prism Element Web Console Guide*.
The **Restore Prism Central** page appears. This page provides you with options to restore the Prism Central instance from a Prism Element cluster or S3-compatible object storage.
3. Select **Restore any Prism Central from S3-compatible object storage** and click **Restore Now**.
The **Restore Prism Central** window appears, showing the services data that will be backed up and those that will not be backed up. For the detailed list, see *Supported Services for Backup and Recovery* and *Unsupported Services for Backup and Recovery* in [Implementation Considerations and Limitations](#) on page 593.
4. Click **Continue**.
5. In the **Restore Prism Central** page under **Connect**, specify the following information according to the source from which you plan to recover:
 - **Endpoint Type**: Select Nutanix Objects.
 - **IP Address or Host Name (https)**: Enter the public IP address or FQDN of the Nutanix Objects object store that hosts the Prism Central backup that you plan to restore.
 - **Bucket Name**: Enter a Nutanix Objects bucket name that hosts the Prism Central backup that you plan to restore.
 - (Optional) **Upload Certificate**: Upload the certificate (including the certificate chain of the object store) present on Nutanix Objects instance.

Note: If you do not upload a certificate, the system deploys self-signed certificates on Nutanix Objects. To replace these self-signed certificates, see [Managing FQDN and SSL Certificate](#) in the *Nutanix Objects User Guide*.

 - **Access Key**: Enter your access key of the Nutanix Objects bucket.
 - **Secret Access Key**: Enter your secret access key of the Nutanix Objects bucket.
 - Click **Next**
6. From the **Source** tab, select a point-in-time backup to restore the Prism Central instance and click **Next**.
The **Source** tab lists the available Prism Central backups in the configured S3 Bucket.
7. From the **Restore Point** tab, select a point-in-time backup to restore the Prism Central instance and click **Next**.
The **Restore Point** tab shows the restore points for the selected point-in-time.
8. From the **Installation** tab, verify the version of the Prism Central instance to be restored and click **Next**.

9. From the **Configuration** tab, specify the following:

- **Network:** From the dropdown menu, select the network where you plan to restore and install Prism Central instances.
- **Subnet Mask:** Enter the subnet mask address in the selected network.
- **Gateway IP:** Enter the gateway IP address in the selected network.
- **DNS Address(es):** Enter the DNS addresses in FQDN, IPv4, or IPv6 format.

For multiple addresses, separate each address by a comma (,).

The **NTP Address(es)**, **Container**, and **Virtual IP** fields show the relevant information associated with the selected network.

The system fetches the **VM name** and **IP address** of the Prism Central instance that you plan to restore.

10. Click **Next**.

11. From the **Summary** tab, review the backup and Prism Central information and click **Restore**.

The Prism Central instance is restored in 60 to 90 minutes, depending on the configuration data of the hosts. The restoration involves the deployment of Prism Central and the restoration of its configuration data from the backup, which can take anywhere between 60 and 120 minutes, depending on the size of the data. The restored Prism Central instance takes an additional 30 to 40 minutes to show all the guest VMs, disks, and metrics. Wait to perform any actions on the restored Prism Central instance until all the recovery tasks are complete on the cluster. You can see the restoration status and the related processes in the **Tasks** window.

What to do next

See [Post Restoration Considerations](#) on page 609.

Post Restoration Considerations

Consider the following after the Prism Central restoration:

- Use the newly restored Prism Central instance only.

If the old Prism Central instance becomes available, shut down or delete the old instance because running the old Prism Central instance can cause data corruption.

Note: If the Prism Central restoration fails, contact Nutanix Support. Do not bring up the old Prism Central instance.

- Reset the credentials.

The Prism Central instance restores with the default credentials. Nutanix recommends changing the default credentials. For information about changing the default credentials, see [Logging Into Prism Central](#) in the *Prism Central Infrastructure Guide*.

Note: If you have both S3-compatible object storage and on-prem clusters configured as backup targets, and you recovered the Prism Central instance through an on-prem cluster, you must reconfigure the S3 bucket credentials after the recovery through the **General > Prism Central Backups** widget in **Settings**.

- Reconfigure the proxy server. For more information, see [Configuring an HTTP Proxy](#) in the *Prism Central Admin Center Guide*

If the old Prism Central instance had a proxy server, reconfigure the proxy server so that the recovered Prism Central instance maps to the correct IP address.

- Reconfigure the fully qualified domain name (FQDN).

If the old Prism Central instance had an FQDN, reconfigure the FQDN so that the recovered Prism Central instance maps to the correct IP address.

- If you had recovery plan jobs (RPJ) in progress, perform the steps in [KB-10962](#).

If the old Prism Central instance had a failover task running (Nutanix Disaster Recovery) or protection policy with entities protected with synchronous replication schedule, perform the steps in [KB 10962](#) to ensure that all the failover tasks stuck in the running state are terminated, and a script is executed for efficient recovery of the Prism Central instance.

- Restore the secret keys (for example, PCKMS) manually through nCLI on the newly restored Prism Central instance. For more information, see [Importing Keys](#) in the *Security Guide*.

First, back up the secret keys present on the old Prism Central instance to a text file:

```
ncli> data-at-rest-encryption backup-software-encryption-keys file-path=/path/textfile.txt password=password
```

Replace `/path/textfile.txt` with the absolute path of the text file wherein you have backed up the secret keys.

Copy the text file to the newly restored Prism Central instance, then restore the secrets manually from the text file:

```
nutanix@pcvm$ mantle_recovery_util -backup_file_path /path/textfile.txt -password <password>
```

Replace `/path/textfile.txt` with the absolute path of the text file wherein you have backed up the secret keys.

Tip: Run the following command to list the secret keys backed up in the text file:

```
nutanix@pcvm$ mantle_recovery_util -backup_file_path /path/textfile.txt -list_key_ids -password <password>
```

Replace `/path/textfile.txt` with the absolute path of the text file wherein you have backed up the secret keys.

- Reconfigure Life Cycle Manager (LCM) if you do not use the Nutanix portal or if you use LCM in a dark site. For more information, see [LCM Settings for Dark Sites - No Internet Connectivity](#). For information on reconfiguration of LCM 3.1, see [KB 17966](#).

Migrating a Prism Central Instance

You can use the Prism Central Backup and Restore (PCBR) capability to migrate a Prism Central instance from one cluster to another.

Before you begin

- Ensure that a backup copy is on the remote cluster before stopping the Prism Central instance.
- Ensure that both the primary cluster (where the Prism Central instance is hosted) and the remote cluster (where you plan to migrate the Prism Central instance) are registered to the Prism Central instance that you want to migrate.
- Read the *Supported Services for PCBR* and *Unsupported Services for PCBR* sections in [Implementation Considerations and Limitations](#) on page 593.

Procedure

1. Configure the remote cluster as the backup location for the Prism Central instance that you plan to migrate.

For more information, see [Backing up Prism Central](#) on page 595.

The backup starts immediately and takes a couple of hours to complete. The backup progress and last backup time are visible on the Prism Central web console (go to [Settings icon](#) and navigate to **General > Prism Central Backups**). For more information, see [Prism Central Settings \(Infrastructure\)](#).

2. Log in to the Prism Central instance that you plan to migrate through SSH as nutanix user.
3. Stop the Prism Central instance that you plan to migrate using the following command, and then type **I agree** as a confirmation:

```
nutanix@pcvm$ cluster stop
```

Note: The Prism Central instance takes 10-15 minutes to stop.

4. Power off the Prism Central instance.
5. Log in to the Prism Element web console on the remote cluster as an administrator.
6. Perform the restore procedure described in [Restoring Prism Central](#) on page 604.

Disabling Prism Central Backup and Restore

The following procedure to disable Prism Central Backup and Restore (PCBR) is the same for continuous and point-in-time backup types.

About this task

To disable PCBR on a Prism Central instance, follow these steps:

Procedure

1. Log in to the Prism Central web console.
2. Click the [Settings icon](#) and from the **Settings** menu, navigate to **General > Prism Central Backups**.

For more information, see [Prism Central Settings \(Infrastructure\)](#) on page 44.

The **Prism Central Backups** page appears. This **Summary** page provides information on the backup status of the Prism Central instance and guides you through configuring backups.

Tip: The **What's New** page appears when you go to the **Prism Central Backups** page for the first time.

3. Perform one of the following:
 - » To disable continuous backup, navigate to **Continuous Backup** and click **Remove** in front of the target cluster.
 - » To disable point-in-time backup, navigate to **Point-in-Time Backup** and click **Remove** in front of the point-in-time backup .

What to do next

When PCBR is disabled, Prism Central resources are unprotected. To ensure recovery of the Prism Central instance in case of a disaster, enable it again after completing the tasks that require PCBR to be disabled.

Intelligent Operations

Intelligent Operations allows you to run your IT operations efficiently using intelligent operational features that include:

- Cluster activity analysis.
- Application discovery in a cluster.
- Application monitoring in a cluster.
- Resource capacity and usage planning.
- Task automation using playbooks.
- Resource and activity related reports generation for Nutanix and non-Nutanix environments

You can enable Intelligent Operations from **Admin Center** application. For more information, see [Enabling Intelligent Operations](#) in *Intelligent Operations Guide*.

Starting from Prism Central version 2023.3, NCM Intelligent Operations (formerly AIOps) now has a new user interface. The following Operations Entities have been moved from Prism Central Infrastructure app to the new user interface created for NCM Intelligent Operations:

- Analysis
- App Discovery
- Monitoring Configurations
- Operations Policies
- Planning
- Playbooks
- Reports
- Settings and Configurations

The documentation for the above-mentioned entities has been moved to the newly created *Intelligent Operations Guide*. For more information, see [Intelligent Operations Guide](#). You can also find the [Intelligent Operations Release Notes](#) under Nutanix Cloud Manager on the documentation portal.

Large Files Upload Using Objects Lite

Objects Lite is miniaturization of Nutanix Objects where the whole objects stack runs as a container in Microservices Architecture.

The Objects Lite service is a compact version of core Nutanix Objects product, specifically designed to operate within the microservices architecture. It is enabled by default on Prism Central for internal usage of Prism Central services. This service enables Prism Central users to upload large objects with streaming and resumability support to other Prism Central services like images, catalog, and LCM.

This method optimizes the interaction between services and eliminates external interaction. While Objects Lite is Amazon S3-compatible, it offers a limited set of APIs specifically tailored to Prism Central services, including images and catalog.

Objects Lite Data Storage

Objects Lite uses one of the Prism Element registered to Prism Central to store its data. If no Prism Element is registered, Objects Lite waits until one becomes available. A data container named `objects1<objects-lite-uuid>` is created for this purpose.

Note: Unregistering Prism Element from Prism Central fails and you need to cleanup Objects Lite to successfully unregister Prism Element from Prism Central. In such a case, Nutanix recommends that you contact Nutanix Support.

Objects Lite Limitations

This section lists the limitations of Objects Lite.

The limitations of Objects Lite are as follows:

- User Roles: Only users with Super Admin or Prism Central Admin roles can upload objects.

Note: If you change your role from Prism Viewer to Prism Admin, subsequent operations might result in an Access Denied error. This happens because the previous role mapping persists in the object cluster for up to two hours, causing the system to still recognize you as a user with Prism Viewer role. To resolve this error, restart the Objects service on Prism Central.

- Upload-only Support: Objects Lite supports only upload operations. Download and list operations are not supported.
- Failure Handling: If an upload fails, Nutanix recommends that you upload the object again. Depending on the method used (PutObject or Multipart), you might need to upload the entire object or just the failed parts. Make sure to keep track of the successfully uploaded parts if you want to resume the `multipart` upload from the point of failure as `ListParts` API for Prism Central user is not supported.
- Concurrency Limits: Objects Lite supports up to eight parallel requests at a time. When concurrency is higher than eight, you might see the Too Many Requests error. In such cases, try again after some time.
- Temporary storage solution: Objects Lite needs to be cleaned before Prism Central Prism Element unregistration for the Prism Element hosting objects-lite's data. Post-compliance backup retention (PCBR) events are also not supported, and all Objects Lite data is lost after a Prism Central's PCBR recovery.

Objects Lite Access

Objects Lite provides an endpoint that allows Admin and Super Admin roles to upload data directly to Objects Lite, bypassing the respective service in the data path. This endpoint is specifically designed for uploads and can be accessed using Prism Central user credentials.

The uploaded data can later be consumed by the corresponding services.

Endpoint URL: The endpoint follows the format,

```
https://<pcip>:9440/api/prism/v4.0/objects/
```

Note: This endpoint is Amazon S3-compatible and works with any S3 SDK or AWS CLI.

To access this endpoint, you must authenticate using your PC credentials.

For more information, see [Accessing Objects Lite Using AWS CLI- Put object](#) on page 613.

Accessing Objects Lite Using AWS CLI- Put object

This section contains the PutObject AWS CLI commands that you can use to upload a file to Objects Lite.

Before you begin

Ensure that you

- Install AWS CLI.
- Configure AWS CLI with Objects Lite using the `pc-ip`, `username` and `password` attributes.

```
access_key: base64(<username>:<password>)
```

```
secret key: base64(<username>:<password>)
aws configure set aws_access_key_id <access-key>
aws configure set aws_secret_access_key <secret-key>
aws configure set endpoint_url https://<pc-ip>:9440/api/prism/v4.0/objects/
```

Note: The access and secret keys for Objects Lite use the base64(<username>:<password>) format.

About this task

Uploading a file using PutObject, Multipart Upload, or High level upload command.

Procedure

Upload the file using the PutObject API with AWS CLI:

```
> aws s3api put-object --bucket <bucket-name> --body <file-path> --key <object-key> --no-verify-ssl
```

- The following list describes the Parameters in the previous command:

- *--bucket*: The name of the bucket where the object is stored.
- *--key*: The key (name) for the object in the bucket.
- *--body*: The path to the file to upload.
- *--no-verify-ssl*: Disables SSL certificate verification.

Note: This command supports uploading files smaller than five GB. For larger files, consider using the multipart upload method or the high-level `aws s3 cp` command.

After the file is uploaded, the command displays the ETag of the uploaded object.

Uploading Files to Objects Lite Using AWS CLI - Multipart Upload

This section contains Multipart upload AWS CLI procedure to upload a file to Objects Lite.

About this task

To upload a file using the Multipart Upload API with AWS CLI, follow these steps:

Procedure

1. Initiate multipart upload.

```
$ aws s3api create-multipart-upload --bucket <bucket-name> --key <object-key> --no-verify-ssl
```

This command returns an Upload ID.

2. Upload parts.

```
> aws s3api upload-part --bucket <bucket-name> --key <object-key> --part-number <part-number> --body <file-path> --upload-id <upload-id> --no-verify-ssl
```

This command returns an ETag for each part.

3. Complete Multipart upload.

```
> aws s3api complete-multipart-upload --bucket <bucket-name> --key <object-key> --upload-id <upload-id> --multipart-upload file://<parts-file> --no-verify-ssl
```

Sample JSON file for completing the upload:

```
{
  "Parts": [
    {
      "ETag": "\"etag1\"",
      "PartNumber": 1
    },
    {
      "ETag": "\"etag2\"",
      "PartNumber": 2
    },
    {
      "ETag": "\"etag3\"",
      "PartNumber": 3
    }
  ]
}
```

The following list describes the parameters in the previous commands:

- *--bucket*: The name of the bucket where the object is stored.
- *--key*: The key (name) for the object in the bucket.
- *--part-number*: The part number for the part being uploaded.
- *--upload-id*: The Upload ID received from `create-multipart-upload`.
- *--multipart-upload*: The JSON file containing the list of parts and their ETags.

Note: Track the successfully uploaded parts to resume the overall upload process in case of a failure and avoid re-uploading the parts that have already been uploaded. The `ListParts` API is not supported.

Uploading Files to Objects Lite Using AWS CLI - Highlevel upload

This section describes the AWS CLI procedure for high level upload that can be used to upload a file to Objects Lite.

About this task

You can copy files from local storage to S3-compatible endpoints using the AWS S3 copy command (`aws s3 cp`). It simplifies the uploading process and supports automatic multipart uploads for large files. It also provides support for concurrent uploads, improving performance.

To upload a file using the high level upload command with AWS CLI, follow these steps:

Procedure

1. Configure the multipart threshold, chunk size, and concurrency.

```
$ aws configure set s3.multipart_threshold <multipart-threshold>
$ aws configure set s3.multipart_chunksize <multipart-chunksize>
$ aws configure set s3.max_concurrent_requests <max-concurrent-requests>
```

The following list describes the parameters in the previous commands:

- *multipart-threshold*: Size threshold for when to use multipart uploads.

- *multipart-chunksize*: Size of each part in a multipart upload.
- *max-concurrent-requests*: Maximum number of concurrent requests for multipart uploads.

2. Upload the file.

```
$ aws s3 cp <file-path> s3://<bucket-name>/<object-key> --no-verify-ssl
```

Note: This command does not handle resumable uploads or retries in case of failures. In that case, consider uploading the file again.

High Availability in Prism Central

Prism Central supports high availability using the following functionalities:

- **VM high availability: Migrates and restarts the Prism Central VM (PCVM) from a failed node to another node.**

When a node that hosts a PCVM fails, Acropolis migrates and restarts the PCVM in another node that has enough data resources and entities required to run the PCVM.

Note: Prism Central supports VM high availability only if your deployment meets the following requirements:

- PCVM runs pc.2024.1 or later
- The cluster (Prism Element) runs AOS 6.8 or later
- AHV version is 20230302.100173 or later.

Both single VM Prism Central and three VM scale-out Prism Central deployments support high availability using the VM high availability feature in Acropolis.

For information on high availability, see [VM High Availability in Acropolis](#) in the *Prism Element Web Console Guide*. For information on how to enable high availability in a PCVM, see [Enabling High Availability for the Cluster](#) in the *Prism Element Web Console Guide*.

- **Scale-out: Scales out a single VM Prism Central deployment to a three-VM deployment with a fault tolerance of one PCVM.**

Prism Central supports high availability by scaling out a single VM Prism Central deployment to a three-VM deployment with a fault tolerance of one PCVM. For information on how to scale out a Prism Central deployment across multiple VMs, see [Expanding \(Scale Out\) Prism Central](#).

In a scale-out Prism Central deployment, if one PCVM fails, the other two PCVMs perform the load balancing for the failed PCVM, and the Prism Central stays active. If two or all three PCVMs fail, then Prism Central is considered inactive.

For information on the configurations, requirements, specifications, and limitations of scalability, see the *Prism Central Scalability* topic in the release notes for the [Prism Central](#) version to be installed.

Finding the Prism Central Version

About this task

You can check the version of your Prism Central instance using the Prism Central web console.

Do the following to check the Prism Central version in the web console.

Procedure

1. Log in to Prism Central.
2. Click the user icon in the main menu and then select the **About Nutanix** option from the dropdown menu.
An **About Nutanix** window appears that includes the Prism Central version number. It also includes a link to Nutanix patent information.
3. Click **Close** to close the window.

Finding the AHV Version on Prism Central

About this task

You can see the installed AHV version in Prism Central.

Procedure

To view the AHV version installed on any host in the clusters managed by Prism Central, perform the following steps:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Hardware** > **Hosts** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#).
The system displays the **List** tab by default with all the hosts across registered clusters.
3. Click the target host name for which you want to see the hypervisor version.
The system displays the [Host Details View](#) with **Summary** tab.

4. Observe the Hypervisor Version field in the **Properties** widget to view the hypervisor version.




Figure 186: Hypervisor Version in Host Detail View

Finding the AOS Version Using Prism Central


To view the Nutanix AOS version running in the Prism Central, do the following:

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#) on page 40, and navigate to **Hardware > Clusters** from the **Navigation Bar**. For information about the **Navigation Bar**, see [Application-specific Navigation Bar](#) on page 61.

The system displays the **List** tab by default with all the registered clusters in **Nutanix** environment.

3. In the **List** tab, the AOS version of the cluster is listed against the cluster name.



Name	AOS Version	Upgrade Status	Hypervisors	Host Count	VM Count	Cluster Runway	Inefficient VMs
auto_cluster_prod_f38293eb9649	6.6	Succeeded	AHV	1	15	302 days	13

Figure 187: AOS Version in a Cluster

Expanding (Scale Out) Prism Central

Before you begin

The following requirements must be met before you can expand Prism Central or add a Prism Central VM:

- The specified gateway must be reachable.
- No duplicate IP addresses can be used.
- The container used for deployment is mounted on the hypervisor hosts.
- For Prism Central versions lower than 2023.3, the hosting cluster should be registered with Prism Central.
- When installing on an ESXi cluster:
 - vCenter and ESXi cluster must be configured properly. For more information about vCenter and ESXi configuration, see the [vCenter Configuration](#) topic in *vSphere Administration Guide for Acropolis*.
 - vCenter must be registered in Prism.
 - DRS must be enabled in vCenter.
 - vCenter is up and reachable during the deployment.

About this task

If Prism Central is just a single VM currently, you can expand it to three VMs. This increases both the capacity and resiliency of Prism Central (at the cost of maintaining two additional VMs). For information on Prism Central scalability, see the *Prism Central Scalability* topic in the release notes for the [Prism Central](#) version to be installed.

To expand this Prism Central instance across multiple VMs, do the following:

Note:

- Scaling out of Prism Central VM not registered with the hosting Prism Element is not supported if any of the portfolio products such as Self-Service or Disaster Recovery is enabled on Prism Central.
- Scaling out of Prism Central is not supported for three-tier non-Nutanix Prism Central deployments.
- All the scale out Prism Central VMs must run on the same cluster. For example, running two VMs in cluster_1 and one VM in cluster_2 is not supported.

- If you have configured application monitoring on a single node Prism Central setup, and you scale out to 3 nodes running Prism Central version PC.2023.4, you must power off and then again power on the newly added nodes.

Procedure

1. Log in to Prism Central.
2. Click the [Settings icon](#) and then select **Prism Central Management** from the **Settings** menu. For more information, see [Prism Central Settings \(Infrastructure\)](#) on page 44.
The **Manage Prism Central** page appears. This page provides information about this Prism Central instance.




Figure 188: Manage Prism Central Page

3. To expand this Prism Central instance from one to three VMs, click **Scale Out PC** to display the **Scale Out PC** page and do the following:

Note: A pop-up window appears explaining that scale out is a one-way process. Click **Continue** to display the **Scale Out PC** page. Once you scale out a Prism Central instance from a single VM to multiple VMs, you cannot revert back. Deleting any of the Prism Central VMs may result in data loss.

Scale Out Prism Central

Prism Central now supports resiliency and higher capacity with scale-out.

auto_cluster_prod_4f046b91b70c
1 VM
2 New

Network
Subnet Mask
Gateway
Virtual IP

VM Name
auto_pc_6405cf2782e14f046b91b70d0
IP
8 vCPUs
32 GiB

VM Name
PC-NameOption-2
IP
8 vCPUs
32 GiB

VM Name
PC-NameOption-3
IP
8 vCPUs
32 GiB

Add PC VM IP

Microservices Infrastructure Configuration
Prism Central Domain Name
Network
Subnet Mask
Gateway
Provided IP Addresses

Scaling out on Microservices Infrastructure requires a minimum of 10 IP addresses. Please provide the remaining IP range for 8 IPs below.
IP Address Range (8 remaining IPs needed)
e.g. 10.0.0.2
e.g. 10.0.0.7
Please ensure that the entered IP address range does not conflict with any reserved IPs (e.g. DHCP IP Pool) in your network.

Expand

Figure 189: Scale Out PC Page

NUTANIX

Prism | Additional Operations | 622

Note: The Prism Central VM image is deployed from the target cluster. If the required image cannot be accessed, typically because there is no Internet access (such as at a dark site), a pop-up error message appears. In this case you have the option to manually download and deploy the image as follows:

1. Log on to the Nutanix customer support portal, click **Downloads > Prism Central**, and click the download link for the target version to save the Prism Central binary .TAR and metadata .JSON

files on your local media. The binary .TAR file is from the Prism Element bundle, used for 1-click deployment.

2. Log on (using SSH) to any Controller VM in the cluster specified in the error message and copy the Prism Central binary .TAR and metadata .JSON files to the Controller VM.
3. Run the following command to deploy the Prism Central image:

```
nutanix@cvm$ ncli software upload software-type=PRISM_CENTRAL_DEPLOY file-path=file_path meta-file-path=metadata_file_path
```

The *file_path* is the full (absolute) path to the .TAR file, and the *metadata_file_path* is the full path to the .JSON file. After this step completes, you can continue the scale out procedure.

- a. Review the **Network**, **Subnet Mask**, and **Gateway** fields, which display the network name, subnet mask value, and gateway IP address applied to this Prism Central instance.

The values in these three fields are read-only and cannot be changed.

Note: The scale-out process uses the IP addresses from the network configured. If you used the default network configuration being the Private Network [default] while enabling microservices infrastructure, Prism Central displays and uses Private Network [default] details in this section.

- b. To specify a virtual IP address for Prism Central, click the **Add Virtual IP** link.

This displays the **Add PC Virtual IP** window. A virtual IP can be used as a single point of access for Prism Central. Enter the IP address in the Virtual IP field, and then click **Update**.

- c. Specify IP addresses for the two new Prism Central VMs.

1. Click the pencil icon for one of the new VMs. (The VM names are set automatically.) The **IP field** opens for editing. Enter the IP address and then click the green check mark icon to save that address.



Figure 190: Edit IP Field

2. Repeat this step for the second new VM.

Note: If proxy is configured on the cluster, ensure to add the IP address of Prism Central and Prism Element to the proxy allowlist.

The vCPU count and memory size in the new VMs are fixed and match the current Prism Central VM values; you cannot change these values.

- d. When all the parameters are correct, click **Expand**.

This starts the process of creating the new VMs and deploying this Prism Central instance. You can monitor progress from the **Tasks** page. For more information, see [Tasks View](#) on page 496.

What to do next

- Set up VM-VM affinity rules between PC VMs. For more information about setting up the VM-VM affinity rules, see [Configuring Legacy VM-VM Anti-Affinity Policy](#).

- In some cases, especially when using NCM Self-Service (formerly known as Calm), a version mismatch with a linked product could occur after expanding Prism Central. To fix this situation, perform a life cycle manager (LCM) inventory after completing the Prism Central expansion procedure. For instructions on how to perform an LCM inventory, see the [LCM documentation](#).

Shutting Down or Starting Up Prism Central VM

About this task

A Prism Central VM (PC VM) is managed like any other VM in a cluster through the Prism Element web console for the cluster in which the Prism Central VM is running. However, shutting down a PC VM requires extra caution.

Note:

Prism Central supports features that could be damaged by shutting down the PC VM abruptly. You must follow the steps carefully to avoid any issues.

Procedure

- To shut down the PC VM in single VM configuration, or to shut down all the three PC VMs in a scale out configuration, see [Shutting down or starting up all the PC VMs in a Prism Central Configuration](#) on page 625.
- To shut down a single PC VM in a scale-out Prism Central configuration, see [Shutting down or starting up a single PC VM in a scale-out PC configuration](#) on page 626.

Shutting down or starting up all the PC VMs in a Prism Central Configuration

About this task

To shut down the PC VM in single VM configuration, or to shut down all the three PC VMs in a scale out configuration, perform the following steps:

Procedure

1. SSH to the IP address of the PC VM that you want to shut down.
2. Stop the cluster.

```
nutanix@pcvm$ cluster stop
```

3. Check the status of the services running on the cluster.

```
nutanix@pcvm$ cluster status
```

The command output should show the following four services still running: Zeus, Scavenger, VipMonitor, and ikatProxy.

4. To shut down a PC VM, log on to the web console of the host and go to the VM dashboard (see [VM Dashboard](#) in Prism Element Web Console Guide). In the VM Table view, select the PC VM you want to shut down (one at a time), and then select **Power Off Actions >> Power Off** or **Power Off Actions >> Guest Shutdown** from the action links.

Alternatively, you can SSH to the IP address of the PC VM and run the command to shut down the PC VM.

```
nutanix@pcvm$ sudo shutdown -h now
```

You must run this command for each PC VM that you want to shut down.

5. To power on a PC VM, log on to the web console of the host and go to the VM Dashboard (see [VM Dashboard](#)). In the VM Table view, select the PC VM you want to power on (one at a time), and then select **Power On**.
6. Once all the PC VMs are powered on, check the status of the services from any one of the PC VMs.

```
nutanix@pcvm$ cluster status
```

In the command output, verify that all the four services (Zeus, Scavenger, VipMonitor, and ikatProxy) are running.

7. Once all the four services start running, start the cluster from any one of the PC VMs.

```
nutanix@pcvm$ cluster start
```

Shutting down or starting up a single PC VM in a scale-out PC configuration

About this task

To shut down or start up a single PC VM in a scale-out PC configuration, perform the following steps:

Procedure

1. SSH to the IP address of the PC VM.
2. Run the following command to gracefully stop all the services:

```
nutanix@pcvm$ genesis stop all
```

3. Now log on to the web console of the host and go to the VM dashboard (see [VM Dashboard](#) in *Prism Element Web Console Guide*).
4. In the VM Table view, select the PC VM you want to shut down, and then select **Power Off Actions >> Power Off** or **Power Off Actions >> Guest Shutdown** from the action links.

Alternatively, you can SSH to the IP address of the PC VM and run the command to shut down the PC VM.

```
nutanix@pcvm$ sudo shutdown -h now
```

5. To power on a PC VM in a scale-out PC configuration, log on to the web console of the host and go to the VM Dashboard (see [VM Dashboard](#)). In the VM Table view, select the PC VM you want to power on (one at a time), and then select **Power On**.
6. Restart the genesis service:

```
nutanix@pcvm$ genesis restart
```

7. Check the status of the services running on the cluster.

```
nutanix@pcvm$ cluster status
```

8. Start the services in the PC VM:

```
nutanix@pcvm$ cluster start
```

Keyboard Shortcuts in Prism Central

You can use the following keyboard shortcuts to invoke important menu options or views in Prism Central:

Table 173: Keyboard Shortcuts for Prism Central Menu Options

Shortcut Key	Menu Option/View
s	Settings Menu
f	Spotlight (search bar)
u	User Menu
h	Help menu (? menu)
p	Recent tasks

You can use the arrow keys to select a particular menu option.

IP Address Reconfiguration

You can use this procedure to reconfigure the IP address and gateway of single Prism Central instances as well as Prism Central VMs in a Scale out Prism Central (clustered PC VMs or PC cluster).

Note: Reconfiguring the IP address and gateway of Prism Central VMs does not require additional steps when using Nutanix Disaster Recovery (formerly Leap).

Preparing to Reconfigure the IP Address and Gateway of PC VMs

Perform the tasks described in this topic before you start the IP address reconfiguration procedure.

About this task

You must perform the following tasks while the Prism Central VMs are still on the existing IP addresses.

Procedure

1. Coordinate Prism Central downtime, because the features and functionality of Prism Central will be unavailable for the entire duration of the IP address reconfiguration procedure.
2. Create a table to map the existing IP addresses with the new IP addresses of the Prism Central VMs for your reference.
3. Use SSH to log on to any running Prism Central VM in the PC cluster or the single PC VM if you do not have a PC cluster.
4. Verify if the PC VM or PC cluster is in a stable state.

```
nutanix@pcvm$ cluster status
```

This command is valid even if you only have a single PC VM and do not have a PC cluster.

5. Run the NCC health checks to make sure that the PC cluster is in a healthy state.
6. If you want to reconfigure the IP addresses of the entities in the AOS clusters (such as IP addresses of CVMs and hypervisor hosts), on which the PC is hosted, reconfigure those IP addresses first before you reconfigure the IP addresses of the PC VMs.
7. Stop the PC VM or PC cluster.

```
nutanix@pcvm$ cluster stop
```

This command is valid even if you only have a single PC VM and do not have a PC cluster.

Wait to proceed until an output showing all the services as DOWN is displayed, except the Zeus and Scavenger services.

Reconfiguring the IP Address and Gateway of Prism Central VMs

You can reconfigure the IP addresses of Prism Central VMs.

Before you begin

Ensure that you have completed the tasks described in [Preparing to Reconfigure the IP Address and Gateway of PC VMs](#) on page 627.

About this task

The procedure to reconfigure the IP addresses of Prism Central VMs differs depending on whether or not IP Address Management (IPAM) is enabled in an AHV cluster. The procedure to reconfigure the IP addresses of Prism Central VMs in an ESXi cluster is the same as the procedure of an AHV cluster in which IPAM is disabled.

For information on IPAM, see [IP Address Management](#) section in the *AHV Administration Guide*.

In this document, managed network refers to an AHV network that has IPAM enabled, and unmanaged network refers to an AHV network that does not use IPAM.

The following conditions apply to Prism Central IP address re-configuration:

- Nutanix Cloud Clusters (NC2) does not support reconfiguration of Prism Central VM IP addresses.
- Do not reconfigure the IP address of the Prism Central VM if any Nutanix Kubernetes Engine (NKE), Nutanix Kubernetes Platform (NKP), or Objects are deployed on the cluster since changing the Prism Central VM IP address is not supported for NKE, NKP, or Objects after deployment.

To reconfigure the IP address of a single Prism Central VM or Prism Central VMs in a Prism Central cluster, follow these steps.

Procedure

1. Log on to a Prism Central VM using SSH with the `nutanix` access, and run the `external_ip_reconfig` script.

```
nutanix@pcvm$ external_ip_reconfig
```

2. Follow the prompts to type the new netmask, gateway, and external IP addresses.

A message similar to the following is displayed when the procedure begins:

```
External IP reconfig started
```

A message similar to the following is displayed if the procedure is completed successfully:

```
External IP reconfig finished successfully. Restart all the CVMs and start the cluster.
```

The message indicates that you must restart the CVMs. In this case, you are reconfiguring the IP addresses of the Prism Central VMs, so you must restart the Prism Central VMs and not the CVMs.

Note:

If your Prism Central VMs are running in a managed network of an AHV cluster, do not start the Prism Central VMs yet but proceed to step 3.

If your Prism Central VMs are in an unmanaged network of an AHV cluster or are in an ESXi cluster, skip steps 3, 4, and 5, and proceed to step 6 directly.

Perform steps 3, 4, and 5 only if the new IP address is in a different IP address range than the previous IP address range and your Prism Central VM is in a managed network of an AHV cluster.

A message similar to the following is displayed if the procedure fails:

```
External IP reconfig Failed
```

If the procedure fails, run the following command to check the log for troubleshooting purposes:

```
nutanix@pcvm$ cat ~/data/logs/ip_reconfig.log
```

3. (Optional) Use SSH to log onto any CVM of the AOS cluster that is hosting the Prism Central cluster or Prism Central VM.

Note: Perform this step only if the new IP address is in a different IP address range than the previous IP address range and your Prism Central VM is in a managed network of an AHV cluster.

4. (Optional) Update the NIC attached to the Prism Central VM with the new IP address of the Prism Central VM using the following command:

Note: Perform this step only if the new IP address is in a different IP address range than the previous IP address range and your Prism Central VM is in a managed network of an AHV cluster.

```
nutanix@cvm$ acli vm.nic_update pc-vm-name pc-vm-mac request_ip=true ip=pc-vm-ip
```

Replace the variables with their appropriate values as follows:

- *pc-vm-name*: Name of the Prism Central VM whose IP address you modified.
- *pc-vm-mac*: MAC address of the Prism Central VM whose IP address you modified (press **Tab** after you type the name of the Prism Central VM to automatically populate the MAC address).
- *pc-vm-ip*: New IP address of the Prism Central VM.

5. (Optional) Log on to the Prism Element web console and modify the IP address range in the IP pool for the Prism Central network to work with the modified IP address of Prism Central VMs. In the Prism Element web console, perform the following:

Note: Perform this step only if the new IP address is in a different IP address range than the previous IP address range and your Prism Central VM is in a managed network of an AHV cluster.

- a. In the **VM** dashboard, click the **Table** view.
 - b. Select a **Prism Central VM** and click **Update**.
 - c. In the **Update VM** dialog box, scroll down to the **Network Adaptors (NIC)** section.
 - d. In the **VLAN Name** column, note the names of the networks attached to your Prism Central VM.
 - e. Click **Close**.
 - f. Click the gear icon in the top-right corner, and under **Settings**, click **Network Configuration**.
 - g. In the **Network Configuration** dialog box, identify the network names you noted in step d.
 - h. Select each network and click the pencil icon next to the network.
 - i. Under **IP Address Pools**, modify the IP address range if the new IP address of the Prism Central VM is in a different IP address range than the previous IP address range.
6. Restart each Prism Central VM in the Prism Central cluster or the single Prism Central VM if you do not have a Prism Central cluster.
 7. Once the Prism Central VMs are powered on, use SSH to log on to each Prism Central VM or the single Prism Central VM if you do not have a Prism Central cluster and verify if the output of the following verification command displays the new IP addresses:

```
nutanix@pcvm$ python cluster/bin/external_ip_reconfig_verify.py
```

8. Start the Prism Central cluster using the following command:

```
nutanix@pcvm$ cluster start
```

9. Re-configure Microservices Infrastructure and the trust communication between Prism Element cluster and Prism Central.

For the steps to re-configure Microservices Infrastructure and trust communication, see [Reconfiguring Microservices Infrastructure and Trust Communications](#) on page 631.

Note: If Prism Element is not registered to Prism Central and you reconfigure the Prism Element or Prism Central IP address, ensure that the trust communication is re-established for the functioning of Microservices (CMSP) infrastructure.

10. Perform the following verification steps:
 - a. Log on to the Prism Element web console of every cluster attached to the Prism Central VM whose IP address you changed.
 - b. On the **Home** page, verify if the **Prism Central Registration** status is displayed as **OK**. This might take a few minutes after the Prism Central VM or Prism Central cluster is powered on.

What to do next

If you need to change the IP address of the Prism Element cluster hosting the Prism Central VMs, you must first complete the steps for changing the IP address of the Prism Central VM and restart the Prism

Central VM. Wait till all the services in Prism Central are restarted and then run the following command on Prism Element to ensure that the new IP address of Prism Central is now associated with the Prism Element cluster:

```
ncli multicluster get-cluster-state
```

If the new IP address of the Prism Central VM is not reflected in the command output of the above command, contact Nutanix Support.

Reconfiguring Microservices Infrastructure and Trust Communications

Reconfigure Microservices Infrastructure and the trust communications between Prism Element cluster and Prism Central.

About this task

To reconfigure Microservices Infrastructure and the trust communications between Prism Element cluster and the Prism Central VM or VMs, follow these steps.

Procedure

1. Log on to a Prism Central VM using SSH with the `nutanix` access.

Complete the following steps on the Prism Central VM:

- a. (Optional) Identify the control plane nodes in a scale-out Prism Central instance.

In a single-VM Prism Central instance, the Prism Central VM is the control plane VM. For scale-out Prism Central instance, run the following command to identify the control plane VM.

```
nutanix@pcvm$ allssh 'sudo systemctl is-enabled kubelet-master'
```

The command output displays all the Prism Central VMs with the control plane **Enabled** status. It displays the message, Failed to get unit file state for kubelet-master.service, for any Prism Central VM that is not a control plane node.

- b. Download the `mssp serviceregistry-cli` and add it to the `/home/nutanix/bin` folder of the control-plane Prism Central VM.

Run this command to download and add the `mssp serviceregistry-cli` file.

```
nutanix@pcvm$ wget https://download.nutanix.com/kbattachments/15469/  
mssp serviceregistry-cli -O /home/nutanix/bin/mssp serviceregistry-cli
```

Note: This step is also not required for Prism Central versions pc.2024.3 and later because the Prism Central binary already includes the `mssp serviceregistry-cli` command interface file.

- c. Configure Microservices Infrastructure for the Eth0 IP address change.

```
nutanix@pcvm$ PYTHONPATH=/home/nutanix/cluster/bin /home/docker/mssp_controller/  
bootstrap/mssp_tools/cmsp-scripts/cmsp_ip_reconfig
```

- d. Find the UUID of the Prism Central VM.

```
nutanix@pcvm:~$ ncli cluster info
```

The Cluster UUID field value is the cluster UUID of Prism Central VM.

2. Log on to a Controller VM of the Prism Element cluster using SSH with the `nutanix` access.

Complete the following steps on the Controller VM:

- a. Reestablish the trust between the Prism Element cluster and scale-out Prism Central.

- » For scale-out Prism Central instance with a virtual IP address, run the following command:

```
nutanix@cvm:~$ python bin/update_zeus_trustsetup.py --cluster_uuid=cluster-  
uuid --cluster_vip=cluster-vip --cluster_fqdn=prism-central-fqdn --  
service_vm_ip_list=comma-seperated-pcvm-ips
```

Replace `cluster_uuid` with UUID of the Prism Central VM cluster, `cluster_vip` with the virtual IP of the Prism Central VM cluster, `service_vm_ip_list` with the IP addresses of the Prism Central VMs.

- » For a single-VM Prism Central without a virtual IP address, run the following commands:

Enable establishing the trust communication without a virtual IP address for Prism Central.

```
nutanix@cvm:~$ sed -i 's/^s*validate_input/ #validate_input/' ~/bin/  
update_zeus_trustsetup.py
```

Establish the trust communication.

```
nutanix@cvm:~$ python ~/bin/update_zeus_trustsetup.py --cluster_uuid=pcvm-  
cluster-uuid --service_vm_ip_list=pcvm-ip
```

3. Log on to a Prism Central VM using SSH with the `nutanix` access.

- a. Verify that Microservices Infrastructure is healthy.

```
nutanix@pcvm$ mspctl cluster health prism-central
```

- b. Check if the Prism Central web console is available on the browser. If the Prism Central web console does not load, restart the `ikat` service.

- » On a single-VM Prism Central instance, run the following command:

```
nutanix@pcvm$ genesis stop ikat_proxy ikat_control_plane && cluster start
```

- » On a scale-out Prism Central instance, run the following command:

```
nutanix@pcvm$ allssh "genesis stop ikat_proxy ikat_control_plane && cluster  
start"
```

- c. If the `ikat` service does not restart, restart `genesis`.

- » On a single-VM Prism Central instance, run the following command:

```
nutanix@pcvm$ genesis restart
```

- » On a scale-out Prism Central instance, run the following command:

```
nutanix@pcvm$ allssh genesis restart
```

Logging Out of Prism Central

Procedure

To log off from Prism Central, select **Sign Out** from the user menu that appears on the far right side of the Prism Central landing page. You are logged out immediately after selecting the option (no prompt or message).

Power Usage

This topic describes the Power Usage feature.

The Power Usage feature provides valuable insights into the power consumption of your IT infrastructure. It enables you to understand and manage energy consumption, which can lead to potential cost savings and supports sustainability efforts. The Power Usage feature displays the accumulated power consumption data as historical line charts over the selected period of time at the cluster level.

By default, the Power Usage widget on the Prism Central Dashboard does not display the power consumption information of the clusters. You must deploy the Power Monitor application from the Nutanix Marketplace and configure the BMC out-of-band credentials for every node of a cluster. After you complete these steps, you can access and analyze detailed power consumption data of your clusters.




Figure 191: Power Usage Widget

The Power Usage feature has the following capabilities:

- Provides real-time and historical line charts of the power usage data for each cluster on the Prism Central (PC) Dashboard for the last hour, last 24 hours, and last week.
- Monitors the power usage alongside key system metrics such as CPU usage, memory usage, disk I/O and more.
- Allows power usage data integration into your management and monitoring tools using v4 APIs.

Power Usage Requirements

It's important to understand the requirements of the Power Usage feature:

- You must be on AOS 7.0 and pc.2024.3.1 version or later.
- You must establish Prism Central BMC OOB network connectivity.
- Only NX, Dell, HPE, Lenovo, and Fujitsu hardware models are supported.

Power Usage Considerations

It's important to understand the software and hardware considerations of the Power Usage feature:

- Only Large and X-Large Prism Central instances are supported.

- Hardware platforms are supported based on the following CPU generations:
 - Intel Cascade Lake
 - Intel Ice Lake
 - Intel Sapphire Rapids
 - AMD Milan
 - AMD Genoa
 - AMD Rome

Activating Power Monitor Workflow

This topic describes the workflow on how to activate the power monitor.

About this task

The following procedure summarizes the Power Monitor activation workflow:

Procedure

1. Log in to Prism Central as an administrator.
2. From the [Application Switcher Function](#), select the **Admin Center** application and from the navigation bar, select **Marketplace**.
3. From the **Nutanix Marketplace** page, deploy the **Power Monitor** application.
For more information, see [Deploying Power Monitor Application](#) on page 634.
4. Locate the **Power Usage** widget on the Prism Central Dashboard and configure the baseboard management controller (BMC) out-of-band credentials.
For more information, see [Configuring Out-of-Band Management Credentials](#) on page 635.
The **Power Usage** widget starts displaying the power usage metrics graph for each cluster. For more information, see [Viewing Power Usage Metrics at a Cluster Level](#) on page 636

Note: The system takes approximately 60 seconds to start displaying the power usage metrics graph.

Deploying Power Monitor Application

Deploy a power monitor application from the Nutanix marketplace.

Before you begin

- Ensure that the Nutanix Marketplace is enabled. For more information, see [Enabling Marketplace](#) in the *Prism Central Admin Center Guide*.
- Review the [Power Usage Requirements](#) and [Power Usage Considerations](#).

About this task

To deploy the power monitor application from the Nutanix Marketplace, follow these steps:

Procedure

1. Log in to Prism Central as an administrator.

2. From the [Application Switcher Function](#), select the **Admin Center** application and from the navigation bar, select **Marketplace**.
3. Under the **Nutanix Apps** section, locate the **Power Monitor** application widget, and click **Get**.
4. On the **Power Monitor** application details page, click **Deploy**.
5. On the **Deploy Power Monitor** confirmation window, click **Deploy**.

The **Power Monitor** application appears on the **My Apps** page.

What to do next

After deploying the power monitor application, you must configure the BMC out-of-band management credentials. For more information, see [Configuring Out-of-Band Management Credentials](#).

Configuring Out-of-Band Management Credentials

Configure the baseboard management controller (BMC) out-of-band (OOB) management credentials for each node of a cluster registered to Prism Central.

Before you begin

Ensure that BMC OOB credentials have Redfish read-only access.

About this task

To configure the BMC OOB management credentials, follow these steps:

Procedure

1. Log in to Prism Central as an administrator.
2. From the [Application Switcher Function](#), select the **Infrastructure** application.
The system displays the **Main Dashboard** page.
3. Locate the **Power Usage** widget and click the **Set up OOB Credentials** link.
The system displays the **Manage Out of Band Management Credentials** page.
4. On the **Manage Out of Band Management Credentials** page, enter the following information:
 - **OOB General Username:** Enter the BMC OOB user name.
 - **OOB Password Selection:**
 - If the **OOB password** is same for every node in the cluster, enable the **Same Password** toggle.
 - If the **OOB password** is different for each node, disable the **Same Password** toggle. This step allows you to set individual passwords for each node.
 - **OOB Password:**
 - If you enabled the **Same Password** toggle, enter the BMC OOB password.
 - If you disabled the **Same Password** toggle, click the down arrow icon next to the **OOB IP Address** and enter the specific **OOB password** for each node.
5. Click **Save & Validate**.
The system activates the **Power Usage** widget on the **Main Dashboard** page and provides historical line charts of power usage data for each cluster.

Viewing Power Usage Metrics at a Cluster Level

View the power usage metrics from the Prism Central Main Dashboard.

About this task

To view the power usage metrics at the cluster level, follow these steps:

Procedure

1. Log in to Prism Central as an administrator.
2. On the Prism Central **Main Dashboard**, under the **Power Usage** widget, click the *cluster_name* to view the power usage mini graph.




Figure 192: Power Usage Widget

3. (Optional) Place the cursor anywhere on the horizontal axis to display the power usage value at that time.
4. (Optional) Select the duration (time interval) from the dropdown list on the right (last one hour, last 24 hours, and last week).
5. (Optional) To see a complete view of the power usage metrics, click **Full View**.
The system opens the cluster's **Metrics Tab**.

Viewing Power Usage Metrics at a Host Level

View the power usage metrics from the **Hosts** details page.

About this task

To view the power usage metrics at the host level, follow these steps:

Procedure

1. Log in to Prism Central as an administrator.

2. From the [Application Switcher Function](#), select the **Infrastructure** application and from the navigation bar, select **Hardware > Hosts**.
By default, the system displays the **List** tab with all the hosts.
3. Click the host name.
4. Click the **Metrics** tab and locate the **Power Usage** metrics chart.
5. (Optional) Place the cursor anywhere on the horizontal axis to display the power usage value at that time.
6. (Optional) Select the duration (time interval) from the dropdown list on the right (last one hour, last 24 hours, and last week).
7. (Optional) Click the **Alert Settings** to configure an alert for the **Power Usage** metric.
For more information, see [Creating Custom Alert Policies](#) in *Prism Central Alerts and Events Reference Guide*.

Updating Out-of-Band Management Credentials

Update the baseboard management controller (BMC) out-of-band (OOB) management credentials.

About this task

To update the BMC OOB credentials, follow these steps:

Procedure

1. Log in to Prism Central as an administrator.
2. From the [Application Switcher Function](#), select the **Infrastructure** application and from the navigation bar, select the **Prism Central Settings**.
3. From the **Settings** navigation bar, select the **OOB Management Credentials**.
4. Select the target cluster checkbox.
5. Click **Manage Credentials**.
6. Update the fields in the **Manage Out of Band Management Credentials** window as needed and click **Save & Validate**.

For more information on the fields available in the **Manage Out of Band Management Credentials** window, see [Configuring Out-of-Band Management Credentials](#) on page 635.

Removing Out-of-Band Management Credentials

Remove the baseboard management controller (BMC) out-of-band (OOB) management credentials.

About this task

To remove the BMC OOB credentials, follow these steps:

Procedure

1. Log in to Prism Central as an administrator.
2. From the [Application Switcher Function](#), select the **Infrastructure** application and from the navigation bar, select the **Prism Central Settings**.
3. From the **Settings** navigation bar, select the **OOB Management Credentials**.
4. Select the target cluster checkbox.

5. Click **Remove Credentials**.
6. On the **Delete Credentials** confirmation window, click **Delete**.

CUSTOMER SUPPORT SERVICES

Nutanix provides customer support services in several ways.

- Nutanix customer support can monitor your clusters and provide assistance when problems occur through the Pulse mechanism. For more information, see [Pulse Health Monitoring](#) information in *Prism Central Admin Center Guide*.
- Nutanix customer support maintains a portal that you can access to request assistance, download various product updates, and view documentation. For more information, see [Accessing the Nutanix Support Portal \(Prism Central\)](#) on page 644.
- If you need help, you can create a support ticket directly from Prism Central. For more information, see [Creating a Support Case](#) on page 640.
- Nutanix technical support can remotely assist with problems in a Prism Central (PC) VM by logging into the faulty PCVM through an SSH connection. For more information, see [Configuring Remote Connection Using CLI](#) on page 639.

Configuring Remote Connection Using CLI

You can enable a remote support connection tunnel for a maximum of 72 hours using CLI.

About this task

If you face an issue in your Prism Central (PC) VM and the Nutanix Support team needs access to it to troubleshoot the issue, you can open a support tunnel for the Nutanix Support team to give remote access to your VM. The Nutanix Support team ensures that the connection to your PCVM through the support tunnel is secure and compliant by consolidating connectivity, authentication, authorization, audit, and recorded sessions.

Use the following procedure to enable the remote support connection tunnel.

Procedure

1. SSH into CVM or PCVM .
2. Start ncli.

```
nutanix@cvm$ ncli  
<ncli>
```

3. Run the cluster start-remote-support command with the duration parameter set as required.
The duration parameter must be set in minutes even if you need a duration of hours.

```
ncli> cluster start-remote-support duration=minutes
```

Add duration that you want to enable the remote support connection tunnel for, in minutes. For example, if you want to keep the connection open for 24 hours, enter **duration=1440**.

Note: You can keep the remote support connection tunnel between 0-72 hours.

Creating a Support Case

About this task

Nutanix customer support maintains a portal where you can get assistance by opening a support case and viewing the status of your open cases (see [Accessing the Nutanix Support Portal \(Prism Central\)](#) on page 644). However, you can also create a support case directly from Prism Central. To create a support case or view information about your open cases, do the following:

Note: A support portal connection is required before you can create a case (see [Prism Licensing](#) on page 54). In addition, at least one cluster must be registered with Prism Central. Also, this feature is available only to customers who get support directly from Nutanix (such as NX and SX model customers), not customers who get initial support from third parties.

Procedure

1. Click the [Help icon](#) in the main menu and select **Create Support Case** from the dropdown menu.

The **Create new support case** page appears. One of the following appears on this page:

- If there are no open cases currently, a blank page with fields for creating a case appears.
- If you have one or more open support cases, summary information about those cases appears (see [Viewing Case Status](#) on page 643). To create a case from this page, click the **Create new support case** button (upper right). The blank page with fields for creating a case will appear.

Create New Support Case

Support Case Details

SUBJECT*

Fan A has stopped or its speed is low

ISSUE CATEGORY

Technical Problem

PRIORITY

P3 - Normal

CLUSTER

BLOCK SERIAL NUMBER ⓘ

PROBLEM DESCRIPTION*

We received the alert below from the CVM of the host:
Fan A has stopped or its speed is low on Controller VM x.x.x.x.

Attachments

Attach log bundle (for last 4 hours)

☒ For PC

☒ Anonymize log bundle

☒ For PE

☒ Anonymize log bundle

Attach NCC summary output

☒ For PC

☒ For PE

Your Contact Details

Tom D

xxx-xxx-xxxx

tomd@example.com

ADDITIONAL USER NOTIFICATIONS

abc@example.com, cba@example.com

Figure 193: Create New Support Case Page

NUTANIX

Prism | Customer Support Services | 641

2. Do the following in the indicated fields:

a. **Subject:** Enter a title that briefly describes the issue.

b. **Issue Category:** Select the type of issue from the dropdown menu.

The types are **Technical Problem, Question, Comment/Feedback, Scheduled Upgrade/Maintenance, RFE** (request for enhancement), **Licensing, Non Technical issue**, and **Technical Problem - Foundation**.

c. **Priority:** Select the priority for this issue from the dropdown menu.

There are four priority levels based on the severity of the problem plus a request for enhancement option. Select the appropriate priority based on the following descriptions.

- **P1 - Emergency.** System is not available and productivity has been halted. Product is unusable in its current state.
- **P2 - Critical.** System is available but experiencing issues which have a direct impact on productivity. Major inconvenience.
- **P3 - Normal.** System is having an occasional issue that has been identified as needing to be resolved, but the issue has not greatly affected productivity. Minor inconvenience.
- **P4 - Low.** Questions about documentation, processes, or procedures. General requests about information.
- **RFE - Request for Enhancement.** Feature requests for the product which would improve the experience or functionality for the customer.

d. **Cluster:** Select the target cluster from the dropdown menu.

e. **Block Serial Number:** Select the serial number of the asset (node) in question from the dropdown menu.

If you need more information, see the "Installed Base" page on the Nutanix support portal for a list of your assets and the corresponding serial numbers.

f. **Tell Us More:** Enter a description of your issue in the text box.

Include any relevant details that might help Nutanix customer support analyze and resolve your issue.

g. **Attach log bundle (for last 4 hours):** Check this box to attach log files from the last four hours to the case. If you want the attached logs to be anonymized (personally identifiable information removed), check the **Anonymize log bundle** box.

- **For PC:** Check this box to attach Prism Central log files from the last four hours to the case. If you want the attached logs to be anonymized (personally identifiable information removed), check the **Anonymize log bundle** box.
- **For PE:** Check this box to attach Prism Element log files from the last four hours to the case. If you want the attached logs to be anonymized (personally identifiable information removed), check the **Anonymize log bundle** box.

Note: If the target cluster is not connected to the Internet, such as in a dark site, a log bundle or NCC summary report (following step) cannot be attached to the case through this method.

h. **Attach NCC summary output**

- **For PC:** When this box is checked, NCC is run, and the summary output of Prism Central is attached. This ensures that the NCC (health checks) results reflect the current state.
- **For PE:** When this box is checked, NCC is run, and the summary output of Prism Element is attached. This ensures that the NCC (health checks) results reflect the current state.

- i. **Your Contact Details:** Enter your name, phone number, and e-mail address in the indicated fields. Case updates are sent to your e-mail address. You can also have the case updates sent to (up to two) additional users by entering their e-mail addresses in the **Additional User Notification** field.

3. When the entered information is completed, click **Submit**.

This submits the case to Nutanix customer support for review and response. The speed of response reflects the priority of the issue, where P1 and P2 issues have a higher priority and more rapid response than less critical issues, but Nutanix customer support strives to handle all cases as quickly as possible.

A Prism Central task is started to create the case. If you selected to attach a log bundle and/or NCC summary output, a task is started for each selected item. The tasks run asynchronously, and you can check the progress (case creation, log collection, and health check tasks) through the tasks dashboard (see [Tasks View](#) on page 496). If a task encounters a problem, you can hover the cursor over the "Failed" status to see a brief explanation. Automatically generated comments (preceded by "Prism Central Automated Note:") are added to the case when the log/health checks collection starts and when the upload completes. When the tasks complete, the output (log bundle and NCC summary report) appear as attachments in the case summary (see [Viewing Case Status](#) on page 643).

Operation Message	Entity	Cluster	Percent	Status	Create Time	Duration
Log collector	Cluster	-	100%	Failed	05/22/19, 11:16:43 AM	29 seconds
Health check	Cluster	-	100%	Succeeded	05/22/19, 11:16:43 AM	44 seconds
Case #00562960 has been updated	Prism Central	-	100%	Succeeded	05/22/19, 11:16:42 AM	49 seconds
Log collector	Cluster	auto_cluster_prod...	100%	Failed	05/22/19, 11:16:42 AM	29 seconds
Health check	Cluster	auto_cluster_prod...	100%	Succeeded	05/22/19, 11:16:42 AM	1 minute 42 ...
Case #00562960 has been updated	Cluster	auto_cluster_prod...	100%	Succeeded	05/22/19, 11:16:41 AM	1 minute 49 ...
Create support case upload	Support case upload	-	100%	Succeeded	05/22/19, 11:16:35 AM	57 seconds

Figure 194: Tasks Dashboard for Prism Central (create case tasks)

19 Total Tasks						1 - 19 of 19 < >	
Operation Message	Entity		Percent	Status	Create Time	Duration	
Log collector	Cluster		100%	Failed	05/22/19, 11:16:42 AM	29 seconds	
Health check	Cluster		100%	Succeeded	05/22/19, 11:16:42 AM	1 minute 42 secon...	
Case #00562960 has been updated	Cluster		100%	Succeeded	05/22/19, 11:16:41 AM	1 minute 49 secon...	
Create support case upload	Support case uplo...		100%	Succeeded	05/22/19, 11:16:35 AM	1 minute 56 secon...	

Figure 195: Tasks Dashboard for Prism Element (create case tasks)

Viewing Case Status

You can view information about open support cases directly from Prism Central by clicking the [Help icon](#) in the main menu and selecting **Create Support Case** from the dropdown menu. This opens the **Create new support case** page, which displays summary information about open cases (if any). When you have one or more open cases, this page displays the following sections:

- **Open Cases** column (on the left). There is an entry (line) for each open case that includes the case name and number; click the desired case entry to displays summary information about that case in the main section of the page.

- Main section (in the middle). This section displays the creation date, description, attachments (if any), and other information related to the selected case. This is an abridged version of the full case details available from the Nutanix support portal (see below).
- **Key Insights** column (on the right). This section includes the following fields:
 - **Status:** Displays the current status of the case.
 - **Created By:** Displays the name of the user who created the case.
 - **Type:** Displays the type (category) of issue the user specified when creating the case.
 - **Case number:** Displays the number assigned to this case.
 - **Serial Number:** Displays the serial number of the asset (node) in question.
 - **Access Portal:** Includes a link to the Nutanix support portal. Clicking the **View Details in Portal** link opens the details page for that case in the support portal in a new tab or window. The details page in the support portal provides the full details and history of that case should you desire more information, and you can submit additional material to the case such as screen shots or other relevant information.




Figure 196: Create New Support Case Page (open cases)

Accessing the Nutanix Support Portal (Prism Central)

About this task

Nutanix provides a variety of support services and materials through its support portal.

Procedure

1. To access the Nutanix support portal from Prism Central, select **Support Portal** from the **User icon** dropdown menu of the main menu.

The login screen for the Nutanix support portal appears in a new tab or window.

2. Enter your support account user name and password.
The Nutanix support portal home page appears.
3. Select the desired service from the screen options.

You can select an option from one of the main menu dropdown menus or search for a topic at the top of the screen, click one of the icons (Documentation, Open Case, View Cases, Downloads) in the middle, or view one of the selections at the bottom such as an announcement or KB article. The following table lists the menu options.

Note: Some options have restricted access and are not available to all users.

Table 174: Main Menu Options

Category	Option	Description
Documentation	Software Documentation	Displays a page from which you can view the Nutanix software manuals.
	Hardware Replacement Documentation	Displays a page from which you can view the Nutanix hardware replacement manuals.
	Knowledge Base	Displays a page from which you can view the knowledge base (KB) articles.
	Solutions Documentation	Displays a page from which you can view documents that describe how to implement the Nutanix platform to solve a variety of business applications.
	EOL Information	Displays a page from which you can view the end of life policy and bulletins.
	Field Advisories	Displays a page from which you can view field advisories.
	Training	Provides a link to the separate Nutanix training portal.
	Security Advisories	Displays a page from which you can view security advisories.
	AOS Upgrade Paths	Displays a page where you can see the supported AOS release upgrade paths.
	Compatibility Matrix	Displays a page from which you can view a compatibility matrix broken down (filtered) by hardware model, AOS version, hypervisor type and version, and feature version (NCC, Foundation, BMC/BIOS).
	Webinar Recordings	Displays a page with links to a selection of Nutanix training webinars.
Support	Open Case	Displays a form to create a support case.
	View Cases	Displays a page from which you can view your current support cases.

Category	Option	Description
Downloads	.NEXT Forums	Provides a link to the (separate) Nutanix Next Community forum.
	Terms & Conditions	Displays a page from which you can view various warranty and terms and conditions documents.
	AOS (NOS)	Displays a page from which you can download AOS releases.
	Acropolis File Services (AFS)	Displays a page from which you can download the Acropolis File Services.
	Acropolis Container Services (ACS)	Displays a page from which you can download the Acropolis Container Services.
	Hypervisor Details	Displays a page from which you can download Acropolis hypervisor versions. You can also download supporting files used when manually upgrading a hypervisor version (AHV, ESXi, or Hyper-V).
	Prism Central	Displays a page from which you can download the Prism Central installation bundle. There are separate bundles for installing on AHV, ESXi, or Hyper-V.
	Phoenix	Displays a page from which you can download Phoenix ISO files.
	Foundation	Displays a page from which you can download Foundation releases.
My Products	Installed Base	Displays a table of your installed Nutanix appliances, including the model type and serial number, location, and support coverage.
	Licenses	Displays a table of your product licenses along with options to add or upgrade licenses for your clusters.




Figure 197: Nutanix Support Portal

Accessing the REST API Explorer

About this task

Nutanix provides a utility with the web console to help you get started with the REST API. The Explorer displays the parameters and format for the API calls that can be included in scripts. Sample API calls can be made to show the type of output you should expect to receive. Only an admin user can access the REST API Explorer.

The v3 API can be viewed in the REST API Explorer.

Procedure

1. Log on to Prism Central, click the [User icon](#) in the upper-right corner of the web console, and click **REST API Explorer**.

The REST API Explorer displays a list of the objects that can be managed by the API.

2. Find the object you want to explore and click the name of the object to expand and show the detailed view of the operations that can be run on this object.

For example, click **alerts**. All the API calls for **alerts** are displayed.

- a. Click **GET** to show the details for this API call.
- b. Click **Try it out!** to test the API call when used with your cluster.
- c. Similarly, you can explore other API calls.

Troubleshooting

For troubleshooting purposes, you can collect Prism Element logs through Prism Central.

Prism Central web console provides the following log operations for troubleshooting:

- You can generate logs for Prism Element clusters (or specific hosts on those clusters) that are connected to the Prism Central instance. For more information, see [Collecting Logs](#) on page 647.
- You can also view the log collection history. For more information, see [Viewing Log Collection History](#) on page 649.

Collecting Logs

You can generate and download the log bundles through Prism Central, for clusters managed through Prism Central.

Before you begin

For log collection through Prism Central, the minimum supported versions are AOS version 7.3 and Prism Central version pc.7.3.

About this task

To generate and download the log bundles through the Prism Central web console, do the following:

Procedure

1. Log in to Prism Central web console.
2. Click the [Help icon](#) in the main menu and from the dropdown menu, select **Troubleshooting > Collect Logs**.

3. In the **Collect Logs** dialog box, in **Entity Selection** section, select the type of entity.
 - » **Prism Central**: Select this check box to generate logs for Prism Central.
 - » **Hosts under a Cluster**: Select this check box to generate logs for hosts under a cluster.
4. If you selected **Hosts under a Cluster**, do the following:
 - a. Select the cluster in **Cluster Name**.
 - b. Click **Add Hosts**.
 - c. Select the hosts for which you want to generate the logs.
 - d. Click **Done**.
5. Click **Next**.
6. In the **Log Settings** section, select the type of logs:
 - » **All logs**: Select this check box to generate all the logs.
 - » **Specific logs (by tags)**: Select this check box to generate specific logs identified by predefined tags. Nutanix recommends using this option to avoid unnecessary load.
7. If you selected **Specific logs (by tags)**, do the following:
 - a. Click **Select Tags**.
 - b. In the **Select Tags** dialog box, select the tags associated with the logs you want to generate.
 - c. Click **Done**.
8. Click **Next**.

9. In the **Output Preferences** section, do the following:

a. To specify the time range for logs, select **Start Date**, **Start Time**, and **Duration**.

Due to space constraints, logs older than 90 days might be overwritten by newer log bundles. Maximum duration allowed is 96 hours or four days.

b. To save the collected logs, select **Destination**.

» **Download Locally:** Logs are collected at the respective hosts. The download option is available only for this destination. To access and download the logs, see [Viewing Log Collection History](#) on page 649.

» **Nutanix Support FTP:** Enter the support case number. Ensure to enter a valid case number to avoid the logs getting lost.

» **Nutanix Support SFTP:** Enter the support case number. Ensure to enter a valid case number to avoid the logs getting lost.

» **Custom Server:** Define a Custom SFTP or FTP server.

c. If you selected **Custom Server**, enter the following information:

- Name or IP address of the server
- Protocol for connecting to the server
- Port number for the server
- User ID for the server
- Password for the server
- Destination path for the log bundle. This destination must be an existing directory on the server.

d. To mask the sensitive information such as IP address, select **Mask sensitive information**.

10. Click **Collect Logs**.

The log-collection (collect log) task is initiated at the selected destination. You can track the progress of this task from **Recent Tasks** list.

After the log-collection task is completed, the web console displays a success notification with a link to download the log, if applicable. You can also download the logs from the log collection history. For more information, see [Viewing Log Collection History](#) on page 649.

Viewing Log Collection History

You can view the log collection history for logs collected through Prism Central.

About this task

Once you initiate the log collection task, the log bundle is retained for up to seven days or until the number of bundles exceeds five—whichever occurs first. On expiry of the retention period or when the bundle count limit is exceeded, the oldest bundles are automatically deleted to maintain compliance.

To view and download (if applicable) the logs collected during the previous seven days:

Procedure

1. Log in to Prism Central web console.
2. Click the [Help icon](#) in the main menu and from the dropdown menu, select **Troubleshooting > View Log Collection History**.

3. In the **Log Collection History** dialogbox, you can view all the log bundles collected during the previous seven days.

In the list, you can view the cluster name and date and time of log collection.

4. Click the download icon to download the log bundle, if applicable.

You can download only the logs that are collected locally on the nodes.

.

HELP RESOURCES

There are several information sources that you can access at any time when you need help:

- Prism Central help documentation (see [Accessing Online Help \(Prism Central\)](#) on page 651).
- Nutanix customer support portal (see [Accessing the Nutanix Support Portal \(Prism Central\)](#) on page 644).
- Nutanix community forum (see [Accessing the Nutanix Next Community](#) on page 652).
- Glossary of terms (see [Nutanix Glossary](#)).

Accessing Online Help (Prism Central)

About this task

Prism Central includes online help documentation that you can access at any time.

Procedure

1. To open the online help, choose one of the following from the [Help icon](#) dropdown menu of the main menu:

- » Select **Help with this page** to display help documentation that describes the current screen.

Note: In a task window click the [Help icon](#) in the upper right to display the help documentation for that window.

- » Select **Online Documentation** to display the **Help Organization** page.

A context-sensitive help page or the **Help Organization** page appears in a new tab or window. (These pages are located on the Nutanix support portal.) The **Help Organization** page provides descriptions of the major help topics with links to the entry page for each major topic. The display includes a breadcrumb at the top to navigate through the help pages.




Figure 198: Help Organization Page

- To select a topic from the table of contents, click the [Navigation](#) icon.

A table of contents pane appears on the left. Click a topic in the table of contents to display that topic.




Figure 199: Table of Contents Pane

- To display all the help contents as a single document, click **epub** or **pdf** in the upper right.

You can view the *Prism Central Infrastructure Guide* in either ePUB or PDF format by selecting the appropriate option. If your browser does not support the selected format, you can download the PDF or ePUB file.

- To search for a topic, click the magnifying glass icon in the main menu bar and enter a search string in the field.

This searches not only the help contents, but also all the documentation, knowledge base articles, and solution briefs. Matching results appear below the search field. Click a topic from the search results to display that topic.




Figure 200: Search Results Example

Accessing the Nutanix Next Community

Nutanix maintains a community forum for customers and partners to facilitate a peer-to-peer exchange of ideas, tips, and information about Nutanix technologies and the rapidly changing landscape of data center IT.

Procedure

- To access the Nutanix next community forum from Prism Central, select **Nutanix Next Community** from the [Help icon](#) dropdown menu of the main menu (see [Prism Central Landing Page](#) on page 35).

The Nutanix Next Community main page appears in a new tab or window. From this page you can search existing posts, ask questions, and provide comments.




Figure 201: Next Community Screen

Glossary

For terms used in this guide, see [Nutanix Glossary](#).

COPYRIGHT

Copyright 2025 Nutanix, Inc.

Nutanix, Inc.

1740 Technology Drive, Suite 150

San Jose, CA 95110

All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. Nutanix and the Nutanix logo are registered trademarks of Nutanix, Inc. in the United States and/or other jurisdictions. All other brand and product names mentioned herein are for identification purposes only and may be trademarks of their respective holders.

Prism Central Admin Center Guide

Prism pc.7.3
October 7, 2025

Contents

About this Publication.....	5
Admin Center Overview.....	6
Application Switcher (Prism Central).....	8
Prism Licensing.....	8
Marketplace.....	9
Enabling Marketplace.....	9
Viewing Marketplace App Details.....	11
Searching or Filtering Marketplace Items.....	11
Nutanix Apps Deployment.....	11
Deploying Self-Service.....	12
Deploying Files.....	12
Deploying Foundation Central.....	13
Deploying Kubernetes Management.....	13
Deploying Objects.....	14
Deploying Database Service.....	15
Deploying Move.....	16
Deploying Power Monitor Application.....	17
Deploying Nutanix Central.....	17
Preferred Partner App Deployment.....	18
OpenShift Overview.....	18
Cisco Intersight Device Connector Overview.....	23
Marketplace Application Update.....	27
Updating Marketplace Applications Using Life Cycle Manager.....	28
My Apps.....	29
Nutanix Apps Management.....	29
Registering Database Service or Move to My Apps.....	30
Preferred Partners Apps Management.....	30
Deleting Cisco Intersight Device Connector.....	30
Upgrading Cisco Intersight Device Connector from My Apps.....	31
Project Management.....	33
Projects Summary View.....	34
Project Details View.....	35
Project Dashboard Tab.....	36
Project Environments Tab.....	39
Project Infrastructure Tab.....	40
Project Policies Tab.....	40
Project Tunnels Tab.....	40
Project Users & Groups Tab.....	40
Project Workloads Tab.....	41
Project Usage Tab.....	42
Creating a Project.....	42

Adding Users to a Project.....	44
Infrastructure in Projects.....	46
Adding Infrastructure in a Self-Service-Enabled Prism Central.....	46
Adding Infrastructure in a Self-Service-Disabled Prism Central.....	48
Modifying a Project.....	49
Deleting a Project.....	50
Environments in Projects.....	50
Configuring Environments in a Project.....	50
Configuring Check Log-in for your Environment.....	66
Support for Multiple Credential.....	67
Adding Credentials to the Environment.....	67
Quota Policy Overview.....	68
Managing Quota Limits for Projects.....	69
Creating a Snapshot Policy.....	71
Identity and Access Management.....	74
Security and User Management (Prism Central).....	74
User Management.....	74
Roles Management.....	75
LCM View.....	77
Category Management.....	78
Creating a Category.....	78
Viewing Category Details.....	79
Assigning a Category to an Entity.....	80
Updating a Category.....	80
Deleting a Category.....	81
Admin Center Settings Options.....	82
Application Log Archive.....	83
Downloading Application Log Archive.....	83
Policy Engine Overview.....	84
Enabling the Policy Engine.....	84
Enabling the Policy Engine at a Dark Site.....	86
Viewing the Policy Engine VM Details.....	87
Disabling the Policy Checks.....	87
Best Practices to Shut Down the Policy Engine.....	88
Backing Up and Restoring the Policy Engine Database.....	88
Pulse Health Monitoring.....	89
Pulse Configuration.....	91
Mask Entity Names and IP Addresses.....	94
Pulse Health Monitoring Data Collection.....	94
Network Configuration for Pulse Health Monitoring.....	100
Prism Central Proxy for Pulse Data.....	100
Prism Central Configuration When a Cluster Uses Proxy Servers.....	101
Who Needs to Use the Allowlist Method.....	101
Configuring an HTTP Proxy.....	103
Allowing Prism Central and Its Managed Clusters (nCLI).....	104
Configuring Name Servers for Prism Central.....	105
Configuring NTP Servers for Prism Central.....	105
Configuring SNMP on Prism Central.....	106

Configuring Alert Emails in Prism Central.....	109
Configuring an SMTP Server for Prism Central.....	111
Configuring Syslog Monitoring.....	112
Syslog Modules.....	113
Setting the Prism Central Default Landing Page.....	116
Changing the Prism Central Language Settings.....	117
Internationalization (i18n) (Prism Central).....	118
Localization (L10n) (Prism Central).....	118
Nutanix Intelligent Virtual Agent Overview.....	118
Turning on NIVA Tips Setting.....	118
Troubleshooting Recommendations for NIVA.....	119
Configuring Prism Central UI Settings.....	119
Configuring the Prism Central Welcome Banner.....	122
Containers and Services Configuration.....	123
Resource Limits for Docker Containers.....	123
Verifying the Container vCPU Limits.....	123
Glossary.....	125
Copyright.....	126

ABOUT THIS PUBLICATION

This document provides information about how to configure and manage administrative tasks using Admin Center in Prism Central.

- For information on all the documents applicable for Prism Central, see [Prism Central Documentation Portfolio](#) in the *Prism Central Infrastructure Guide*.
- To access other Nutanix documents, see [Nutanix support portal](#).

ADMIN CENTER OVERVIEW

Admin Center facilitates you to configure and manage common administrative tasks that are applicable to the platform and various Nutanix apps.




Figure 1: Admin Center

Admin Center Entities

You can use Admin Center to manage the following configurations:

- **Define Projects and Various App Management Features Within Projects**

Projects provide logical groupings of user roles for managing resource utilization within your organization. The project construct helps you define users and groups for marketplace app deployment, networks to use for app deployment, VM specifications and deployment options, quota policies, snapshot policies, and so on.

The project configuration options differ depending on whether you have enabled Self-Service in your Prism Central or not. In a Prism Central where Self-Service is disabled, you can configure projects with basic setups such as add users, add local Nutanix account, and manage usage and workloads. Application management features such as environment configuration, policy configuration, and cloud infrastructure setups are possible only when you deploy Self-Service in your Prism Central. For detailed information on projects, see [Project Management](#) on page 33.

- **Deploy Nutanix Apps and Preferred Partner Apps**

Nutanix Marketplace facilitates provisioning of Nutanix Apps such as Self-Service, Files, Objects, Foundation Central, Move, Database Service, and Kubernetes Management and Preferred Partner Apps, such as OpenShift and Cisco Intersight Device Connector.

Important: Starting with Prism Central version pc.2024.3.1, you can access and deploy custom apps, hybrid cloud apps, and runbooks (previously listed as **Other Apps** on the Nutanix Marketplace and My Apps page) from NCM Marketplace in Self-Service. For more information, see the [Marketplace in Self-Service](#) section in the *Self-Service Administration and Operations Guide*.

For detailed information on Marketplace, see [Marketplace](#) on page 9.

- **Manage Deployed Apps from a Common Workspace**

The My Apps entity of Admin Center provides a single workspace to manage all the apps that you deployed from the Marketplace.

Important: Starting with Prism Central version pc.2024.3.1, you can manage custom apps and hybrid cloud apps (previously listed as **Other Apps** on the Nutanix Marketplace and My Apps page) from the Applications page in Self-Service. For more information, see the [Applications in Self-Service](#) section in the *Self-Service Administration and Operations Guide*.

For detailed information on My Apps, see [My Apps](#) on page 29.

- **Manage Users, Roles, and Authentication Mechanisms**

Identity and Access Management (IAM) facilitates you to view details of local and directory-specified users, view system-defined roles and custom roles, and create custom roles. The [Security Management Using Prism Central](#) section in the *Security Guide* covers different aspects of users, roles, and authentication management, including SAML/IDP supported providers.

- **Perform inventory and Software Updates**

The life cycle manager (LCM) facilitates you to view information about the current inventory and lets you update the versions as needed. For detailed information on LCM, see the [Life Cycle Manager Guide](#).

- **Manage Licensing**

You can apply licenses to enable a variety of features. For detailed information, see the [Nutanix License Manager Guide](#).

- **Manage Global Prism Central Settings Applicable Across Prism Central**

The setting options in Admin Center are common across multiple Nutanix apps that you deploy and access using Prism Central. You can enable policy engine and configure various app management, networking, alerts, and appearance related settings. For more information on Admin Center settings, see [Admin Center Settings Options](#) on page 82.

Admin Center Navigation Bar

You get the Admin Center navigation bar when you select Admin Center in the Application Switcher. The Admin Center navigation bar has the following entities.

- My Apps
- Marketplace
- Projects
- IAM
- LCM
- Licensing
- Settings

You can click an entity in the navigation bar to view the summary or manage configurations associated with that entity. For example, you can click **Projects** to create a project, add users, add infrastructure, and configure various app management components.

You can also lock the navigation bar so that it remains open when you switch between different entities within the Admin Center.

Application Switcher (Prism Central)

Prism Central provides a centralized environment to procure multiple Nutanix Apps such as Self-Service, Files, Database Service, Objects, Kubernetes Management, and Foundation Central and manage these apps from a single workspace. The Application Switcher facilitates you to seamlessly switch between the Nutanix Apps that you deployed and that you have access to in Prism Central.

For more information, see [Application Switcher Function](#) in the *Prism Central Infrastructure Guide*.

Access to Apps in the Application Switcher

In a new deployment of Prism Central, the Application Switcher displays the Admin Center, Infrastructure, Apps and Marketplace, Cost Governance, and Security Central for users with the Nutanix admin role. You can view other Nutanix Apps in the Application Switcher after enabling them from the Marketplace. For more information, see [Marketplace](#) on page 9.

If you have upgraded to the current version of Prism Central, you can view all the Nutanix Apps that you enabled in the previous version along with the Admin Center, Infrastructure, Apps and Marketplace, Cost Governance, and Security Central.

For users with the non-admin role, Application Switcher displays the apps based on the access policies. For example, if you a user with system-defined consumer role, then you can view only Infrastructure, Apps and Marketplace, and any other Nutanix Apps that you have access to.

Prism Licensing

Nutanix provides licenses you can apply to enable a variety of features.

The Prism Central web console and Nutanix Support Portal provide the most current information on your licenses. For more information on licenses, see the [License Manager Guide](#).

MARKETPLACE

Nutanix Marketplace facilitates deployment and management of Nutanix Apps and Preferred Partner Apps.

Important:

- Starting with Prism Central version pc.2024.3.1, you can access and deploy custom apps, hybrid cloud apps, and runbooks (previously listed as **Other Apps** on the Nutanix Marketplace and My Apps page) from NCM Marketplace in Self-Service. For more information, see the [Marketplace in Self-Service](#) section in the *Self-Service Administration and Operations Guide*.
- Only a Prism Admin can view the Nutanix App and Preferred Partner App details or deploy the apps from the Marketplace.
- The IAM interface has been revamped to provide fine-grained RBAC so that custom roles can be created with granular operations or permissions. For information on how to create roles and add authorization policy, see the [Security Management using Prism Central \(PC\)](#) section of the *Security Guide*.

As a Prism Central customer, you do not require any additional license to use the new Marketplace capabilities and access Nutanix apps.

The first step to use Marketplace for app deployment is to enable the Marketplace in your Prism Central instance. For more information, see [Enabling Marketplace](#) on page 9.

Marketplace App Categories

The following table list the different categories of apps in Nutanix Marketplace.

Table 1: Marketplace Items

Marketplace Item	Description
Nutanix Apps	These apps are developed and certified by Nutanix. These apps include Self-Service, Files, Objects, Foundation Central, Move, Database Service, Kubernetes Management, and so on. These apps run on the Nutanix platform. You might have to apply for necessary licenses to use these apps. For more information, see the Nutanix License Manager Guide. For information on how to deploy Nutanix apps, see Nutanix Apps Deployment on page 11.
Preferred Partner Apps	These apps are developed and certified by Nutanix in collaboration with the preferred partners. These apps are pre-seeded for deployment on the Nutanix platform. For example, OpenShift. For information on how to deploy Preferred Partner apps, see Preferred Partner App Deployment on page 18.

Enabling Marketplace

Enabling Marketplace is a prerequisite for app deployment from Nutanix Marketplace or app management in My Apps. Only a Prism Admin can enable Marketplace.

Before you begin

- Marketplace enablement requires an additional 2 GB of memory for a small Prism Central deployment and 4 GB of memory for a large Prism Central deployment. Ensure that your Prism Central has required resources to enable Marketplace in the same cluster.
- A unique data service IP address should be configured in the Prism Element web console cluster that is running on Prism Central. For information on configuring data service IP address, see the [Modifying Cluster Details](#) in the *Prism Element Web Console Guide*.
- Nutanix recommends you to observe the impacts before you change the virtual IP address and Data Services IP address. For more information, see the [Virtual IP Address Impact](#) and [iSCSI Data Services IP Address Impact](#) sections in the *Prism Element Web Console Guide*. See the product-specific documents for detailed information on how to handle the changes for virtual IP address and Data Services IP address.

Note:

- Marketplace is supported only on AHV and ESXi hypervisors on a Nutanix cluster and is not supported on Hyper-V hypervisor.
- Marketplace is not supported on vSphere Essentials edition because vSphere Essentials edition does not support hot-pluggable virtual hardware.
- Marketplace is not supported on x-small Prism Central deployments.

About this task

Note: Once you enable Marketplace after an upgrade, you must perform the LCM inventory to see the correct version of Nutanix Apps on the My Apps page.

Procedure

1. Log in to your Prism Central instance as an administrator.
2. From the Application Switcher, select **Admin Center**.
3. In the navigation bar, click **Marketplace**.
4. Click **Enable Marketplace**.




Figure 2: Enable Marketplace

Viewing Marketplace App Details

You can view the Marketplace app overview, version, and the associated actions on the app details page.

Procedure

1. Log in to your Prism Central instance.
2. From the Application Switcher, select **Admin Center**.
3. In the navigation bar, click **Marketplace**.
4. Click **Get** for the application to view details.
The app details page opens.
5. On the app details page, view details such as the overview, version, and actions.

Searching or Filtering Marketplace Items

You can search or filter Marketplace to get a list of items based on your search keywords or the selected filter categories.

Procedure

1. Log in to your Prism Central instance.
2. From the Application Switcher, select **Admin Center**.
3. In the navigation bar, click **Marketplace**.
4. (Optional) To search an app, in the **Search Marketplace** field, enter the name of the app.
The Marketplace page filters the apps based on the keywords you entered in the **Search Marketplace** field.
5. (Optional) To filter an app, click **Filters** and select **Nutanix**, **Preferred_Partners**, or both.
The Marketplace page filter apps based on the categories you select.

Nutanix Apps Deployment

As a Prism Admin, you can enable the following Nutanix apps from Admin Center.

- Self-Service. For information on the Self-Service deployment, see [Deploying Self-Service](#) on page 12
- Files. For information on the Files deployment, see [Deploying Files](#) on page 12
- Foundation Central. For information on the Foundation Central deployment, see [Deploying Foundation Central](#) on page 13
- Kubernetes Management. For information on the Kubernetes Management deployment, see [Deploying Kubernetes Management](#) on page 13
- Objects. For information on the Objects deployment, see [Deploying Objects](#) on page 14
- Database Service. For information on the Database Service deployment, see [Deploying Database Service](#) on page 15.
- Move. For information on the Move deployment, see [Deploying Move](#) on page 16.
- Power Monitor. For information on the Power Monitor deployment, see [Deploying Power Monitor Application](#) on page 17.

- Nutanix Central On-Premises. For information on the Nutanix Central on-premises deployment, see [Deploying Nutanix Central](#) on page 17.

Deploying Self-Service

About this task

Self-Service is integrated into Prism Central and does not require you to deploy any additional VMs. To start using Self-Service, you have to deploy the Self-Service app from the Marketplace.

Before you begin

- Ensure that you have enabled Marketplace. For more information, see [Enabling Marketplace](#) on page 9.
- Ensure that you meet the prerequisites to enable Self-Service. For more information, see the [Prerequisites to Deploy Self-Service](#) section in the *Self-Service Administration and Operations Guide*.

Note:

If the Prism Element web console is not registered from a Prism Central and the app blueprints have a subnet, image, or VMs on the Prism Element web console, the Self-Service functionality is impacted.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Marketplace** in the navigation bar.
4. In the Nutanix Apps section, click **Get** for the Self-Service app.
5. On the app details page of Self-Service, click **Deploy**.
6. In the Deploy Self-Service confirmation window, click **Deploy**.

After deployment, the Self-Service app appears on the **My Apps** page. You can select **Self-Service** in the Application Switcher to access and use the app. See the [Self-Service Administration and Operations Guide](#) for information on configuring and using Self-Service.

Deploying Files

About this task

Nutanix Files provides file services to clients for file sharing across user work stations from a centralized and protected location.

Before you begin

- Ensure that you have enabled Marketplace. For more information, see [Enabling Marketplace](#) on page 9.
- For deployment prerequisites, see the [Deploying Files on the Files Manager](#) section in the *Files Manager User Guide*.

Procedure

1. Log in to Prism Central as an administrator.

2. Select **Admin Center** in the Application Switcher.
3. Click **Marketplace** in the navigation bar.
4. In the Nutanix Apps section, click **Get** for the Files app.
5. On the app details page, click **Deploy**.
6. In the Deploy Files confirmation window, click **Deploy**.

After deployment, the Files app appears on the **My Apps** page. You can select **Files** in the Application Switcher to access and use the app. See the [Nutanix Files Manager Guide](#) for information on configuring and using Files in Prism Central.

Deploying Foundation Central

About this task

Foundation Central can manage several Foundation instances from a single pane of glass, allowing you to create clusters of remote nodes without needing to configure each of them individually.

Before you begin

- Ensure that you have enabled Marketplace. For more information, see [Enabling Marketplace](#) on page 9.
- For deployment prerequisites, see the [Set Up Foundation Central](#) section in the *Foundation Central Guide*.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Marketplace** in the navigation bar.
4. In the Nutanix Apps section, click **Get** for the Foundation Central app.
5. On the app details page of Foundation Central, click **Deploy**.
6. In the Deploy Foundation Central confirmation window, click **Deploy**.

After Foundation Central is enabled, the Foundation Central app appears on the **My Apps** page. You can select **Foundation Central** in the Application Switcher to access and use the app. See the [Foundation Central Documentation](#) for information on configuring and using Foundation Central.

Deploying Kubernetes Management

About this task

Nutanix Kubernetes Management is a curated turnkey offering that provides simplified provisioning and operations of Kubernetes clusters. Kubernetes is an open-source container orchestration system for deploying and managing container-based applications.

Before you begin

- Ensure that you have enabled Marketplace. For more information, see [Enabling Marketplace](#) on page 9.
- For more information about the deployment requirements, see the [Requirements](#) section in the *Nutanix Kubernetes Engine Guide*.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Marketplace** in the navigation bar.
4. In the Nutanix Apps section, click **Get** for the Kubernetes Management app.
5. On the app details page, click **Deploy**.
6. In the Deploy Kubernetes Management confirmation window, click **Deploy**.

After deployment, the Kubernetes Management app appears on the **My Apps** page. You can select **Kubernetes Management** in the Application Switcher to access and use the app. See the [Nutanix Kubernetes Engine Guide](#) for information about configuring and using Kubernetes.

Deploying Objects

About this task

Nutanix Objects allows you to create and manage object stores in a cluster.

Before you begin

- Ensure that you have enabled Marketplace. For more information, see [Enabling Marketplace](#) on page 9.
- For deployment prerequisites, see the [Deployment and Network Prerequisites](#) section in the *Objects User Guide*.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Marketplace** in the navigation bar.
4. In the Nutanix Apps section, click **Get** for the Objects app.
5. On the app details page, click **Deploy**.
6. In the Deploy Objects confirmation window, click **Deploy**.
7. Select **Objects** in the Application Switcher.
A welcome page appears with the prerequisite details to create an object store.
8. Click **Download Creation Checklist** to download the list of prerequisites for deploying an object store.

9. (Only for ESXi clusters) Click **vCenter Registration** before deploying object store on an ESXi cluster. To deploy object stores on the ESXi clusters, you need to provide the vCenter credentials and configure the IPAM for the ESXi networks. For more information, see the [Managing vCenter for Object Service](#) section in the *Objects User Guide*.

Note: A non-admin user can perform vCenter Registration only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions:

- Create Object Store
- View Object Store

Administrators and non-admin users are Prism Central users with the following roles:

- Administrator—A Super Admin or a Prism Admin in Prism Central.
- Non-admin user—A Prism Central user without any administrator privileges.

For information on the built-in roles in Prism Central, see the [Built-in Role Management](#) section in the *Security Guide*.

10. Click **Next**.

11. Click **Create Object Store** to start creating the first object store.

For information on creating the object store, see the [Creating or Deploying an Object Store \(Prism Central\)](#) section in the *Objects User Guide*.

After deployment, the Objects app appears on the **My Apps** page. You can also select **Objects** in the Application Switcher to access and use the app. See the [Nutanix Objects Guide](#) for information on using object stores.

Deploying Database Service

About this task

Nutanix Database Service (NDB) automates and simplifies database administration, bringing one-click simplicity and invisible operations to database provisioning and life-cycle management. You can deploy multiple instances of Database Service.

Before you begin

- Ensure that you have enabled Marketplace. For more information, see [Enabling Marketplace](#) on page 9.
- For the network requirements for the Database Service, see the [NDB Network Requirements](#) section in the *Nutanix Database Service Administration Guide*.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Marketplace** in the navigation bar.
4. In the Nutanix Apps section, click **Get** for the Database Service app.
5. On the app details page, click **Deploy**.

6. In the Deploy Database Service confirmation window, click **Deploy**.
7. Enter the **Application Name** and **Application Description**.
8. In the **EraService** section, keep the default values to deploy Database Service with a non-HA configuration.
For other deployment configurations, see the [NDB Control Plane Configuration and Scalability](#) section in the *Nutanix Database Service Administration Guide*.
9. In the **NETWORK ADAPTERS (NICs)** section, select an appropriate NIC from the dropdown menu.
10. Click **Deploy**.

After deployment, the Database Service app appears on the **My Apps** page. You can select **Database Service** in the Application Switcher to access and use the app. See the [Nutanix Database Service Administration Guide](#) for information on configuring and using Database Service. You can perform these steps again to deploy another instance of Database Service.

Deploying Move

About this task

Nutanix Move (Move) is a cross-hypervisor mobility solution to migrate virtual machines (VMs) with minimal downtime. You can deploy multiple instances of Move.

Before you begin

- Ensure that you have enabled Marketplace. For more information, see [Enabling Marketplace](#) on page 9.
- Ensure that you have enabled DHCP in the network adapter.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Marketplace** in the navigation bar.
4. In the Nutanix Apps section, click **Get** for the Move app.

Note: The Move version that is available to you for deployment is the one that comes bundled with this Prism Central version. This bundled version may not necessarily be the latest Move version.

5. On the app details page, click **Deploy**.
6. In the Deploy Move confirmation window, click **Deploy**.
The Deploy Move page appears.
7. Enter the **Application Name** and **Application Description**.

8. In the **MoveService** section, do the following:
 - a. Update the value of **vCPUs** to 2.
 - b. Update the value of **Cores per vCPU** to 2.
 - c. Update the value of **Memory (GiB)** to 8.
 - d. Click **+** in the **DISKS** section.
 - e. For **Disk 2: SCSI**, enter the disk size as 50 in the **Size (GiB)** field.
 - f. In the **NETWORK ADAPTERS (NICs)** section, select an appropriate NIC from the dropdown menu.

Note: You must select an appropriate NIC for a successful deployment.

9. Click **Deploy**.

After deployment, the Move app appears on the **My Apps** page. You can select **Move** in the Application Switcher to access and use the app. See the [Move User Guide](#) for information on configuring and using Move. You can perform these steps again to deploy another instance of Move.

Deploying Power Monitor Application

Deploy a power monitor application from the Nutanix marketplace.

Before you begin

- Ensure that the Nutanix Marketplace is enabled. For more information, see [Enabling Marketplace](#) in the *Prism Central Admin Center Guide*.
- Review the [Power Usage Requirements](#) and [Power Usage Considerations](#).

About this task

To deploy the power monitor application from the Nutanix Marketplace, follow these steps:

Procedure

1. Log in to Prism Central as an administrator.
2. From the [Application Switcher Function](#), select the **Admin Center** application and from the navigation bar, select **Marketplace**.
3. Under the **Nutanix Apps** section, locate the **Power Monitor** application widget, and click **Get**.
4. On the **Power Monitor** application details page, click **Deploy**.
5. On the **Deploy Power Monitor** confirmation window, click **Deploy**.

The **Power Monitor** application appears on the **My Apps** page.

What to do next

After deploying the power monitor application, you must configure the BMC out-of-band management credentials. For more information, see [Configuring Out-of-Band Management Credentials](#).

Deploying Nutanix Central

Use this procedure to deploy Nutanix Central on-premises application from the Prism Central Marketplace.

Before you begin

Complete the prerequisites. For more information, see [Nutanix Central On-Premises Prerequisites](#).

About this task

To deploy Nutanix Central on-premises, follow these steps:

Procedure

1. Log on to Prism Central web console as an administrator.
2. From the [Application Switcher Function](#), select **Admin Center**.
3. In the navigation bar, click **LCM > Prism Central Cluster**.
4. Perform inventory.
For more information, see [Performing Inventory with the Life Cycle Manager](#) in the *Life Cycle Manager Guide*.
5. Update the **Marketplace Bundle** component version.
For more information, see [Performing Firmware and Software Updates in a Connected Site Setup](#) in the *Life Cycle Manager Guide*.
6. In the navigation bar, click **Marketplace**.
7. Under **Nutanix Apps**, in the **Nutanix Central** tile, click **Get**.
The **Nutanix Central** application page appears.
8. Click **Deploy**.
The **Deploy Nutanix Central** dialog box appears.
9. Click **Deploy**.
The **Nutanix Central** deployment page appears.

What to do next

For further deployment steps, see [Deploying Nutanix Central On-Premises](#).

Preferred Partner App Deployment

A Preferred Partner app is developed and certified by Nutanix in collaboration with the preferred partner. The app is pre-seeded for deployment on the Nutanix platform. For example, OpenShift.

Similar to Nutanix apps, the Preferred Partner app is available in the Marketplace for deployment immediately after you enable Marketplace in Prism Central.

Ensure that you meet all the resource requirements specific to the app before you deploy the Preferred Partner app.

OpenShift Overview

The OpenShift app in the Nutanix Marketplace is developed and certified by Nutanix in collaboration with Red Hat.

Red Hat OpenShift brings together tested and trusted services to reduce the friction of developing, modernizing, deploying, running, and managing applications. Built on Kubernetes, OpenShift delivers a consistent experience across public cloud, on-premises, hybrid cloud, or edge architecture.

You can deploy OpenShift from the Nutanix Marketplace. For more information, see [Deploying OpenShift](#) on page 19. Ensure that you meet all prerequisites before you deploy OpenShift from the Marketplace. For more information, see .

Prerequisites to Deploy OpenShift

Before you deploy OpenShift from the Nutanix Marketplace, ensure that you meet these prerequisites:

- Your Prism Central instance must have the required resources to enable OpenShift. For more information, see the *Red Hat Documented Cluster Resource Requirements* section in the *OpenShift Documentation*.
- You must have an internet connection with or without proxy. If you use a firewall, allow these URLs in the firewall to ensure that your Prism Central and Prism Element clusters have access to these URLs:
 - *.docker.io
 - *.production.cloudflare.docker.com
 - *.nutanix.github.io
 - download.nutanix.com
 - .quay.io
 - .cloudfront.net
 - *.us-west-2.amazonaws.com

For URLs specific to OpenShift, see the *OpenShift Documentation*.

- You must create DNS records for two static IP addresses in the appropriate DNS server for the Nutanix instance that hosts your OpenShift Container Platform cluster. In each record, `<cluster_name>` is the cluster name and `<base_domain>` is the cluster base domain that you specified during the cluster installation. A complete DNS record has the `<component>.<cluster_name>.<base_domain>` format. Create the following DNS records:
 - API VIP: This DNS record must point to the load balancer for the control plane machines. The clients that are outside the cluster and the nodes within the cluster must resolve this DNS record. The record format for this component is `api.<cluster_name>.<base_domain>`.
 - Ingress VIP: This DNS record is a wildcard record that points to the load balancer targeting the machines that run ingress router pods. The clients that are outside the cluster and the nodes within the cluster must resolve this DNS record. The record format for this component is `*.apps.<cluster_name>.<base_domain>`.

Note: If your OpenShift Container Platform (OCP) cluster includes a Windows worker node, add another DNS record in the `api-int.<cluster_name>.<base_domain>` format for internal communication within the cluster.

Deploying OpenShift

The Nutanix Marketplace provides a pre-configured OpenShift application to simplify the deployment process so that you can quickly get a production-ready cluster running with just a few clicks.

About this task

Note:

- You can deploy OpenShift using the Internal Project only. For information on the Internal Project, see [Project Management](#) on page 33.
- OpenShift deployment from the Nutanix Marketplace requires an AHV hypervisor. For Nutanix clusters with an ESXi hypervisor, you must perform VMware Installer Provisioned Infrastructure (IPI) or User Provisioned Infrastructure (UPI) based installation.

Before you begin

Ensure that you meet all the OpenShift deployment prerequisites. For more information, see [Prerequisites to Deploy OpenShift](#) on page 19.

Procedure

To deploy OpenShift, follow these steps:

1. Log in to your Prism Central instance.
2. From the Application Switcher, select **Admin Center**.
3. In the navigation bar, click **Marketplace**.
4. Under Preferred Partners, click **Get** for the OpenShift app.
5. Click **Deploy**.
6. In the **Application Name** field, enter the name of the application.
Following are the rules for naming convention.
 - Start the name of the blueprint with an alphanumeric character or an underscore.
 - Use only space, underscore, and dash as special characters.
 - Ensure that you have at least one character in the name.
 - Do not end the name with a dash.

7. Under Profile Variables, follow these steps.
 - a. From the **OpenShift minor release version** and **OpenShift patch release version** dropdown menus, select the required versions.
 - b. In the **OpenShift cluster name**, enter the name of the cluster.
 - c. In the **OpenShift base domain** field, enter the base domain of your environment to create routes to your OpenShift Container Platform cluster components.

The full DNS name for your cluster is a combination of the base domain and cluster name parameter values that uses the `cluster_name.base_domain` format.
 - d. In the **Prism Central FQDN** field, enter the Prism Central IP address or FQDN.

This FQDN must be in the SAN properties of the Prism Central SSL certificate.
 - e. In the **Prism Element FQDN** field, enter the Prism Element IP address or FQDN.

This FQDN must point to the cluster virtual IP address where you intend to deploy the provisioner VM. Enter the value of **Prism Element FQDN** in the `your.prismcentral.domainname` format.
 - f. In the **Prism Element port** field, enter the Prism Element port.

The default value for the Prism Element console port is 9440.
 - g. In the **OpenShift Machine Network** field, enter the IP address blocks for virtual machines.

The block must not overlap with cluster and service networks. Enter the value in the 10.0.0.0/16 format.
 - h. In the **OpenShift API VIP** field, enter the virtual IP (VIP) address that you configured for control plane API access.

The VIP address must be in the same subnet as the machine network. Enter the value of this field in the XXX.XXX.XXX.XXX format.
 - i. In the **OpenShift Ingress VIP** field, enter the virtual IP (VIP) address that you configured for cluster ingress.

The VIP address must be in the same subnet as the machine network. Enter the value of this field in the XXX.XXX.XXX.XXX format.
 - j. In the **OpenShift Pod Network** field, enter the IP address block for Kubernetes pods.

The default value for this field is 10.128.0.0/14 with a host prefix of /23.
 - k. In the **OpenShift Pod addresses per host** field, enter the subnet prefix length to assign to each individual node.

For example, if host prefix is set to 23 then each node is assigned a /23 subnet out of the given CIDR. A host prefix value of 23 provides 510 pod IP addresses.
 - l. Enter the IP address block for Kubernetes services in the **OpenShift Service Network** field.

The default value is 172.30.0.0/16.
 - m. Enter the number of CPUs per control plane host in the **Control Plane CPU** field (a minimum of 4 vCPUs and a maximum of 128 vCPUs).
 - n. Enter the memory (MiB) per control plane host in the **Control Plane Memory** field (a minimum of 16384 MiB and a maximum of 1048576 MiB).
 - o. Enter the number of compute machines (worker machines) to provision in the **Compute replicas** field.
 - p. Enter the pull secret in the **OpenShift Pull Secret** field.

You can get a pull secret from the Red Hat OpenShift Cluster Manager to authenticate downloading container images for OpenShift Container Platform components from services, such as *Quay.io*. <https://console.redhat.com/openshift/install/pull-secret>.

- q. (Optional) Enter the details of the HTTP proxy server in the **(Optional) Internet proxy for HTTP** field.
The format for defining an HTTP proxy typically includes the proxy server address, port, and sometimes authentication credentials (when required). Enter the HTTP proxy server details in the following format.
- Without Authentication: The format must be `http://<ip>:<port>`, where `<ip>` is the domain name or IP address of the proxy server and `<port>` is the port number of the proxy server.
 - With Authentication: The format must be `http://<username>:<pswd>@<ip>:<port>`, where `<username>` is the username for authenticating with the proxy and `<pswd>` is the corresponding password for the username.
- r. (Optional) Enter the details of the HTTPS proxy server in the **(Optional) Internet proxy for HTTPS** field.
The format for defining an HTTPS proxy typically includes the proxy server address, port, and sometimes authentication credentials (when required). Enter the HTTPS proxy server details in the following format.
- Without Authentication: The format must be `https://<ip>:<port>`, where `<ip>` is the domain name or IP address of the proxy server and `<port>` is the port number of the proxy server.
 - With Authentication: The format must be `https://<username>:<pswd>@<ip>:<port>`, where `<username>` is the username for authenticating with the proxy and `<pswd>` is the corresponding password for the username.
- s. (Optional) Enter the destination domain names, IP addresses, or other network CIDRs to exclude from the proxy in the **(Optional) Bypass the proxy** field.
The details you enter in the **(Optional) Bypass the proxy** field must be a comma-separated list in the following format.
- example.com,.example.com,192.168.1.1,192.168.0.0/24
- t. To diagnose and resolve issues when something goes wrong in the deployment process, select **True** in the **DEBUG** dropdown menu.

8. In the Provisioner section, follow these steps:

- a. Click **NETWORK ADAPTERS**.
- b. From the **NIC** dropdown menu, select a NIC.

The NIC dropdown menu options are based on the cluster configuration.

Note: Cluster details are optional and are inferred from the NIC you select.

9. In the Credentials section, follow these steps:

- a. Expand the credentials.
- b. In the **Password** field, specify a password for the credential.
- c. In the **SSH Private Key** field, specify a private key for the credential.

Note: The user must be a local account with the User Admin role in Prism Central and Prism Element. Nutanix recommends you to create a service account in Prism Central and Prism Element.

10. Click **Deploy**.

Cisco Intersight Device Connector Overview

The Cisco Intersight Device Connector application enables Nutanix Prism Central to securely connect to the Cisco Intersight SaaS platform or on-premises virtual appliance. This connection allows IT teams to monitor, manage, and automate their infrastructure using the advanced features of Cisco Intersight while running Nutanix clusters.

By integrating this device connector with Prism Central, Nutanix provides a secure, streamlined method for endpoints to communicate with the Cisco Intersight portal over the internet. IT teams can claim (register) and unclaim (unregister) their Nutanix infrastructure directly through Intersight, using the well-established security protocols of Cisco.

Prism Central integrates with Cisco Intersight using the Cisco Intersight Device Connector application to periodically collect the following data:

- Inventory data, such as cluster configuration information, node information, capacity used in the system, and so on
- Alerts
- Licensing and alarms (supports up to 5000 alarms)
- Disks and networks
- Clusters and hosts statistics
- Virtual machines and cluster-level GPUs

Note:

- The Cisco Intersight Device Connector is qualified only on AHV and ESXi hypervisors.
- Deployment of Cisco Intersight Device Connector creates the associated service account, role, and authorization policy. You can view these entities in the following sections of the IAM page:
 - Service account: Appears in the **Imported Users** section of the **Identities** tab.
 - Role: On the **Roles** tab.
 - Authorization Policy: On the **Authorization Policies** tab.

Do not modify or delete the service account, role, or authorization policy associated with the Cisco Intersight Device Connector.

- The current version of Cisco Intersight Device Connector supports a maximum of 25 clusters with up to 8 nodes.

You deploy the Cisco Intersight Device Connector application from the Nutanix Marketplace. For information on deploying the device connector, see [Deploying Cisco Intersight Device Connector](#) on page 24.

You must consider the following points before deploying the Cisco Intersight Device Connector.

- You deploy the Cisco Intersight Device Connector application only if your Prism Central instance is managing Cisco-server-based clusters.
- You must have the Prism Admin or Super Admin role to deploy, open, and manage the Cisco Intersight Device Connector application from the Marketplace. For more information on roles, see the [Security Guide](#).
- You must have the Microservices Infrastructure (sometimes referred to as MSP) enabled and running for your Prism Central instance. For more information on MSP, see [Microservices Infrastructure](#) in the *Prism Central Infrastructure Guide*.

When the number of inventory items (clusters and VMs) reaches its threshold level, you can manually resize Prism Central resources and the Cisco Intersight Device Connector container for proper functioning of the Cisco Intersight Device Connector application. For more information on resizing, see [KB-18989](#).

Deploying Cisco Intersight Device Connector

As a Prism Admin or Super Admin, you can deploy the Cisco Intersight Device Connector application from the Marketplace.

About this task

For more information on Cisco Intersight Device Connector, see [Cisco Intersight Device Connector Overview](#) on page 23.

For proxy configuration, see [Service Manager for Installing Services and Applications](#) in the *Prism Central Infrastructure Guide*.

Before you begin

- Ensure that you enabled Marketplace. For more information, see [Enabling Marketplace](#) on page 9.
- Ensure that you are running Prism Central pc.2024.3 or later versions.

Procedure

To deploy Cisco Intersight Device Connector, perform the following steps.

1. Log in to your Prism Central instance.
2. From the Application Switcher, select **Admin Center**.
3. In the navigation bar, click **Marketplace**.
4. Under Preferred Partners, on the **Cisco Intersight Device Connector** application tile, click **Get**.
5. For information on the device connector resource requirements, scroll the page and click the **Cisco Intersight Device Connector Resource Requirements** link.
6. To read the End User License Agreement for the device connector, click the **Cisco End User License Agreement** link.
7. On the Cisco Intersight Device Connector application details page, click **Deploy**.
8. In the confirmation window, click **Deploy**.
The deployment process creates a tile for the application in **My Apps** from where you can view and manage the application.
9. (Optional) To view the progress of the deployment, click the **Audit** tab.
10. (Optional) To download a collective audit log report, click **Generate Audit Logs**.
You can also click the Download icon next to a task to download individual script logs.

What to do next

After the Cisco Intersight Device Connector application displays the **Running** state, open the application either from the **Marketplace** or **My Apps** to access the device connector on a new browser tab.

You can then claim your Prism Central instance to Cisco Intersight. For more information, see [Claiming your Prism Central Instance to Cisco Intersight](#) on page 25 and [Claiming your Prism Central Instance to Cisco Intersight Virtual Appliance](#) on page 26.

Deploying Cisco Intersight Device Connector in a Dark Site Environment

You can deploy Cisco Intersight Device Connector on your Prism Central instance that is in a dark site environment.

Before you begin

- Ensure that you deployed the Prism Central MSP Apps Bundle in the dark site environment. For more information, see [Service Manager for Installing Services and Applications](#) in the *Prism Central Infrastructure Guide* and [Enabling Security Dashboard on Dark Sites](#) in the *Security Guide*.
- Ensure that you upgraded Prism Central in the dark site environment to version pc.7.3 or later. For more information, see [Firmware and Software Updates Management in a Dark Site \(No-Internet Connectivity\) Setup](#) in the *Life Cycle Manager Guide*.

Procedure

To deploy Cisco Intersight Device Connector in the dark site environment, follow these steps:

1. Download the Cisco Intersight Device Connector installation bundle version 7.3.0 or later from the [Nutanix Downloads](#) page.
2. Place the Cisco Intersight Device Connector installation bundle on the local web server that you created to upgrade your Prism Central instance in the dark site environment.
For more information on placing the bundle on a web server, see [Placing the Firmware and Software LCM Bundles on a Web Server](#) in the *Life Cycle Manager Guide*.
3. Follow these steps to verify that your Prism Central instance is pointing to the URL of the local web server:
 - a. Log on to your Prism Central instance.
 - b. From the Application Switcher, select **Admin Center**.
 - c. In the navigation bar, select **LCM**.
 - d. On the **LCM** page, click the **Prism Central Cluster** tab.
 - e. Select the Prism Central cluster in the list.
 - f. In the right pane, open the **Settings** tab.
 - g. In the **URL** field, verify that the URL of the local web server is correct.
4. Deploy Cisco Intersight Device Connector from the Marketplace.
For more information, see [Deploying Cisco Intersight Device Connector](#) on page 24.

Claiming your Prism Central Instance to Cisco Intersight

After you deploy the Cisco Intersight Device Connector from the Marketplace, claim (register) your Prism Central instance to Cisco Intersight for data collection.

Before you begin

- You must have outbound internet access to connect to Cisco Intersight.
- You must have the Prism Admin or Super Admin role to claim your Prism Central instance to Cisco Intersight.
- The Device Connector must communicate with svc.intersight.com on port 443, and outbound connectivity must be enabled for each Prism Central VM.

Procedure

To claim your Prism Central instance to the public Cisco Intersight cloud, follow these steps:

1. Log on to Cisco Intersight using your Cisco credentials.

2. From the dropdown menu, select **System**.
3. In the left pane, click **Targets**.
4. Click **Claim a New Target**.
5. Select **Hyperconverged**, then select the **Nutanix Prism Central** tile.
6. Click **Start**.
7. Open a new browser tab or window and log on to your Prism Central instance.
8. From the Application Switcher, select **Admin Center**.
9. In the navigation bar, click **My Apps**.
10. On the Cisco Intersight Device Connector tile, click **Open**.
The **Device Connector** page opens.
11. On the **Device Connector** page, click the copy icon in the **Device ID** section to copy the ID.
12. On the Cisco Intersight user interface, paste the device ID in the **Device ID** field.
13. On the **Device Connector** page, click the copy icon in the **Claim Code** section to copy the code.
14. On the Cisco Intersight user interface, paste the claim code in the **Claim Code** field.
15. (Optional) Select the resource groups that this target should be added to.
16. Click **Claim**.

Claiming your Prism Central Instance to Cisco Intersight Virtual Appliance

After you deploy the Cisco Intersight Device Connector from the Marketplace, claim (register) your Prism Central instance to Cisco Intersight for data collection.

Before you begin

- Ensure that you deploy and configure Cisco Intersight Virtual Appliance in your on-premises environment.
- Ensure that your Prism Central instance has access to the Cisco Intersight Virtual Appliance within your local network.
- Ensure that you have the IP address (or hostname) and credentials of your Prism Central instance.
- Ensure that you have the Prism Admin or Super Admin role to claim your Prism Central instance.

Procedure

Perform the following steps to claim your Prism Central instance to Cisco Intersight Virtual Appliance.

1. Log on to Cisco Intersight Virtual Appliance using your Cisco credentials.
You access the Cisco Intersight Virtual Appliance local management interface using the IP address or hostname.
2. From the dropdown menu, select **System**.
3. In the left pane, click **Targets**.
4. Click **Claim a New Target**.
5. Select **Hyperconverged**, then select the **Nutanix Prism Central** tile.

6. Click **Start**.
The Claim Nutanix Prism Central Target page opens.
7. In the **Hostname/IP Address** field, enter the hostname or IP address.
8. In the **Username** field, enter the username.
9. In the **Password** field, enter the password.
10. (Optional) Select the resource groups that this target should be added to.
11. Click **Claim**.

Upgrading Cisco Intersight Device Connector using Life Cycle Manager

You can upgrade Cisco Intersight Device Connector using Life Cycle Manager (LCM).

Before you begin

Ensure that you deployed Cisco Intersight Device Connector.

Procedure

To upgrade Cisco Intersight Device Connector using LCM, follow these steps:

1. Log on to your Prism Central instance.
2. From the Application Switcher, select **Admin Center**.
3. In the navigation bar, select **LCM**.
4. Select the Prism Central cluster in the list.
5. Click **Perform Inventory**.
The LCM tool retrieves the latest firmware and software LCM bundles and lists them on the **Inventory** page. You can view the status on the **Operation** tab in the right pane.
6. In the **Action** column, click **Upgrade**.
The **Select Upgrades** page displays the installed version and upgrade version of Cisco Intersight Device Connector.
7. Select the **Cisco Intersight Device Connector** checkbox and click **Next**.
8. Click **Apply Upgrade Plan**.
You can view the upgrade status on the **Operation** tab and verify the upgraded version on the **Software** tab in the right pane.

Marketplace Application Update

Update Marketplace applications to their latest version using Life Cycle Manager (LCM) so that you can deploy the latest version of these applications after the update.

Nutanix provides Marketplace application updates as bundles through LCM. These bundles might include any of these components:

- The latest versions of Nutanix applications or Preferred Partner applications
- Any new applications that is not currently available in the Marketplace

The Marketplace update simplifies application management in the following ways:

- Application updates can occur independently of the Prism Central release cycle, which provides more flexibility and faster access to the latest versions.
- The Marketplace page notifies you whenever a new version of a Marketplace bundle is available in LCM after you perform an inventory.

Updating Marketplace Applications Using Life Cycle Manager

Update Marketplace applications to their latest version using Life Cycle Manager (LCM).

Before you begin

Ensure that you enabled Marketplace. For more information, see [Enabling Marketplace](#) on page 9.

Procedure

1. Log on to your Prism Central instance.
2. From the Application Switcher, select **Admin Center**.
3. In the navigation bar, select **LCM > Prism Central Cluster**.

4. Perform inventory.

For more information, see [Performing Inventory with the Life Cycle Manager](#) in the *Life Cycle Manager Guide*.

5. Update the **Marketplace Bundle** component.

For more information, see [Performing Firmware and Software Updates in a Connected Site Setup](#) in the *Life Cycle Manager Guide*.

Note: To update the **Marketplace Bundle** component in the dark site environment, see [Performing Firmware and Software Updates in a Dark Site Setup](#) in the *Life Cycle Manager Guide*.

MY APPS

My Apps provides a single workspace to manage all Nutanix apps and Preferred Partner apps that you deployed from the Nutanix Marketplace. For more information on the app categories, see [Marketplace](#) on page 9.

The My Apps page includes the following sections:

- The Nutanix Apps section lists the Nutanix Apps, such as Self-Service, Objects, and so on, that you deployed from the Marketplace. For information on managing Nutanix Apps, see [Nutanix Apps Management](#) on page 29.
- The Preferred Partner Apps section lists the Preferred Partner apps, such as OpenShift or Cisco Intersight Device Connector, that you deployed from the Marketplace. For information on managing Preferred Partner Apps, see [Preferred Partners Apps Management](#) on page 30.

Note:

- You can access and manage apps using the My Apps page only after your Prism Admin enables the Marketplace. For more information, see [Enabling Marketplace](#) on page 9.
- Starting with Prism Central version pc.2024.3.1, you can manage custom apps and hybrid cloud apps (previously listed as **Other Apps** on the Nutanix Marketplace and My Apps page) from the Applications page in Self-Service. For more information, see the [Applications in Self-Service](#) section in the *Self-Service Administration and Operations Guide*.

Nutanix Apps Management

The App Details page of a Nutanix App provides the app summary and the audit trail of its deployment. The Overview tab on the App Details page displays the app UUID, name of the app with its version, and so on. The Audit tab displays the audit logs of the actions performed on the app.

The App Details page also provides the option to upgrade the app. For more information on upgrading Nutanix Apps, see [Life Cycle Manager Guide](#).

The My Apps page also has the option to register existing instances of Database Service or Move so that you can manage them on the My Apps page. For more information, see [Registering Database Service or Move to My Apps](#) on page 30.

Note:

- When you upgrade Nutanix apps through LCM, you might notice the deployed app version on the My Apps page differs from the respective Marketplace Item version on the Marketplace Page. The Marketplace Item version refers to the version that is bundled with the Prism Central version you have upgraded to. Whereas, the app version on the My Apps page shows the current version the app is running.
- When you upgrade from a previous version of Prism Central, you must enable Marketplace and perform the LCM inventory to see the correct versions of Nutanix Apps on the My Apps page.

Nutanix Apps Permissions

The following table lists the different operations for Nutanix Apps and permissions for different users to perform those operations in My Apps:

Table 2: My Apps Operations for Nutanix Apps

Operation	Prism Central Admin	Project Admin	Developer	Consumer	Operator
View or open Nutanix Apps	Yes	Yes	Yes	Yes	Yes
Search or Filter Nutanix Apps	Yes	Yes	Yes	No	No
Upgrade Nutanix Apps	Yes	No	No	No	No

Registering Database Service or Move to My Apps

About this task

You can register existing instances of Database Service or Move so that you can manage them through My Apps.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **My Apps** in the navigation bar.
4. On the My Apps page, do the following:
 - » To register an existing instance of Database Service, click **Register Database Service**.
 - » To register an existing instance of Move, click **Register Move**.
5. On the Register page, select the VM of the instance from the list.
You can also search the VMs on the basis of name, cluster name, or IP address.
6. Click **Register**.
The registered instance appears on the My Apps page.

Preferred Partners Apps Management

You can manage the preferred partners app through Prism Central.

Preferred Partner apps are applications developed by Nutanix partners that are curated and made available for easy deployment directly from the Nutanix Marketplace. These applications are validated by Nutanix to ensure they integrate well with the platform.

This section contains topics to describe management of preferred partners apps.

Deleting Cisco Intersight Device Connector

After deployment, the Cisco Intersight Device Connector application appears in the **My Apps** section of the Admin Center and provides you the options to manage the application.

Before you begin

You must have the Prism Admin or Super Admin role to delete the Cisco Intersight Device Connector application.

About this task

Warning: Deleting the Cisco Intersight Device Connector application impacts the connectivity between Cisco Intersight and Prism Central and lead to inconsistencies and errors. Ensure that you have analyzed the outcome before deleting the application.

Procedure

To delete the Cisco Intersight Device Connector application, follow these steps:

1. Log in to your Prism Central instance.
2. In the Application Switcher, select **Admin Center**.
3. In the navigation bar, click **My Apps**.
4. On the Cisco Intersight Device Connector tile, click **Manage**.
5. On the **Overview** tab of the Cisco Intersight Device Connector application page, click **Delete**. The **Delete Application** window opens.
6. In the confirmation field, type DELETE.
7. click **Delete**.

You can view the progress of the app deletion on the **Audit** tab.

Upgrading Cisco Intersight Device Connector from My Apps

You can upgrade Cisco Intersight Device Connector application to the latest version from **My Apps**.

About this task

You must have the Prism Admin or Super Admin role to manage or upgrade the Cisco Intersight Device Connector application.

You can also upgrade Cisco Intersight Device Connector directly from the LCM page. For more information, see [Upgrading Cisco Intersight Device Connector using Life Cycle Manager](#) on page 27.

Procedure

1. Log on to your Prism Central instance.
2. From the Application Switcher, select **Admin Center**.
3. In the navigation bar, select **My Apps**.
4. On the Cisco Intersight Device Connector tile, click **Manage**. The **Overview** tab opens.
5. Click **Upgrade**.

The confirmation window opens and shows the **Upgrade** or **Redirect to LCM** option depending on whether you performed inventory in LCM.

6. If you did not perform inventory in LCM, click **Redirect to LCM** and follow these steps:
 - a. On the **LCM** page, click the **Prism Central Cluster** tab.
 - b. Select the Prism Central cluster in the list.
 - c. Click **Perform Inventory**.

The LCM tool retrieves the latest firmware and software LCM bundles and lists them on the **Inventory** page. You can view the status on the **Operation** tab in the right pane.
 - d. In the **Action** column, click **Upgrade**.

The **Select Upgrades** page displays the installed version and upgrade version of Cisco Intersight Device Connector.
 - e. Select the **Cisco Intersight Device Connector** checkbox, and click **Next**.
 - f. Click **Apply Upgrade Plan**.

You can view the status on the **Operation** tab and verify the upgraded version on the **Software** tab in the right pane.
7. If you performed inventory, click **Upgrade** to upgrade Cisco Intersight Device Connector.

PROJECT MANAGEMENT

A project is a set of users with a common set of requirements or a common structure and function, such as a team of engineers collaborating on an engineering project.

Projects provide logical groupings of user roles for managing resource usage within your organization. You use a directory or identity provider configured in your Prism Central to assign different roles to users or groups in a project.

The project construct helps you define various settings, such as:

- Permissions, such as the user accounts and groups who can deploy a marketplace application.
- The networks to use while deploying an application.
- Default VM specifications and deployment options, such as vCPUs, vRAM, storage, base images, Cloud-Init, or Sysprep specs.
- The quota and snapshot policy definition.
- Credentials.

Note: For information about the maximum tested and supported values for projects in Prism Central, see *Prism Central Scalability* section in the [Prism Central Release Notes](#).

A project has user roles. Each role has predefined functions that a user to whom the role is assigned can perform in a project. For information on the user roles, see [Roles Management](#) on page 75 and [Role-based Access Control in Self-Service](#) in the *Self-Service Administration and Operations Guide*.

Projects provide a consistent experience when you access them from Prism Central or from Self-Service. However, the project features and configurations differ depending on whether you have deployed Self-Service in your Prism Central or not.

The following table lists the project features that are available with Prism Central when Self-Service is disabled and when Self-Service is enabled.

Table 3: Project Feature Availability in Prism Central

Project features	When Self-Service is disabled	When Self-Service is enabled
Allowing multiple AHV clusters and subnets	Yes	Yes
Viewing workloads per project	Yes (for VMs)	Yes (for VMs and Applications)
Viewing usage statistics (vCPU, Memory, Usage)	Yes	Yes
Internal project	Yes	Yes
Managing project-level quotas	Yes	Yes
Allowing VPC Overlay subnets	Yes	Yes
Adding other private or public cloud accounts (ESXi, AWS, Azure or GCP)	No	Yes

Project features	When Self-Service is disabled	When Self-Service is enabled
Managing account or cluster quotas & quotas for ESXi (new quotas with policy engine)	No	Yes
Managing environments	No	Yes
Managing tunnels for VPC or SaaS in Projects	No	Yes
Managing snapshot policies	No	Yes

Upgrade to the Refactored Project

You upgrade to the refactored project when you upgrade to Prism Central version 2022.6 or later. With this upgrade, all existing entities in Prism Central that are assigned to the default project are migrated to an Internal Project. The Internal Project is a system-defined project that encompasses all underlying infrastructure.

After the upgrade when you create VMs as an admin, the Internal Project is automatically assigned to the VMs. You can later assign the VM to a different project using the **Manage Ownership** option. For more information, see [VM Management](#) in the *Prism Central Infrastructure Guide*.

Note: With the refactored project, Self-Service no longer provides an out-of-the-box Default Project when you enable it. The Default project is visible only when you upgrade your Prism Central from a previous version to version 2022.6 or later to retain existing Self-Service-specific entities (if any) using the Default project. The Default Project does not appear for any new deployments of Prism Central version 2022.6 or later.

Projects Summary View

The Projects Summary page allows you to view the list of existing projects and their definitions in an at-a-glance view. You can click a project name to view and manage the detailed configuration of the project on the Project Details page. To access the Projects Summary page, select **Admin Center** in the Application Switcher, and click **Projects** in the navigation bar.

The following table describes the definitions that appear for a project on the Projects Summary page.

Table 4: Projects List Fields

Definition	Description
Name	Displays the project name.
Status	Displays the status of the project.
Infrastructure	Displays the number of on-prem and cloud accounts added to the project.
Users	Displays the number of users in the project.
Apps	Displays the number of applications deployed using the project.
VMs	Displays the number of on-prem AHV cluster VMs in the project.
vCPU	Displays the number of vCPUs used by the on-prem AHV clusters within the project.
Memory	Displays the amount of memory (in GiB) used by the on-prem AHV clusters within the project.

Definition	Description
Storage	Displays the amount of storage (in GiB) used by the on-prem AHV clusters within the project.
Cost	Displays the resource cost of your applications for the last 30 days using Showback. Showback is supported only for Nutanix and VMware through vCenter platforms.

You can filter the projects list based on different parameter values. The following table describes the filter options available when you open the **Filter** pane. To apply a filter, select a parameter and select the box of the desired value (or multiple values) you want to use as a filter. You can apply filters across multiple parameters.

Table 5: Filter Pane Fields

Parameter	Description
Name	Filters based on the project name. Select a condition from the dropdown list and enter a keyword in the field to get a list of projects that satisfy the condition and keyword combination. You can select Contains , Does not contain , Starts with , Ends with , or Equals to condition.
vCPU Usage	Filters based on the number of vCPUs used. Enter a range in the From <low> to <high> field to get a list of projects within the range of vCPU usage that you specify.
Memory Usage	Filters based on the amount of memory used. Enter a range in the From <low> to <high> GiB fields to get a list of projects within the range of memory usage that you specify.
Storage Usage	Filters based on the amount of storage space used. Enter a range in the From <low> to <high> GiB field to get a list of projects within the range of storage usage that you specify.
VM Count	Filters based on the number of VMs. Enter a range in the From <low> to <high> field to get a list of projects within that range of the total number of VMs.

You can click the Settings icon to select, deselect, or reorder the columns that you want to view on the Projects Summary page. You can click **Export** to export the data in a .csv file.

The Projects Summary page also includes a **Create Project** to create projects. For more information, see [Creating a Project](#) on page 42.

The **Actions** menu appears when you select a project on the Projects Summary page. The menu allows you to delete the selected project.

To view the details of a project, click the project name to open the Project Details page. For more information, see [Project Details View](#) on page 35.

Project Details View

The Project Details page allows you to view the entire setup of a project in different tabs and helps you manage their configuration. To access the Project Details page, click the project name on the Project Summary page. For information on the Project Summary page, see [Projects Summary View](#) on page 34.

A few tabs on the Project Details page appear only when you enable Self-Service in your Prism Central instance. The Project Details page has the following tabs:

- Dashboard
- Usage
- Workloads
- Users & Groups
- Infrastructure
- Environments (appears when you enable Self-Service in your Prism Central instance)
- Policies (appears when you enable Self-Service in your Prism Central instance)
- Tunnels (appears when you enable Self-Service in your Prism Central instance)

Project Dashboard Tab

The **Dashboard** tab appears when you first open the Project Details page. The project dashboard contains tiles to display various configurations and usage summary of the project at one place. The availability of tiles and the appearance of the information on the tiles differ based on Self-Service enablement (see [Deploying Self-Service](#) on page 12) and the overall project configuration and usage data.

The dashboard has the following tiles:

Table 6: Project Dashboard

Tiles	Description
Overview	The Overview tile has the following sections. • Description: Displays the description you provided for the project. You can hover your mouse on the tile and click Edit to add or modify the description of the project. • Summary: Displays the number of users, groups, roles, infrastructure (accounts), and environments configured for the project. You can hover your mouse and click the Edit icon to open the respective tab and edit the configuration. For more information about the infrastructure (accounts), see Infrastructure in Projects on page 46. For more information about environments, see Environments in Projects on page 50.

Tiles	Description
Workload Summary	Provides information related to the overall workloads associated with the project. When Self-Service is disabled, the tile only displays the total number of VMs that are up and the number of VMs that are in the active, stopped, or error state within the project. When you hover your mouse, you can also view the number of local and remote VMs. The local VM count indicates the number of VMs associated with the local AHV clusters added to the project. The remote VM count indicates the number of VMs associated with the remote PC (remote clusters) added to the project. When Self-Service is enabled, the tile also displays the total number of applications and the number of applications that are in the active, stopped, or error state within the project. You can click View Workloads to navigate to the Workloads tab and view the details of the VMs or applications. For more information, see Project Workloads Tab on page 41.
Top Workloads	Provides information related to the workloads that consume the maximum resources. When Self-Service is disabled, the tile displays the VMs that consume the maximum resources within the project. When Self-Service is enabled, the tile also displays the applications (on-prem AHV and ESXi) that consume the maximum resources within the project. You can use the switcher at the top of the tile to view the list of VMs or applications. You can filter the list of VMs or applications based on the vCPU, memory, or storage consumption. You can click View all VMs or View all Apps to navigate to the Workloads tab and view the state, owner, and other details of the VMs or applications. For more information, see Project Workloads Tab on page 41.

Tiles	Description
More Project Features	Displays recommendations and options to enable additional features for your projects. You can click the options available on the tile to enable Self-Service and configure these additional features for your project. • Enable Policies: Policies in a project allows you to control the resources that are consumed in your project. With policies, you can configure resource quotas, approvals, and other policy types. For more information, see Policy Engine Overview in the <i>Self-Service Administration and Operations Guide</i>. • Deploy workloads on cloud accounts: In addition to the on-prem private clouds with AHV or ESXi, you can also deploy your workloads on public clouds when you enable Self-Service. You can configure your public clouds such as AWS, Azure, or GCP in Self-Service. For more information, see Provider Account Settings in Self-Service in the <i>Self-Service Administration and Operations Guide</i>. • Provision, scale, and manage multi-VM applications: You can design complex blueprints that span multiple environments and teams using Nutanix Self-Service. For more information, see Self-Service Blueprints Overview in the <i>Self-Service Administration and Operations Guide</i>. • Further bucket your infrastructure: You can configure environments to segregate your infrastructure for cases such as development, testing, production, disaster recovery, and so on. For more information, see Environments in Projects on page 50.
Top Users	Display the users that consume the maximum resources within the project. You can filter the users based on the vCPU, memory, or storage consumption using the resource switcher on the tile. You can click View all Users to navigate to the Users & Groups tab and view the user details such as the group that the user belongs and the resources that the user is consuming. For more information, see Project Users & Groups Tab on page 40.
Quota Consumption	Displays the percentage of resource quotas (vCPU, memory, and storage) that you have defined for the Nutanix accounts associated with the project. You can click Update Quotas to navigate to the Infrastructure tab and edit the quota values for the account. When Self-Service and policy engine are enabled, you can also define quota values for any VMware accounts that you add to your project. You can click View Quota Utilization Report to view the resource utilization, quota utilization, and application quota utilization at the project level. To get better quota consumption data with the upgraded quota capabilities, it is recommended to enable Self-Service and then the policy engine. You do not need to license Self-Service to enable policy engine. For more information about policy engine, see Policy Engine Overview in the <i>Self-Service Administration and Operations Guide</i>.

Tiles	Description
Project Expenditure Trends	Display a graph to report the expenditure trend within the project from an hour, day, or month viewpoint when Self-Service is enabled in your Prism Central instance. Expenditure trend is generated when you enable Showback in Self-Service. For more information, see Showback in the <i>Self-Service Administration and Operations Guide</i>. You can use the Generate Report option to generate a showback report. The showback report contains the cost data of the active applications within the project over the last 90 days. You can also use the Generate Report option to download the generated report.
Environments	Displays the environment usage within the project when you have configured environments in your project and have launched applications using those environments. You can sort the environment usage on the basis of compute, memory, or usage using the resource switcher on the tile. You can click View all Environments to navigate to the Environments tab and view or update your environment configuration. For more information on environment configuration, see Environments in Projects on page 50.
Tunnels	Displays the number of tunnels inherited from the account and the VPCs associated with the project.

Project Environments Tab

The **Environments** tab appears when you enable Self-Service in your Prism Central instance. The tab displays all the environments that you configured within the project.

An environment is a subset of a project that you configure to use during blueprint creation or application launch. For more information, see [Environments in Projects](#) on page 50.

You can use the icons at the top of the page to view the environment details in a list or tile mode. You can use the Search field to search environments on the page.

The tab displays the following details about an environment.

Table 7: Environments

Parameter	Description
Name	Displays the name of the environment. You can click the environment name to navigate to the Environment Configuration page where you can view or manage the configuration.
Created By	Displays the user (role) who created the environment.
Accounts	Displays the accounts associated with the environment.
Credentials	Displays the number of credentials associated with the environment.
Ready For	Displays the readiness of the environment for blueprint creation or application launch.

You can click the Settings icon to select, deselect, or reorder the columns in the list mode.

You can click the **Create Environment** button to create and configure an environment. For more information, see [Configuring Environments in a Project](#) on page 50.

You can click the vertical ellipsis next to an environment to update or delete the environment.

Project Infrastructure Tab

The **Infrastructure** tab allows you to add accounts to your project and displays the details of the added accounts. The tab also allows you to remove accounts from the project. The left pane on the tab lists all the accounts that you added. When you click an account, the right pane displays the account details.

You can click **Add Infrastructure** to add accounts to your project.

When Self-Service is disabled in your Prism Central instance, you can add your local Nutanix account and select clusters and subnets for the project. For more information, see [Adding Infrastructure in a Self-Service-Disabled Prism Central](#) on page 48.

When Self-Service is enabled, you can add your on-prem private clouds with AHV or ESXi or any public clouds (such as AWS, Azure, or GCP) that you configured in Self-Service. For more information, see [Adding Infrastructure in a Self-Service-Enabled Prism Central](#) on page 46.

For more information about account configuration in Self-Service, see [Provider Account Settings in Self-Service](#) in the *Self-Service Administration and Operations Guide*.

Project Policies Tab

The **Policies** tab appears when you enable Self-Service in your Prism Central instance. You must enable policy engine to view or manage policy configurations on the **Policies** tab.

The **Quotas** tab in the left pane provides a unified view of the resource quota limits that you defined for the project and the accounts within the project. For information about Quota Policy, see [Quota Policy Overview](#) on page 68 or [Managing Quota Limits for Projects](#) on page 69.

The **Snapshot** tab in the left pane provides the details of all snapshot policies you created within the project. For information on snapshot policy creation, see [Creating a Snapshot Policy](#) on page 71.

Project Tunnels Tab

The **Tunnels** tab appears when you enable Self-Service in your Prism Central instance and displays the number of tunnels inherited from the account and the VPCs associated with the project. For more information about tunnels, see [Tunnels for Orchestration within a VPC](#) in the *Self-Service Administration and Operations Guide*.

Project Users & Groups Tab

The **Users & Groups** tab displays a list of users and groups assigned to the project.

Table 8: Users & Groups

Parameter	Description
Name	Displays the Active Directory name (typically in the form of <i>name@domain</i>) associated with the project.
User group	Displays the group because of which the user is added to the project.
Total VMs	Displays the number of VMs owned by the user.
vCPU	Displays the number of vCPUs consumed by the user.

Parameter	Description
Memory	Displays the amount of memory (in GiB) consumed by the user.
Storage	Displays the amount of storage (in GiB) consumed by the user.

You can click **Add/Edit Users & Groups** to add or edit users or group of users to your project. For more information, see [Adding Users to a Project](#) on page 44.

Project Workloads Tab

The **Workloads** tab displays the overall workloads associated with the project. The Workloads tab has the following tabs that you can access in the left pane.

Table 9: Tabs Within the Workloads Tab

Tabs	Description
VMs	The VMs tab shows the total number of VMs and the number of local and remote VMs associated with the project. The local VM count indicates the number of VMs associated with the local AHV clusters that you add to the project. The remote VM count indicates the number of VMs associated with the remote PC (remote clusters) that you add to the project. The VMs tab displays the details of only those VMs that are associated with the local AHV clusters. You can view the following details: • Name: Displays the name of the VM. You can click the VM name to navigate to the VM Details page where you can view and update the VM configuration. • Cluster Name: The name of the cluster to which the VM is running. • State: Displays whether the VM is currently On (green circle) or off (red circle). • Owner: Displays the user who owns the VM. • vCPU: Displays the number of vCPUs allocated to the VM. • Memory: Displays the amount of memory (in GiB) allocated to the VM. • Storage: Displays the amount of disk space (in GiB) allocated to the VM.

Tabs	Description
Apps	This tab appears only when Self-Service is enabled in your Prism Central instance and you have launched applications ((on-prem AHV and ESXi) using the blueprints that are configured within the project. You can view the following details of an application on this tab. • Name: Displays the name of the application. You can click the application name to navigate to the Application Details page where you can view or manage the application configuration. • Source Blueprint: Displays the name of the blueprint that was used to launch the application. You can click the blueprint name to navigate to the Blueprint details page where you can view or manage the blueprint configuration. • State: Displays the state of the application. The state can be Running, Stopped, or Error. • Owner: Displays the user who owns the application. • Created on: Displays the date and timestamp of the application creation. • Last Updated at: Displays when the application was updated the last time. • Cost: Displays the service cost of the application incurred in the last 30 days. The information appears for Nutanix and VMware when the Showback is enabled in Self-Service.

Project Usage Tab

The **Usage** tab provides the graphical representation of the resource usage within the project. You can select the time period for the graphs by selecting the 1 day or 1 week option in the **Show** list. The tab displays the following graphs:

- The **vCPU Usage** graph displays a rolling time interval monitor of project vCPU usage. Placing the cursor anywhere on the horizontal axis displays the value at that time.
- The **Memory Usage** graph displays a rolling time interval monitor of project memory usage.
- The **Storage Usage** graph displays a rolling time interval monitor of project storage usage.

The graphs appears for on-prem AHV clusters only.

Creating a Project

You create a project to provide logical groupings of user roles for managing resource usage within your organization.

About this task

Use this procedure to define the basic setup of your project. After you define the basic setup, you can further add users to your project or configure other application management features such as accounts, environments, and policies.

Note: You must enable Self-Service in your Prism Central instance to configure application management features.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Projects** in the navigation bar.
4. Click **+ Create Project** on the Project Summary page to create the project.
5. Type a name for the project in the **Project Name** field.
6. Type a description for the project in the **Description** field.
7. Select the **Enable Directory or Provider Shortlist** checkbox to shortlist Active Directories (ADs) or Identity Providers (IDPs) in the project and then select ADs or IDPs in the **Directory or Provider** dropdown menu.

After you shortlist ADs or IDPs in the project, project admins can add users or groups only from the shortlisted ADs or IDPs. By default, the option to shortlist ADs and IDPs remains disabled.

You can use this feature to ensure that the project do not show all ADs or IDPs during multi-tenancy.

8. Select an admin for the project from the **Project Admin** dropdown menu. This step is optional.
Type the first few letters of the project admin to get the list of matching results in the dropdown menu.

Note: The project automatically adds you as a Project Admin when you create it. You can add other users after you configure the basic setup of the project. For more information, see [Adding Users to a Project](#) on page 44.

If you have multiple directories or identity providers configured, ensure that the directory or identity provider from which you want to add the project admin is selected. Do the following:

- Click the Settings icon to open the Search Directories window.
- Select the option for the directory service or identity provider that you want to use to add an admin.
- Click **Save**.

Note: Local users are not supported in a project. You can only add an admin from your configured directory or identity provider.

9. Check the **Allow Collaboration** checkbox to allow project users to collaboratively manage VMs and applications within the project.
The **Allow Collaboration** checkbox appears when you add your first user to the project. By default, the **Allow Collaboration** checkbox is selected and enables a project user to view and manage VMs and applications of other users in the same project. If you clear the **Allow Collaboration** checkbox, project users can manage only the VMs and applications that they create. You cannot change this configuration after you add the first user and save the project. To change the configuration, you need to remove all users from the project.
10. Click **Create**.
The **Dashboard** tab for the project appears.
11. Add users or groups to your project on the **Users & Groups** tab. For more information, see [Adding Users to a Project](#) on page 44.
12. Add infrastructure to your project on the **Infrastructure** tab. For more information, see [Adding Infrastructure in a Self-Service-Disabled Prism Central](#) on page 48.

13. If you have enabled Self-Service in your Prism Central instance, you can configure infrastructure, environments, and policies in your project. Do the following:
 - a. Use the **Infrastructure** tab to add different accounts to your project. For more information, see [Adding Infrastructure in a Self-Service-Enabled Prism Central](#) on page 46.
 - b. Use the **Environments** tab to add credentials and configure VMs for the provider accounts that you selected for your project. For more information, see [Configuring Environments in a Project](#) on page 50.
 - c. Use the **Policies** tab to define your quota and snapshot policies. For more information, see [Quota Policy Overview](#) on page 68 and [Creating a Snapshot Policy](#) on page 71.
14. Click **Save** on the Project Details page.

Adding Users to a Project

Use projects to assign roles to users or user groups. Based on the role assigned, users are allowed to configure and manage resources within the organization.

About this task

For information on the user roles, see [Roles Management](#) on page 75 and [Role-Based Access Control in Self-Service](#) in the *Self-Service Administration and Operations Guide*.

Note: Local users are not supported in a project. You can only add users from your configured directory or identity provider.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Projects** in the navigation bar.
4. Do one of the following:
 - » Click **+ Create Project** to create a new project and add users to the project. For information on creating a project, see [Creating a Project](#) on page 42.
 - » Click a project name in the list of existing projects to add users to that project.
5. Do the following to shortlist Active Directories (ADs) or Identity Providers (IDPs) in a project.
 - a. On the **Users & Groups** tab, click **Directory Settings**.
 - b. On the Directory Settings window, select the **Enable Directory or Provider Shortlist** checkbox.
By default, the option to shortlist ADs and IDPs remains disabled.
 - c. From the **Directories or Providers** dropdown menu, select the ADs or IDPs that you want to shortlist in the Project.
You can select multiple ADs or IDPs to shortlist. Shortlisting ADs or IDPs in the project allows project admins to add only those users or user groups to the project that belong to the shortlisted ADs or IDPs. This feature is specifically helpful when you want to ensure that the project do not show all ADs or IDPs during multi-tenancy.
 - d. Click **Save**.

6. On the **Users & Groups** tab, click **Add/Edit Users & Groups**.

The **Add/Edit Users & Groups** window appears.

7. Click the **Directory and Provider** settings icon.

The **Search Directories** window appears with options to select the AD or IDP from which you want to add the user or user group to the project.

You can add users and user groups from multiple ADs or IDPs configured in your Prism Central and shortlisted in the project. After you add users or user groups from one AD or IDP, you can select another AD or IDP in the Search Directories window and add users or user groups belonging to the that AD or IDP.

8. Select the radio button for the directory service or identity provider that you want to use to add users or user groups.
9. To save the changes and close the **Search Directories** window, click **Save**.
10. To add a user or group with a role in the project, do the following.

- a. Click **+ Add User**.

A blank row is added with the **Name** and the **Role** columns.

- b. In the **Name** column, enter the directory name or identify provider name of the user or a user group belonging to the AD or IDP you selected.

Note:

- Type the name of the AD user to search and add the user.
- For IDP user names, enter the complete IDP user name in the field and press Enter.
- The **Name** column accepts user group names. If you have created an IDP user group with a name, type the user group name in the **Name** column to add the group to the project.

- c. In the **Role** column, select a user role from the list.

The default value in the **Role** column is **Project Admin**. You can select a value from the list to change the user role.

The **Allow Collaboration** checkbox appears when you add the first user to your project. By default, the **Allow Collaboration** checkbox is selected and enables a project user to view and manage VMs and applications of other users in the same project. If you clear the **Allow Collaboration** checkbox, project users can manage only the VMs and applications that they create. You cannot change this configuration after you add the first user and save the project. To change the configuration, you need to remove all users from the project.

- d. Click **Save Users and Project**.

The **Users, Groups & Roles** section on the Overview tile displays the number of users you added to the project.

Note:

- If you add a group to a project, users in the group might not appear in the project members list until they log in.
- Nested groups (groups within a group) are not supported. For example, if a selected group (Group1) includes a nested group (Group1.1) along with individual names, only individual names are added to the project. The group members of Group 1.1 are not added to the project.

Infrastructure in Projects

Infrastructure allows you to add accounts to your project. When Self-Service is disabled, you can add only the local Nutanix account to your project. For more information, see [Adding Infrastructure in a Self-Service-Disabled Prism Central](#) on page 48). However, you can enable Self-Service and add multiple accounts of the same provider or different providers you configured in Self-Service to your projects. For more information, see [Adding Infrastructure in a Self-Service-Enabled Prism Central](#) on page 46).

Infrastructure supports accounts of the following providers:

- Nutanix
- VMware
- AWS
- Azure
- GCP
- Kubernetes

When you add Nutanix accounts to your project, you can allow clusters and their corresponding VLAN subnets. A VLAN subnet is bound to a Prism Element cluster. When you use the VLAN subnet to provision a VM, the VM is placed on that Prism Element.

You can also allow clusters and their corresponding overlay subnets. An overlay subnet can span from a few clusters to all the clusters of your Prism Central. Therefore, when you configure your Nutanix account in your project, you must allow clusters before allowing the subnets.

Note:

- Allowing clusters before their subnets enables you to have projects where you can allow VPCs and their corresponding subnets without allowing any VLAN subnets.
- You can also define resource quota limits for the Nutanix and VMware accounts within your project. However, to define resource quota limits for the accounts, you must enable policy engine from the Admin Center Settings. For more information, see [Enabling the Policy Engine](#) on page 84 and [Managing Quota Limits for Projects](#) on page 69.

Adding Infrastructure in a Self-Service-Enabled Prism Central

When Self-Service is enabled in your Prism Central, you can add multiple provider accounts or credential provider accounts that you configured in Self-Service to your project. You can also define resource quota limits for quota checks for Nutanix and VMware accounts within the project if you enable policy engine in your Prism Central instance.

About this task

Use this procedure to add provider accounts or credential provider accounts to your project.

Before you begin

- Ensure that you have configured your provider accounts or credential provider accounts in Self-Service. For information on configuring a provider account, see [Provider Account Settings in Self-Service](#) in the *Self-Service Administration and Operations Guide*. For information on configuring a credential provider account, see [Credential Settings in Self-Service](#) in the *Self-Service Administration and Operations Guide*.
- For resource quota limit definition, ensure that you have enabled the policy engine in the Admin Center Settings. For information on enabling policy engine, see [Enabling the Policy Engine](#) on page 84.

Procedure

1. Log in to Prism Central as an administrator.
 2. Select **Admin Center** in the Application Switcher.
 3. Click **Projects** in the navigation bar.
 4. Do one of the following:
 - » Click **+Create Project** to create a new project and add providers to the project. For information on creating a project, see [Creating a Project](#) on page 42.
 - » Click a project name in the list of existing projects to add providers to that project.
 5. Click the **Infrastructure** tab.
 6. Click **Add Infrastructure** and select a provider account or a credential provider account from the list.
 7. If you selected a credential provider account, then go to step 11.
 8. If you selected a Nutanix account, then do the following to configure resources in the project.
 - a. Click **Configure Resources**.
 - b. On the **Select Clusters and VLANs** tab, select the cluster that you want to allow in the project from the **Select clusters to be added to this project** list.

Selecting a cluster for the account is mandatory. You must select a cluster irrespective of whether you want to allow VLAN subnets or not.
 - c. Under Select VLANs for the above clusters section, click **Select VLANs** to view and select the VLANs that you want to allow in the project. This step is optional.
 - d. Click **Select VPCs & Subnets**.

The **Select VPCs & Subnets** button or tab appears only on VPC-enabled setups.
 - e. On the **Select VPCs & Subnets** tab, select the VPC from the **Select VPCs to view overlay subnets below** list to view the associated subnets. This step is optional.

The Select overlay subnets section displays the overlay subnets associated with the VPC you selected.

If tunnel is not configured for the selected VPC, you can perform only basic operations, such as VM provisioning, on the VPC. To perform check log-in and orchestration, ensure that you create a tunnel for the VPC. For more information, see [Creating VPC Tunnels](#) in the *Self-Service Administration and Operations Guide*.
 - f. Under the Select overlay subnets section, select the overlay subnets that you want to allow in the project. This step is optional.
 - g. Do one of the following:
 - » For the local Nutanix account, click **Confirm and Select Default**, select a default VLAN subnet, view the configuration summary, and then click **Confirm**.
 - » For a remote PC account, click **Confirm**, view the configuration summary, and then click **Confirm**.
- The default VLAN subnet is used when you create a virtual machine in Prism Central.

Note: You can select one or more AHV clusters. You can allow one or more subnets per cluster. When you configure a blueprint, only the allowed networks and clusters appear for the user to select during network configuration. If the network selection is a runtime attribute, only the allowed networks and clusters are available to update while launching a blueprint.

9. If you selected a Nutanix account, then do the following to configure services in the project.

A Nutanix service refers to a specific software component or functionality that Nutanix provides to enhance the capabilities, performance, and management of the Nutanix infrastructure and software stack. Currently, you can connect a Nutanix Database Service (NDB) account to your Nutanix account in Self-Service to leverage the database management capabilities of NDB with the application management workflows. For more information, see [Connecting NDB to an Account](#) in the *Self-Service Administration and Operations Guide*.

 - a. In the Nutanix Services section, click **Configure Services**.
 - b. Under Add or Remote Accounts, select the NDB accounts that you want to allow in the project.

The list displays the name and the Server IP addresses of the NDB accounts that you connected to the selected Nutanix account.
 - c. Click **Next**.
 - d. View the list of selected accounts, and click **Confirm**.

The **NDB** section on the **Infrastructure** tab displays the list of selected NDB accounts with their Server IP addresses.
10. If you selected a Nutanix or VMware account, you can define resource quota limits.

The options to define resource quota limits appear only when you enable policy engine in the Admin Center Settings of Prism Central. For information on enabling policy engine, see [Enabling the Policy Engine](#) on page 84.

 - a. Select the **Quotas**.

The **vCPU**, **Memory**, and **Disk** fields are enabled.
 - b. Enter quota values for **vCPU**, **Memory**, and **Disk** for the selected clusters.

The **Available/Total** row shows the available resources quota and the total quota allocated to the provider. The **Physical Capacity** row shows the used and total physical capacity. Use these details while defining resource quota limits.
11. Click **Save**.

The Overview tile on the **Dashboard** tab displays the number of accounts you added to the project.

Adding Infrastructure in a Self-Service-Disabled Prism Central

When Self-Service is disabled in your Prism Central instance, you can add only the local Nutanix account and allow clusters and subnets in the project. You can also define resource quota limits for quota checks within the project if you enable policy engine in your Prism Central instance.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Projects** in the navigation bar.
4. Do one of the following:
 - » Click the **+Create Project** button to create a new project and add the Nutanix account to the project. For more information on creating a project, see [Creating a Project](#) on page 42.
 - » Click a project name in the list of existing projects to add the Nutanix account to that project.
5. Click the **Infrastructure** tab.
6. Click **Add Infrastructure** and select the local Nutanix account.

7. Do the following.
 - a. Click **Configure Resources**.
 - b. On the **Select Clusters and VLANs** tab, select the cluster that you want to allow in the project from the **Select clusters to be added to this project** dropdown menu.

Selecting a cluster for the account is mandatory. You must select a cluster irrespective of whether you want to allow VLAN subnets or not.
 - c. Under Select VLANs for the above clusters section, click the cluster and then **Select VLANs** to view and select the VLANs that you want to allow in the project.
 - d. Click **Select VPCs & Subnets**.

The **Select VPCs & Subnets** button or tab appears only on VPC-enabled setups.
 - e. On the **Select VPCs & Subnets** tab, select the VPC from the **Select VPCs to view overlay subnets below** list to view the associated subnets. This step is optional.

The Select overlay subnets section displays the overlay subnets associated with the VPC you selected.

If tunnel is not configured for the selected VPC, you can perform only basic operations, such as VM provisioning, on the VPC. To perform check log-in and orchestration, ensure that you create a tunnel for the VPC. For more information, see [Creating VPC Tunnel](#) in the *Self-Service Administration and Operations Guide*.
 - f. Under the Select overlay subnets section, select the overlay subnets that you want to allow in the project. This step is optional.
 - g. Click **Confirm and Select Default**.
 - h. Select a default subnet.

Default subnets are used when you create a VM from Prism Central.
 - i. Click **Confirm**.
8. To define resource quota limits for the Nutanix account, under the Quotas section, enter quota values for **vCPU**, **Storage**, and **Memory**.

The options to define resource quota limits appear only when you enable policy engine in the Admin Center Settings of Prism Central. For information on enabling policy engine, see [Enabling the Policy Engine](#) on page 84.
9. Click **Save**.

The Overview tile on the **Dashboard** tab displays the number of accounts you added to the project.

Modifying a Project

You can modify the description, users, infrastructure, or any application management specific details of your project.

About this task

Use this procedure to modify an existing project.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Projects** in the navigation bar.

4. On the Project Summary page, click the project that you want to modify.
The Project Details page appears. The page includes the same tabs and fields that appeared while creating the project you want to modify.
5. Update the field values as required and then click **Save**.

Deleting a Project

You can delete a project that is not associated with any application or blueprint. If the project is already used to create any applications or blueprints, you cannot delete the project. In such cases, a window appears displaying the association of the project with different applications or blueprints.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Projects** in the navigation bar.
4. On the Project Summary page, select the checkbox adjacent to the project that you want to delete.
5. From the **Action** dropdown menu, select **Delete**.
Prism Central checks the association of the project with entities such as applications, blueprints, runbooks, endpoints, and so on. In case the project is associated with any entities, the details of the associated entities are displayed in a window. You cannot delete the project until you remove the project from these entities. In case the project is not associated with the entities, a confirmation window appears where you can confirm to delete the project.
6. In the confirmation window, click **Delete**.

Environments in Projects

If Self-Service is enabled in your Prism Central instance, you can configure environments as the subset of your project.

To configure environments, you must add the accounts that you configured in Self-Service to your project and then configure environments using those accounts. Environment configuration involves configuring VMs and adding credentials for the accounts that you added to your project. For more information, see [Configuring Environments in a Project](#) on page 50.

You use a configured environment either during your blueprint creation or during an application launch. When you select a configured environment while launching an application from the marketplace, the values for the application launch are picked up from the selected environment. For more information, see [Launching a Blueprint from the Marketplace](#).

You can add multiple accounts of the same provider or accounts of different providers to your project. You can also configure multiple environments in a project and set one of the environments as the default environment for the project. You can optionally select one environment for each application profile in the blueprint.

Note: Environment is not supported for Kubernetes.

Configuring Environments in a Project

You configure environments as a part of your project creation so that you can use the configured environments when you create blueprints in Self-Service or launch marketplace applications in the Admin Center. You can configure multiple environments in your project.

Before you begin

- Ensure that you have enabled Self-Service in your Prism Central instance. For more information, see [Deploying Self-Service](#) on page 12.
- Ensure that you have configured a project and have added the required accounts to your project. For more information, see [Creating a Project](#) on page 42 and [Infrastructure in Projects](#) on page 46.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Projects** in the navigation bar.
4. On the Projects Summary page, click the project in which you want to configure the environment.
5. Click the **Environments** tab.
6. Click **Create Environment**.
7. On the **General** tab, do the following:
 - a. Type a name for the environment in the **Name** field.
 - b. Provide a description for the environment in the **Description** field.
 - c. Select **Set as default environment** checkbox to use the environment as a default environment to launch applications for the project.

The **Set as default environment** checkbox is selected by default for the first environment you configure in your project.
 - d. Click **Next**.
8. On the **Infrastructure** tab, do the following:
 - a. Click the **Select Infrastructure** dropdown menu, and select an account.

The **Select Infrastructure** dropdown menu shows the accounts that you added to your project. For more information, see [Infrastructure in Projects](#) on page 46.

If you have already selected an account, then you can click **+ Add Infrastructure** to select and add more than one provider accounts to the environment.
 - b. Expand the **VM Configuration** section to configure the virtual machine details for the environment.
 - For information on how to configure a VM for Nutanix, see [Configuring Nutanix Environment](#) on page 52.
 - For information on how to configure a VM for AWS, see [Configuring AWS Environment](#) on page 58.
 - For information on how to configure a VM for VMware, see [Configuring VMware Environment](#) on page 54.
 - For information on how to configure a VM for GCP, see [Configuring GCP Environment](#) on page 64.
 - For information on how to configure a VM for Azure, see [Configuring Azure Environment](#) on page 60.
 - c. Click **Next**.

9. On the **Credentials** tab, add credentials for your environment. For more information, see [Adding Credentials to the Environment](#) on page 67.
10. Click **Save Environment & Project**.
The **Overview** tile on the **Dashboard** tab displays the number of environments you configured for the project.
You can go to the **Environments** tab to view the details of each environment you configured for your project.

Configuring Nutanix Environment

Environment configuration involves configuring VMs and adding credentials for the accounts that you added to your project.

About this task

Use this procedure to configure environment variables with a Nutanix account.

Before you begin

Ensure that you have configured a project and selected a Nutanix account. For more information, see [Creating a Project](#) on page 42 and [Configuring Environments in a Project](#) on page 50.

Procedure

1. On the **Infrastructure** tab of the Create Environment page, click the **Select Infrastructure** dropdown menu, and add a Nutanix account.
2. Select the account you added in the left pane.
The Resource Configuration section displays the cluster, VLAN subnets, and overlay subnets you selected in the project for the account.
When you configure the environment for the first time, the cluster and subnets that you selected for the project is selected for the environment by default. However, you can click **Configure Resources** to configure resources specific to the environment. You can select additional subnets for the environment or remove any subnets that you have already selected for the project.
3. In the right pane, expand the **VM Configuration** section and select either **Windows** or **Linux** as the operating system for the VM.
4. In the **Cluster** dropdown menu, select the cluster where you want to place the VM.
The **Cluster** dropdown menu displays the clusters that you allowed in the project.
The VLAN subnets have direct association with the cluster. When you select a VLAN subnet under the Network Adapters section, the associated cluster is auto-populated in the **Cluster** dropdown menu. However, if you intend to use overlay subnets, you must select the cluster in the list.
5. Enter a name of the VM in the **VM Name** field.
You can use Self-Service macros to provide a unique name to the VM. For example, `vm-@{calm_time}@`. For more information on Self-Service macros, see [Macros Overview](#).
6. Configure the processing unit of the VM by entering the number of vCPU, cores of each vCPU, and total memory in GB of the VM in the **vCPU**, **cores per vCPU**, and **Memory (GiB)** fields.

7. If you want to customize the default OS properties of the VM, select the **Guest Customization**.
Guest customization allows you to modify the properties of the VM operating system. You can prevent conflicts that might result due to the deployment of virtual machines with identical settings, such as duplicate VM names or same SID. You can also change the computer name or network settings by using a custom script.
 - a. Select **Cloud-init** for Linux or **SysPrep** for Windows, and enter or upload the script in the **Script** panel.
For Sysprep, you must use double back slash for all escape characters. For example, \\v.
 - b. For Sysprep script, click **Join a Domain** checkbox and configure the following fields.
 - Enter the domain name of the Windows server in the **Domain Name** field.
 - Select a credential for the Windows VM in the **Credentials** dropdown menu. You can also add new credentials.
 - Enter the IP address of the DNS server in the **DNS IP** field.
 - Enter the DNS search path for the domain in the **DNS Search Path** field.
8. To add a virtual disk to the VM, click the **+** icon next to the **DISKS** section and do the following.
 - a. Select the device for the image from the **Device Type** dropdown menu.
You can select either **CD-ROM** or **Disk**.
 - b. Select the device bus from the **Device Bus** dropdown menu.
You can select **IDE** or **SATA** for CD-ROM and **SCSI, IDE, PCI**, or **SATA** for DISK.
 - c. From the **Operations** dropdown menu, select one of the following.
 - » To allocate the disk memory from the storage container, select **Allocate on Storage Container**.
 - » To clone an image from the disk, select **Clone from Image Service**.
 - d. If you selected **Allocate on Storage Container**, enter the disk size in GB in the **Size (GiB)** field.
 - e. If you selected **Clone from Image Service**, select the image you want to add to the disk in the **Image** field.
All the images that you uploaded to Prism Central are available for selection. For more information about image configuration, see the [Image Management](#) section in the *Prism Central Infrastructure Guide*.
 - f. Select the **Bootable** checkbox for the image that you want to use to start the VM.

Note: You can add more than one disk and select the disk with which you want to boot up the VM.
9. Under the **Boot Configuration** section, select a firmware type to boot the VM.
 - » To boot the VM with legacy BIOS firmware, select **Legacy BIOS**.
 - » To boot the VM with UEFI firmware, select **UEFI**. UEFI firmware supports larger hard drives, faster boot time, and provides more security features.
10. (For GPU-enabled clusters only) To configure a vGPU, click the **+** icon under the **vGPUS** section and do the following:
 - a. From the **Vendor** dropdown menu, select the GPU vendor.
 - b. From the **Device ID** dropdown menu, select the device ID of the GPU.
 - c. From the **Mode** dropdown menu, select the GPU mode.

11. Under the **Categories** section, select a category in the **Key: Value** dropdown menu.
Use this option to tag your VM to a defined category in Prism Central. The list options are available based on your Prism Central configuration. If you want to protect your application by a protection policy, select the category defined for the policy in your Prism Central. Categories list is available only for Nutanix.
12. To add a network adapter, click the **+** icon next to the Network Adapters (NICS) field.
The **NIC** dropdown menu shows all the VLAN and overlay subnets. The VLAN subnets have direct association with the cluster. Therefore, when you select a VLAN subnet, the associated cluster is auto-populated in the **Cluster** dropdown menu.
The NICs of a VM can either use VLAN subnets or overlay subnets. For example, if you select an overlay subnet in NIC 1 and then add NIC 2, the NIC 2 list displays only the overlay subnets.
If you select a VLAN subnet in NIC 1, any subsequent VLAN subnets belong to the same cluster. Similarly, if you select an overlay subnet, all subsequent overlay subnets belong to the same VPC.
13. Configure the connection in your environment. For more information, see [Configuring Check Log-in for your Environment](#) on page 66.
14. Click **Next**.
15. Add credentials for the environment. For more information, see [Adding Credentials to the Environment](#) on page 67. This step is optional.
16. Click **Save Environment & Project**.

What to do next

You can use the environment details while configuring a blueprint for Nutanix or launching a blueprint.

Configuring VMware Environment

Environment configuration involves configuring VMs and adding credentials for the accounts that you added to your project.

About this task

Use this procedure to configure environment variables for VMware.

Before you begin

Ensure that you have configured a project and selected a VMware account. For more information, see [Creating a Project](#) on page 42 and [Configuring Environments in a Project](#) on page 50.

Procedure

1. On the **Infrastructure** tab of the Create Environment page, click the **Select Infrastructure** dropdown menu, and add a VMware account.
2. Select the account you added in the left pane.
3. In the right pane, expand the **VM Configuration** section, and select either **Windows** or **Linux** as the operating system for the VM.

4. Select the **Compute DRS Mode** checkbox to enable load sharing and automatic VM placement.
Distributed Resource Scheduler (DRS) is a utility that balances computing workloads with available resources in a virtualized environment. For more information about DRS mode, see the *VMware documentation*.
 - » If you selected **Compute DRS Mode**, then select the cluster where you want to host your VM from the **Cluster** dropdown menu.
 - » If you have not selected **Compute DRS Mode**, then select the host name of the VM from the **Host** dropdown menu.

5. Do one of the following:

- » Select **VM Templates** and then select a template from the **Template** dropdown menu.

Templates allow you to create multiple virtual machines with the same characteristics, such as resources allocated to CPU and memory or the type of virtual hardware. Templates save time and avoid errors when configuring settings and other parameters to create VMs. The VM template retrieves the list options from the configured vCenter.

Note:

- Install the VMware Tools on the Windows templates. For Linux VMs, install *Open-vm-tools* or *VMware-tools* and configure the *Vmtoolsd* service for automatic start-up.
- Support for *Open-vm-tools* is available. When using *Open-vm-tools*, install Perl for the template.
- Do not use SysPrepped as the Windows template image.
- If you select a template that has unsupported version of VMware Tools, then a warning appears stating VMware tool or version is unsupported and could lead to VM issues.
- You can also edit the NIC type when you use a template.

For more information, refer to VMware KB articles.

- » Select **Content Library**, a content library in the **Content Library** dropdown menu, and then select an OVF template or VM template from the content library.

A content library stores and manages content (VMs, vApp templates, and other types of files) in the form of library items. A single library item can consist of one file or multiple files. For more information about the vCenter content library, see the *VMware Documentation*.

Caution: Content Library support is currently a technical preview feature in Self-Service. Do not use any technical preview features in a production environment.

6. If you want to use the storage DRS mode, then select the **Storage DRS Mode** checkbox and a datastore cluster from the **Datastore Cluster** dropdown menu.
The datastore clusters are referred as storage pod in vCenter. A datastore cluster is a collection of datastores with shared resources and a shared management interface.
7. If you do not want to use storage DRS mode, then do not select the **Storage DRS Mode** checkbox, and select a datastore from the **Datastore** dropdown menu.

8. In the **VM Location** field, specify the location of the folder in which the VM must be created when you deploy the blueprint. Ensure that you specify a valid folder name already created in your VMware account.

To create a subfolder in the location you specified, select the **Create a folder/directory structure here** and specify a folder name in the **Folder/Directory Name** field.

Note: Self-Service prefers the VM location specified in the environment you select while launching an application. For example, you specify a subfolder structure as the VM location in the blueprint and the top-level folder in the environment. When you select this environment while launching your application, Self-Service considers the VM location you specified in the environment and creates the VM at the top-level folder.

Select the **Delete empty folder** checkbox to delete the subfolder you create within the specified location, in case the folder does not contain any VM resources. This option helps you to keep a clean folder structure.

9. Enter the instance name of the VM in the **Instance Name** field.
This field is pre-populated with a macro as suffix to ensure name uniqueness. The service provider uses this name as the VM name.
10. Select the **CPU Hot Add** checkbox if you want to increase the VCPU count of a running VM.
Support for **CPU Hot Add** depends on the Guest OS of the VM.
11. Update the **vCPUs** and **Cores Per Socket** count.
The number of sockets is calculated by dividing the total number of vCPUs by Cores Per Socket. For example, if the number of vCPUs is 4 and the Cores Per Socket is 2, then the number of sockets will be 2.
12. Select the **Memory Hot Plug** checkbox if you want to increase the memory of a running VM.
Support for **Memory Hot Plug** depends on the Guest OS of the VM.
13. Update the memory in the **Memory** field.
14. Under **Controller**, click **+** to add the type of controller.
You can select either SCSI or SATA controller. You can add up to three SCSI and four SATA controllers.

15. Under the **Disks** section, click the **+** icon to add vDisks and do the following:
 - a. Select the device type from the **Device Type** dropdown menu.
You can either select **CD-ROM** or **DISK**.
 - b. Select the adapter type from the **Adapter Type** dropdown menu.
You can select **IDE** for CD-ROM.
You can select **SCSI**, **IDE**, or **SATA** for DISK.
 - c. Enter the size of the disk in GiB.
 - d. In the **Location** field, select the disk location.
 - e. If you want to add a controller to the vDisk, select the type of controller in the **Controller** dropdown menu to attach to the disk.

Note: You can add either SCSI or SATA controllers. The available options depend on the adapter type.

- f. In the **Disk mode** dropdown menu, select the type of the disk mode. Your options are:
 - » **Dependent:** Dependent disk mode is the default disk mode for the vDisk.
 - » **Independent - Persistent:** Disks in persistent mode behave like conventional disks on your physical computer. All data written to a disk in persistent mode are written permanently to the disk.
 - » **Independent - Nonpersistent:** Changes to disks in nonpersistent mode are discarded when you shut down or reset the virtual machine. With nonpersistent mode, you can restart the virtual machine with a virtual disk in the same state every time. Changes to the disk are written to and read from a redo log file that is deleted when you shut down or reset.
16. Under the **Tags** section, select tags from the **Category: Tag pairs** field.
You can assign tags to your VMs so you can view the objects associated with your VMs in your VMware account. For example, you can create a tag for a specific environment and assign the tag to multiple VMs. You can then view all the VMs that are associated with the tag.
17. If you want to customize the default OS properties of the VM, then click **Enable** under **VM Guest Customization** and select a customization from the **Predefined Guest Customization** dropdown menu.

18. If you do not have any predefined customization available, select **None** and do the following.
 - a. Select **Cloud-init** or **Custom Spec**.
 - b. If you selected **Cloud-init**, enter or upload the script in the **Script** field.
 - c. If you have selected **Custom Spec**, enter the network details for the VM in the following fields:
 - Enter the hostname in the **Hostname** field.
 - Enter the domain in the **Domain** field.
 - Select timezone from the **Timezone** dropdown menu.
 - Select **Hardware clock UTC** to enable hardware clock UTC.
 - Click the **+** icon to add network settings.
 - To automatically configure DHCP server, enable **Use DHCP** and then skip to the **DNS Setting** section.
 - Enter a name for the network configuration you are adding to the VM in the **Setting name** field. Settings name is the saved configuration of your network that you want to connect to your VM.
 - Enter values in the **IP Address**, **Subnet Mask**, **Default Gateway**, and **Alternative Gateway** fields.
 - Under the **DNS Settings** section, enter values in the **DNS Primary**, **DNS Secondary**, **DNS Tertiary**, and **DNS Search Path**.
19. Configure the connection in your environment. For more information, see [Configuring Check Log-in for your Environment](#) on page 66.
20. Click **Next**.
21. Add credentials for the environment. For more information, see [Adding Credentials to the Environment](#) on page 67. This step is optional.
22. Click **Save Environment & Project**.

What to do next

You can use the environment details while configuring a blueprint for VMware or launching a blueprint.

Configuring AWS Environment

Environment configuration involves configuring VMs and adding credentials for the accounts that you added to your project.

About this task

Use this procedure to configure environment variables with an AWS account.

Before you begin

Ensure that you have configured a project and selected an AWS account. For more information, see [Creating a Project](#) on page 42 and [Configuring Environments in a Project](#) on page 50.

Procedure

1. On the **Infrastructure** tab of the Create Environment page, click the **Select Infrastructure** dropdown menu, and add an AWS account.
2. Select the account you added in the left pane.

3. In the right pane, expand the **VM Configuration** section and select either **Windows** or **Linux** as the operating system for the VM.
4. Enter the name of the instance in the **Instance Name** field.

This field is pre-populated with a macro as suffix to ensure name uniqueness. The service provider uses this name as the VM name.
5. Select the **Associate Public IP Address** checkbox to associate a public IP address with your AWS instance.

If you do not select the **Associate Public IP Address** checkbox, ensure that the AWS account and Self-Service are on the same network for the scripts to run.
6. Select an AWS instance type from the **Instance Type** dropdown menu.

Instance types comprise varying combinations of CPU, memory, storage, and networking capacity and give you the flexibility to choose the appropriate mix of resources for your applications. Each instance type includes one or more instance sizes, allowing you to scale your resources to the requirements of your target workload. The dropdown menu displays the instances that are available in the AWS account. For more information, see *AWS documentation*.
7. Select the region from the **Region** dropdown menu and configure the following.

Note: The dropdown menu displays the regions which are selected while configuring the AWS setting.

- a. Select the availability zone from the **Availability Zone** dropdown menu.

An Availability Zone is one or more discrete data centers with redundant power, networking, and connectivity in an AWS region. Availability Zones allow you to operate production applications and databases that are more highly available, fault tolerant, and scalable than would be possible from a single data center.
 - b. Select the machine image from the **Machine Image** dropdown menu.

An Amazon Machine Image is a special type of virtual appliance that is used to create a virtual machine within the Amazon Elastic Compute Cloud. It serves as the basic unit of deployment for services delivered using EC2.
 - c. Select the IAM role from the **IAM Role** dropdown menu.

An IAM role is an AWS Identity and Access Management entity with permissions to make AWS service requests.
 - d. Select the key pairs from the **Key Pairs** dropdown menu.

A key pair (consisting of a private key and a public key) is a set of security credentials that you use to prove your identity when connecting to an instance.
 - e. Select the VPC from the **VPC** dropdown menu.

Amazon Virtual Private Cloud (Amazon VPC) allows you to provision a logically isolated section of the AWS cloud where you can launch AWS resources in your defined virtual network.

 - Select the **Include Classic Security Group** checkbox to enable security group rules.
 - Select security groups from the **Security Groups** dropdown menu.
8. Enter or upload the AWS user data in the **User Data** field.
 9. Enter the AWS tags in the **AWS Tags** field.

AWS tags are key and value pair to manage, identify, organize, search for, and filter resources. You can create tags to categorize resources by purpose, owner, environment, or other criteria.

10. Under the **Storage** section, configure the following to boot the AWS instance with the selected image.
 - a. In the **Device** field, select the device to boot the AWS instance. The available options are based on the image you have selected.
 - b. In the **Size (GiB)** field, enter the required size for the bootable device.
 - c. In the **Volume Type** dropdown menu, select the volume type. You can select either **General Purpose SSD**, **Provisioned IOPS SSD**, and **EBS Magnetic HDD**.
For more information on the volume types, see *AWS documentation*.
 - d. Select the **Delete on termination** checkbox to delete the storage when the instance is terminated. This step is optional.You can also add more secondary storages by clicking the + icon next to the **Storage** section.
11. Configure the connection in your environment. For more information, see [Configuring Check Log-in for your Environment](#) on page 66.
12. Click **Next**.
13. Add credentials for the environment. For more information, see [Adding Credentials to the Environment](#) on page 67. This step is optional.
14. Click **Save Environment & Project**.
The **Environments** tile on the **Overview** tab displays the number of environments you configured for the project. You can go to the **Environments** tab to view the details of the environment you configured for your project.

What to do next

You can use the environment details while configuring a blueprint for AWS or launching a blueprint.

Configuring Azure Environment

Environment configuration involves configuring VMs and adding credentials for the accounts that you added to your project.

About this task

Use this procedure to configure environment variables for Azure.

Before you begin

- Ensure that the following entities are already configured in the Azure account.
 - Resource group
 - Availability set
 - Network security group
 - Virtual network
 - Vault certificates
- Ensure that you have configured a project and selected an Azure account. For more information, see [Creating a Project](#) on page 42 and [Configuring Environments in a Project](#) on page 50.

Procedure

1. On the **Infrastructure** tab of the Create Environment page, click the **Select Infrastructure** dropdown menu, and add a Azure account.
2. Select the account you added in the left pane.
3. In the right pane, expand the **VM Configuration** section, and select either **Windows** or **Linux** as the operating system for the VM.
4. Under **VM Configuration**, enter the instance name of the VM in the **Instance Name** field.
This field is pre-populated with a macro as suffix to ensure name uniqueness. The service provider uses this name as the VM name.
5. Select a resource group from the **Resource Group** dropdown menu or select the **Create Resource Group** checkbox to create a resource group.
Each resource in Azure must belong to a resource group. A resource group is simply a logical construct that groups multiple resources together so you can manage the resources as a single entity. For example, you can create or delete resources as a group that share a similar life cycle, such as the resources for an n-tier application.
The **Resource Group** dropdown menu displays the resource groups that are associated with the subscriptions you selected in your Azure account. In case you have not selected any subscriptions, Self-Service considers all the subscriptions that are available in the Azure service principal to display the resource groups. Each resource group in the dropdown menu also displays the associated subscription.
6. If you selected a resource group from the **Resource Group** dropdown menu, then do the following:
 - a. Select the geographical location of the datacenter from the **Location** dropdown menu.
 - b. Select **Availability Sets** or **Availability Zones** from the **Availability Option** dropdown menu.
You can then select an availability set or availability zone. An availability set is a logical grouping capability to ensure that the VM resources are isolated from each other to provide High Availability if deployed within an Azure datacenter. An availability zone allows you to deploy your VM into different datacenters within the same region.
 - c. Select the hardware profile as per your hardware requirements from the **Hardware Profile** dropdown menu.
The number of data disks and NICs depends upon the selected hardware profile. For information about the sizes of Windows and Linux VMs, see Windows and Linux Documentation.
7. If you selected the **Create Resource Group** checkbox to create a resource group, then do the following:
 - a. Select a subscription associated to your Azure account in the **Subscription** field.
 - b. Enter a unique name for the resource group in the **Name** field.
 - c. Select the geographical location of the datacenter that you want to add to the resource group in the **Location** dropdown menu.
 - d. Under **Tags**, enter a key and value pair in the **Key** and **Value** fields respectively.
Tags are key and value pairs that enable you to categorize resources. You can apply a tag to multiple resource groups.
 - e. If you want to automatically delete a resource group that has empty resources while deleting an application, click the **Delete Empty Resource Group** checkbox.
 - f. Specify the location and hardware profile.

8. Under the **Secrets** section, click the **+** icon and do the following:
 - a. Enter a unique vault ID in the **Vault ID** field.
These certificates are installed on the VM.
 - b. Under **Certificates**, click the **+** icon.
 - c. Enter the URL of the configuration certificate in the **URL** field.
The URL of the certificate is uploaded to key vault as a secret.
 - d. Enter the certificate store for the VM in the **Store** field.
 - For Windows VMs, specify the certificate store on the virtual machine to which the certificate is added. The specified certificate store is implicitly created in the LocalMachine account.
 - For Linux VMs, the certificate file is placed under the /var/lib/waagent directory with the file name <UppercaseThumbprint>.crt for the X509 certificate file and <UppercaseThumbpring>.prv for private key. Both of these files are .pem formatted.
9. (For Windows) Select the **Provision Windows Guest Agent** checkbox.
This option indicates whether or not to provision the virtual machine agent on the virtual machine. When this property is not specified in the request body, the default behavior is to set it to true. This ensures that the VM Agent is installed on the VM, and the extensions can be added to the VM later.
10. (For Windows) To indicate that the VM is enabled for automatic updates, select the **Automatic OS Upgrades** checkbox.
11. Under the **Additional Unattended Content** section, click the **+** icon and do the following:
 - a. Select a setting from the **Setting Name** dropdown menu.
You can select **Auto Logon** or **First Logon Commands**.

Note: Guest customization is applicable only on images that allows or support guest customization.
 - b. Enter or upload the XML content. See [Sample Auto Logon and First Logon Scripts](#).
12. Under the **WinRM Listeners** section, click the **+** icon and do the following:
 - a. Select the protocol from the **Protocol** dropdown menu.
You can select **HTTP** or **HTTPS**.
 - b. If you selected HTTPS, then select the certificate URL from the **Certificate URL** dropdown menu.
13. Under the **Storage Profile** section, select the **Use Custom Image** checkbox to use a custom VM image created in your subscription.
You can then select a custom image or publisher-offer-SKU-version from the **Custom Image** dropdown menu.
14. Under the **VM Image Details** section, select an image type in the **Source Image Type** dropdown menu.
You can select **Marketplace**, **Subscription**, or **Shared Image Gallery**.
Do one of the following:
 - » If you selected **Marketplace**, then specify the publisher, offer, SKU, and version for the image.
 - » If you selected **Subscription**, then select the custom image.
 - » If you selected **Shared Image Gallery**, then select the gallery and the image.

15. Under the **OS Disk Details** section, do the following:
 - a. Select the storage type from the **Storage Type** dropdown menu.
You can select **Standard HDD**, **Standard SSD**, or **Premium SSD**.
 - b. Select a disk storage account from the **Disk Storage** dropdown menu.
This field is available only when the **Use Custom Image** is enabled.
 - c. Select disk caching type from the **Disk Caching Type** dropdown menu.
You can select **None**, **Read-only**, or **Read write**.
 - d. Select disk create option from the **Disk Create Option** dropdown menu.
You can select **Attach**, **Empty**, or **From Image**.
16. Under the **Network Profile** section, add NICs as per your requirement and do the following for each NIC:
 - a. Select a security group from the **Security Group** dropdown menu.
 - b. Select application security groups from the **ASGs** dropdown menu.
Unlike Network Security Groups that are defined at the network level, Application Security Groups (ASGs) are logical entities that are defined at the application level. ASGs help to manage the VM security by grouping the VMs according to the applications that run on them and allow application-centric use of Network Security Groups. You can select multiple ASGs in a NIC.

Note:

 - ASGs are supported when you clone your applications and are also supported in snapshot and restore.
 - ASGs are not supported in VM configuration updates.
 - c. Select application security groups from the **ASGs** list.
Unlike Network Security Groups that are defined at the network level, Application Security Groups (ASGs) are logical entities that are defined at the application level. ASGs help to manage the VM security by grouping the VMs according to the applications that run on them and allow application-centric use of Network Security Groups. You can select multiple ASGs in a NIC.

Note:

 - ASGs are supported when you clone your applications and are also supported in snapshot and restore.
 - ASGs are not supported in VM configuration updates.
 - d. Select a virtual network from the **Virtual Network** dropdown menu.
 - e. Under **Public IP Config**, enter a name and select an allocation method.
 - f. Under **Private IP Config**, select an allocation method.
If you selected **Static** as the allocation method, then enter the private IP address in the **IP Address** field.
17. Enter tags in the **Tags** field. This step is optional.
18. Configure the connection in your environment. For more information, see [Configuring Check Log-in for your Environment](#) on page 66.
19. Click **Next**.

20. Add credentials for the environment. For more information, see [Adding Credentials to the Environment](#) on page 67. This step is optional.
21. Click **Save Environment & Project**.

What to do next

You can use the environment details while configuring a blueprint for Azure or launching a blueprint.

Configuring GCP Environment

Environment configuration involves configuring VMs and adding credentials for the accounts that you added to your project.

About this task

Use this procedure to configure environment variables for GCP.

Before you begin

Ensure that you have configured a project and selected a GCP account. For more information, see [Creating a Project](#) on page 42 and [Configuring Environments in a Project](#) on page 50.

Procedure

1. On the **Infrastructure** tab of the Create Environment page, click the **Select Infrastructure** dropdown menu, and add a GCP account.
2. Select the account you added in the left pane.
3. In the right pane, expand the **VM Configuration** section, and select either **Windows** or **Linux** as the operating system for the VM.
4. Under **VM Configuration**, enter the instance name of the VM in the **Instance Name** field. This field is pre-populated with macro as suffix to ensure name uniqueness.
5. Select the zone from the **Zone** dropdown menu.
Zone is a physical location where you can host the VM.
6. Select machine type from the **Machine Type** dropdown menu.
The machine types are available based on your zone. A machine type is a set of virtualized hardware resources available to a virtual machine (VM) instance, including the system memory size, virtual CPU (vCPU) count, and persistent disk limits. In Compute Engine, machine types are grouped and curated by families for different workloads.
7. Under the **Disks** section, click the **+** icon to add a disk.
You can also mark the added vDisks runtime editable so you can add, delete, or edit the vDisks while launching the blueprint. For more information about runtime editable attributes, see [Runtime Variable Overview](#).
8. To use an existing disk configuration, select the **Use existing disk** checkbox, and then select the persistent disk from the **Disk** dropdown menu.

9. If you have not selected the **Use existing disk** checkbox, then do the following:
 - a. Select the type of storage from the **Storage Type** dropdown menu. The available options are as follows.
 - » **pd-balanced**: Use this option as an alternative to SSD persistent disks with a balanced performance and cost.
 - » **pd-extreme**: Use this option to use SSD drives for high-end database workloads. This option has higher maximum IOPS and throughput and allows you to provision IOPS and capacity separately.
 - » **pd-ssd**: Use this option to use SSD drives as your persistent disk.
 - » **pd-standard**: Use this option to use HDD drives as your persistent disk.

The persistent disk types are durable network storage devices that your instances can access like physical disks in a desktop or a server. The data on each disk is distributed across several physical disks.
 - b. Select the image source from the **Source Image** dropdown menu.

The images available for your selection are based on the selected zone.
 - c. Enter the size of the disk in GB in the **Size in GB** field.
 - d. To delete the disk configuration after the instance is deleted, select the **Delete when instance is deleted** checkbox under the **Disks** section.
10. To add a blank disk, click the **+** icon under the **Blank Disks** section and configure the blank disk.
11. To add networking details to the VM, click the **+** icon under the **Networking** section.
12. To configure a public IP address, select the **Associate Public IP address** checkbox and configure the following fields.
 - a. Select the network from the **Network** dropdown menu and the sub network from the **Subnetwork** dropdown menu.
 - b. Enter a name of the network in the **Access configuration Name** field and select the access configuration type from the **Access configuration type** dropdown menu.

These fields appear when you select the **Associate public IP Address** checkbox.
13. Under the **SSH Key** section, click the **+** icon and enter or upload the username key data in the **Username** field.
14. Select **Block project-wide SSH Keys** checkbox to enable blocking project-wide SSH keys.
15. Under the **Management** section, do the following:
 - a. Enter the metadata in the **Metadata** field.
 - b. Select the security group from the **Network Tags** dropdown menu.

Network tags are text attributes you can add to VM instances. These tags allow you to make firewall rules and routes applicable to specific VM instances.
 - c. Enter the key-value pair in the **Labels** field.

A label is a key-value pair that helps you organize the VMs created with GCP as the provider. You can attach a label to each resource, then filter the resources based on their labels.
16. Under the **API Access** section, do the following:
 - a. Specify the service account in the **Service Account** field.
 - b. Under Scopes, select **Default Access** or **Full Access**.

17. Configure the connection in your environment. For more information, see [Configuring Check Log-in for your Environment](#) on page 66.
18. Click **Next**.
19. Add credentials for the environment. For more information, see [Adding Credentials to the Environment](#) on page 67. This step is optional.
20. Click **Save Environment & Project**.

What to do next

You can use the environment details while configuring a blueprint for GCP or launching a blueprint.

Configuring Check Log-in for your Environment

You configure a check log-in task to check whether you are able to SSH into the VM provisioned.

Before you begin

The network adapter must be specified in the VM configuration before you configure check log-in for your environment.

Procedure

1. On the **Infrastructure** tab of the Create Environment page, under **VM Configuration**, expand the **Connection** section.
2. Select the **Check log-in upon create** checkbox.
3. In the **Credential** list, select **Add New Credential** to add a new credential and do the following:
 - a. Type a name of the credential in the **Credential Name** field.
 - b. Type the user name in the **Username** field.
 - c. Select the secret type from the **Secret Type** list.
You can either select **Password** or **SSH Private Key**.
 - d. Do one of the following.
 - » If you selected password, type the password in the **Password** field.
 - » If you selected SSH Private Key, enter or upload the SSH private key in the **SSH Private Key** field.
If the private key is password protected, click **+Add Passphrase** to provide the passphrase. This step is optional.
 - e. If you want this credential as your default credential, select the **Use as default** checkbox.
 - f. Click **Done**.
4. Select address from the **Address** dropdown menu.
You can either select the public IP address or private IP address of a NIC.
5. Select the connection from the **Connection Type** dropdown menu.
Select **SSH** for Linux or **Windows (Powershell)** for Windows.
The **Connection Port** field is automatically populated depending upon the selected **Connection Type**. For SSH, the connection port is 22 and for PowerShell the connection port is 5985 for HTTP and 5986 for HTTPS.

6. If you selected **Windows (Powershell)**, then select the protocol from the **Connection Protocol** list. You can select **HTTP** or **HTTPS**.
7. Enter the delay in seconds in the **Delay** field.

Delay timer defines the time period when the check login script is run after the VM starts. It allows you to configure the delay time to allow guest customization script, IP, and all other services to come up before running the check login script.
8. In the **Retries** field, enter the number of log-on attempts the system must perform after each log on failure.
9. Click **Next** to add credentials to your environment or save the environment and the project.

For information on adding credentials, see [Adding Credentials to the Environment](#) on page 67.

Support for Multiple Credential

Credentials help in abstracting identity settings while connecting to an external system. You can configure multiple credentials of the same type (either SSH key or password) under the **Environments** tab. You can use the configured credentials during the launch of an application blueprint.

Adding Credentials to the Environment

Credentials are used to authenticate a user to access various services in Self-Service. Self-Service supports key-based and password-based authentication method.

About this task

Use this procedure to add credentials.

Before you begin

Ensure that you have configured a project and created environments for the project. For more information, see [Creating a Project](#) on page 42 and [Configuring Environments in a Project](#) on page 50.

Procedure

1. On the **Credentials** tab of the Create Environment page, click **+ Add Credentials**.
2. Enter a name of the credential in the **Credential Name** field.
3. Enter a username in the **Username** field.
4. Select a secret type from the **Secret Type** dropdown menu.

You can either select **Password** or **SSH Private Key**.
5. Do one of the following.
 - » If you selected password, enter the password in the **Password** field.
 - » If you selected SSH Private Key, enter or upload the SSH private key in the **SSH Private Key** field.

Optionally, if the private key is password protected, click **+Add Passphrase** to provide the passphrase. The type of SSH key supported is RSA. For information on how to generate a private key, see [Generating SSH Key on a Linux VM](#) or [Generating SSH Key on a Windows VM](#).
6. Click **Save Environment & Project**.

Quota Policy Overview

Quota policies enforce a usage limit on an infrastructure resource for projects and restrict project members to use more than the specified quota limits. Quotas ensure that a single project or a few projects do not overrun the infrastructures. If the cluster runs out of a resource, project members cannot use the resource even if the project has not reached its specified limit.

Note:

Starting with version pc.2024.1, Prism Central discontinued supporting legacy quotas and enforcing any legacy quotas that you defined in your projects. To enforce quotas in projects, you must enable the policy engine. For more information, see [KB-15343](#).

When you upgrade your Prism Central instance to pc.2024.1 or later versions and enable the policy engine, your Prism Central instance migrates any legacy project quotas to policy-engine-based quotas. You cannot reverse the migration of legacy quotas.

Quota policies also enforce a usage limit on an infrastructure resource at the provider account level to ensure that the resource consumption is within the specified quota limits across all projects of that provider account.

Note: Quotas do not reserve any specific amount of infrastructure resources.

Quota Allocation

Quotas are allocated at the account and project levels. Enforcement of resource quota depends on these factors:

- The status of the policy engine.
You must enable the policy engine to enforce resource quota policies. For more information, see [Enabling the Policy Engine](#) on page 84.
- The resource quotas you allocate to the Nutanix and VMware provider accounts. For more information, see [Allocating Resource Quota to an Account](#) in the *Self-Service Administration and Operations Guide*.
- The resource quotas you allocate for a project at the project level. For more information, see [Managing Quota Limits for Projects](#) on page 69.
- The resource quotas you allocate to the Nutanix and VMware accounts within a project. For more information, see [Adding Infrastructure in a Self-Service-Enabled Prism Central](#) on page 46.

Project-Level Quota Allocation

You can define resource quota limits at the project level for the projects that you create in your Prism Central instance. The **Policies** tab of a project provides a unified view of all the resource quota limits that you defined for the project and the accounts within the project. You can manage all your project-specific quota definition on the **Policies** tab.

You must follow these conditions while allocating quotas at the project level.

- The quota limits you define at the project level cannot be less than the sum of quotas you allocated to different accounts within the project. You can, however, increase the project-level quota limits.
- You can define resource quota limits for a project without defining quota limits for the associated Nutanix or VMware accounts.
- The project-level quota limits cannot be more than the sum of quotas allocated at the account level to the associated providers. For example, if your project has a Nutanix account and a VMware account, the project-level quota limits cannot be more than the sum of quotas allocated globally to the associated Nutanix and the VMware accounts.

- When you disable quota at the project level, the policy engine does not perform quota utilization checks for any actions that you perform within the project from the time you disabled the quota.

Quota Checks

The policy engine performs quota checks for every resource provisioning request. Quota check happens for multi-VM applications, single-VM applications, and the VMs that you created from your Prism Central instance within a project.

- For multi-VM applications, quota check happens when you launch a blueprint, update an application, or perform a scale-out action to increase the number of replicas of a service deployment.
- For single VM applications, quota check happens when you launch a blueprint or update an application.
- For successful launch of a blueprint, application update, or scale-out, the requested resource must be within the quota limit allocated to the associated provider and project.
- In case of quota violation, your Prism Central instance displays a notification with details such as the associated project, associated account if applicable, and the reasons for violation.
- In case of a scale down or application delete action, the policy engine releases the consumed resources back to the total available quota.

Reporting

The Prism Admin and Project Admin can view project-wise usage of infrastructure resources for each cluster. For more information, see [Managing Quota Limits for Projects](#) on page 69.

For information on quota utilization report, see [Viewing Quota Utilization Report](#) in the *Self-Service Administration and Operations Guide*.

Managing Quota Limits for Projects

The **Policies** tab of a project provides you the options to define and manage quota limits for the project.

About this task

Quotas in projects work in the following ways:

- When Self-Service is disabled in your Prism Central instance, you can define and manage quotas only for Nutanix (AHV).
- When Self-Service is enabled in your Prism Central instance, you can define and manage quotas for Nutanix (AHV) and VMware accounts.

You can do the following as part of the quota limit management at the project level:

- Enable or disable project-level quota checks.
- Assign quota limits for the project.
- Edit quota limits for the associated Nutanix and VMware accounts within the project. For information on allocating resource quota limits to the associated accounts, see [Adding Infrastructure in a Self-Service-Enabled Prism Central](#) on page 46.
- View quota utilization report.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.

3. Click **Projects** in the navigation bar.
4. Do one of the following:
 - » Click **+Create Project** to create a new project and define quota limits. For information on creating a project, see [Creating a Project](#) on page 42.
 - » Click a project name in the list of existing projects to define quota limits for that project.
5. Configure your project with users, accounts, and environments. For more information, see [Adding Users to a Project](#) on page 44, [Infrastructure in Projects](#) on page 46, and [Environments in Projects](#) on page 50.
6. Click the **Policies** tab of the project.
7. Ensure that the **Quotas** tab is selected in the left pane.
8. To enable quota checks at the project level, enable the **Quotas** toggle button in the right pane.
9. To assign resource quota limits at the project level, enter quota values for **vCPU**, **Memory**, and **Disk** for the project.
 - The quota limits you define at the project level for vCPU, Memory, and Disk must be equal to or more than the sum of quotas you allocated to different accounts within the project.
 - The **Project/Global Quota** row shows the total quota limit defined for the different associated accounts within the project and globally at the account levels. The project-level quota limits must be equal to or less than the sum of the quota allocated at the account level to the associated providers globally.
 - If you hover your mouse over the status bar of a resource in the **Quota Utilization** row, you can view the resources consumed and the resources allocated to the project.
10. To manage provider quota limits within the project, expand the provider account in the Provider Quotas section and do the following:

Note: The options to define and manage provider quota limits are available only when Self-Service is enabled in your Prism Central instance.

- a. View the quota limits (vCPU, memory, and disk) allocated to the provider account within the project.
- b. Click **Edit** to enable and add the quota limit or modify the existing quota limit.

The Edit Account window opens where you can enable quotas for the account, add quota limits, or modify the existing limits.

The **Available/Total** row shows the available resources quota and the total quota allocated to the provider. The **Physical Capacity** row shows the used and total physical capacity. Use these details while allocating resource quotas to the cluster.

11. To review quota utilization at the project level, click **Quota Utilization Report**.

Note: The Quota Utilization Report is available only when Self-Service is enabled in your Prism Central instance.

You can use the Quota Utilization Report to review the following details:

- Use the **Resource Utilization** tab to review the total utilization of vCPU, Memory, and Disk at the project level. You can also review the utilization of infrastructure resources by each cluster of the accounts associated with the project.
 - Use the **Quota Utilization** tab to review detailed utilization data at the level of each cluster of the associated accounts. You can review the quota allocated and quota used by each account. You can also review the quota allocated for each cluster and the percentage of quota utilization at the cluster level.
 - Use the **App Quota Utilization** tab to review the resource quota utilization in terms of applications in Self-Service.
12. To disable quota checks for the actions that are performed within the project, disable the **Quotas** toggle button.

Creating a Snapshot Policy

A snapshot policy allows you to define rules to create and manage snapshots of application VMs that run on a Nutanix or VMware platform. The policy determines the overall intent of the snapshot creation process. For applications VMs that run on Nutanix platform, you can create rules to manage your snapshots on a local cluster, on a remote cluster, or both and also configure snapshot expiration. For applications VMs that run on VMware platform, you can select allowed clusters and associated hosts for the local snapshots.

About this task

Perform the following steps to create your snapshot policy. For information on the snapshot configuration and creation, see [Snapshot and Restore Overview](#).

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Projects** in the navigation bar.
4. Do one of the following:
 - » Click **+Create Project** to create a new project and create snapshot policies within the project. For information on creating a project, see [Creating a Project](#) on page 42.
 - » Click a project name in the list of existing projects to create snapshot policies within the project.
5. Ensure that you have your project configured with users, accounts, and environments. For more information, see [Adding Users to a Project](#) on page 44, [Infrastructure in Projects](#) on page 46, and [Environments in Projects](#) on page 50.
6. On the **Policies** tab, click **Snapshot** in the left pane.
7. Click **+ Create Snapshot Policy**.
8. In the **Policy Name** field, type a name for the snapshot policy.
9. In the **Policy Description** field, type a description for the snapshot policy.

10. Select the **Set as default snapshot policy** checkbox to make this your default snapshot policy.
11. Under Primary Site, select a primary environment.
You can select an environment you configured with your Nutanix or VMware account.
12. If you selected a Nutanix environment, then do the following to configure local snapshots in your policy.
 - a. Select an account in the primary environment to which you want to associate the snapshot policy.
Selecting an account in the **Infrastructure** dropdown menu enables **Local Snapshots**. You can view all the clusters that you allowed for the account in the **Primary Cluster** column under Snapshot Rules.
 - b. Under Snapshot Rules, in the **Snapshot Expiry** field for each cluster, specify the number of days after which the snapshot should expire.

Note: The storage cost of the snapshot depends on the days of expiration you specify for the cluster. The longer the days of expiration, the higher the storage cost. The default value is zero days, which indicates that the snapshot will never expire.

If you do not want to include an allowed cluster in the policy, click the **Delete** icon next to the cluster. You can use **+ Add Rule** to include a deleted cluster to the policy.

13. If you selected a VMware environment, then do the following to configure local snapshots in your policy.

Note: The snapshot policy rules are created for VMware clusters and hosts. The Snapshot action is enabled only for the clusters or hosts that you select in the policy.

- a. Select an account in the primary environment to which you want to associate the snapshot policy.
Selecting an account in the **Infrastructure** dropdown menu enables **Local Snapshots**. You can view all the clusters that you allowed for the account in the **Cluster** column under Snapshot Rules. The snapshot rule is applicable to the clusters that appear in the list.
 - b. To remove an allowed cluster from the snapshot policy, click the Delete icon next to the cluster.
 - c. To select a host for the snapshot policy, click **Add Rule**, and then select a host in the **Select Host** dropdown menu.
You can select multiple hosts for the snapshot policy.
14. If you selected a Nutanix environment, then do the following to enable remote snapshots in your policy.
 - a. Enable the toggle button next to **Remote Snapshots**.
You can enable remote snapshots if your target environment and the associated account has multiple allowed clusters, and you want to use one of the clusters to store snapshots. Remote snapshots are particularly useful when your Prism Central has a computer-intensive cluster managing workloads and a storage-intensive cluster managing your data, snapshots, and so on.
You can anytime use the toggle button to enable or disable remote snapshots in your policy.
 - b. Click **+ Add Rule**.
 - c. From the **Primary Cluster** dropdown menu, select the primary cluster where the snapshots of the VMs are taken.
 - d. From the **Target Cluster** dropdown menu, select the target cluster where you want to store the snapshots.
 - e. From the **VM Categories** dropdown menu, select the VM category.
 - f. For each cluster, in the **Snapshot Expiry** field for each cluster, specify the number of days after which the snapshot should expire.
 - g. Click **+ Add Rule** and repeat the steps to add more primary and target clusters to the rule.

15. Click **Save Snapshot Policy**.

IDENTITY AND ACCESS MANAGEMENT

Identity and Access Management (IAM) is an authentication and authorization feature that uses attribute-based access control (ABAC). Starting from version pc.2024.1, you can find multiple improvements that focus on making the IAM feature better optimized and more autonomous, such as:

- Introduction of new system-defined roles to facilitate easier assignment of permissions (granular RBAC) and simplify the authorization process.
- Consolidation of role mapping and role assignment into the new Authorization Policies workflow.
- Simplified CAC authentication workflow.

Note: Prism Central welcome banner reverts to default after a series of disabling, upgrading and re-enabling IAM. For more information, see [KB-17644](#).

For information on IAM prerequisites, considerations, and enablement, see the [Identity and Access Management \(IAM\)](#) section in the *Security Guide*.

For IAM Software Support, see the [Prism Central Release Notes](#).

Security and User Management (Prism Central)

Prism Central provides several mechanisms to manage security and control user access. For more information, see the following topics in the [Security Management using Prism Central \(PC\)](#) chapter of the *Security Guide*.

- Set the user authentication method to local, directory service, or both. For more information, see [Configuring Authentication](#).
- Create authorization policies to assign roles (built-in or custom) to identities (users or user groups) for a global or customized scope (of allowed actions). For more information, see [Authorization Policies](#).
- Add or edit local user accounts. For more information, see [Managing Local User Accounts](#).
 - Update your account. For more information, see [Updating My Account](#).
- Assign roles to users. For more information, see [Controlling User Access \(RBAC\)](#).

Note: For recovery plan RBAC, see [Nutanix Disaster Recovery Guide](#).

- Import SSL certificates. For more information, see [Importing an SSL Certificate](#).
- Control SSH access to Prism Central. For more information, see [Controlling Remote \(SSH\) Access](#).
- Create and apply security-related policies. For more information, see [Security Policies](#) in the *Flow Network Security Next-Gen Guide*.

User Management

Use the **Identities** tab on the Identify and Access Management page to view a list of local users, imported users, and user groups. The tab also provides you the option to add a local user. To access the **Identities** tab in Admin Center, click **IAM** in the navigation bar, and then click the **Identities** tab.

For information on creating and managing local user accounts, see the [Managing Local User Accounts](#) section in the *Security Guide*.

The following table lists the actions that you can perform on the **Identities** tab and their associated options.

Table 10: Identities Tab - Options

Action	Option
Add a local user	To add a local user. For more information, see Managing Local User Accounts .
Edit a local user	To edit or disable a local user. For more information, see Editing or Disabling a Local User Account .
Order the user list alphabetically (A-Z or Z-A)	Click (toggle) the Name , Username , Type , Email Address , Last Log In , or Modified On column header.
View imported users	Click the Imported Users tab.
View Users Groups	Click the User Groups tab.
Filter by name or username	Type the name or username and press Enter. For User Groups, you can only filter by name.

Roles Management

The **Roles** tab on the Identify and Access Management page allows you to view summary information about built-in (default) and custom (created) roles and to access detailed information about each role. To access the list of roles in Admin Center, click **IAM** in the navigation bar, and then click the **Roles** tab.

For information on how to create, manage, and apply roles, see [Controlling User Access \(RBAC\)](#) in the *Security Guide*.

By default, you can view the following information about the roles on the **Roles** tab.

Table 11: Roles List Fields

Parameter	Description
Role	Displays the role name.
Description	Describes the role (if a description was provided when the role was created or updated).
Accessible Services	The services that are accessible to the role. In case multiple services are accessible to the role, you can hover the mouse to view the list of services the role has access to.
Accessible Entity Types	The entity types that are accessible to the role. In case multiple entity types are accessible to the role, you can hover the mouse to view the list of entity types the role has access to.
Modified On	Displays the date and time when the role was modified the last time.

You can filter the list based on the role name.

You can use the **Create Role** button to create a custom role. For more information, see [Creating a Custom Role](#) in the *Security Guide*.

The **Actions** menu appears when you select a role. You can use the Actions menu to one or more roles are selected. It allows you to duplicate roles, to add authorization policy, to update a custom role, and to delete a custom role.

Note: You cannot delete or update the System roles.

To access details of a role, click the role name on the **Roles** tab. The Role page has the following action buttons.

- **Add Authorization Policy:** Use this button to add an authorization policy to the role. For information on authorization policies, see [Authorization Policies](#) in the *Security Guide*.
- **Duplicate:** Use this button to duplicate the role. It opens the role configuration screen preconfigured with the same permissions as this role. For more information, see [Creating a Custom Role](#) in the *Security Guide*.
- **Delete:** Use the button to delete the role. You can only delete a custom role.
- **Update:** Use this button to update the role permissions. You can only update a custom role. For more information, see [Modifying a Custom Role](#) in the *Security Guide*.

The role page also displays the entity types accessible to the role and their associated operations. You can expand the entity types to view more details about the operations.

LCM VIEW

The life cycle manager (LCM) tracks software and firmware versions of the various components in a cluster.

View information about the current inventory and update the versions as needed.

To access the LCM dashboard, select **Admin Center** from the Application Switcher and then click **LCM** in the navigation bar. You must have the Prism Admin role to access LCM.

For more information on life cycle manager, see the [Life Cycle Manager Guide](#).

CATEGORY MANAGEMENT

A category groups entities into key-value pairs. You assign new entities to a category based on specific criteria and link policies to the entities that belong to a specific category value.

For example, suppose that you have a category for departments that includes engineering, finance, and HR. You create one backup policy that applies to engineering and HR and a separate, stricter backup policy for finance. You then use categories to implement different policies across groups and quickly view established relationships.

The following table shows the relationship between four policies (Backupbasic, Engenvironment, Hourly alerts, and Daily backup) associated with three departments (Engineering, Finance, and HR). The policies apply to more than 30 virtual machines in each department.

Table 12: Dept Details View

Policy	Values	Entities
Backupbasic (DR)	Department: Engineering	VMname1, VMname2, VMname3, and 34 others
	Department: HR	VMname1, VMname2, VMname3, and 34 others
Engenvironment (Security)	Department: Engineering	VMname1, VMname2, VMname3, and 34 others
Hourly Alerts (Affinity)	Department: Finance	VMname1, VMname2, VMname3, and 34 others
	Department: HR	VMname1, VMname2, VMname3, and 34 others
Daily Backup (DR)	Department: Finance	VMname1, VMname2, VMname3, and 34 others

Important:

The **Storage: \$Default** category is used to assign the default storage policy to an entity such as a VM or a VG. Nutanix recommends that you do not associate any other policy with the **Storage: \$Default** category or make any other changes to this category.

For more information on the default storage policy, see [Default Storage Policy](#).

Creating a Category

You create a category to assign new entities based on specific criteria, and then link policies to the entities belonging to that particular category value.

Procedure

To create a category, follow these steps:

1. Log in to your Prism Central instance.
2. From the Application Switcher, select **Admin Center**.

3. In the navigation bar, click **Categories**.
The **Categories List** page displays a list of all existing system-defined and user-defined categories.
4. Click **New Category**.
The **Create Category** page appears.
5. In the **Name** field, enter a name for the category.
Category names are attribute types to filter entities, for example, Department or Location.
6. (Optional) In the **Purpose** field, enter the purpose of the category.
7. In the **Values** field, enter the category values.
Category values are attribute instances to filter entities. For example, if your category name is Department, the category values can be Engineering, HR, Sales, Marketing and so on.
When determining duplicates, the **Values** field considers the character case of the entered values. For example, the field considers **Sales** and **sales** as two unique values and accepts these values.
8. Click **Save**.
The new category appears on the **Categories List** page.

What to do next

You can assign the category to different entities. For more information, see [Assigning a Category to an Entity](#) on page 80.

Viewing Category Details

You can view the summary or details of the categories from the Prism Central Admin Center.

Procedure

To view the details of a category, follow these steps:

1. Log in to your Prism Central instance.
2. From the Application Switcher, select **Admin Center**.
3. In the navigation bar, click **Categories**.
The **Categories List** page displays a list of all existing system-defined and user-defined categories.
4. View the value, associated entities, and associated policies.
The Associated Entities column lists the number of entities assigned to a category, and the Associated Policies column lists the number of policies assigned to the category. A dash (-) in any column indicates that the value for the field is not available or not applicable.

Note: For a non-admin user, your Prism Central instance displays the list of associated entities when you expand the row and does not display the cumulative list of associated entities.

5. In the list of categories, click a category name.
The **Summary** page for the category opens and displays the values defined for the category, the associated policies, and the associated entities.

Note: The list does not include the policies that are associated with multiple categories.

6. Place your cursor over a policy, value, or entity.
The page highlights the association among the three parameters.

Assigning a Category to an Entity

You can assign a category value to clusters, VMs, hosts, volume groups, images, and subnets.

About this task

Note: Each entity supports only 64 categories, so you can assign maximum 64 categories to any entity.

Procedure

To assign a category value to one or more entities, follow these steps:

1. Log in to your Prism Central instance.
2. From the Application Switcher, select **Infrastructure**.
3. Navigate to the dashboard of the entity to which you need to assign the category.
4. On the dashboard, select the entities that you need to tag with the same category value.
You can select clusters, VMs, hosts, volume groups, subnets, or images from the **List** tab of their respective pages.
5. From the **Actions** dropdown menu, select **Manage Categories**.
The **Manage Categories** page for the entity opens.
6. From the **Set Categories** dropdown menu, select the target values.
The **Set Categories** dropdown menu provides a list of matching categories when you enter the keywords to search. Policies that are associated with the category you select appear in the **Possible Associated Policies** field. Repeat this step to assign multiple category values to the same set of entities.
7. Click **Save**.

Updating a Category

You can update the name, purpose, or the attribute values of a user-defined category.

About this task

You cannot update a system-defined category.

Procedure

To update a user-defined category, follow these steps:

1. Log in to your Prism Central instance.
2. From the Application Switcher, select **Admin Center**.
3. In the navigation bar, click **Categories**.
The **Categories List** page displays a list of all existing system-defined and user-defined categories.
4. Click a user-defined category name.
The **Summary** page of the category opens.
5. Click **Update**.
The **Update Category** page opens.

6. Update the **Name**, **Purpose**, and **Values** of the category.

For more information on how to update the **Name**, **Purpose**, and **Values** fields, see [Creating a Category](#) on page 78.

7. Click **Save**.

Deleting a Category

You can delete a category and its associated values when the category is not used in any existing policy.

About this task

You cannot delete a system-defined category.

Procedure

To delete a user-defined category, follow these steps:

1. Log in to your Prism Central instance.
2. From the Application Switcher, select **Admin Center**.
3. In the navigation bar, click **Categories**.
The **Categories List** page displays a list of all existing system-defined and user-defined categories.
4. To delete a category from the **Categories List** page, follow these steps:
 - a. Select the category name from the list of categories.
 - b. From the **Actions** dropdown menu, click **Delete**.
The **Deleting Category** window appears.
 - c. Click **Delete**.

ADMIN CENTER SETTINGS OPTIONS

You can use the Admin Center settings to configure various application management, networking, alerts, and appearance related settings. These administrative settings are common across multiple Nutanix applications that you deploy and access using Prism Central or affect the platform. You must log in as an administrator to configure these settings.

You have the following categories of settings in Admin Center.

- **App Management**

Nutanix Central is a unified cloud management console where you can access multiple Prism Centrals. Registration is a two-way API-key process. Nutanix Central uses role-based access control (RBAC) to manage access of the registered Prism Central domains.

For information on Nutanix Central and its registration process, see [Nutanix Central Admin Guide](#).

Prism Central automatically archives run logs of the deleted applications and custom actions that are older than one month to clean up resources. These archives can be downloaded and viewed by an administrator to get valuable insights within 7 days of archive creation. For more information, see [Application Log Archive](#) on page 83.

- **Governance Policies**

The policy engine enforces resource quota policy, orchestrates applications through tunnels on a virtual private network (VPC) on Nutanix accounts, enforces approval policies, and schedules app actions and runbook executions. For more information on policy engine, see [Policy Engine Overview](#) on page 84.

- **Setup**

Nutanix Cloud Manager Intelligent Operations (formerly Prism Pro/Ultimate) combines aspects of administration, reporting, and intelligent automation for IT Operations. Intelligent Operations empowers you to plan intelligently and optimize for capacity, proactively detect performance anomalies, and automate the IT operations. For information on Intelligent Operations enablement, see [Enabling Intelligent Operations](#) in the *Intelligent Operations Guide*.

The Pulse feature provides diagnostic system data about Prism Central to Nutanix to deliver proactive, context-aware support for Nutanix solutions. When enabled, Pulse sends a summary email of cluster configuration to a Nutanix server daily. With this information, Nutanix can apply advanced analytics to optimize your implementation and to address potential problems.

- **Network**

You can configure the following network settings:

- **HTTP Proxy:** A server to communicate with the Nutanix Service Center if Prism Central cannot send traffic to a Nutanix Service Center directly.
- **Name Server:** A server that hosts a network service to provide responses to queries against a directory service, such as a DNS server.
- **NTP Server:** A server to which Prism Central must connect to synchronize the system clock.
- **SNMP Server:** A server to facilitate the exchange of management information between network devices.

- **Security**

The Cluster Lockdown feature allows you to delete (or add) public authentication keys used for SSH access into Prism Central. Removing all public keys locks down Prism Central from external access. For more information, see [Controlling Remote \(SSH\) Access](#) in the *Security Guide*. You can use the SSL Certificate option to create a self-signed certificate. For more information, see [Installing an SSL Certificate](#) in the *Security Guide*.

- **Alerts and Notifications**

You can configure the alert messages that are sent from Prism Central, define rules, and the email content. As part of alert and notification settings, you can also configure the following servers:

- SMTP server: To allow Prism Central to send notifications.
- Syslog server: To allow syslog monitoring to forward system logs (API Audit, Audit, Security Policy Hitlogs, and Flow Service Logs) of the registered clusters to an external syslog server.

- **Appearance**

You can configure settings related to accessibility and appearance, such as:

- Default landing page: The page that appears when you or other users first log in to Prism Central.
- Language: The language and locale for your Prism Central instance.
- UI Settings: The theme and timeout duration for admin and non-admin users.
- Banner page: The Welcome page appearance and content.

Application Log Archive

Prism Central automatically archives run logs of the deleted applications and custom actions that are older than 1 month. You can download the archives within 7 days from the time of archive creation.

For a running application, data is not archived for the system-generated Create actions.

You can get the following information for Start, Restart, Stop, Delete, and Soft Delete system-generated actions and user-created actions.

- Started by
- Run by
- Status

Prism Central archives all action details of a deleted application.

Only an administrator can view and download the application log archive. For more information, see [Downloading Application Log Archive](#) on page 83.

Downloading Application Log Archive

Prism Central periodically archives application logs to clear resources. You can download the archived application logs from the Admin Center settings.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Settings** in the navigation bar.
4. Under App Management, click **App Log Archive**.
The right pane displays the timestamp and the validity of the archive.
5. Click **Download** and navigate to the location on your computer where you want to save the tar.gz file.

Policy Engine Overview

The policy engine is a single-VM setup for the single or scale-out Prism Central. When you enable the policy engine for your Prism Central instance, a new VM is created and deployed for the policy engine. All you need is an available IP address that belongs to the same network as that of your Prism Central VM for the policy engine VM.

As an administrator, you can enable the policy engine to:

- Enforce resource quota policy for the infrastructure resources (compute, memory, and storage) on Nutanix and VMware. The quota policy enforcement allows better governance on resources across infrastructures at the provider and project levels. For more information, see [Quota Policy Overview](#) on page 68.
- Orchestrate apps through tunnels on a virtual private network (VPC) on Nutanix accounts. For more information, see [Tunnels for Orchestration within a VPC](#) in the *Self-Service Administration and Operations Guide*.
- Enforce approval policies to manage resources and control actions in your environment. For more information, see [Approval Policy Overview](#) in the *Self-Service Administration and Operations Guide*.
- Schedule app actions and runbook executions. For more information, see [Scheduler Overview](#) in the *Self-Service Administration and Operations Guide*.

Note:

- In versions earlier than pc.2024.1, the policy engine enablement was possible only when Self-Service was deployed and licensed in your Prism Central instance. Starting from Prism Central version pc.2024.1, you can enable the policy engine from the Admin Center irrespective of whether Self-Service is deployed or not.
- Nutanix recommends that you enable policy engine immediately after you upgrade your Prism Central to version pc.2024.1. Until the policy engine is enabled, Prism Central cannot perform any quota checks expected from the existing project quotas (legacy project quotas) or any validation on the resource consumption. The Quota section in your Projects pages also remains unavailable until you enable the policy engine.
- When you upgrade your Prism Central instance to pc.2024.1 or later versions and enable the policy engine, your Prism Central instance migrates any existing project quotas (legacy project quotas) to the new version of quotas designed as per the policy engine. You cannot reverse the migration of legacy quotas. For information on policy engine quotas, see [Quota Policy Overview](#) on page 68.
- If you delete your policy engine VM and deploy a new VM, the quotas set earlier do not persist with the new policy engine VM. Therefore, you must back up the policy engine database so you can restore the database whenever required. For more information, see [Backing Up and Restoring the Policy Engine Database](#) on page 88.

Enabling the Policy Engine

The policy engine is a single-VM setup for the single or scale-out Prism Central. You can enable the policy engine from the Admin Center Settings of your Prism Central instance.

Before you begin

- Ensure that you have an available IP address that belongs to the same network as that of your Prism Central VM. Prism Central uses the IP address to create and deploy the policy engine VM. If an HTTP proxy is configured, ensure that you add the policy engine IP address to the HTTP-proxy allowlist on Prism Central. For information on configuring HTTP Proxy for Prism Central, see [Configuring an HTTP Proxy](#).
- Ensure that you enable Marketplace before enabling the policy engine. For information on enabling marketplace, see [Enabling Marketplace](#) on page 9.

About this task

- When you enable the policy engine for your Prism Central instance, a new VM with 4 vCPUs, 6 GiB memory, and 40 GiB disk space is created and deployed for the policy engine.
- When you upgrade Prism Central to the latest version and if the policy engine was enabled in the previous version, you must also upgrade the policy engine to the latest version using LCM.

Note:

- In versions earlier than pc.2024.1, the policy engine enablement was possible only when Self-Service was deployed and licensed in your Prism Central instance. Starting from Prism Central version pc.2024.1, you can enable the policy engine from the Admin Center irrespective of whether Self-Service is deployed or not.
- When you upgrade your Prism Central instance to pc.2024.1 or later versions and enable the policy engine, your Prism Central instance migrates any existing project quotas (legacy project quotas) to the new version of quotas designed as per the policy engine. You cannot reverse the migration of legacy quotas. For information on the policy engine quotas, see [Quota Policy Overview](#) on page 68.

To enable the policy engine, follow these steps:

Procedure

1. Log in to your Prism Central instance as an administrator.
2. From the Application Switcher, select **Admin Center**.
3. In the navigation bar, click **Settings**.
4. In the **Governance Policies** section, click **Policy Engine**.
5. If your Prism Central instance is on AHV, follow these steps:
 - a. In the **IP Address for Policy Engine** field, enter an IP address for the policy engine VM.
 - b. Click **Enable**.
 - c. Click **Confirm** to enable the policy engine.
6. If your Prism Central instance is on ESXi and Self-Service is disabled, follow these steps to enter the details of your vCenter server where the Prism Central VM resides.

Prism Central requires the vCenter server details to establish communication with the server prior to creating the policy engine VM. Once you enter the server details, Prism Central validates the credentials and enables the **IP Address for Policy Engine** field so that you can enter the IP address of the policy engine VM.

 - a. In the **Server** field, enter the IP address of the vCenter server.
 - b. In the **Username** field, enter the user name of the vCenter account.

If a domain is part of the username, use the <username>@<domain> username syntax.
 - c. In the **Password** field, enter the password of the account.
 - d. In the **Port** field, enter the port number as 443.
 - e. In the **IP Address for Policy Engine** field, enter an IP address for the policy engine VM.
 - f. Click **Enable**.
 - g. Click **Confirm** to enable the policy engine.

7. If your Prism Central instance is on ESXi and Self-Service is enabled, follow these steps:
 - a. Ensure that you added the VMware provider account for the vCenter that manages the host where the Prism Central VM resides to your Self-Service.
For more information, see [Configuring a VMware Account](#) in the *Self-Service Administration and Operations Guide*.
 - b. In the **IP Address for Policy Engine** field, enter an IP address for the policy engine VM.
 - c. Click **Enable**.
 - d. Click **Confirm** to enable the policy engine.

What to do next

- Allocate resource quotas to provider accounts. For more information, see [Allocating Resource Quotas to an Account](#) in the *Self-Service Administration and Operations Guide*.
- Allocate resource quotas to projects. For more information, see [Managing Quota Limits for Projects](#).
- Create VPC tunnels on Nutanix accounts. For more information, see [Creating VPC Tunnels](#) in the *Self-Service Administration and Operations Guide*.
- Create approval policies for runbook executions, app launch, or app day-2 operations. For more information, see [Creating an Approval Policy](#) in the *Self-Service Administration and Operations Guide*.
- Schedule app actions and runbook executions. For more information, see [Creating a Scheduler Job](#) in the *Self-Service Administration and Operations Guide*.

Enabling the Policy Engine at a Dark Site

The policy engine is a single-VM setup for the single or scale-out Prism Central instance. You can enable the policy engine at a dark site.

Before you begin

If your Prism Central instance is on ESXi, add the VMware provider account for the vCenter that manages the host where the Prism Central VM resides to your Prism Central instance.

Procedure

1. Download the policy engine image from the [Support & Insights Portal](#).
When you download the image, ensure that the image version is compatible with your Prism Central version.
2. If your Prism Central is on AHV, upload the image on your Prism Central instance with the following name:
<Self-Service version number>-CalmPolicyVM.qcow2

3. If your Prism Central is on ESXi, follow these steps to upload and convert the image to a VM template:
 - a. Manually upload the image as a VM template on the vCenter host where the Prism Central VM resides with the following name:
<Self-Service version number>-CalmPolicyVM.ova

You must upload the OVA file as an OVF template in the same container where the Prism Central VM is located.
 - b. In the vCenter inventory, locate the generated VM to convert the VM into a template.
 - c. Right-click the VM (<Self-Service version number>-CalmPolicyVM.ova) and click **Template > Convert to Template**.
 - d. Verify that the VM template name is <Self-Service version number>-CalmPolicyVM.ova.
You must edit the VM template name in case the name is other than <Self-Service version number>-CalmPolicyVM.ova.
4. After uploading the image, enable the policy engine from the Admin Center Settings page. For information on enabling the policy engine, see [Enabling the Policy Engine](#) on page 84.

Viewing the Policy Engine VM Details

After enabling the policy engine, you can review the policy engine VM configuration, network configuration, and cluster information on the Admin Center Settings page. For example, you can view the power status, protection status, or cluster name of the policy engine VM.

Before you begin

Ensure that you have enabled the policy engine for your Prism Central instance. For information on enabling the policy engine, see [Enabling the Policy Engine](#) on page 84.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Settings** in the navigation bar.
4. Under Governance Policies in the left pane, click **Policy Engine**.
5. Expand the **Policy Engine VM Details** section to view the policy engine VM details.

Disabling the Policy Checks

Disable the policy checks for your Prism Central instance if the policy engine VM encounters any connectivity issues or the policy engine VM does not respond.

About this task

Note: Disabling policy checks disables approval policies, suspends quota policy enforcement, and pauses scheduled jobs.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.

3. Click **Settings** in the navigation bar.
4. Under Governance Policies in the left pane, click **Policy Engine**.
5. Select the **Skip Policy Checks** checkbox to disable the policy enforcement.
6. Click **Confirm**.

Best Practices to Shut Down the Policy Engine

Consider the following points before you shut down your policy engine.

- Make sure that no VM related action such as VM create, update, delete, or so on is running in Prism Central.
- Make sure that no actions are running in your Self-Service (if enabled).

You can restart the policy engine VM from the Prism Central VM list page.

Backing Up and Restoring the Policy Engine Database

Enabling policy engine creates and deploys a new VM for the policy engine in your Prism Central network. After the deployment of the policy engine VM, you can anytime create a backup of your policy engine database. You can use the backup to restore the policy engine to the earlier state on your existing policy engine VM or on a new policy engine VM.

About this task

You must run the backup and restore commands from your Prism Central instance.

Procedure

1. Get the IP address of the Policy Engine VM.

For information on how to get the Policy Engine IP address, see [Viewing the Policy Engine VM Details](#) on page 87.

2. Create a backup of your policy engine.

```
$ ssh nutanix@<policy_vm_ip> /home/nutanix/scripts/backup.sh
```

Where *<policy_vm_ip>* is the IP address of the policy engine VM.

This command creates a local backup of the policy engine database on the policy engine VM at /home/nutanix/data/backups/.

Note: When you run the command to create the backup, the policy engine remains unavailable until the backup is created.

3. View all the backups that are available on your policy engine VM.

```
$ ssh nutanix@<policy_vm_ip> /home/nutanix/scripts/restore.sh --list
```

4. Follow these steps to restore the policy engine to a new policy engine VM:

- a. Copy the backup to Prism Central and then to the new policy engine VM.

```
nutanix@pcvm$ scp nutanix@<policy_vm_ip>:/home/nutanix/data/backups/<Backup File Name> ~/tmp/
```

```
nutanix@pcvm$ scp ~/tmp/<Backup File Name> nutanix@<new_policy_vm_ip>:/home/nutanix/data/backups/<Backup File Name>
```

- b. Restore policy engine to its earlier state.

```
$ ssh nutanix@<policy_vm_ip> /home/nutanix/scripts/restore.sh -f=<backup_name>
```

Where *<policy_vm_ip>* is the IP address of the policy engine VM and *<backup_name>* is the local backup file available on the policy engine VM.

Note: When you run the command to restore your policy engine on the existing policy engine VM, the policy engine remains unavailable until it is restored.

5. Follow one of these steps:

- » If the IP address of the policy engine VM remains the same, restore the tunnel firewall rules for existing tunnel VMs on the Prism Central VM.

```
$ ssh <Prism Central IP address>
nutanix@pcvm docker exec -ti domain_manager bash
activate
update_policy_vm_host_data.py
```

- » If the IP address of the policy engine VM changes, delete the existing tunnel VMs and create new tunnel VMs.

Pulse Health Monitoring

The Pulse feature provides diagnostic system data about Prism Central or a Prism Element cluster to Nutanix Support to help deliver proactive, context-aware support for Nutanix solutions. When you enable Pulse, Pulse periodically sends data to Nutanix Insights and the Nutanix Support team for troubleshooting. Pulse is the underlying technology that securely transmits system-level diagnostic data to the Insights platform, enabling predictive health and context-aware support automation workflows. Nutanix Insights is an integrated service that utilizes this data to augment product support, reducing customer case volume and expediting issue resolution time.

Pulse collects data automatically without affecting the system performance and shares only the basic system-level information required for monitoring the health and status of Prism Central or a Prism Element cluster. This information includes the following items.

- System alerts
- System tasks
- System logs
- System configuration
- Performance metrics
- Current Nutanix software version
- Nutanix processes and Controller VM information
- Hypervisor details such as type and version
- Cumulative data about the monitored clusters

All the information in the above categories is specific to the internal processes of Nutanix and does not contain information regarding customer workloads or applications.

Nutanix processes data that Pulse sends consistent with your agreement with Nutanix and, where applicable, by the Nutanix Privacy Statement. For more information, see [Nutanix Privacy Statement](#). Some of the Nutanix products require Pulse enablement for functionality and features. See the Nutanix Privacy Statement and applicable product documentation for more details.

Pulse frequently collects important data, like system-level statistics and configuration information, to automatically detect issues and help simplify troubleshooting. With this information, Nutanix Support can apply advanced analytics to optimize your implementation and address potential problems. Nutanix can use pulse data to trend the adoption of versions, features, and configurations.

Pulse sends messages through HTTPS (port 443) to [insights.nutanix.com:443](#) using TLS 1.2. The HTTPS request uses certificate authentication to validate that Pulse has established communication with the Nutanix Remote Diagnostics service. The TLS 1.2 protocol uses public key cryptography and server authentication to provide confidentiality, message integrity, and authentication for traffic passed over the Internet. For the complete list of required ports, see [Ports and Protocols](#).

You can enable or disable Pulse in Prism Central or in a cluster registered to Prism Central. For more information, see [Pulse Configuration](#) on page 91.

Pulse Transport Methods

Nutanix recommends that you configure one of the following Pulse transport methods (in the order of preference):

1. Enable Pulse and use Prism Central as a proxy for the Pulse data transmitted by each node (for clusters registered with Prism Central). For more information, see [Prism Central Proxy for Pulse Data](#) on page 100.

Advantages: The configuration is automatic and no new firewall configurations are required when you add a node to the cluster or remove a node from the cluster.

2. Enable Pulse and configure an HTTP proxy server. For information on how to configure an HTTP proxy server, see [Configuring an HTTP Proxy](#) on page 103.

Advantage: No new firewall configurations are required when you add a node to the cluster or remove a node from the cluster.

3. Enable Pulse and configure your firewall. Enable Pulse by using each Controller VM IP address in each managed cluster. For information on enabling Pulse in a cluster, see [Enabling Pulse](#) in the *Prism Element Web Console Guide*. For information on configuring the firewall, see the *Pulse Access Requirements* section in the [Pulse Health Monitoring](#) topic of the *Prism Element Web Console Guide*.

Disadvantage: New firewall configurations are required when you add a node to the cluster or remove a node from the cluster.

Remote Diagnostics

Remote Diagnostics enables Nutanix Support to request granular diagnostic information from Pulse-enabled clusters. Pulse streams a collection of configuration data, metrics, alerts, events, and select logs to Nutanix Insights, providing a high-level representation of the cluster state. If the Pulse data stream is not detailed enough to diagnose a specific issue, Nutanix Support might need to collect more diagnostic data from the cluster. Remote Diagnostics allows Nutanix to remotely collect the following data only.

- Logs
 - Nutanix services logs
 - Custom gflags set for any Nutanix service
 - Activity traces for Nutanix services
 - Hypervisor logs
 - Hypervisor configuration
 - Cluster configuration
 - System statistics like memory usage
- Nutanix NCC health check reports
- Nutanix log summary report (Panacea)
- A curated set of read-only commands

Each time Remote Diagnostics triggers a collection, it adds an entry to the audit trail of the cluster. There are always two entries, the start (initiation) and finish (termination) of the diagnostics collection.

By default, every Pulse-enabled cluster has Remote Diagnostics enabled. If your security policy or other consideration does not allow cluster access to Nutanix Support for remote diagnostics collection, you can disable Remote Diagnostics without disabling Pulse. Nutanix Support continues to provide seamless and proactive support based on the Pulse data.

- To check the Remote Diagnostics status, log on to a Controller VM through SSH and run the following command.

```
nutanix@cvm$ zcat /appliance/logical/nusights/collectors/kCommand/override_config
```

Note: This command prints the Remote Diagnostics status only if the Remote Diagnostics status is set explicitly. The command does not print anything if the status is the default status.

- To disable Remote Diagnostics, log on to a Controller VM through SSH and run the following command.

```
nutanix@cvm$ /home/nutanix/ncc/bin/nusights/set_remote_diagnostics_status --  
enable=false --reason="text"
```

Replace *text* with a string describing the reason for disabling Remote Diagnostics. The `--reason` parameter is optional.

- To enable Remote Diagnostics, log on to a Controller VM through SSH and run the following command.

```
nutanix@cvm$ /home/nutanix/ncc/bin/nusights/set_remote_diagnostics_status --  
enable=true --reason="text"
```

Replace *text* with a string describing the reason for disabling Remote Diagnostics. The `--reason` parameter is optional.

Note: To enable remote diagnostics, clusters must run NCC 3.7.0.1 or later versions.

Pulse Configuration

When you log in to Prism Central for the first time after an installation or upgrade, the system checks whether Pulse is enabled. If it is not enabled, a pop-up window appears recommending you enable Pulse.

- To enable Pulse in Prism Central and the clusters registered to Prism Central, click **Continue** in the message and follow the prompts.

- To continue without enabling Pulse in Prism Central and the clusters registered to Prism Central, select the **Disable Pulse (not recommended)** checkbox and click **Continue**. For subsequent enablement of Pulse, see [Enabling Pulse](#) on page 92.
- To enable Pulse in Prism Central but not in the clusters registered to Prism Central, select the **Do not enable Pulse for registered clusters. (not recommended)** checkbox and then click **Continue**.
- For Pulse configuration recommendations, see [Pulse Transport Methods](#).

Enabling Pulse

About this task

Perform the following steps to enable Pulse in Prism Central and the clusters registered to Prism Central.

Note:

- To enable Pulse in an individual Prism Element cluster, see [Enabling Pulse](#) in the *Prism Element Web Console Guide*.
- Nutanix recommends that you enable Pulse to allow Nutanix Support to receive cluster data and deliver proactive and context-aware support.
- Nutanix does not collect any personally identifiable information (PII) through Pulse.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Settings** in the navigation bar.
4. Under Setup, click **Pulse**.
The **Enable Insights powered by Pulse** window appears.
5. Click **Enable Pulse**.
The system displays the **Enable Pulse** window.
6. Click **Proceed**.

Note: To prevent enabling Pulse in the clusters registered to Prism Central, clear the **Enable Pulse automatically for registered clusters. (recommended)** checkbox. By default, this checkbox is selected.

Pulse is enabled in Prism Central and for the clusters registered to Prism Central, and the following message is displayed.

Pulse has been enabled for this Prism Central and enablement is in progress for registered Clusters.

Important:

- Enabling Pulse in Prism Central enables Pulse for all the existing clusters registered to Prism Central and those registered to Prism Central in the future. To prevent enabling Pulse in the clusters registered in the future, disable the **Enable Pulse automatically for all existing Clusters and Clusters registered to this Prism Central in the future. (recommended)** toggle in **Pulse Settings**. By default, this toggle is enabled. However, the existing registered clusters continue to have Pulse enabled in them.

- Nutanix recommends enabling Pulse to allow Nutanix Support to receive system data and deliver proactive and context-aware support.
 - You can enable or disable the Pulse settings in an individual registered cluster through the Prism Element web console of the cluster. Changing the Pulse settings in a registered cluster through the Prism Element web console takes precedence over the Pulse setting changes done previously for that cluster through Prism Central.
For example, if a user enables Pulse through Prism Central on all the clusters registered to Prism Central and then disables Pulse on one of the clusters through the Prism Element web console of the cluster, Pulse remains disabled in the cluster, as the changes done through the Prism Element web console takes precedence.
 - To reset the Pulse settings for all the registered clusters, you must disable the **Enable Pulse automatically for all existing Clusters and Clusters registered to this Prism Central in the future. (recommended)** toggle in the related Prism Central.
7. Verify the **Pulse Connection Status for Prism Central** section for connection information. You can view the following information.
- **Status:** Displays the transport mechanism status, which is **SUCCESS** when the HTTP client can access the Pulse server successfully or **FAILED** when it cannot (or Unknown when the status is not known).
 - **Last Checked Time:** The last time Pulse sent data and checked the connectivity with the Pulse server.

Disabling Pulse

About this task

Perform the following steps to disable Pulse from Prism Central.

Caution: Nutanix recommends that you enable Pulse to allow Nutanix Support to receive cluster data and deliver proactive and context-aware support. If Pulse is disabled, the clusters do not send alerts to Nutanix Support when problems occur.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Settings** in the navigation bar.
4. Under Setup, click **Pulse**.
The **Enable Insights powered by Pulse** window appears.
5. Click **Disable Pulse**.
The system displays the **Disable Pulse** window.
6. Click **Disable**.

Note: Disabling Pulse from Prism Central does not disable Pulse from the clusters registered to Prism Central. To disable Pulse from the registered clusters, log in to the Prism Element Web console of each cluster and disable Pulse. For more information, see [Disabling Pulse](#) in the *Prism Element Web Console Guide*.

Pulse is disabled from Prism Central and the following message is displayed.

Pulse has been disabled for this Prism Central.

Mask Entity Names and IP Addresses

Nutanix processes data that Pulse sends in accordance with your agreement with Nutanix and, where applicable, by the Nutanix Privacy Statement. For more information, see [Nutanix Privacy Statement](#).

For the purpose of Pulse data, you can mask the entity names and IP addresses that are not masked by default. The entity names and IP addresses link to Nutanix entities, such as cluster names, and not to individuals.

To check masking settings on Prism Central, run these commands.

- Check if partial scrubbing is enabled.

```
nutanix@cvm$ curl -H "Content-Type: application/json" -X GET -H "X-Nutanix-Preauth-User:admin" http://localhost:9080/PrismGateway/services/rest/v1/pulse
```

If the output contains the string "identificationInfoScrubbingLevel":"PARTIAL" or "identificationInfoScrubbingLevel":"AUTO", partial scrubbing is enabled. By default, the scrubbing level is set to AUTO. If the output contains the string "identificationInfoScrubbingLevel":"ALL", it means partial scrubbing is not enabled, and the **Open In Prism** button on the Insights Portal remains deactivated.

- Mask the entity names and IP addresses and update the personally identifiable information (PII) scrub level for Prism Central.

```
nutanix@cvm$ curl -H "Content-Type: application/json" -X PUT -H "X-Nutanix-Preauth-User:admin" --data '{"identificationInfoScrubbingLevel": "ALL"}' http://localhost:9080/PrismGateway/services/rest/v1/pulse
```

To check masking settings for all registered clusters, run these commands on Prism Central.

- Check if partial scrubbing is enabled for all the registered clusters.

```
nutanix@cvm$ curl -H "Content-Type: application/json" -X GET -H "X-Nutanix-Preauth-User:admin" http://localhost:9080/PrismGateway/services/rest/v1/pulse?proxyClusterUuid=all_clusters
```

If the output contains the string "identificationInfoScrubbingLevel":"PARTIAL" or "identificationInfoScrubbingLevel":"AUTO", partial scrubbing is enabled. By default, the scrubbing level is set to AUTO. If the output contains the string "identificationInfoScrubbingLevel":"ALL", identificationInfoScrubbingLevel is still set to ALL, and the **Open In Prism** button on the Insights Portal remains deactivated.

- Mask the entity names and IP addresses and update the PII scrub level for all registered clusters.

```
nutanix@cvm$ curl -H "Content-Type: application/json" -X PUT -H "X-Nutanix-Preauth-User:admin" --data '{"identificationInfoScrubbingLevel": "ALL"}' http://localhost:9080/PrismGateway/services/rest/v1/pulse?proxyClusterUuid=all_clusters
```

Pulse Health Monitoring Data Collection

The following categories serve as a general overview of the types of information that Pulse gathers from the clusters. Note that some of this information may be anonymized depending on your settings. This list is not exhaustive; for more details about the Pulse information your clusters send to Nutanix, contact Nutanix Support.

Table 13: Pulse Data Collection

Entity	Data Collected
Cluster	• Cluster name (may be anonymized) • Uptime • AOS version • Cluster ID • Block serial number • HW model • Cluster IOPS • Cluster latency • Cluster memory

Entity	Data Collected
Hardware	- Model number - Serial number - Part number - Block number - Node UUID - Type - Size - Version - Name (may be anonymized) - Manufacturer - Status - Memory (size) - Hypervisor type - Hypervisor version - Firmware version - Disk type - Disk model - Disk capacity - Node temperature - Network interface model - SATADOM firmware - PSU status - Node location - IPMI version - Fan RPM - Component location - DIMM bank connection - Clock speed - DIMM temperature - BIOS release date - BIOS ROM size - CPU signature - CPU core count - CPU cores enabled

Note: In this context, hardware can include nodes, blocks, boards, disks, BMCs, fans, DIMMs, BIOS, CPUs, NICs, storage controllers, and power supplies.

Entity	Data Collected
Storage Pool	• Name (may be anonymized) • Capacity (logical used capacity and total capacity) • IOPS and latency
Container	• Container name (may be anonymized) • Capacity (logical used and total) • IOPS and latency • Replication factor • Compression ratio • Deduplication ratio • Inline or post-process compression • Inline deduplication • Post-process deduplication • Space available • Space used • Erasure coding and savings
Controller VM (CVM)	• Details of logs, attributes, and configurations of services on each CVM • CVM memory • vCPU usage • Uptime • Network statistics • IP addresses (may be anonymized)

Entity	Data Collected
VM	• Name (may be anonymized) • VM state • vCPU • Memory • Disk space available • Disk space used • Number of vDisks • Name of the container that contains the VM (may be anonymized) • VM operating system • IOPS • Latency • VM protection status • Management VM (yes or no) • I/O pattern (read, read/write, random, sequential) • IP address (may be anonymized)
Disk Status	• Performance stats • Usage
Hypervisor	• Hypervisor software and version • Uptime • Installed VMs • Memory usage • Attached datastore
Datastore	• Usage • Capacity • Name
Protection Domains	• Name (may be anonymized) • Count and names of VMs in each protection domain

Entity	Data Collected
Gflags	• Key and value • State (set) • Node ID • Service name • Time of modification
Feature	• Feature ID • Name • State (enabled or disabled) • Mode
License	License type (Starter, Ultimate, or Pro)
Alerts	• Alert ID • Type • Severity • Resolution status • Acknowledgement status • Impact type • Message • Creation time • Modification time
Tasks	• Task ID • Operation type • Status • Entities • Message Completion percentage • Creation time • Modification time
Logs	• Component • Timestamp • Source file name • Line number • Message

Entity	Data Collected
Nutanix Services	Service-specific metrics

Network Configuration for Pulse Health Monitoring

When Pulse Health Monitoring is enabled, all telemetry data streams to the Insights service hosted at insights.nutanix.com over port 443. Therefore, telemetry traffic must be allowed to reach this destination.

The source of Pulse data varies depending on the deployment:

- If Prism Central is deployed, Pulse routes all the information from every node in a managed cluster through this Prism Central. You can opt for setting up a Prism Central Proxy and allowing traffic from Proxy Server IPs to insights.nutanix.com over port 443.
- If Prism Central is not deployed, Pulse routes this information from each CVM for every node in the cluster. You can opt for setting up a Prism Element Proxy and allowing traffic from Proxy Server IPs to insights.nutanix.com over port 443.

Prism Central Proxy for Pulse Data

Prism Central can automatically act as a proxy for Pulse data transmitted by each node in a Prism Element cluster registered to that Prism Central instance.

How Do I Enable Prism Central Proxy for Pulse Data?

You do not have to explicitly enable this feature. It depends on your Prism Central and Prism Element configuration.

Pulse data from Prism Element nodes is automatically routed through Prism Central and then sent to Nutanix Support if you satisfy these requirements:

- You enable Pulse on each registered Prism Element cluster.
- Prism Central and each Prism Element cluster node are running NCC 3.5.2 or later.
- The Prism Element clusters do not have direct internet connectivity, and you have not configured an HTTP proxy on Prism Element clusters registered to this Prism Central instance. If your Prism Element clusters are configured to use an HTTP proxy or direct internet connectivity, the cluster nodes bypass the Prism Central Pulse proxy and transmit Pulse data to Nutanix Support.

If your Prism Central deployment is not available, the cluster nodes bypass the Prism Central Pulse proxy and transmit Pulse data to Nutanix Support.

For a Prism Central scale-out deployment, each Prism Element node selects a Prism Central VM at random to act as its proxy.

Can I Use This Feature If I Have Configured an HTTP Proxy on Prism Central and Prism Element?

If you have configured an HTTP proxy on Prism Central, you can use this feature automatically by satisfying these requirements:

- You have configured an HTTP proxy on Prism Central but it does not require basic authentication (a user name and password). For information on setting up an HTTP proxy on Prism Central, see [Configuring an HTTP Proxy](#) on page 103.
- Your Prism Element clusters are not configured to use a proxy. If your Prism Element clusters are configured to use a proxy, data is transmitted from each node to Nutanix support, bypassing the Prism Central Pulse proxy for pulse data.

Prism Central Configuration When a Cluster Uses Proxy Servers

Note:

- Prism Central and its managed clusters are not supported in environments deploying Network Address Translation (NAT).
- If you are planning to use Prism Central as a proxy for Pulse data transmitted by each node in a Prism Element cluster managed by that Prism Central instance, see [Prism Central Proxy for Pulse Data](#) on page 100.

To communicate with the Nutanix service center directly, you need to configure the HTTP proxy. For information on how to configure the HTTP proxy through the Prism Central web console, see [Configuring an HTTP Proxy](#) on page 103. After the HTTP proxy is configured, all the HTTP or HTTPS communication initiated by Prism Element or the Prism Central is routed through the proxy server.

The communication from Prism is routed through the proxy server until an allowed target entry indicates otherwise. For information on who needs to allow the IP addresses, see [Who Needs to Use the Allowlist Method](#) on page 101. To overcome a communication failure between a Prism Element and the registered Prism Central, you need to manually allow the target entries from the Prism Element web console or nCLI. These options enable you to add Prism Central and its managed/registered clusters to an allowlist, where any HTTP proxy settings are ignored. This configuration allows network traffic between them, bypassing any proxy servers configured in the cluster. The allowlist also enables you to register new clusters with Prism Central successfully where clusters are using an HTTP proxy.

For information on how to configure the allowlist IP address through the Prism Central web console, see [Configuring an HTTP Proxy](#) on page 103.

Alternatively, you can use the `ncli http-proxy add-to-whitelist` and `ncli http-proxy delete-from-whitelist` nCLI command options. For information on how to configure the allowlist IP address through the nCLI, see [Allowing Prism Central and Its Managed Clusters \(nCLI\)](#) on page 104.

- You can add or delete one allowlist entry at a time.
- Each allowlist entry cannot exceed 253 characters.
- A maximum of 1000 allowlist entries are supported.
- When deleting an entry from an allowlist, delete the target, not the target type.
- The commands do not support the IPv4 network mask network prefix or * (asterisk) prefix notation. You can specify a subnet and netmask, such as 10.0.0.0/255.0.0.0 or 192.168.1.0/255.255.255.0. This results in adding the specified subnet to the allowlist.
- When applying an allowlist, domain names like contoso.com are not processed the same as www.contoso.com and are treated as separate, distinct entities.
- Use fully qualified domain names to allow.

Who Needs to Use the Allowlist Method

SSL port 9440 needs to be open in both directions between the Prism Central VM and any registered clusters or clusters to be registered. For the complete list of required ports, see [Port Reference](#).

If you are implementing a proxy server in your cluster environment with this port open as shown in this simple graphic, you do not need to allow Prism Central and its managed/registered clusters.




Figure 3: Prism Central, Proxy Server with SSL Port 9440 Open

If you are implementing a proxy server in your cluster environment with this port closed as shown in this simple graphic, you must allow direct communication between Prism Central and its managed/registered clusters. For more information, see [Configuring an HTTP Proxy](#) on page 103 and [Allowing Prism Central and Its Managed Clusters \(nCLI\)](#) on page 104.




Figure 4: Prism Central, Proxy Server with SSL Port 9440 Closed and with Allowlisting

Configuring an HTTP Proxy

About this task

Note:

- To use Prism Central as a proxy for Pulse data transmitted by each node in a Prism Element cluster managed by that Prism Central instance, see [Prism Central Proxy for Pulse Data](#) on page 100.

About this task

If Prism Central cannot send traffic to a Nutanix Service Center directly, you need to configure an HTTP proxy. Use the following steps to configure an HTTP proxy.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Settings** in the navigation bar.
4. Under Network, click **HTTP Proxy**.
5. To add an HTTP proxy, click **New Proxy** and configure the following fields:

Note: You can configure only one HTTP proxy at a time. If an HTTP proxy exists, you must delete the existing proxy before creating a new one.

- a. **Name:** Enter a proxy server name.
- b. **Address:** Enter an IP address or host name for the proxy server.
- c. **Port:** Enter the port number to use.
- d. **Username:** Enter a user name.
- e. **Password:** Enter a password.
- f. **Protocols:** Select the protocol to use.
 - » Select **HTTP** to route all HTTP requests through the proxy.
 - » Select **HTTPS** to route all HTTPS requests through the proxy.
 - » Select **HTTP** and **HTTPS** to route both HTTP and HTTPS requests through the proxy.
- g. Click **Save**.

After you save the configuration, the dialog box displays the new HTTP proxy entry in the list.

6. To edit an HTTP proxy entry, click the pencil icon for that entry, update entries, and then click **Save**.
Ensure that you add or update the allowlist entry in the correct format. For example, `10.0.0.0/255.255.255.0`

The **Update HTTP Proxy** window appears with the same fields as the **Create HTTP Proxy** window with an additional option to add allowlist entries.

To configure HTTP proxy allowlist entries, do the following:

- To add an allowlist target, click the **+ Create** link. This step opens a line to enter a target address. Enter the target IP address or the network address and the subnet mask and then click the **Save** link in that field.
Prism Central sends traffic to the allowlist IP addresses directly rather than through the HTTP or HTTPS proxy.
- To edit an allowlist target, click the pencil icon for that target and update as needed.
- To delete an allowlist target, click the X icon for that target.

Note: Where proxy allowlist configuration is required, it is recommended that you add the cluster virtual IP and all external CVM IPs of each registered cluster to the proxy allowlist for Prism Central. The virtual IP of Prism Central and the IP of each Prism Central VM (PCVM) should be added to the proxy allowlist on the registered Prism Element clusters. For easier management, you can add the subnets instead of adding the individual IPs.

7. To delete an HTTP proxy entry, click the X icon for that entry, and then click **OK** in the confirmation window.

Allowing Prism Central and Its Managed Clusters (nCLI)

About this task

Using this example, you can bypass a proxy server used by a managed Prism Element cluster and allow network traffic between Prism Central and the cluster. This example helps if you attempted to register a cluster that implemented a proxy server, and the registration failed.

Procedure

1. Open an SSH session to any Controller VM in the cluster to be managed by Prism Central.
2. In this example, add the Prism Central VM IP address to the allowlist, then ensure the Prism Central VM IP address was added to the allowlist.

```
nutanix@cvm$ ncli http-proxy add-to-whitelist target-type=ipv4_address  
target=10.4.52.40
```

```
nutanix@cvm$ ncli http-proxy get-whitelist  
Target Type : IPV4_ADDRESS  
Target : 10.4.52.40
```

Note: Repeat this step for additional Prism Central VM IP addresses and the virtual IP address in the case of scale-out Prism Central.

3. Open an SSH session to the Prism Central VM managing the cluster where you just modified the HTTP allowlist.
4. Add the cluster virtual IP address to the allowlist, then ensure the IP address was added to the allowlist.

```
nutanix@pcvm$ ncli http-proxy add-to-whitelist target-type=ipv4_address  
target=10.4.52.10
```

```
nutanix@cvm$ ncli http-proxy get-whitelist  
Target Type : IPV4_ADDRESS  
Target : 10.4.52.10
```

- Repeat this step for individual CVM IP addresses.

- For large clusters or for multiple clusters in a network it can be more efficient to add a network to the http-proxy allowlist. For example,

The following is a sample command to add a network to the http-proxy allowlist.

```
nutanix@pcvm$ ncli http-proxy add-to-whitelist target-type=ipv4_network_mask
target=10.4.52.0/255.255.255.0
```

The following is sample output of the get-whitelist command, which displays the allowed network added above.

```
nutanix@cvm$ ncli http-proxy get-whitelist
```

Target Type : IPV4_NETWORK_MASK Target : 10.4.52.0/255.255.255.0

In this case, Prism Central and its managed cluster can communicate, with network traffic bypassing any proxy servers configured in the cluster.

Configuring Name Servers for Prism Central

About this task

Name servers are computers that host a network service for providing responses to queries against a directory service, such as a DNS server.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Settings** in the navigation bar.
4. Under Network, click **Name Servers**.
5. To add a name server, enter the server IP address in IPv4 or IPv6 format in the **Server IP** field and then click **Add** to the right of that field.

The server is added to the **IP Address** list.

Note: Changes in the name server configuration might take up to 5 minutes to take effect. Functions that rely on DNS might not work properly during this time. You can configure a maximum of three name servers.

6. To delete a name server entry, click the X icon for that server in the **IP Address** list and then click **OK** in the confirmation window.

Configuring NTP Servers for Prism Central

About this task

Network Time Protocol (NTP) is a protocol for clock synchronization between computers, and Prism Central must connect to an NTP server to synchronize the system clock.

Note: If Prism Central is running on Hyper-V, you must specify the IP address of the Active Directory Domain Controller server, not the hostname. Do not use DNS hostnames or external NTP servers.

Procedure

1. Log in to Prism Central as an administrator.

2. Select **Admin Center** in the Application Switcher.
3. Click **Settings** in the navigation bar.
4. Under Network, click **NTP Servers**.
5. To add an NTP server, enter the fully-qualified host name or IP address in IPv4 or IPv6 format in the **NTP Server** field and then click **Add**.
The name or address is added to the **Hostname or IP Address** list.
6. To delete an NTP server entry, click the delete icon for that server in the **Hostname or IP Address** list and then click **OK** in the confirmation window.

Configuring SNMP on Prism Central

About this task

The Simple Network Management Protocol (SNMP) is an application layer protocol that facilitates the exchange of management information between network devices.

Note:

- Prism Element (individual cluster) supports both the SNMP service (agent) and SNMP traps, but Prism Central supports only SNMP traps. The SNMP capability of Prism Central is limited to just sending traps. For information on configuring SNMP for an individual cluster and for details about the Nutanix MIB, see the [Prism Element Web Console Guide](#).
- The Net-SNMP package version 5.7.2 does not support 256-bit AES encryption.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Settings** in the navigation bar.
4. Under Network, click **SNMP**.
5. To enable SNMP for Nutanix Objects, select the **Enable for Nutanix Objects** checkbox.
For information on Nutanix Objects, see the [Nutanix Objects Guide](#).

6. To add an SNMP user entry, click the **Users** tab and the **New User** button and then do the following in the indicated fields:
 - a. **Username:** Enter a user name.
 - b. **Priv Type:** Select **AES** (Advanced Encryption Standard) as the privacy encryption type from the dropdown menu.
 - c. **Priv Key:** Enter a privacy key phrase (password) in the field.
 - The key phrase is AES encrypted when the user is created.
 - The password you choose must meet the following complexity requirements:
 - Contain a minimum 8 characters and a maximum 64 characters.
 - Contain one upper case letter (A-Z).
 - Contain one lower case letter (a-z).
 - Contain one digit (0-9).
 - Contain only the following special characters: ~ @ # \$ % * - = _ + { } | [] ? , . : &.
 - d. **Auth Type:** Select **SHA** (Secure Hash Algorithm) as the authentication hash function type from the dropdown menu.
 - e. **Auth Key:** Enter an authentication key phrase (password) in the field.
 - The key phrase is SHA-1 encrypted when the user is created.
 - The password you choose must meet the following complexity requirements:
 - Contain a minimum 8 characters and a maximum 64 characters.
 - Contain one upper case letter (A-Z).
 - Contain one lower case letter (a-z).
 - Contain one digit (0-9).
 - Contain only the following special characters: ~ @ # \$ % * - = _ + { } | [] ? , . : &.
 - f. Click **Save**.

After you save the configuration, the page displays the new user entry in the list.
7. To view the Nutanix MIB (NUTANIX-MIB.txt), click **View MIB**. To download NUTANIX-MIB.txt, right-click and select the appropriate download action for your browser.
8. To add an SNMP trap receiver, click the **Traps** tab.

9. Click **New Trap Receiver** and then do the following in the indicated fields:

- a. **Receiver Name:** Enter the receiver name.
- b. **SNMP Version:** Select the SNMP version, either v3 or v2c. For SNMP v2c, Nutanix supports only SNMP TRAP and not SNMP GET.
- c. The following fields appear based on the SNMP Version you select:
 - **Trap Username:** This field appears when you select v3 as the SNMP Version. Select a user from the list.
 - **Community:** This field appears when you select v2c as the SNMP Version. The default value for v2c trap community is Public. You can enter any other value of your choice.

All users added previously (see step 7) appear in the list. You cannot add a trap receiver entry until you add at least one user.
- d. **Address:** Enter the target address.

An SNMP target address specifies the destination and user that receives outgoing notifications, such as trap messages. SNMP target address names must be unique within the managed device.
- e. **Port:** Enter the port number to use.

The SNMP trap receiver uses UDP port number 162. For the complete list of required ports, see [Port Reference](#).
- f. **Engine ID:** Enter an engine identifier value, which must be a hexadecimal string between 5 and 32 characters long. This step is optional.

If you do not specify an engine ID, an engine ID is generated for you for use with the receiver. Every SNMP v3 agent has an engine ID that serves as a unique identifier for the agent. The engine ID is used with a hashing function to generate keys for authentication and encryption of SNMP v3 messages.
- g. **Inform:** Select **True** from the dropdown menu to use inform requests as the SNMP notification method; select **False** to use traps as the SNMP notification method.

SNMP notifications can be sent as traps or inform requests. Traps are one-way transmissions; they do not require an acknowledgment from the receiver. Informs expect a response. If the sender never receives a response, the inform request can be sent again. Therefore, informs are more reliable than traps. However, informs consume more resources. Unlike a trap, which is discarded as soon as it is sent, an inform request must be held in memory until a response is received or the request times out. Also, traps are sent only once, while an inform may be retried several times. The retries increase traffic and add overhead on the network. Thus, traps and inform requests provide a trade-off between reliability and resources.

Note: The SNMP server is blocked if it does not get a response to the inform traps in five consecutive attempts.
- h. **Transport Protocol:** Select the protocol to use from the list.

The options are **TCP**, **TCP6**, **UDP**, and **UDP6**.
- i. Click **Save**.

After you save the configuration, the page displays the new trap entry in the list.
- j. To test all configured SNMP traps, click the **Traps** tab, and then click **Test All**.

The Nutanix cluster sends test alerts to all the SNMP trap receivers configured on the cluster.

10. To edit a user or trap receiver entry, click the appropriate tab (**Users** or **Traps**) and then click the pencil icon for that entry in the list.

An edit window appears for that user or trap receiver entry with the same fields as the add window. (Transport entries cannot be edited.) Enter the new information in the appropriate fields and then click **Save**.

11. To delete an SNMP entry, click the appropriate tab (**Users** or **Traps**), click the X icon for that entry in the list, and then click **OK** in the confirmation box.

Configuring Alert Emails in Prism Central

You can configure the alert messages that are sent from Prism Central. You can configure alert settings, reporting rules, and message templates.

About this task

Alert emails sent from Prism Central are in addition to any alert emails you might have configured on individual clusters through the Prism Element web console. The system sends you email from both the entities in this case. Prism Central alert emailing is not enabled by default; you must explicitly enable it and specify the recipients. If you enable alerts through Prism Central and do not want to receive duplicate email notifications for the same alert, disable customer email notification for those alerts on the individual clusters through Prism Element web console, however keep the email notification support enabled for Nutanix customer support.

Note: Prism Central requires an SMTP server to send alert email messages. For information on how to configure an SMTP server, see [Configuring an SMTP Server for Prism Central](#).

Procedure

To configure alert emails in Prism Central, perform the following steps:

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** from the [Application Switcher Function](#), click **Settings** in the navigation bar, and navigate to **Alerts and Notifications > Alert Email Configuration** or directly click the [Settings icon](#).

For information on the **Navigation Bar**, see [Application-specific Navigation Bar](#) information in *Prism Central Infrastructure Guide*.

You can also select the **Infrastructure** application from the [Application Switcher Function](#), and perform the following steps to access the **Alert Email Configuration** window:

- a. Navigate to **Activity > Alerts** from the **Navigation Bar**.
The system displays the **Alerts** page.
- b. Click **Email Configuration**.
The system displays the **Alert Email Configuration** window.

3. Click the **Settings** tab, and perform the following actions:

a. Under **Email Preference**, set the following fields for alert emails:

- Select the **Every Single Alert** checkbox to configure Prism Central to send an email whenever the event occurs.
- Select the **Daily Digest** checkbox to configure Prism Central to send a cumulative (24 hour) list of alerts once a day. Set the time when you want the alert to be sent.
- Select the **Skip the daily digest email if there are no alerts generated on a given day** checkbox to configure Prism Central to send an email only when there are any alerts. Prism Central does not send emails for the days when no alerts occur.

If you do not select any checkboxes, no alert emails are sent from Prism Central.

b. To send alert notifications to others, enter their email addresses as a comma-separated list in the **Email Recipients** field.

c. Click **Save**.

Note: The **Tunnel Connection** section displays mail transport status information.

If you want to configure rules or an email message template, perform the following steps before you click **Save**.

4. Click the **Custom Settings** tab to create a custom alert email setting.

5. Under **Create New Setting** set the following fields:

a. Specify the conditions to generate the alert:

- **Alert Severity:** Select one or more of the severities from the list (Critical, Warning, Info, or All).
- **Impact Type:** Select one or more of the categories from the list (Availability, Capacity, Configuration, Performance, System Indicator, or All).
- **Cluster:** Select one or more of the clusters from the list (*cluster_name* or All).
- **Alert Contains:** Enter a key phrase or word that should generate an email notification whenever the alert contains that phrase. For example, if you want to get an email notification when an alert contains the phrase low memory, enter low memory in the field.

b. Specify who should receive the alert email by entering recipient email addresses as a comma-separated list in the **Send Email To** field.

c. Select **Enable** checkbox to enable the new setting. If you have created custom rules (settings) in a version prior to Prism Central 2023.3, those rules are enabled by default after upgrade to Prism Central 2023.3.

Note: Custom settings override the **Email Preference** configured under **Settings** tab, and are honoured even if you do not specify **Email Preference**.

d. Click **Save** to save the setting.

e. Repeat these steps to create more custom settings.

6. Click the **Email Content** tab to create a template for the email messages.

Enter the following field information in the **Email Content** tab:

- **Prepend Subject:** Enter the desired text.

This text appears at the beginning of the subject field in each alert email. If the field is left blank, no prepended text appears in the subject.

- **Append Body:** Enter the desired text.

This text appears at the end of the message body in each alert email. If the field is left blank, no appended text appears in the message body.

7. Click **Save**.

Configuring an SMTP Server for Prism Central

About this task

Simple Mail Transport Protocol (SMTP) is an Internet standard protocol for electronic mail transmission across Internet Protocol (IP) networks. Prism Central uses SMTP to send alert emails and to exchange emails. Use the following steps to configure an SMTP server entry.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Settings** in the navigation bar.
4. Under Alerts and Notifications, click **SMTP Server**.
5. Do the following in the indicated fields:
 - a. **Host Name or IP Address:** Enter the IP address or fully-qualified domain name for the SMTP server.
 - b. **Port:** Enter the port number to use.

The standard SMTP ports are 25 (unencrypted), 587 (TLS), and 465 (SSL). For the complete list of required ports, see [Port Reference](#).
 - c. **Security Mode:** Enter the desired security mode from the dropdown menu.

The options are **NONE** (unencrypted), **STARTTLS** (use TLS encryption), and **SSL** (use SSL encryption).
 - d. **User:** Enter a user name.

The **User** field appears only when a secure option (**STARTTLS** or **SSL**) is selected. The user name might need to include the domain (*user@domain*) depending on the authentication process.
 - e. **Password:** Enter the user password.

The **Password** field appears only when a secure option (**STARTTLS** or **SSL**) is selected.
 - f. **From Email Address:** Enter an e-mail address that appears as the sender address. This step is optional.

By default, alert and status information e-mails display "cluster@nutanix.com" as the sender address. You have the option to replace that address with a custom address by entering a sender address in this field.
6. Click **Save**.

Configuring Syslog Monitoring

You can configure syslog monitoring to forward system logs (API Audit, Audit, Security Policy Hitlogs, and Flow Service Logs) of the registered clusters to an external syslog server using Prism Central.

About this task

The Prism Central method of syslog monitoring configuration propagates the configuration to the Prism Element clusters. If you do not wish the configuration to be propagated to the clusters, you must use Nutanix command-line interface (nCLI) for syslog monitoring configuration. For information on configuring remote syslog monitoring using nCLI, see [Send Logs to Remote Syslog Server](#) in the *Acropolis Advanced Administration Guide*.

Note:

Configure only one syslog server. Nutanix does not support sending logs to multiple syslog servers.

To configure syslog monitoring through Prism Central, follow these steps:

Before you begin

- You must have the IP address of the syslog server that is deployed in your environment.
- For forwarding Flow logs, the Flow feature must be enabled.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Settings** in the navigation bar.
4. Under Alerts and Notifications, click **Syslog Server**.
5. Click **Add Syslog Server**.
6. On the Syslog Server tab, follow these steps:
 - a. In **Server Name**, enter a descriptive name for the server.
 - b. Enter the **IP Address** and **Port**.
 - c. Select the **UDP** or **TCP** transport protocol.
 - d. Select the **Enable RELP (Reliable Event Logging Protocol)** checkbox to enable RELP. This step is optional.

Reliable Event Logging Protocol (RELP) is a networking protocol for data logging in computer networks. It extends the functionality of the syslog protocol to provide reliable delivery of event messages.
 - e. Click **Next**.

7. On the **Data Sources** tab, follow these steps:

a. Select one or more log modules from the following log types. For more information, see [Syslog Modules](#) on page 113.

- **API Audit**
- **Audit**
- **Security Policy Hit Logs** (policy hit log files logs)
- **Flow Service Logs** (Flow processes logs)
- **Calm**
- **Epsilon**

Caution:

Nutanix recommends that you only configure the modules listed in [Syslog Modules](#) on page 113. Modifications to the modules you make using Prism Central web console overwrite the configuration of other modules using nCLI commands.

b. Select the **Severity Level** for the data sources you selected. For more information, see [Syslog Modules](#) on page 113.

c. Click **Save**.

Syslog Modules

Refer to the following table to understand the information that is sent to the syslog server based on the selected log modules and severity levels.

Note: For the list of modules and module names for use in nCLI commands, see the [Configuring the Remote Syslog Server Settings](#) topic in the *Acropolis Advanced Administration Guide*.

Table 14: Log Modules and Severity Levels

Module Name	Severity Level	Description
API Audit	0-7	API Audit logs send information about REST API endpoints that are called and who called them from both Prism Central and Prism Element. All log level settings send the same type of information. Changing the log level changes the received log level at the syslog server, but not the content that is sent.
Audit	0-7	Audit logs send information about VM, Category, and Security Policy creation, update, and delete operations. All log level settings send the same type of information. Changing the log level changes the received log level at the syslog server, but not the content that is sent. See the <i>Syslog Module: AUDIT Fields</i> table below for information on Syslog Audit fields.

Module Name	Severity Level	Description
Security Policy Hit Logs	All Levels	Security Policy Hit Logs send information about network flows that are acted on by a security policy. This information includes the source and destination IP address, protocol, and port. It also includes the action taken by the policy such as allow, monitor, or drop. Network statistics are also included in each message for bytes sent and received in the flow. Messages are always sent to the syslog server as level 6 - informational.
Flow Service Logs		Flow service logs contain detailed debug information used for troubleshooting the backend processes used by Flow.
		Note: Do not enable these unless instructed to do so by Nutanix support for troubleshooting purposes. Unlike the other data sources, the syslog level selected for Flow Service Logs has an impact on the types of data sent to the syslog server.
	0 - 5	Critical service failure messages as well as log rotation indications. Sent at the selected log level to syslog.
	6 - 7	Detailed log messages from the various components that make up the Flow product on both Prism Central and Prism Element.
Calm	0-7	Nucalm service logs are sent to the syslog server based on the log level you select. These service logs contain information about the module from the Nucalm container (such as styx, hercules, jove, helios, iris, and so on) based on the selected log-level priority. For example, if Warning (level-number: 4) is selected, all the logs that belong to the selected log level are stored on the syslog server.
Epsilon	0-7	Epsilon service logs are sent to the syslog server based on the log level you select. These service logs contain information about the module from the Epsilon container (such as jove, narad, indra, karan, arjun, durga, and so on) based on the selected log-level priority. For example, if Warning (level-number: 4) is selected, all the logs that belong to the selected log level are stored on the syslog server.

The following table provides information on the syslog severity levels.

Table 15: Syslog Severity Levels

Value	Severity	Keyword	Description
0	Emergency	emerg	System is unusable.
1	Alert	alert	Should be corrected immediately.
2	Critical	crit	Critical conditions.
3	Error	err	Error Conditions.
4	Warning	warning	Indication that an error might occur if an action is not taken.
5	Notice	notice	Events that are unusual, but not error conditions.

Value	Severity	Keyword	Description
6	Informational	info	Normal operational messages that require no action.
7	Debug	debug	Information useful to developers for debugging the application.

The following table provides information on the AUDIT fields of the Syslog Module.

Table 16: Syslog Module: AUDIT Fields

AUDIT Field	Description
entityType	Indicates the resource type of the entity that is associated with the audit record. Example. The entity type is <i>vm</i> for operations on virtual machines.
name	Indicates the name of the entity that is associated with the audit record. Example. Name of the virtual machine for audits or operations on virtual machines.
uuid	Indicates the unique identifier of the entity that is associated with the audit. uuid is also used to fetch the details of the entity instance.
alertUid	Indicates the name of the audit record. Example. <i>VmMigrateAudit</i> and <i>VmCreateAudit</i> .
creationTimestampUsecs	Indicates the time (in microseconds) from Epoch when the audit record was created.
defaultMsg	Provides a brief description of the current state of the operation that is captured by the audit.
opStartTimestampUsecs	Indicates the time (in microseconds) from Epoch when the operation that is captured by the audit begins.
opEndTimestampUsecs	Indicates the time (in microseconds) from Epoch when the operation that is captured by the audit is completed.
operationType	Indicates the category of the audit record. Example. <i>Create</i> , <i>Delete</i> , and <i>Migrate</i> .
originatingClusterUuid	Indicates the cluster (if applicable) where the operation of is performed.
params	Additional key value pairs that capture information pertaining to the operation that is captured by the audit record.
userName	Indicates the name of the user who triggered the operation.

AUDIT Field	Description
userId	Indicates the unique identifier of the user who triggered the operation.

The following table provides information on the different fields of the Security Policy Hit Logs.

Table 17: Security Policy Hit Logs Fields

Field	Description
<Timestamp>	Indicates the time of the hit log capture.
<Policy UUID>	Indicates the unique identifier of the policy that is associated with the hit log.
<Policy Name>	Indicates the name of the policy as configured by the user.
<Session Information> Create, Update, or Destroy	Indicates the session information. • Create - New session started. TCP handshake or UDP bidirectional packets. • Update - Sent every 30 seconds for active sessions. • Destroy - Session ended by TCP handshake or UDP timeout.
SRC	Indicates the source IP address.
DST	Indicates the destination IP address.
PROTO	Indicates the transport protocol.
ACTION	Indicates the action taken on traffic by the applied policy. • Allow - traffic allowed by policy • Drop - traffic dropped by policy • Monitor - traffic allowed by policy in monitor mode
<Packet Stats>	Indicates the packet statistics.

Setting the Prism Central Default Landing Page

About this task

The default landing page is the page that appears when you first log in to Prism Central. You can set any of the deployed Nutanix apps, Infrastructure page, or the Admin Center page as your default landing page. As an administrator, you can also set the system default landing page that determines the default landing page for all other users. Use the following steps to set the default landing page.

Procedure

1. Log in to Prism Central as an administrator.

2. Select **Admin Center** in the Application Switcher.
3. Click **Settings** in the navigation bar.
4. Under Appearance, click **Default Landing Page**.
5. To set your default landing page, select an option from the **My default landing page** dropdown menu.
6. To set the system default landing page (for all other users), select an option from the **System default landing page** dropdown menu.

Note: Users can override the system default landing page and can set their own default landing page.

7. Click **Save**.

Changing the Prism Central Language Settings

You can change the language setting in Prism Central to Simplified Chinese, Japanese, or Korean (supported only in Self-Service).

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Settings** in the navigation bar.
4. Under Appearance, click **Language Settings**.
5. To change the language setting of the cluster, select a language from the **Language** dropdown menu. You have the following options:

- » **Simplified Chinese**
- » **Japanese**
- » **Korean** (supported only in Self-Service)

By default, the **English** language is selected.

6. To change the locale settings (date, time, and calendar), select the appropriate region from the **Region** dropdown menu.

By default, the locale is set to the language setting that you have set in the **Language** dropdown menu. However, you can change the **Region** to display the date, time, or calendar in some other format. This format for date, time, and calendar is applied for the entire cluster.

7. Click **Save**.

The language and locale settings (date, time, and calendar) is changed according to the selection. For example, in the below image, once you click **Save** the language setting for the cluster is changed to Chinese and locale setting is changed to Russian. For information on the entities that are supported in Simplified Chinese, see [Internationalization \(i18n\) \(Prism Central\)](#) on page 118. Also, the user interface is localized according to the selection. For information on localization, see [Localization \(L10n\) \(Prism Central\)](#) on page 118.

Note: If you are logged on as a Domain administrator and the language settings changes are not saved, you might need to update the *User logon name* for the Domain administrator account. For more information, see [KB10166](#).

Internationalization (i18n) (Prism Central)

The following table lists all the supported and unsupported entities in UTF-8 encoding.

Table 18: Internationalization Support for Prism Central

Supported Entities	Unsupported Entities
User-defined dashboard name	Password fields
Custom widget name	Static dashboard name
First and last name under Update Profile	Static widget name
User name, first name, and last name under User Management	
Chart name	

Localization (L10n) (Prism Central)

Nutanix localizes the user interface in Simplified Chinese, Japanese, and Korean (supported only in Self-Service) language. All the static screens are translated to the selected locale language.

You have an option to change the language settings of the cluster from English (default) to Simplified Chinese, Japanese, or Korean (supported only in Self-Service). For more information, see [Changing the Prism Central Language Settings](#) on page 117.

If the Prism Element instance is launched from Prism Central, language settings of Prism Central takes precedence over Prism Element.

The dashboards (including tool tips) and menus of Prism Central are localized.

Guidelines and Limitations

- Logical entities that do not have a contextual translation available in the localized language are not translated.
- The AOS generated alerts and events are not localized to the selected locale language.
- Following strings are not localized: VM, CPU, vCPU, Language Settings, licensing details page, hardware names, storage denominations (GB, TB), About Nutanix page, EULA, service names (SNMP, SMTP), hypervisor types.

Nutanix Intelligent Virtual Agent Overview

Nutanix Intelligent Virtual Agent (NIVA) provides contextual tips, recommendations, and help resources for the Self-Service and Intelligent Operations features.

These contextual tips, recommendations, and help resources guide you to use these features efficiently in your Prism Central instance. For more information on NIVA, see the [Nutanix Intelligent Virtual Agent Guide](#).

Important: Starting with Prism Central version pc.7.3.1, Nutanix rebranded Intelligent Assist to Nutanix Intelligent Virtual Agent (NIVA). While the user interface for Prism Central version pc.7.3 still shows the name Intelligent Assist, all current documentation and resources use the updated name, either Nutanix Intelligent Virtual Agent or NIVA.

Turning on NIVA Tips Setting

You turn on the NIVA **Tips** setting for your Prism Central instance to view NIVA contextual tips and recommendations.

About this task

For more information on NIVA, see the [Nutanix Intelligent Virtual Agent Guide](#).

Before you begin

- Ensure that NIVA is connected to My Nutanix.
- Ensure that you either have the Prism Admin or Super Admin role.

Procedure

To turn on the NIVA **Tips** setting, follow these steps:

1. Log on to your Prism Central instance.
2. From the Application Switcher, select **Admin Center**.
3. In the navigation bar, click **Settings**.
4. In the **Appearance** section, click **NIVA**.
5. Turn on the **Tips** toggle.

Troubleshooting Recommendations for NIVA

Troubleshooting guidance for NIVA includes probable causes for different issues and recommended resolutions.

For more information on troubleshooting recommendations for NIVA, see the [Nutanix Intelligent Virtual Agent Guide](#).

Configuring Prism Central UI Settings

About this task

The Prism Central login page includes background animation by default, and users are logged out automatically after being idle for 15 minutes. You can change one or both of these settings.

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Settings** in the navigation bar.
4. Under Appearance, click **UI Settings**.

5. To set different variations of the Prism UI background themes, select any of the following options from the **Prism Themes** dropdown menu. The UI changes cosmetically to reflect the selection.

Important:

- The Prism themes feature is currently in technical preview. You may encounter visual anomalies in few settings or views where the prism themes are not applied. For example, Licensing Settings. Nutanix recommends that you register a case on the support portal to report any anomalies.
 - After you change the background theme, save your selection and refresh any open Prism tabs for the changes to reflect.
 - The theme setting is not common between Prism Central and Prism Element web console. You have to configure background themes in Prism Central and Prism Element web console separately.
- Select **Light Theme** for the light background with high-contrast view. This is the default Prism background theme.
 - Select **Dark Theme** for the dark background with high-contrast view, and then click **Continue** to proceed.
 - Select **Auto (OS defined)** to apply background themes defined for the operating system, and then click **Continue** to proceed. For example, if you have defined **Dark** mode in the operating system setting, the Prism UI uses the setting and sets a dark background theme for Prism Central.

6. To customize the theme, background color, title text, or blurb text on the logon page, do the following:




Figure 5: UI Settings - Theme, Title Text, and Blurb Text

- Simultaneously press the Alt key on the Windows system (Option key on the MAC system) and click on the top bar (as shown in the image) of the **UI Settings** page. The options to customize the theme, title text, and blurb text appear on the page.
 - Select the theme from the options displayed for **Theme**. You can change the HEX codes to create your own custom gradient background color for the logon page.
 - In the **Title Text** field, enter the text to create your custom title.
 - In the **Blurb Text** field, enter the text to create your custom blurb text. This text is displayed below the password field.
7. To configure session timeout, do the following under **Security Settings**:
- Select the session idle timeout for the current user from the **Session Idle Timeout For Current User** dropdown menu.
 - Select the default session idle timeout for all non-administrative users from the **Default Session Idle Timeout For Non-Admin Users** dropdown menu.
 - Select the appropriate option from the **Session Idle Override For Non-Admin Users** dropdown menu to override the session timeout for non-administrative users.

Note: The Idle timeout interval for an administrator cannot be set for more than an hour.

8. Clear the **Disable 2048 game** checkbox to disable the 2048 game.
9. Click **Save**.

Configuring the Prism Central Welcome Banner

About this task

The welcome banner is the first screen that appears you attempt to log in to Prism Central. As an administrator, you can configure the content of the banner page by adding a custom message or any graphics. Use the following steps to configure the banner page.

Note: Prism Central welcome banner reverts to default after a series of disabling, upgrading and re-enabling IAM. For more information, see [KB-17644](#).

Procedure

1. Log in to Prism Central as an administrator.
2. Select **Admin Center** in the Application Switcher.
3. Click **Settings** in the navigation bar.
4. Under Appearance, click **Welcome Banner**.
5. Enter or paste the desired content in HTML format in the pane on the left.

Only "safe" HTML tags are supported. Inline event handlers, scripts, and externally-sourced graphics are not allowed.

Welcome Banner ?

☐ Show in Login Screen

Only safe HTML tags are allowed (e.g. `<ul>`, `<li>`, `<div>`, `<img>`, `<h1>-<h6>`, etc). Inline event handlers and scripts are not allowed.

```
<h1 style="text-align:center, color:#dd0000"> Welcome to this Site
</h1>
<h3 style="text-align:center, color:blue"> Terms and Conditions<
h3>
<p>Read the following terms and then click the bar at the bottom
to accept them.</p>
```

Terms and Conditions

Read the following terms and then click the bar at the bottom to accept them.

Save

Figure 6: Welcome Banner

6. Click **Preview** to display the banner in the pane on the right.
7. In case the banner is not correct, update the HTML code until the preview pane displays the desired message.
8. To enable the banner, select the **Enable Banner** checkbox and then click **Save**.

A live banner page includes an "Accept terms and conditions" bar at the bottom. Clicking the bar sends the user to the login page.

To disable the banner, clear the **Enable Banner** checkbox.

CONTAINERS AND SERVICES CONFIGURATION

You can configure and manage the containers and services meant for end-users through Admin Center module of Prism Central.

The topics in this section cover configurations for different Admin Center components, such as resource management for containers.

Resource Limits for Docker Containers

The Domain Manager, Nucalm, and Epsilon are the Docker containers that run on the Prism Central VM and manage services for Projects, Nutanix Marketplace, and Self-Service.

Nutanix implemented control group (cgroup) vCPU limits on these containers to prevent any single container from consuming an excessive amount of compute resources, which can negatively impact the performance of other containers running on the same host. This approach not only optimizes resource allocation but also enhances overall system efficiency, allowing multiple services to run concurrently without interference.

The following table shows the default vCPU limits set for each container based on the Prism Central configuration:

Table 19: Default vCPU Limits for Domain Manager, Nucalm, and Epsilon

Prism Central Configuration	vCPU Limit for Domain Manager	vCPU Limit for Nucalm	vCPU Limit for Epsilon
Tiny Prism Central	0.5	Not Applicable	Not Applicable
Small Prism Central	1	1	1.5
Large Prism Central	2	2	3
Extra Large Prism Central	2	2	3

You can also verify the resource limits of your containers. For more information, see [Verifying the Container vCPU Limits](#) on page 123.

If you need to modify the default vCPU limits for these Docker containers, contact Nutanix Support. For more information, see [KB-18581](#).

Verifying the Container vCPU Limits

You can verify the control group (cgroup) vCPU limits of your Domain Manager, Nucalm, and Epsilon containers.

Procedure

To verify the resource limits, follow these steps:

1. Log on to your Prism Central instance using the SSH session as a `nutanix` user.
2. Verify the resource limit:

```
nutanix@pcvm$ for container in nucalm domain_manager epsilon; do allssh 'docker inspect \
```

```
["$container" | grep -iE "CpuQuota|CpuPeriod"; done
```

Example output:

```
=====10.44.76.70=====
"CpuPeriod": 100000,
"CpuQuota": 200000,
=====10.44.76.70=====
"CpuPeriod": 100000,
"CpuQuota": 200000,
=====10.44.76.70=====
"CpuPeriod": 100000,
"CpuQuota": 300000,
```

GLOSSARY

For terms used in this guide, see [Nutanix Glossary](#).

COPYRIGHT

Copyright 2025 Nutanix, Inc.

Nutanix, Inc.

1740 Technology Drive, Suite 150

San Jose, CA 95110

All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. Nutanix and the Nutanix logo are registered trademarks of Nutanix, Inc. in the United States and/or other jurisdictions. All other brand and product names mentioned herein are for identification purposes only and may be trademarks of their respective holders.

Flow Virtual Networking Guide

Flow Virtual Networking pc.2024.2

October 30, 2025

Contents

Purpose.....	5
Related Documentation.....	6
Flow Virtual Networking Overview.....	7
Flow Virtual Networking Architecture.....	9
Essential Concepts.....	12
Requirements and Limitations of Flow Virtual Networking.....	21
Flow Virtual Networking Configurations.....	26
Enabling the Network Controller.....	28
Disabling the Network Controller.....	31
Disabling Network Controller to Unregister a Prism Element Cluster.....	31
Upgrading the Network Controller.....	32
Dark Site Installation and Upgrade.....	33
Deploying the Network Controller at a Dark Site.....	34
Upgrading the Network Controller at a Dark Site.....	35
Control User Access in Flow Virtual Networking (RBAC).....	36
Flow Virtual Networking Roles and Permissions.....	36
Flow Virtual Networking Operational Authorizations.....	37
Network Types.....	40
Changing the Default VLAN Type.....	42
Troubleshooting Tips.....	43
Network Gateway Upgrades.....	45
Identifying the Gateway Version.....	45
Detecting Upgrades for Gateways.....	45
Upgrading the Network Gateway.....	46
Installing or Upgrading the Network Gateway in a Dark Site.....	46
Network and Security Entities.....	48
Subnets.....	48
Subnets Summary View.....	49
Subnet Details View.....	51
Virtual Private Clouds.....	52
Viewing the Virtual Private Cloud Summary.....	52
Viewing the Virtual Private Cloud Details.....	54
Floating IPs.....	59
Floating IPs Summary View.....	59
Connectivity.....	60
Gateways Summary View.....	60
Gateway Details View.....	62

VPN Connections Summary View.....	63
VPN Connection Details View.....	65
Subnet Extensions Summary View.....	68
Subnet Extension Details View.....	69
BGP Sessions Summary View.....	72
BGP Session Details View.....	73
Security Policies.....	75
Security Dashboard.....	75
Virtual Private Cloud Management.....	76
VM IP Address Management.....	77
Creating Secondary IP Addresses.....	78
Assigning Secondary IP Addresses to Interfaces.....	79
Assigning Secondary IP Addresses to Floating IPs.....	79
VM and Network Migration.....	80
Migration of VMs between VLAN Basic Subnet and VPC Subnets.....	80
Migration of VLAN Basic Subnets.....	85
VPC Management.....	88
Creating a Virtual Private Cloud.....	89
Requesting Floating IPs.....	92
Creating a Subnet.....	93
Attaching a Subnet to a Virtual Machine.....	96
Creating a Policy.....	97
Creating Static Routes.....	102
Updating a Virtual Private Cloud.....	103
Updating a Subnet.....	105
Updating a Policy.....	105
Updating Static Routes.....	106
Deleting a Virtual Private Cloud.....	106
Deleting Subnets, Policies or Routes.....	107
Connections Management.....	108
Network Gateway Management.....	108
Creating a Network Gateway.....	108
Updating a Network Gateway.....	115
Deleting a Network Gateway.....	116
Virtual Private Network Connections.....	116
VPN Workflow.....	117
Prerequisites for VPN Configurations.....	118
Creating a VPN Connection.....	120
Updating VPN Connection.....	122
Deleting a VPN Connection.....	122
VPN Connection within Same Prism Central.....	122
Layer 2 Network Extension.....	124
Subnet Extension Workflow.....	126
Layer 2 Network Extension Over VPN.....	127
Layer 2 Network Extension Over VTEP.....	132
PBR-based Tromboning in L2 Extended Subnet.....	140
Updating an Extended Subnet.....	142
Removing an Extended Subnet.....	142
Border Gateway Protocol Sessions.....	143
Creating a BGP session.....	144
Updating a BGP session.....	146
Deleting a BGP session.....	147

Copyright.....148

PURPOSE

This *Flow Virtual Networking* Guide describes how to enable and deploy Nutanix Flow Virtual Networking on Prism Central.

RELATED DOCUMENTATION

The [Nutanix Support Portal](#) provides software download pages, documentation, compatibility, and other information.

Documentation	Description
Release Notes Flow Virtual Networking	Flow Virtual Networking Release Notes
Port And Protocols	Port Reference: See this page for details of ports that must be open in the firewalls to enable Flow Virtual Networking to function.
Nutanix Security Guide	Prism Element and Prism Central security, cluster hardening, and authentication.
Flow Network Security Next Gen	Flow Network Security Next-Gen is the next-generation Nutanix microsegmentation solution with an enhanced policy model, advance policy operation, and enterprise readiness features. FNG
AOS guides and release notes	Covers AOS Administration, Hyper-V Administration for Acropolis, Command Reference, Powershell Cmdlets Reference, AOS Family Release Notes, and AOS release-specific Release Notes
Life Cycle Manager Guides	How to upgrade core and other Nutanix software.
AHV guides and release notes	Administration and release information about AHV.
Prism Central and Web Console guides and release notes	Administration and release information about Prism Central and Prism Element.

FLOW VIRTUAL NETWORKING OVERVIEW

Flow Virtual Networking, powered by Network Controller, drives network virtualization to offer a seamless network experience with enhanced security.

Flow Virtual Networking is a software-defined networking solution that provides multi-tenant isolation, self-service provisioning, and IP address preservation using VPCs, subnets, and other virtual components that are separate from the physical network, for the AHV clusters. It integrates tools to deploy networking features like Virtual LANs, Virtual Private Cloud (VPC), Virtual Private Network (VPN), Layer 2 Virtual Network Extension using VPN or Virtual Tunnel End Point (VTEP), Border Gateway Protocol sessions to support flexible app-driven networking that focuses on VMs and applications.

Flow Virtual Networking deploys the following components to manage software-defined network virtualization:

Network Controller

The Network Controller is the networking component of Prism Central that manages and controls configuration, monitoring and optimization of network resources for Flow Virtual Networking VPCs and VLAN subnets. It provides programmability, automation, and centralized control for configuring and managing network flows.

Network Controller is necessary to use centralized VLAN management, Flow Virtual Networking and Flow Network Security Next Generation.

Network Gateway

The network gateway is used to create VPN, VTEP, or BGP gateways to connect subnets using VPN connections, Layer 2 subnet extensions over VPN or VTEP, or over BGP sessions. The network gateway appliance is available along with the Network Controller when you install Prism Central. Network gateway VMs are used to create VPN, VTEP, or BGP gateways to connect subnets using VPN connections, Layer 2 subnet extensions over VPN or VTEP, or BGP sessions.

Flow Virtual Networking comprises of the following features:

- **Centralized Agile Management**

Prism Central helps you enable the Network Controller that provides Flow Virtual Networking (application-driven network virtualization) as well as centralized VLAN management. Flow Virtual Networking leverages the Network Controller and optionally, network gateway appliance to help you manage network configuration changes with speed and agility. It delivers a centralized network management solution with multi-tenant networking, self-service network provisioning, and a multi-cluster network control plane.

Prism Central provides the centralized network management plane that helps you manage the control plane provided by the Network Controller. The Network Controller as the control plane, deploys network virtualization. The Open vSwitch (OVS) infrastructure on the AHV hosts provide the data plane. For more information on the architecture of Flow Virtual Networking, see [Flow Virtual Networking Architecture](#) on page 9.

- **Programmability with Context and Visibility**

Flow Virtual Networking helps you directly program network features and configure network resources quickly and easily through automated services on Prism Central. The Network Controller allows you to design and configure self-service networks using the Prism Central user interface and REST APIs. Flow Virtual Networking enables you to manage networks and network lifecycles easily, to accommodate the increasing demand for network services, without impacting the overall network.

You can view the networks, connection endpoints, and the traffic parameters. This helps you easily redirect the traffic to improve service delivery, reduce service disruptions for your customers and increase network responsiveness, thus helping you deliver a seamless customer experience.

- **Secure Multi-tenancy Solution**

Flow Virtual Networking allows per-tenant isolation using VPC-based network segmentation and namespace isolation. These isolated virtual networks provide security by default.

You can apply policy based routing using the Network Controller to improve the security of the networks by redirecting traffic through security VMs within the VPC. Flow Virtual Networking, with the Network Controller and network gateway, allows you to manage cloud networking by abstracting and unifying cloud resources effectively. The Network Controller uses Virtual Private Cloud (VPC) networks that are abstracts of the underlying network to unify multi-cluster based resources (managed by the Prism Central) into isolated network spaces (VPCs). Secure egress of traffic to the underlying VLAN network is managed using SNAT, Floating IP addresses, or routing with support for static and BGP advertisement.

- **Interoperable Secure Connectivity Solution**

With Flow Virtual Networking, you can use VPN, VTEP or BGP gateway-based configurations for multiple sites, with automated network gateway appliance upgrades. You can also extend subnets across sites using Layer 2 virtual subnet extensions (VPN or VTEP based) for connectivity without using physical gateways, in a vendor-neutral environment.

- **NAT-based Secure Egress**

The Network Controller allows you to configure NAT based traffic egress routes to external networks, with IP address retention and policy-based routing. You can also use a no-NAT, or the routed, option for external networks. For more information on NAT, see [Essential Concepts](#) on page 12.

- **Enhanced Networking for Disaster Recovery**

The Network Controller supports Nutanix [Disaster Recovery](#) solutions.

Note: Prism Central Backup and Restore (PCBR) supports Flow Virtual Networking. For more information, see [Prism Central Backup, Restore, and Migration](#) documentation.

Deployment Workflow

The Flow Virtual Networking Network Controller is auto-enabled when you install an X-Large Prism Central instance or upgrade the version of your existing X-Large Prism Central instance to pc.2023.3 or later. On Small and Large Prism Central instances, you need to enable the Network Controller. Flow Virtual Networking is not supported on X-Small Prism Central instances.

On Small and Large Prism Central instances, you need to enable the Network Controller. Flow Virtual Networking is not supported on X-Small Prism Central instances.

For steps to enable the Network Controller, see [Enabling the Network Controller](#) on page 28.

When Flow Virtual Networking is enabled, the Network Controller and the network gateway appliance are installed. The Network Controller is a collection of containerized services that run directly on the Prism Central VM(s). The Network Controller orchestrates all the virtual networking operations.

- You can deploy Flow Virtual Networking Network Controller in a dark site (a site that does not have Internet access) environment. For more information, see [Deploying the Network Controller at a Dark Site](#) on page 34.
- You can upgrade the Network Controller. Nutanix releases an upgrade for the Network Controller with Prism Central releases. For more information, see [Upgrading the Network Controller](#) on page 32.
- You can create and manage virtual private clouds (VPCs) and overlay subnets to leverage the underlying physical networks that connect clusters and datacenters. For more information, see [Virtual Private Cloud Management](#) on page 76.

You can also upgrade the network gateway version. For more information, see [Network Gateway Upgrades](#) on page 45.

Flow Virtual Networking Architecture

Flow Virtual Networking lets you create integrated software-defined networks and virtual private cloud capabilities and provides software-defined networking with multi-tenant isolation, self-service provisioning, and IP address preservation. The Flow Virtual Networking architecture uses a three-plane approach to simplify network virtualization.

Flow Virtual Networking provides a software-defined networking (SDN) solution with the three-plane architecture that SDN is built with. Flow Virtual Networking as SDN follows the following three-plane architecture:

- **The Management Plane**

The Management Plane provides the interface between you and the configuration and management interfaces. Primarily, it allows you to configure, manage, and monitor the virtual network resources such as IP addresses, subnets, routes and protocols.

Prism Central provides the Management Plane for Flow Virtual Networking. The **Network & Security** entity provides the Flow Virtual Networking components like **Subnets, Virtual Private Clouds, Floating IPs, and Connectivity** (which encompasses network gateways, connections like VPN or VTEP or BGP sessions.) Prism Central also allows you to control access to these virtual networking components on Prism Central using Role-based access control (RBAC).

- **The Control Plane:**

This plane is defined by the SDN controller. This plane is essentially decoupled from the data plane. In other words, the appliance that houses the SDN controller is a different and separate entity from the one that houses the data transport network or the Data Plane.

In Flow Virtual Networking, the Control Plane is defined by the network controller. Prism Central enables the Microservices Infrastructure when you deploy Prism Central. The network controller is enabled on the Prism Central as containerized services using Microservices Infrastructure.

The network controller that allows you to create a virtual overlay network as an abstraction of the complex underlay network infrastructure. The network controller manages the network services and direct packet traffic throughout the network. The network controller, with the network gateway appliance, helps you manage the networks, connections (such as VPN and VTEP connections, and BGP sessions) and devices with ease.

In x-Large Prism Central, the network controller is automatically enabled when Prism Central is deployed. In Small and Large Prism Central deployments, you must enable the network controller manually. For more information, see [Enabling the Network Controller](#) on page 28.

- **The Data Plane**

The Open vSwitch (OVS) deploys a collection of bridges within the AHV hosts. The traffic flows through these bridges between the AHV hosts. To configure and manage these bridges, AHV allows you to deploy virtual switches. AHV deploys a default virtual switch vs0 during the installation process. The default virtual switch manages the bridges br0 on all the AHV hosts in the cluster. For more information on the virtual switches, see [About Virtual Switch](#).

This OVS infrastructure on the AHV hosts provides the Data Plane for Flow Virtual Networking. For more information on OVS, see [About Open vSwitch](#).

This architecture provides a foundation for Flow Virtual Networking as depicted in the following chart.




Figure 1: Flow Virtual Networking Architecture

Implementation Constructs of Flow Virtual Networking

Flow Virtual Networking provides the following virtual constructs to provide a complete networking solution:

- Virtual Private Clouds or VPCs:
- Subnets as VLAN or Overlay Subnets
- Routes
- Policies for routing.
- External Networks such as:
 - NAT based external networks
 - Routed (or NoNAT) external networks
 - Multiple networks or a set of networks with both NAT and NoNAT external networks
- Network gateways such as:
 - Layer 3 Virtual Private Network or VPN
 - Layer 2 Network Extensions with VxLAN or Virtual Tunnel End Point (VTEP)
 - Border Gateway Protocol based gateways and sessions.

For more information on these constructs, see [Essential Concepts](#) on page 12.

Flow Virtual Networking Operation

Each VPC, as an isolated network namespace with a virtual router instance, connects all of the subnets inside the VPC. The VPCs are created in Prism Central that manages all the nodes and clusters that the VPCs span across.

Each VPC can have one or more subnets and all the subnets are connected to the same VPC virtual router. A VPC uses Geneve encapsulation to tunnel traffic between the AHV hosts. When two VMs in a VPC on two different hosts send traffic to each other, the packets are encapsulated in Geneve on the first host, sent to the other host where the packets are decapsulated, and sent to the destination VM.

When you select a NIC for a VM, place that NIC in an overlay subnet, or a VLAN Basic Subnet (VLAN on AHV networking stack). When you choose an overlay subnet, you are also choosing the VPC that the subnet is a part of. Each VM can be placed inside only a single VPC. You cannot connect a VM to both a VPC and a VLAN (AHV-based VLAN Basic Subnet or Network Controller-based VLAN Subnet) at the same time, or to two different VPCs at the same time.

Every VPC contains a single virtual router and different types of routes like External networks, direct connections, remote connections. The virtual router acts as a control point for traffic inside a VPC. An External Network is the primary way traffic enters and exits a VPC. External Networks are created in Prism Central and exist on only a single Prism Element cluster. This network defines the VLAN, the default gateway, the IP address pool, and the NAT type for all the VPCs using it. One External Network can be used by many VPCs.

Direct and remote connections can be established using network gateways in one-to-one (VPN, VTEP or BGP) or one-to-many (VTEP) connections. All connections require network gateways. For example, a VPN connection requires a local gateway and a remote gateway. While the VPN and VTEP gateways are a part of the data plane, BGP gateways are part of the control plane.

You can apply simple stateless policies here, and the traffic that flows through the router is evaluated by the policies. Policies do not apply to traffic from one VM to another VM inside the same subnet. Inside a VPC, policies are evaluated in priority order from highest (1,000) to lowest (10). Once traffic is matched a policy can take one of the following actions:

- Permit
- Deny
- Reroute including Redirect traffic to another /32 IPv4 address in another subnet.

Stateless policies require separate rules defined in both the forward and reverse direction if a Permit rule is overriding a Drop rule. Otherwise, return traffic would be denied by the Drop rule. Use similar priorities to group these matching forward and reverse entries.

Thus, Flow Virtual Networking allows you to create completely isolated virtual networks that are separated from the physical network. These isolated virtual networks provide security by default.

Deployment Scale

Flow Virtual Networking supports the scale provided on the [Nutanix Configuration Maximums](#) page.

Note: For information on the algorithms supported by Flow Virtual Networking (Network Controller and network gateway) APIs, see [Nutanix Networking Versioned APIs \(4.0.1-alpha-1\)](#).

Supported Third Party Appliances

Nutanix has validated that the following the network gateway appliances work in Flow Virtual Networking VPCs:

- AWS
- CheckPoint
- Cisco ASA
- Fortinet
- Juniper SRX
- PaloAlto
- SonicWall NSv
- VyOS

Essential Concepts

Network Controller

The Network Controller is defined as networking component of Prism Central that manages and controls configuration, monitoring and optimization of network resources for Flow Virtual Networking. It provides programmability, automation, and centralized control for configuring and managing network flows.

Network Controller is necessary to use centralized VLAN management, Flow Virtual Networking and Flow Network Security Next-Gen.

VPC

A Virtual Private Cloud (VPC) is an independent and isolated IP address space that functions as a logically isolated virtual network. A VPC could be made up of one or more subnets that are connected through a logical or virtual router. VPCs allow you to manage the isolated and secure virtual network with enhanced automation and scaling. The isolation is done using network namespace techniques like IP-based subnets or VLAN based networking.

The IP addresses within a VPC must be unique. However, IP addresses may overlap across VPCs, in other words, the IP addresses inside of one VPC to overlap with any other VPC, or even with the physical network. As VPCs are provisioned on top of another IP-based infrastructure (connecting AHV nodes), they are often referred to as the overlay networks. Tenants may spin up VMs and connect them to one or more subnets within a VPC. Virtual Private Cloud (VPC) is a virtualized network of resources that are specifically isolated from the rest of the resource pool. A VPC can expand to include any cluster managed by the same Prism Central. A VPC might exist within a single AHV cluster, or within clusters in the same availability zone.

The default VPC type that is referred to as *VPC* in this documentation is the one you create to isolate selected subnets of connected VMs. This VPC is also called as user VPC or guest VPC, but generally referred to as VPC.

The other VPC type that Flow Virtual Networking supports is *transit VPC*. For more information, see the *Transit VPC* section below. You need a minimum Prism Central version of pc.2024.1 to deploy transit VPCs.

Shared VPC Connections

Shared VPC connections involve connecting VPCs such that you can route traffic between them using private IP addresses. The VPCs can, then, communicate as if they are in the same network. You can connect a VPC to another VPC either directly or through a transit VPC to achieve shared connections. For information on transit VPCs, see the *Transit VPC* section below.

VPC Subnets

You can use IP address-based subnets to network virtual machines within a VPC. A VPC may use multiple subnets. VPC subnets use private IP address ranges. IP addresses within a single VPC must be unique, in other words, IP addresses inside the same VPC cannot be repeated. However, IP addresses can overlap across multiple VPCs. The following figure shows two VPCs named Blue and Green. Each VPC has two subnets, 192.168.1.0/24 and 192.168.2.0/24, that are connected by a logical router. Each subnet has a VM with an IP address assigned. The subnets and VM IP addresses overlap between the two VPCs.




Figure 2: VPC Subnet

The communication between VMs in the same subnets or different subnets in the same VPC (also called East-West communication) is enabled using Generic Network Virtualization Encapsulation (Geneve). If a Prism Central manages multiple clusters, then the VMs that belong to the same VPC could be deployed across different clusters. The virtual switch on the AHV nodes provide distributed virtual switching and distributed virtual routing for all VPCs.

The communication from a VM in a VPC to an endpoint outside the VPC (called external communication or North-South communication) is enabled by an external network connection. Such a connection may be secured using VPN.

Note: You must configure the default route (0.0.0.0/0) to the external subnet as the next hop for connectivity outside the cluster (north-south connectivity).

The following figure shows the logical connectivity of the VPCs to the external network, and subsequently to the Internet.




Figure 3: External Communication

Transit VPC

For external connectivity, connect a user VPC to a transit VPC or an Overlay External Subnet with external connectivity. You could use a maximum of one NAT and one No-NAT external network for a given VPC.

Transit VPC use a hub-and-spoke architecture. Transit VPCs are useful in the following cases.

- Transit VPCs simplify and scale routing configuration (for North-South traffic) for large number of VPCs by introducing a hub VPC in the path. This minimizes the need for dynamic routing advertisement to infrastructure routers or configuring infrastructure router statically.
- Transit VPCs enable you to route traffic between user VPCs using private IPv4 addresses (using Externally Routable Prefix or ERP routes), thus allowing user VPCs to access resources you have in one of your regular VPCs. An added advantage is that traffic does not need to be routed on the physical infrastructure.
- Transit VPCs enable hosting shared services among VPCs (by hosting these services on overlay subnets under a transit VPC).
- Transit VPCs allow a logical separation between provider (transit VPC) and tenant (user VPC) network in a multi-tiered model. Multi-tiered models allow for layers of access control where each tenant controls their own routing and security policies, whereas transit VPCs allow the administrators to control the routing and security policies in the layer above the tenant layer.
- Transit VPCs allow for routing and policy control over cross-tenant communication without touching the physical infrastructure.

Conditions applicable to transit VPCs:

- Use VLAN subnets with external connectivity for North-bound connections of a transit VPC.
- Use Overlay subnets with external connectivity (Overlay external subnet) for South-bound connections from a transit VPC to non-transit or user VPCs. Overlay subnets with external connectivity can only connect to transit VPCs.
- Use the Overlay subnet without external connectivity to connect a transit VPC with entities such as VMs.
- Configure Externally Routable Prefixes (ERPs) on the VPCs to ensure that the transit VPC has a route to the Overlay subnets for the VPCs.
- When you connect a transit VPC to a VLAN backed No-NAT external network, deploy a Border Gateway Protocol (BGP) gateway to advertise the networks that connect through the transit VPC. Scale out the No-NAT gateways to provide maximum connectivity. For more information on scaling out No-NAT gateways, see *No-NAT Gateways* section in the following pages.
- Floating IP addresses supported for Recovery Plans in Disaster Recovery do not work if the floating IP addresses are configured for transit VPCs.

Note: When multiple regular VPCs with externally routable prefixes (ERPs) are connected by a transit VPC, while adding ERPs for the transit VPC, add the ERPs of the regular VPCs that are required to be advertised by the BGP Gateway of the transit VPC.

The BGP gateway of the transit VPC services only the transit VPCs and not the regular VPCs connected by the transit VPC. Therefore, in a transit VPC that provides external connectivity to the regular VPCs connected to it, the ERPs of the regular VPCs must be added as ERPs of the transit VPC.

If an ERP of a regular VPC connected to the transit VPC is not added to the list of ERPs of the transit VPC, the BGP gateway of the transit VPC does not advertise the ERP. Prism Central also raises an alert with alert ID, **802007**. For information on alerts, see [Prism Central Alerts and Events Reference Guide](#).

For example, two VPCs connect to the virtual router of the transit VPC through a No-NAT Overlay External Subnet. The transit VPC connects to the network infrastructure through a No-NAT External Network (that may be a No-NAT VLAN External Network). The VPCs also connect to the transit VPC through a NAT Overlay External Subnet. The transit VPC connects to the internet through a NAT External Subnet (that may be a NAT VLAN External Network).




Figure 4: Transit VPC in a network

External Subnets

Subnets that provide external connectivity to a VPC are external subnets. External subnets may be subnets within the deployment but not included in a specific VPC. External subnets may also be subnets that connect to the endpoints outside the deployment such as another deployment or site.

External subnets can be deployed with NAT or without NAT. You can add a maximum of two external subnets - one external subnet with NAT and one external subnet without NAT to a VPC. Both external subnets cannot be of the same type. For example, you cannot add two external subnets, both with NAT.

You can deploy VLAN subnets (Network Controller based VLANs) or Overlay subnets as external subnets. However, an Overlay subnet deployed as an external subnet (Overlay external subnet) can be attached to only a transit VPC. You cannot attach an Overlay External subnet to a regular, non-transit VPC.

Primary and Secondary IP Addresses for VMs

For information on Primary and Secondary IP Addresses, see [VM IP Address Management](#) on page 77.

SNAT and Floating IP Address

SNAT and Floating IP addresses are used only when you use NAT for an external subnet.

In Source Network Address Translation (SNAT), the NAT router modifies the IP address of the sender in IP packets. SNAT is commonly used to enable hosts with private addresses to communicate with servers on the public Internet.

For VMs within the VPC to communicate with the rest of the deployment, the VPC must be associated with an external network. In such a case, the VPC is assigned a unique IP address, called the SNAT IP, from the subnet prefix of the external network. When the traffic from a VM needs to be transmitted outside the VPC, the source IP address of the VM, which is a private IP address, is translated to the SNAT IP address. The reverse translation from SNAT IP to private IP address occurs for the return traffic. Since the SNAT IP is shared by multiple VMs within a VPC, only the VMs within the VPC can initiate connections to endpoints outside the VPC. The NAT gateway allows the return traffic for these connections only. Endpoints outside the VPC cannot initiate connections to VMs within a VPC.

In addition to the SNAT IP address, you can also request a Floating IP address — an IP from the external subnet prefix that is assigned to a VM via the VPC that manages the network of the VM. Unless the floating IP address is assigned to the private IP address (primary or secondary IP address) of the VM, the floating IP address is not reachable. When the VM transmits packets outside the VPC, the private IP of the VM is modified to the Floating IP. The reverse translation occurs on the return traffic. As the VM uses the Floating IP address, an endpoint outside the VPC can also initiate a connection to the VM with the floating IP address.

The translation of the private IP addresses to Floating IP or SNAT IP address, and vice versa, is performed in the hypervisor virtual switch. Therefore, the VM is not aware of this translation. Floating IP translation may be performed on the hypervisor that hosts the VM to which the floating IP is assigned to. However, SNAT translation is typically performed in a centralized manner on a specific host.

Network Address Translation

Network Address Translation (NAT) provides a method to map the IP addresses of an internal or private subnet to a public IP address that can communicate with the internet or other subnets. It is a process for modifying the source or destination addresses in the headers of an IP packet when the packet is put in transit. In general, the sender and receiver applications are not aware that the IP packets are being manipulated.

For example, consider the following scenario:




Figure 5: NAT

When VPC#1 and VPC #2 need access to a common segment of the overall organization's network, there would be conflicts with overlapping IP addresses in the common segment, VPC#1, and VPC#2 subnets. Using a NAT external subnet in this scenario eliminates the conflicts and connectivity issues. When the two VPCs (#1 and #2) communicate with each other as well, conflicting IP address would lead to connectivity issues. Especially while connecting to unknown subnets or the Internet, NAT provides security and masking.




Figure 6: NAT Access to Common Segment




Figure 7: NAT Access between VPCs

NAT Gateways are used only when you use Network Address Translation (NAT) for an external subnet.

NAT Gateway

A NAT Gateway service provides the entities inside an internal network with connectivity to the Internet without exposing the internal network and its entities. It performs the process of Network Address Translation as a service.

A NAT Gateway service works as follows:

- A NAT Gateway service is deployed as an AHV host. You need an AHV host to implement a NAT Gateway service because NAT gateway services involve and require operations like load balancing and routing that are automatically performed by Flow Virtual Networking. One of the AHV hosts in a cluster (that also hosts the Prism Central AZ) is deployed as the NAT Gateway.
- A NAT Gateway service is connected to the internal network with an internal subnet IP address and to the external network with an externally-routable IP address.

The externally-routable IP address is an IP address selected from IP address pool of the external subnet configured for the VPC.

No-NAT Gateway

Like the NAT Gateway service, the No-NAT gateway service also provide external connectivity. However, it does not perform Network Address Translation.

The No-NAT gateway service selects an AHV host or node from the Prism Element cluster to act as the gateway and route the external traffic. You can deploy scale-out gateway services with up to four AHV hosts acting as gateways, when you create a VPC with a VLAN subnet providing external connectivity. The external (North South) traffic for the VPC is distributed across the number of AHV hosts or nodes selected for the VPC.

For information on setting External connectivity for a VPC, see *External Connectivity* in the table in [Creating a Virtual Private Cloud](#) on page 89.

The following considerations apply to No-NAT scale-out gateway services providing up to four No-NAT gateways:

- You can deploy a scale-out No-NAT gateway only if you attach a No-NAT VLAN external subnet (not an Overlay external subnet).
- The externally-routable IP address may be an IP address from a private IP address space or a private network (RFC1918) address.
- The No-NAT gateway IP address can be manually selected or chosen dynamically from the IP pool of the external subnet.

Static IP Address

A static IP address is a fixed IP address that is manually assigned to an interface in a network. Static IP addresses provide stable routes that do not have to be updated frequently in the routing table since the static routes generated using static IP addresses do not need to be updated.

Usually in a large IP-based network (a network that uses IP addresses), a Dynamic Host Configuration Protocol or DHCP server assigns IP addresses to interfaces of an entity (using DHCP client service on the entity). However, some entities may require a static IP address that can be reached (manual remote access or via VPN) quickly. A static IP address can be reached quickly because the IP address is fixed, assigned manually and is stored in the routing table for a long duration. For example, a printer in an internal network would need a static IP address so that it can be connected reliably. Static IP addresses can be used to generate static routes which remain unchanged in routing tables, thus providing stable long-term connectivity to the entity that has the static IP address assigned.

Virtual IP Address

Any IP address in a VPC subnet, that is assigned, manually or otherwise, to an entity like a VM may be termed as a virtual IP address.

Do not confuse this virtual IP address with the virtual IP addresses assigned to Prism Central or Prism Element cluster.

Static Route

Static routes are fixed routes that are created manually by the network administrator. Static routes are more suited for small networks or subnets. Irrespective of the size of a network, static routes may be required in a variety of cases. For example, in VPCs where you use virtual private networks (VPNs) or Virtual Tunnel End Point (VTEP) over VxLAN transport connections to manage secure connections, you could use static routes for specific connections such as site-to-site connections for disaster recovery. In such a case it is necessary to have a known reliable route over which the disaster recovery operations can be performed smoothly. Static routes are primarily used for:

- Facilitating the easy maintenance of the routing table in small networks that are not expected to grow.
- Routing to and from other internal route or stub networks. A stub network or an internal route network is a network accessed using a single route and the router has only one neighbor.
- Use as a default or backup route. Such a route is not expected to specifically match any other route in the routing table.

In a network that is not constantly changing, static routes can provide faster and more reliable services by avoiding the network overheads like route advertisement and routing table updates for specific routes.

Reroute Policy

The network controller supports traffic rerouting through one service IP address for both directions or two separate IP addresses for incoming and outgoing traffic.

You can set a **Fallback Action** for the reroute policy. The **Fallback Action** is initiated when the service VM IP address is not reachable. You can configure a **Fallback Action** from the drop down menu. Flow Virtual Networking allows you to configure **Pass-through, Drop, Allow** or **No Action**. Select the Re-route option to configure the traffic routing for entities like High Availability (HA) firewall VMs with single-legged or 2-legged firewall configurations.

For example, when you want to persist the IP address assigned to an entity like a firewall VM to ensure that traffic is sent to a specific IP address irrespective of the entity it is assigned to, create a Reroute policy with the **No Action** option. Select **Re-route** in **Actions**. Do not select **Configure separate reroute IP for incoming and outgoing traffic** and enter the **Reroute IP Address (Incoming and Outgoing traffic)** for the **No Action** selection in **Fallback Action**. This configuration works only for a single-legged Firewall VM configuration. The single Re-route IP address for both incoming and outgoing traffic leads to looping of traffic.

For a two-legged firewall VM configuration, select **Configure separate reroute IP for incoming and outgoing traffic** and configure a **Reroute IP Address (Incoming Traffic)** for the inside interface and a **Reroute IP Address (Outgoing Traffic)** for the outside interface of the two-legged deployment.

For a three-legged firewall design that includes a Demilitarized Zone (perimeter network), select **Configure separate reroute IP for incoming and outgoing traffic** and configure a **Reroute IP Address (Incoming Traffic)** for the inside interface, a for the outside interface and the **Destination** IP address for the perimeter network interface.

For more information on Re-route configurations, see [Creating a Policy](#) on page 97.

VLAN Basic Subnets (or Basic VLANs)

VLAN Basic Subnets refer to the AHV networking based VLANs that Acropolis creates while creating the AHV clusters (VLAN0 - default VLAN that is used to network the CVMs and AHV hosts) or the VLANs that you create to network the guest VMs using the **Network Configuration** page in Prism Element Web Console.

These traditional AHV VLAN with or without IP management (VLAN Basic Subnets networks with or without IPAM) are managed by Acropolis. Therefore, you can create or manage these VLAN Basic Subnets in the Prism Element Web Console and in Prism Central.

For information on VLAN Basic Subnets, see [AHV Administration Guide](#).

VLAN Subnets (or VLANs)

You create or manage the Network Controller VLANs (or just VLANs) using the network controller. You can only create or manage these VLANs in Prism Central. You cannot use Prism Element Web Console to create or manage these VLANs. The Network controller does not drop unicast traffic when it is specifically supported in VLAN Subnets (Network Controller based VLAN).

If you need to use Network Controller VLANs (VLANs) to the latest networking and network security features such as Flow Network Security Next-Gen.

You cannot migrate the VLANs to Basic VLANs. For information on migration of networks and VMs, see [VM and Network Migration](#) on page 80.

For information on the requirements and limitations of VLANs, see [Network Types](#) on page 40.

Overlay subnets

You can create an IP-based Overlay subnet for a VPC. An Overlay subnet is a virtualized network that is configured on top of an underlying virtual or physical network. A peer-to-peer network or a VPN are

examples of Overlay subnets. An important assumption for the underlying network is connected such that the set of AHV hosts using the same VPCs must have layer 3 connectivity.

There are two types of Overlay subnets and their conditions are:

- Overlay subnets without external connectivity or regular Overlay subnets:
 - Overlay subnets are regular IP-based subnets *without* external connectivity.
 - You can attach an Overlay subnet to regular VPCs or transit VPCs to connect the VPC or transit VPC to VMs or workload entities.
- Overlay subnets with external connectivity or Overlay external networks:
 - You can attach an Overlay external network **only** to a transit VPC.
 - You can connect only VPCs to an Overlay external network. You cannot connect VMs or workload entities to an Overlay external network.
 - You can configure an Overlay external network of either the NAT or the No-NAT type. The No-NAT Overlay external subnet does not support No-NAT gateway scale-out. For information on No-NAT gateway scale-out, see **No-NAT Gateway** in this section.

Traffic Behavior

Broadcast Traffic

Flow Virtual Networking forwards the broadcast traffic to all the guest VMs in the same subnet, irrespective of which AHV hosts these VMs are running on.

Unicast Traffic

Flow Virtual Networking transmits unicast traffic based on the configured networking policies.

Unknown Unicast Traffic

Flow Virtual Networking drops unknown unicast traffic. It is not transmitted to any guest VM within or outside the source AHV host.

Multicast Traffic

Inside a Flow Virtual Networking VPC, multicast traffic is forwarded only within a subnet and to all VMs in that subnet. Currently there is no IGMP snooping within VPCs.

REQUIREMENTS AND LIMITATIONS OF FLOW VIRTUAL NETWORKING

Before enabling the Network Controller for Flow Virtual Networking, make sure the Nutanix environment fulfils these requirements and you are aware of the limitations.

Requirements

Ensure that the following requirements are met before you enable the Flow Virtual Networking Network Controller (or the Network Controller) on Prism Central.

- **Prism Central, AHV Hosts and Preferred Site**

Nutanix strongly recommends that you deploy a three-node scale-out Prism Central for production deployments, although Flow Virtual Networking might be enabled on a single-node Prism Central. The availability of Flow Virtual Networking services in Prism Central is critical for performing operations on VMs that are connected to Overlay or VLAN Subnets. A three-node scale-out Prism Central ensures that Flow Virtual Networking continues to run even if one of the nodes with a Prism Central VM fails.

Interruptions to Network Controller services can cause loss of connectivity upon live migration of guest VMs networked in Overlay Subnets or Network Controller-backed VLAN Subnets. When the Network Controller is down or connectivity between Prism Central and connected AHV clusters is interrupted, and any VMs networked in Overlay subnets or VLAN Subnets are migrated, the migrated VMs might become unreachable until the Network Controller service and connectivity is restored.

Flow Virtual Networking VPC subnets and Network Controller-backed VLAN Subnets require reliable connectivity between the Prism Central Network Controller and registered AHV clusters. Nutanix recommends that all AHV clusters reside at the same site or data center as their registered Prism Central instance when using VPC subnets or Network Controller-backed VLAN Subnets to avoid network control plane interruption. Each site must have a local Prism Central when using Flow Virtual Networking VPCs. You can exclude specific AHV clusters from Flow Virtual Networking VPCs using CLI configurations.

- **User Role**

Ensure that you log on to Prism Central with **Prism Admin** role.

You can log on as a local account user or a non-local account (IDP-based) user. Your user account must be configured with the **Prism Admin** role privileges to use Flow Virtual Networking.

If you log on to Prism Central without the **Prism Admin** role privileges, then Prism Central does not allow you to enable or use Flow Virtual Networking. The task is reported as **Failed** with a User Denied Access message.

For more information on the user access, roles, and permissions required to enable and operate Flow Virtual Networking, see [Control User Access in Flow Virtual Networking \(RBAC\)](#) on page 36.

- **Ports and Protocols**

Nutanix deploys a number of ports and protocols in its software. These ports must be open in the firewalls to enable Flow Virtual Networking to function. For information on the ports and protocols used for Flow Virtual Networking, see [Ports and Protocols](#).

- **Software Versions**

Ensure that the Prism Central running Flow Virtual Networking Network Controller is hosted on an AOS cluster running AHV.

The Network Controller has a dependency only on the AHV and Prism Central versions. Ensure that the nodes in all the clusters managed by the same Prism Central are running the same compatible AHV version. For

information on compatible Network Controller, AHV, and Prism Central versions, see *Software Compatibility* in the Flow Virtual Networking [Release Notes](#).

When you deploy Prism Central with the Network Controller, the prechecks that are run include a check of the AOS and AHV versions. An incompatible version of AHV, Prism Central creates the Network Controller but issues an alert (Failed to configure host for Atlas networking) during Network Controller enablement. For more information on the alert, see [Prism Central Alerts and Events Reference Guide](#).

Upgrade the AOS and AHV versions, as applicable, to the compatible versions.

- **Microservices Infrastructure**

Microservices Infrastructure is enabled by default on a Prism Central that is running pc.2022.9 or later version. For more information, see [Microservices Infrastructure](#) in the *Prism Central Infrastructure Guide*.

- **Prism Central Size Supported**

Small, Large and X-Large Prism Central deployments support Flow Virtual Networking. Flow Virtual Networking is not supported on X-Small Prism Central instances.

The Network Controller is auto-enabled when you install or upgrade an X-Large Prism Central to pc.2023.3 or later versions.

Note: Before you enable the Network Controller on a Small or Large Prism Central, ensure that the Prism Central instance is registered to the same Prism Element cluster that hosts the Prism Central VM(s).

- **Resource Requirement Per Prism Central VM**

When you enable the Network Controller on a Small or Large Prism Central deployment, the deployment requires resources per Prism Central VM in addition to the resource requirement of the Small or Large Prism Central. Ensure that the additional resources are available to the Prism Central deployment before enabling the Network Controller.

- For Flow Virtual Networking on a small Prism Central: Every Prism Central VM requires additional 3 GB memory and 2 vCPUs.
- For Flow Virtual Networking on a large Prism Central: Every Prism Central VM requires additional 4 GB memory and 3 vCPUs.

If the additional resources are not available on the hosting nodes, then Network Controller is not enabled.

- **Resource Requirement Per AHV Host**

When you enable the Network Controller on a Prism Central deployment, the deployment requires 2 GB of memory per AHV host.

- **Connectivity**

Flow Virtual Networking requires reliable connectivity between Prism Central and registered AHV clusters. Ensure that all AHV clusters reside at the same site or data center as their registered Prism Central instance. Do not register AHV clusters to a Prism Central at a remote site when using Flow Virtual Networking Virtual Private

Clouds (VPCs). Each site requires a local Prism Central when using Flow Virtual Networking. You can exclude specific AHV clusters from Flow Virtual Networking using CLI configurations.

Ensure connectivity:

- Between Prism Central and its managed Prism Element clusters.
- To the Internet for connectivity (not required for dark site) to:
 - ECR for Docker images
 - S3 storage for LCM portal

Note: For dark site deployments, Nutanix provides a dark site bundle, which has the Docker images (normally hosted on ECR) and the Network Controller package (normally hosted on LCM portal). These dark site bundles can be downloaded using an internet-connected system outside the dark site.

- **Prism Central VM registration**

You cannot unregister the Prism Element cluster that is hosting the Prism Central deployment where you have enabled Flow Virtual Networking. You can unregister other clusters being managed by this Prism Central deployment.

- Ensure that you have created a virtual IP address (VIP) for Prism Central. Once set, do not change this address.
- **MTU Settings**

Nutanix recommends increasing the MTU to 9000 bytes on the virtual switch vs0 and ensure that the physical networking infrastructure supports higher MTU values (jumbo frame support). Nutanix recommends configuring the MTU value in the range of 1500 to 9000 bytes.

Note:

If you try to configure an MTU value that does not fall within the range of 1500 to 9000 bytes on the default virtual switch vs0, Prism Central displays an error and fails to apply the configuration.

By default, the Nutanix Controller VMs use the standard Ethernet MTU (maximum transmission unit) of 1,500 bytes for all the network interfaces. The system advertises the MTU of 1442 bytes to guest VMs using DHCP to account for the extra 58 bytes used by Generic Network Virtualization Encapsulation (Geneve). However, some VMs ignore the MTU advertisements in the DHCP response. Therefore, to ensure that Flow Virtual Networking

functions properly with such VMs, enable jumbo frame support on the physical network and the default virtual switch vs0.

If you cannot increase the MTU of the physical network, decrease the MTU of every VM in a VPC to 1442 bytes in the guest VM console.

Note: Do not change the MTU of the CVM.




Figure 8: Sample Configurations with and without Higher MTU: VS0, CVM and UVMs

Table 1: Flow Virtual Networking MTUs

Feature	MTU (Overhead Calculation)
VPC	Regular Geneve = 1442 (1500 - 58 bytes Geneve)
VPC + Subnet Extension	Geneve + VXLAN = 1392 (1500 - 58 bytes Geneve - 50 bytes VXLAN)
VPC + VPN	Geneve + IPsec = 1356 (1500 - 58 bytes Geneve - 86 bytes IPsec)
VPC + VTEP + VPN	Geneve + VXLAN VTEP + IPsec = 1306 (1500 - 58 bytes Geneve - 86 bytes IPsec - 50 bytes VXLAN)

Requirements for Upgrades

The following applies to Network Controller (**Advanced Networking** in **Prism Central Settings**) upgrades:

- Compatible AHV Versions**

Ensure that the AHV hosts in the Prism Element clusters managed by a Prism Central that has Network Controller enabled are running an AHV version compatible with the Network Controller **upgrade version**. The Network Controller is upgraded but not enabled, if any of the AHV hosts is running an incompatible version.

Important: Before you upgrade the Prism Central version to upgrade the Network Controller, upgrade the AHV version on the hosts with incompatible AHV versions using LCM to the AHV version compatible with the Network Controller upgrade version.

For information on compatible AHV versions, see the [Release Notes](#). For information on Prism Central, AHV and AOS version compatibility, see the [Compatibility and Interoperability Matrix](#).

Note: When the Network Controller is deployed with a compatible Prism Central deployment package but with incompatible AHV package, the Network Controller is deployed with Prism Central, but not enabled.

- Ensure that all the AHV hosts in the AOS cluster are running the version compatible with the Network Controller upgrade version.

Limitations

The following are the limitations of Flow Virtual Networking:

- The Flow Virtual Networking Network Controller is supported only on AHV clusters. It is not supported on ESXi or Hyper-V clusters.
- Flow Virtual Networking is not supported in clusters with Compute-only nodes.
- The Flow Virtual Networking Network Controller is not enabled by default on a new Prism Element cluster registered with the Flow Virtual Networking-enabled Prism Central if the Prism Element cluster has nodes with incompatible AHV versions.
- The Flow Virtual Networking Network Controller does not support updating a VLAN Basic Subnet that is configured as an external subnet.

You cannot enable the external connectivity option in the **Update Subnet** dialog box. Therefore, you cannot modify an existing VLAN-backed subnet to add external connectivity.

VLAN Basic Subnets for external connectivity are managed by the Flow Virtual Networking control plane. Traditional AHV VLAN IPAM networks are managed by Acropolis.

Note: Do not configure the same VLAN as both a Flow Virtual Networking external network and an AHV IPAM network, as this can lead to IP address conflicts.

- You cannot disable the Flow Virtual Networking Network Controller if any external subnets and VPCs are in use. Delete the external subnets and VPCs and then disable Flow Virtual Networking.

FLOW VIRTUAL NETWORKING CONFIGURATIONS

The Flow Virtual Networking Network Controller is auto-enabled when you install an X-Large Prism Central instance or upgrade the version of your existing X-Large Prism Central instance to pc.2023.3 or later. On Small and Large Prism Central instances, you need to enable the Network Controller. Flow Virtual Networking is not supported on X-Small Prism Central instances.

When you select **Subnets** (see step 2 in [Subnets Summary View](#) on page 49) for the first time, a dialog box, indicating that Flow Virtual Networking is *auto enabled*, is displayed:

Network Controller Settings View

- Log in to Prism Central.
- Click **Prism Central Settings** from the **Navigation Bar** of the **Infrastructure** application.
- On the **Prism Central Settings** page, click **Network Controller**.

The **Network Controller (formerly Advance Networking)** page opens.

☆

Network Controller (formerly Advanced Networking)



Service Status

● Enabled

Health Status

● Good

Version

3.0.0

Check for Updates

Disable Network Controller

Network Controller for VLAN Management

Enable for centralised VLAN management and latest networking & security features including Flow Network Security 2.0. Set this as default for easy subnet creation, can be changed during subnet creation.

☐ Set as default

Unchecking this will default subnet creation to VLAN Basic. [Learn More](#)

☒ Use the [VLAN migrate workflow](#) to convert VLAN Basic subnets to Network Controller managed VLAN Subnets.

Clusters and Compatibility

Cluster :	AOS :	AHV :	Compatibility :
Q			
auto_clus...	6.6.1	el7.nutanl...	● Upgrade AHV
auto_clus...	6.7	el8.nutanl...	● Supported

Figure 9: Network Controller (formerly Advance Networking)

The **Network Controller (formerly Advanced Networking)** page displays the following:

- **Service Status**—This section displays the status as **Enabled** when the Network Controller is enabled.
- **Health Status**—This section displays **Good** for a healthy Network Controller.
- **Version**—The version of Network Controller such as 4.0.0.
- **Check for Updates**—This link helps you check for available Network Controller upgrades through the LCM page.
- **Disable Network Controller**—This link helps you disable the Network Controller. For more information, see [Disabling the Network Controller](#) on page 31.
- **Network Controller for VLAN Management**—This section provides the **Set as default** check box is clear by default to ensure that VLAN Basic (AHV based VLANs) is the default VLAN type. Select the **Set as default** check box to make VLAN Subnets (Network Controller based VLANs) the default VLAN type.

This section also provides information about migration of VLAN Basic Subnets to VLAN Subnets. For information on the types of networks that Flow Virtual Networking creates and manages, see [Network Types](#) on page 40.

- **Clusters and Compatibility**—This section displays a table with information on the AOS and AHV versions of the cluster and hosts in the clusters that the Network Controller spans over. If any cluster requires an upgrade, this status is indicated in the **Compatibility** column of the table.

For information on migrating VMs from AHV-based VLANs or VLAN Basic Subnets to Overlay subnets, see [Migration of VMs between VLAN Basic Subnet and VPC Subnets](#) on page 80.

For information on converting the AHV-based VLANs or VLAN Basic Subnets to Network Controller based VLANs or VLAN Subnets, see [Migration of VLAN Basic Subnets](#) on page 85.

Enabling the Network Controller

If you have a Small or Large Prism Central deployment, you need to manually enable the Network Controller.

About this task

Before you proceed to enable the Network Controller by clicking the **Network Controller** option on the **Prism Central Settings** page, see [Requirements and Limitations of Flow Virtual Networking](#) on page 21.

To enable the Network Controller, follow these steps:


Procedure

1. Log in to Prism Central.
2. Click **Prism Central Settings** from the **Navigation Bar** of the **Infrastructure** application. The **Prism Central Settings** page opens.
3. Click **Network Controller**.

4. In the **Network Controller (formerly Advanced Networking)** pane, click **Enable**.

Ensure that the prerequisites specified on the pane are fulfilled.

Network Controller (formerly Advanced Networking)



Network Controller is a component of Prism Central that manages and controls configuration, monitoring and optimization of Network resources. It provides programmability, automation, and centralized control for configuring and managing network flows.

Network Controller is necessary to use centralized VLAN management, Flow Network Security Next-Gen or Flow Virtual Networking.

Requirements

Network Controller requires 2 vCPUs and 3GB RAM for small Prism Central VMs and 3 vCPUs and 4GB RAM for large Prism Central VMs. The VMs will be resized automatically upon enabling Network Controller.

To enable Network Controller ensure that

- 1 Prism Central can access download.nutanix.com
- 2 Prism Central has [Microservice Infrastructure](#) configured.

After enabling,

- Network Controller will be available for all AHV Clusters running AOS 6.1 or higher.
- The default configuration when creating new subnets will be VLAN Basic. You can change the subnet type during creation under the Advanced Configuration option or change the default to set Network Controller managed subnet.
- VLAN Basic subnets can be migrated to Network Controller managed VLAN subnets.

Enable

Figure 10: Enabling Network Controller

Prism Central displays the deployment progress and completion.

Disabling the Network Controller

You can disable the Flow Virtual Networking Network Controller.

About this task

If any external subnets and VPCs are in use, you cannot disable the Network Controller. Delete the external subnets and VPCs, and then disable the Network Controller.

Procedure

1. Log in to Prism Central.
2. Click **Prism Central Settings** from the **Navigation Bar** of the **Infrastructure** application. The **Prism Central Settings** page opens.
3. Click **Network Controller**.
4. On the **Network Controller (formerly Advance Networking)** page, click **Disable Network Controller**.
5. On the confirmation message box, click **Confirm** to confirm disablement.

To exit without disabling the **Network Controller**, click **Cancel**.

Disabling Network Controller to Unregister a Prism Element Cluster

Before unregistering a Prism Central from the Prism Element cluster, disable the Network Controller on that Prism Element cluster.

About this task

When Flow Virtual Networking is enabled on a Prism Central, it propagates the capability to participate in VPC networking to all the registered Prism Elements that are running the required AHV version.

In cases where there are VMs on the Prism Element attached to the VPC network, or if the Prism Element is used to host one or more of the external VLAN networks attached to a VPC, Prism Central alerts you with a prompt. When being alerted about the aforementioned conditions, close the CLI and make adequate configuration to resolve the condition (for example, select a different cluster for the external VLAN network and delete the VMs attached to the VPC network running on the Prism Element). After making such configurations, execute the network controller CLI to disable Flow Virtual Networking. If the command goes through successfully, it is safe to unregister the Prism Element.

For example, in a deployment of three Prism Elements - PE1, PE2 and PE3 - registered to the Flow Virtual Networking-enabled Prism Central, you need to unregister PE3 from the Prism Central. You must first disable Flow Virtual Networking using the steps in [Disabling the Network Controller](#) on page 31 or the following steps:

Procedure

1. SSH to PE3.
2. Run the `ncli cluster info` or `ncli cluster get-params` command to display the cluster parameters.
Copy the cluster UUID (For example: 017457d3-1012-465c-9c54-aa145f2da7d9) from the displayed cluster parameters.
3. SSH to the Prism Central VM.

4. Open the network controller console by executing the `atlas_cli` command.

```
nutanix@pcvm$ atlas_cli  
<atlas>
```

5. Execute the `config.add_to_excluded_clusters <cluster uuid>` command, providing the cluster UUID that you copied earlier.

An example of the Prism Central alert, for the condition that PE3 VM is attached to an external network, is as follows:

```
<atlas> config.add_to_excluded_clusters 0005bf8d-2a7f-3b2e-0310-d8e34995511e  
Cluster 0005bf8d-2a7f-3b2e-0310-d8e34995511e has 1 external subnet,  
which will lose connectivity. Are you sure? (yes/no)
```

Note: To enable Flow Virtual Networking on the cluster, execute the `config.remove_from_excluded_clusters <cluster uuid>` command, providing the cluster UUID.

What to do next

To verify if the Flow Virtual Networking Network Controller is disabled, SSH to PE3 and run the `acli atlas_config.get` command.

The output displays the `enable_atlas_networking` parameter as `False` if Flow Virtual Networking is disabled and as `True` if Flow Virtual Networking is enabled on the Prism Element.

```
nutanix@cvm$ acli atlas_config.get  
config {  
  anc_domain_name_server_list: "10.10.10.10"  
  enable_atlas_networking: False  
  logical_timestamp: 19  
  minimum_ahv_version: "20190916.101588"  
  ovn_cacert_path: "/home/certs/OvnController/ca.pem"  
  ovn_certificate_path: "/home/certs/OvnController/OvnController.crt"  
  ovn_privkey_path: "/home/certs/OvnController/OvnController.key"  
  ovn_remote_address: "ssl:anc-ovn-external.default.anc.aj.domain:6652"  
}
```

You can now unregister the Prism Central from the Prism Element cluster. For steps to unregister a Prism Central from a Prism Element cluster, see [Unregistering a cluster from Prism Central](#)

Upgrading the Network Controller

You can upgrade the Flow Virtual Networking controller (*Advanced Networking Controller* in *Prism Central Settings*) using Life Cycle Manager (LCM) on Prism Central.

Before you begin

See [Requirements and Limitations of Flow Virtual Networking](#) on page 21.

In case of upgrading the Flow Virtual Networking controller in a dark site, ensure that LCM is configured to reach the local web server that hosts the dark site upgrade bundles.

Note:

The network controller upgrade fails to start after the pre-check if one or more clusters have Flow Virtual Networking enabled and are running an AHV version incompatible with the new network controller upgrade version.

About this task

To upgrade the network controller using LCM, do the following.

Procedure

1. Log in to Prism Central.
2. Select the **Admin Center** application from the [Application Switcher Function](#), and click **LCM** from the **Navigation Bar**.
The **LCM** page opens displaying the **Best Practices** tab.

3. Click the **Inventory** tab.

4. Click **Perform Inventory**.

When you click **Perform Inventory**, the system scans the registered Prism Central cluster for software versions that are running currently. Then it checks for any available upgrades and displays the information on the LCM page under the **Updates** tab.

5. Click the **Updates** tab.

The **Updates** page opens displaying the available software updates.

6. Select the check box associated with **Networking Controller** and click **View Upgrade Plan**.
The **Review Update Plan** window opens.

7. Click **Apply 1 Updates**.

Dark Site Installation and Upgrade

Dark sites are primarily on-premises installations which do not have access to the internet. Such sites are disconnected from the internet for a range of reasons including security. To install or upgrade the Network Controller at such dark sites, you need to deploy the Network Controller bundle at the site.

This dark site deployment procedures include downloading and deploying the LCM dark site server bundles, downloading and deploying the Nutanix Compatibility bundle to ensure that the latest product meta data is available, and the network controller bundles.

See [Requirements and Limitations of Flow Virtual Networking](#) on page 21.

Prerequisite steps

You need access to the Nutanix Portal from an Internet-connected device to complete these steps.

Note: For dark site deployments, Nutanix provides a dark site bundle, which has the Docker images (normally hosted on ECR) and the Network Controller package (normally hosted on LCM portal). These dark site bundles can be downloaded using an internet-connected system outside the dark site.

Do the following before you install or upgrade the Network Controller:

- Update the LCM Framework. For more information, see [Updating the LCM Framework Using a Web Server](#) in the *Life Cycle Manager Dark Site Guide*.
- Install and prepare the LCM Dark Site server. For more information, see [Setting up a Local Web Server](#) in the *Life Cycle Manager Dark Site Guide*.

Take note of the FQDN or IP address of the LCM Dark Site server (Local Web Server). For example, in this documentation, `<LCM-web-server-ip>` is used to indicate the IP address of the LCM Dark Site server and `~/release` is the path of the dark site server folder.

- Ensure that you have configured the **Dark Site (Local Web Server)** settings on the **LCM > Settings** page.

- Update the firmware specific to the installed platform hardware (including the Nutanix Compatibility bundle). For more information, see [Fetching the Firmware Update Bundle Using a Web Server](#).

Note: After you have downloaded the Nutanix Compatibility bundle tar.gz file, verify if the contents match the following output:

```
[root@<LCM-web-server-ip> ~]$ tar -tvf
nutanix_compatibility_bundle.tar.gz

-rw-r--r-- jenkins/jenkins nutanix_compatibility.tgz
-rw-r--r-- jenkins/jenkins nutanix_compatibility.tgz.sign
-rw-rw-r-- jenkins/jenkins nutanix_compatibility.tgz.v2.sign
-rw-rw-r-- jenkins/jenkins lcm_cert_v2.crt
-rw-rw-r-- jenkins/jenkins lcm_intermediate_v2.crt

nutanix@cvm$
```

- On the [Flow Virtual Networking Downloads](#) page, ensure that **Network Controller (formerly ANC)** is selected in the component selection dropdown menu. Download the Network Controller bundle: Copy the **Md5** value for the bundle.

Deploying the Network Controller at a Dark Site

Before you begin

See the prerequisites provided in [Requirements and Limitations of Flow Virtual Networking](#) on page 21.

Complete the *Prerequisite steps* provided in [Dark Site Installation and Upgrade](#) on page 33.

About this task

When you deploy Prism Central in a dark site, the Network Controller bundle needs to be separately downloaded for deployment by Prism Central.

In x-Large Prism Central deployments, the Network Controller is automatically enabled.

In small and large Prism Central deployments, you must manually enable the Network Controller. See [Enabling the Network Controller](#) on page 28.

To upgrade the installed Network Controller, see [Upgrading the Network Controller at a Dark Site](#) on page 35.

Procedure

1. Log on to the LCM Dark Site server (Local Web Server) with root privileges.
2. Verify that the contents of the Network Controller bundle is similar to the following sample output for the Network Controller 3.0.0 bundle:

```
[root@<LCM-web-server-ip> ~]$ tar -tzf 3.0.0.tar.gz

builds/
builds/atlas-controller/
builds/atlas-controller/3.0.0/
builds/atlas-controller/3.0.0/atlas_network_controller.tar.gz
builds/atlas-controller/3.0.0/metadata.sign
builds/atlas-controller/3.0.0/metadata.json
```

3. Extract the Network Controller bundle to `~/release`

The following is a sample of the command to extract the Network Controller bundle.

```
[root@<LCM-web-server-ip> ~]$ sudo tar -zxvf 3.0.0.tar.gz -C ~/release/
```

4. Run the following command after unpacking to ensure that the file permissions are not disrupted during the unpacking:

```
chmod -R +r builds
```

5. In Prism Central, navigate to **Admin Center > LCM > Settings**.

- Select **Source > Dark Site (Local Web Server)**
- Enter the `http://<LCM-web-server-ip>/release` in **URL**.

6. SSH into the Prism Central VM as an admin user and run the following commands.

```
admin@pcvm$ mspctl controller airgap enable --url=http://<LCM-web-server-ip>/release
```

```
admin@pcvm$ mspctl controller airgap get
```

7. Verify that the source for deployment is configured as the dark site server.

Log on to the Prism Central VM through an SSH session as a `nutanix` user, and run the following command.

```
nutanix@pcvm$ configure_lcm --print | grep -i "msp\|atlas\|dark"
```

The following sample output shows that `is_darksite` is `True`.

```
msp: {"url": "BASE_URL/msp-builds/", "flags": [], "component": "msp", "tags": []}
atlas_controller: {"url": "BASE_URL/atlas-controller/", "flags": [], "component":
"atlas_controller", "tags": []}
is_darksite: True
enable_https_darksite: False

nutanix@NTNX-10-19-57-54-A-PCVM:~$
```

Where `BASE_URL` is the source location for the bundles. This should match `http://<LCM-web-server-ip>/release`.

8. Enable **Network Controller**. For more information, see [Enabling the Network Controller](#) on page 28.

Upgrading the Network Controller at a Dark Site

This procedure lets you upgrade the Network Controller in a dark site.

About this task

The procedure to upgrade the Network controller in a dark site consists of all the steps in the [Deploying the Network Controller at a Dark Site](#) on page 34 procedure up to the step that verifies that the source for upgrades is configured as the dark site server.

After the verification step, perform the following steps.

Procedure

1. In Prism Central, navigate to **Admin Center > LCM > Inventory** and click **Perform Inventory**.
The **LCM > Updates** tab displays the **Networking Controller** upgrade version bundle.
2. Select the **Networking Controller** component.

3. Run **Pre-Upgrade > Upgrade Prechecks**.

- a. On the **Initiate Precheck?** window, click **Continue**.
LCM runs the prechecks for upgrade.
- b. When the **Precheck successful!** message is displayed, click **Return to Updates** to return to the **Updates** page.

4. Upgrade **Networking Controller**.

- a. Click **View Upgrade Plan**.
- b. On the **Review Upgrade Plan** page, click **Apply _ Updates**.
- c. Click **Return to Updates** after the upgrade is complete.

Control User Access in Flow Virtual Networking (RBAC)

Flow Virtual Networking supports role-based access control (RBAC) that you can configure to provide customized access permissions for users based on their assigned roles. The roles dashboard allows you to view information about all defined roles and the users and groups assigned to those roles.

For more information on configuring RBAC for Flow Virtual Networking, see [Controlling User Access \(RBAC\)](#) in the *Nutanix Security Guide*.

Flow Virtual Networking Roles and Permissions

Flow Virtual Networking provides certain pre-configured roles and permissions with those roles.

Prism Central provides two roles for Flow Virtual Networking management:

- **VPC Admin** which has 41 permissions pre-configured to manage Overlay or VPC networking including create, update, and delete networks.
- **Network Infra Admin** which has 24 permissions pre-configured to manage the network infrastructure (underlay) on the AHV network stack.

The table provides the list of permissions that are pre-configured for the two roles.

Entity	Permissions	VPC Admin	Network Infra Admin
Virtual Switch (DVS)	View, Create, Update and Delete	View Only	View Only
	Migrate	No	Yes
VLAN Subnets	View, Create, Update, Delete, Migrate and IP Reservation	No	Yes
IPFix	View, Create, Update and Delete	Super Admin role can only perform this operation.	
VLAN External Subnets	View	Yes	Yes
	Create, Update and Delete	No	Yes
Advanced Network Controller deployment	Super Admin and Prism Admin Roles have permissions to perform this operation. The permissions for this operation are not provided on the Roles page.		
VPC	View and View_NS_stats	Yes	No

Entity	Permissions	VPC Admin	Network Infra Admin
Overlay Subnets	Create	Yes	No
	Update	Yes	No
	Delete	Yes	No
	View, Create, Update and Delete	Yes	No
Overlay External Subnets	View	Yes	No
	Create, Update and Delete	No	No
Floating IP Addresses	View, Create, Update and Delete	Yes	No
Policy Based Routing (PBR)	View, Create, Update, Delete and Clear Containers	Yes	No
Routes	View, Update	Yes	No
Network Gateways (VPN, VTEP and BGP)	View, Create, Update and Delete	Yes	Yes
VPN Connections	View, Create, Update, Delete and Configuration Download	Yes	Yes
Layer 2 Stretch	View, Create, Update and Delete	Yes	Yes
BGP Sessions	View, Create, Update and Delete	Yes	No
VM (Not networking objects)	View	Yes	Yes
	Create, Update and Delete	No	No
Cluster (Not networking objects)	View	No	Yes
Cluster Networking capabilities	View	Yes	Yes
Uplink Bonds	View	No	Yes
Status of schedulable Nodes	View	Yes	Yes
VPC Virtual Switch mappings	View and Update	Yes	No
Layer 2 Stretch related entities	View	Yes	Yes
Availability Zones (AZs)	View	Yes	No

Flow Virtual Networking Operational Authorizations

Flow Virtual Networking requires you to have certain permissions or authorizations to complete certain tasks.

The table provides the list of permissions you need to perform various operations. Ensure that the necessary permissions are set for your role, to perform the necessary operations.

Note: This table provides a sample list of operations and necessary authorizations for Flow Virtual Networking. This list may not be a complete or extensive list.

Operations	Authorizations You need
Enable or deploy the Network Controller	- Permission to Create Network Controller. - Permission to View the VLAN Subnet on which the Network Controller is deployed.
Create Virtual Switch	- Permission to Create virtual switch - Permission to View the Host on which the virtual switch is instantiated.
Create VLAN Subnet	Permission to Full Access for the VLAN Subnet entity
Create vNIC on a VLAN Subnet or Overlay Subnet	- Necessary permissions for VM operations. Set these for the VM entity, including View Overlay Subnet and View Subnet View permission for VLAN Subnet or Overlay subnet.
Create or Update a VPC	Create VPC, Update VPC
Attach the VPC to an external subnet	View External Subnet
Create Overlay Subnet, PBR or Route	Create Overlay Subnet, Create Routing Policy for PBR or Update VPC Route Table
Update Route table	View VPC to view the VPC to which the Overlay subnet, PBR or route table is attached
Create Network Gateway on VPC	Create Network Gateway or Update Network Gateway View VPC to view the VPC
Create Network Gateway on a VLAN subnet	Create Network Gateway or Update Network Gateway View Subnet to view the VLAN subnet
Create or Update BGP Gateway on VPC (Optionally with serviced VPC specified).	Create and Update permissions for BGP Gateways. View VPC to view the VPC subnet View permissions for the serviced VPC if specified.

Operations	Authorizations You need
Create or Update BGP on VLAN (Optionally with serviced VPC specified).	• Create or Update BGP gateway. • View Subnet • View VPC to view the serviced VPC only if serviced VPC is specified. • View Network Gateway to view the local and remote network gateways
Create VPN Connection	• Create VPN Connection or Update VPN Connection • View Subnet to view the VLAN subnet
Create a Direct Connect on a VPC (For Nutanix Cloud Cluster on AWS)	• Create permission for Direct Connect. • View permission for the VPC.
Create Direct Connect Virtual Interfaces (VIFs)	• Create permission for Direct Connect. • View permission to view the Direct Connect VIF entities.
Create or Update Floating IP Address on an external subnet	• Create Floating IP or Update Floating IP • View External Subnet to view the external subnet
Attach Floating IP to vNIC of a VM	• Update Floating IP • View VM to view the vNIC of the VM
Attach Floating IP to Private IP of a VPC	• Update Floating IP • View VPC to view the VPC properties
Create a Layer 2 Network Extension on a subnet through a VPN connection	• Create permission for Layer 2 Network Extension • View permissions for the subnet that needs to be extended. • View permission for the VPN connection to be used for the extension.
Create a Layer 2 Network Extension on a subnet through a VTEP connection	• Create permission for Layer 2 Network Extension • View permissions for the subnet that needs to be extended. • View permission for the VTEP gateways to be used for the extension.

Operations	Authorizations You need
Create BGP Session	• Create BGP Session • View Network Gateway permissions for local and remote gateways.

Network Types

Flow Virtual Networking Network Controller supports Overlay and VLAN type networks.

Overlay networks

You can create an IP-based Overlay subnet for a VPC. An Overlay network is a virtualized network that is configured on top of an underlying virtual or physical network. Examples of Overlay networks are:

- You can create an Overlay subnet with external connectivity (Overlay external subnet) to connect a transit VPC to other regular VPCs.
- You can create a special purpose multicast network as an Overlay network within an existing network.
- A peer-to-peer network or a VPN.

An important assumption for an Overlay network is that the underlying network is fully connected. Nutanix provides the capability to create Overlay network-based VPCs.

For more information, see *Overlay networks* in [Essential Concepts](#) on page 12.

VLAN networks

Starting with Prism Central pc.2023.3 with AOS 6.7 and AHV 20230302.198, Network Controller 3.0.0 and later versions support the creation of VLANs (VLAN Subnets) on the Flow Virtual Networking Network Controller. The Network Controller also supports migration of VLAN Basic Subnets to VLAN Subnets subject to support and limitations information provided in the *VLAN Subnets Support* section.

For information on migration of VLAN networks, see [VM and Network Migration](#) on page 80.

VLAN Basic Subnets (or Basic VLANs)

VLAN Basic subnets are not managed by the network controller in Prism Central, and are instead managed by the Acropolis leader of their Prism Element cluster. VLAN Basic Subnets refer to the AHV networking based VLANs that Acropolis creates while creating the AHV clusters (VLAN0 - default VLAN that is used to network the CVMs and AHV hosts) or the VLANs that you create to network the guest VMs using the **Network Configuration** page in Prism Element Web Console.

These traditional AHV VLAN with or without IP management (VLAN Basic Subnets networks with or without IPAM) are managed by Acropolis. Therefore, you can create or manage these VLAN Basic Subnets in the Prism Element Web Console and in Prism Central. For more information, see *AHV Networks* in the [AHV Administration Guide](#).

You can only use Prism Central to migrate these VLAN Basic Subnets to Network Controller-based VLANs that you can manage in Prism Central (see [Migration of VLAN Basic Subnets](#) on page 85).

VLAN Subnets (VLANs)

Create or manage the VLAN Subnets (VLANs) or Network Controller managed VLANs using the Flow Virtual Networking Network Controller. You can only create or manage these VLAN Subnets in Prism Central. You cannot use Prism Element Web Console to create or manage these VLAN Subnets.

Note: Clusters with CO nodes do not support the creation of VLAN Subnets.

For more information, see *VLANs (or VLAN Subnets)* in [Essential Concepts](#) on page 12.

vNIC Subnet Change

You cannot change the VLAN Subnet associated with the VM vNIC. Instead, delete the VM vNIC and create a new vNIC and associate this to the new VLAN Subnet.

You can change the VLAN Basic Subnet associated with the VM vNIC to another VLAN Basic subnet.

You cannot update the VLAN ID of a VLAN Subnet or a VLAN Basic Subnet.

VLAN Subnets Support

VLAN Subnets (VLANs) supports the following:

- **IGMP Snooping**

For more information on IGMP snooping in Nutanix networks, see the [IGMP Snooping](#) documentation.

- **vNIC creation with Access VLAN mode**

Network Controller VLAN Subnets support only access mode, and do not support VLAN trunk mode.

- **vNIC Scale**

The Network Controller only supports VMs with vNIC associated with either the AHV networking stack or the Network Controller stack.

- **DHCP options on managed VLAN Subnets**

VLAN subnets that are managed networks (networks which use IPAM managed IP addresses) support DHCP options.

- **Traffic Mirroring**

VLAN Subnets support Traffic Mirroring. For information on Traffic Mirroring, see [Traffic Mirroring on AHV Hosts](#) in *AHV Administration Guide* and [Traffic Mirroring](#) in *Prism Central Infrastructure Guide*.

- **Traffic Support**

VLAN Subnets support broadcast, unicast including unknown unicast, and multicast traffic.

- **IPFIX Exporter**

VLAN Subnets support IPFIX Exporter.

- **TFTP Server IP Address**

If you need to configure a TFTP server for a managed network, use the IP address of the TFTP server instead of the FQDN.

VLAN Subnet (VLAN) does not support the following:

1. Nutanix Files
2. Trunk mode
3. Virtual NICs in kDirect mode
4. You cannot update a vNIC on a VLAN Subnet that was created by migrating a VLAN Basic Subnet. Delete the vNIC that needs to be updated and create a new vNIC with the updated parameters.
5. Unknown unicast flooding and disabling port security. Any VMs or workloads that depend on unknown unicast traffic are impacted during the subnet migration workflow. When the VMs are migrated to the Flow Virtual Networking network controller, port security is enabled and unknown unicast stops working.
6. Service chaining (see *Service Chain* in [Essential Concepts](#) on page 12).
7. Remote Office Branch Office (ROBO) deployments
8. Clusters with Compute-only nodes

Changing the Default VLAN Type

With a minimum Prism Central version of PC.2023.3 that deploys Network Controller 3.0.0, you can change the default VLAN creation from VLAN Subnet (VLAN) type to VLAN Basic Subnet type and vice versa.

About this task

To change the default VLAN type created using the [Creating a Subnet](#) workflow, do the following:

Procedure

1. Log on to Prism Central.
2. Select the **Infrastructure** application from the [Application Switcher](#) function.
3. To change the default VLAN type in **Prism Central Settings**, do the following.
 - a. Click **Prism Central Settings** from the **Navigation Bar** of the **Infrastructure** application.
For more information on the **Navigation Bar** of Prism Central applications, see [Application-specific Navigation Bar](#) in the Prism Central Infrastructure Guide).
The Prism Central Settings page opens.
 - b. Click **Network Controller**.
 - c. On the **Network Controller (formerly Advanced Networking)** page, under
 - » Clear the **Set as default** check box to set the default VLAN type as VLAN Basic Networking. The **Set as default** is cleared by default ensuring that the VLAN Basic Networking is the default VLAN type when you deploy or upgrade Prism Central.
 - » Select the **Set as default** check box to set the default VLAN type as VLAN Subnets (or Network Controller VLANs).
4. To change the default VLAN type when you navigate to **Network & Security > Subnets** for the first time after the Prism Central is deployed or upgraded.
The **Network Controller for VLAN Management** page opens.
 - a. On the **Network Controller for VLAN Management**, click **Change Default Settings**.
The **Network Controller (formerly Advanced Networking)** page in **Prism Central Settings** opens.
 - b. On the **Network Controller (formerly Advanced Networking)** page, under
 - » Clear the **Set as default** check box to set the default VLAN type as VLAN Basic Networking. The **Set as default** is cleared by default ensuring that the VLAN Basic Networking is the default VLAN type when you deploy or upgrade Prism Central.
 - » Select the **Set as default** check box to set the default VLAN type as VLAN Subnets (or Network Controller VLANs).

5. To change the default VLAN type to VLAN Basic Subnet while [Creating a Subnet](#) on page 93, do the following.
 - a. Navigate to **Network & Security > Subnets**
 - b. Click **Create Subnet** to create a VLAN network.
Provide the necessary configuration details for the VLAN Basic Subnet.
 - c. Under **Advanced Configuration**, select the **Advanced Configuration** check box under **VLAN Basic Networking**.For more information on creating a subnet, see [Creating a Subnet](#) on page 93.

Troubleshooting Tips

This section provides information to assist troubleshooting of Flow Virtual Networking deployments. This is in addition to the information that the *Prism Central Infrastructure Guide* provides.

Audit Logs

Prism Central generates audit logs for all the flow networking activities like it does for other activities on Prism Central. For more information, see [Audit Summary View](#) in the *Prism Central Infrastructure Guide*.

Support Bundle Collection

To support troubleshooting for Flow Virtual Networking, you can collect logs.

To collect the logs, run the following commands on the Prism Central VM console:

```
nutanix@cvm$ logbay collect -t msp,anc
```

An example of the command is as follows:

```
nutanix@cvm$ logbay collect -t msp,anc -0  
msp_pod=true,msp_systemd=true,kubectl_cmds=true,persistent=true --  
duration=-48h0m0s
```

Where:

- **-t** flag indicates the tags to collect
 - **msp** tag will collect logs from the services running on MSP pods and persistent log volumes (application-level logs)
 - **anc** tag will collect the support bundle, which includes database dumps and OVN state
- **-0** flag adds tag-level options
 - **msp_pod=true** collects logs from MSP service pods
On the PC, these logs can be found under /var/log/containers.
 - **persistent=true** collects persistent log volumes (application-level logs for ANC)
On the PC, these can be found under /var/log/ctrlog
 - **kubectl_cmds=true** runs **kubectl** commands to get the Kubernetes resource state
- **--duration** sets the duration from the present to collect

The command run generates a zip file at a location, for example: /home/nutanix/data/logbay/bundles/<filename>.zip

Unzip the bundle and you'll find the **anc** logs under a directory specific to your MSP cluster, the worker VM where the pod is running, and the logging persistent volume of that pod. For example:

./msp/f9684be8-b4e8-4524-74b4-076ed53ca1fd/10.48.128.185__worker_master_etcd/persistent/default/ovn/anc-ovn_StatefulSet/

For more information on the task run, see the text file that the command generates at a location, for example: /home/nutanix/data/logbay/taskdata/<taskID>/collection_result.txt

For more information on the logbay collect command, see the *Logbay Log Collection (Command Line)* topic in the *Nutanix Cluster Check Guide* (NCC Guide).

Layer 2 Virtual Subnet Extension Alert

The L2StretchLocalIfConflict alert (Alert with Check ID - 801109) may occur while performing Layer 2 virtual subnet extensions. For more information, see KB-10395 for more information about its resolution.

NETWORK GATEWAY UPGRADES

Nutanix deployment can detect and install upgrades for the on-premises Network Gateways. Network Gateways may be deployed for Virtual Private Networks (VPNs) connections, Virtual Tunnel End Point (VTEP) connections and Border Gateway Protocol (BGP) sessions.

For information on identifying the current Nutanix Gateway version, see [Identifying the Gateway Version](#) on page 45.

For on-premises Network Gateways, the upgrades must be detected and installed on the respective Prism Central on which each Network Gateway is installed. For more information, see [Detecting Upgrades for Gateways](#) on page 45.

For more information on the upgrade procedure, see [Upgrading the Network Gateway](#) on page 46.

Note: Upgrading the VPN appliance causes disruption of traffic for the duration of the upgrade operation.

Identifying the Gateway Version

About this task

To identify the current Nutanix Gateway version, do the following:

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from Application Switcher Function, and navigate to **Network & Security > Connectivity** from the **Navigation Bar**.
The **Connectivity** page opens displaying the **Gateways** tab.

3. Click the Gateway name link text to open the Gateway details page.

In the Gateway table, the Gateway name is a clickable link text.

The **Gateway Version** is listed in the **Properties** widget.

Detecting Upgrades for Gateways

About this task

Prism Central can detect whether new Gateway upgrades are available, or not, for Nutanix Gateways using LCM. You can then install the upgrade.

Procedure

- Log in to Prism Central.
- Select the **Admin Center** application from the [Application Switcher Function](#), and click **LCM** from the **Navigation Bar**.
The **LCM** page opens displaying the **Best Practices** tab.
- Click the **Inventory** tab.

- Click **Perform Inventory**.

Note:

Nutanix recommends that you select **Enable LCM Auto Inventory** in the **LCM** page in Prism Central to continuously detect new Gateway upgrades as soon as they are available.

The upgrade notification banner is displayed on the **Gateways** page.

Upgrading the Network Gateway

About this task

Perform upgrades of the Network Gateway using the respective Prism Central on which the Gateway is deployed.

Note: Upgrading the VPN appliance causes disruption of traffic for the duration of the upgrade operation.

To upgrade the Network Gateway, perform the following steps.

Procedure

1. Log in to Prism Central as an admin user.
2. Select the **Admin Center** application from the [Application Switcher Function](#), and click **LCM** from the **Navigation Bar**.
The **LCM** page opens displaying the **Best Practices** tab.

3. Click the **Inventory** tab.

4. Click **Perform Inventory**.
The **Perform Inventory** window opens.

5. Click **Proceed**.

When you click **Proceed**, the system scans the registered Prism Central cluster for software versions that are running currently. Then it checks for any available upgrades and displays the information on the **LCM** page under **Software**.

Note: Skip this step if you have enabled auto-inventory in the **LCM** page in Prism Central.

6. Click the **Updates** tab.
The **Updates** page opens displaying the available software updates.
7. Select the checkbox associated with the Gateway version you want to upgrade and click **View Upgrade Plan**.
The **Review Upgrade Plan** window opens.
8. Click **Apply 1 Updates**.

LCM upgrades the gateway version. This process takes some time.

Installing or Upgrading the Network Gateway in a Dark Site

Dark sites are primarily on-premises installations which do not have access to the internet. Such sites are disconnected from the internet for a range of reasons including security.

Before you begin

The Network Gateway is deployed for three types of connections:

- Virtual Private Networks as VPN Gateways
- Virtual Tunnel End Points as VTEP Gateways
- Border gateway Protocol sessions as BGP Gateways

Ensure that you complete the following tasks before you upgrade the Network Gateway in a dark site.

- Upgrade the LCM framework.
- Ensure that you have installed and prepared the LCM Dark Site server. For more information, see [Setting up a Local Web Server](#) in the *Life Cycle Manager Dark Site Guide*.
- Ensure that you have configured the **Dark Site (Local Web Server)** settings on the **LCM > Settings** page.
- See the *Release Notes* for the Network Gateway version compatible with the Network Controller and Prism Central version. To access the complete set of documentation, including the *Release Notes*, log on to [Flow Virtual Networking](#).
- On the [Flow Virtual Networking Downloads](#) page, select **Network Gateway** in the component selection dropdown menu. Download the compatible Network Gateway bundle that you ascertained in the preceding task. Copy the **SHA256** value for the bundle.
- Place the extracted vyos_<version>.qcow2 image and vyos_<version>.metadata.json files in the LCM Dark Site server. (See [Setting up a Local Web Server](#) in the *Life Cycle Manager Dark Site Guide*.)
- **Perform Inventory** on the **LCM** page in Prism Central **Admin Center** application.

Tip:

To go to the **LCM** page, select the **Admin Center** application from the [Application Switcher Function](#), and click **LCM** from the **Navigation Bar**.

In the **Updates** tab of the **LCM** page, **Network Gateway** now appears as an available update.

See the [Life Cycle Manager Dark Site Guide](#) for more information about **Perform Inventory** and the **Updates** tab.

About this task

To install or upgrade the Network Gateway at such dark sites, you need to deploy the Network Gateway bundle at the site.

Procedure

- See [KB-12393](#) and contact Nutanix Support to complete the Network Gateway version upgrade in the dark site.

NETWORK AND SECURITY ENTITIES

You can access the following networking and security entity items from the **Network and Security** entity of the **Infrastructure** application. For information on how to access the entity items available in **Network and Security** entity, see [Application-specific Navigation Bar](#) in the *Prism Central Infrastructure Guide*.

- **Subnets:** This page displays the subnets and the operations you can perform on subnets. For more information, see [Subnets](#) on page 48.
- **Virtual Private Clouds:** This page displays the VPCs and the operations you can perform on VPCs. For more information, see [Virtual Private Clouds Summary View](#) on page 53.
- **Floating IPs:** This page displays a list of floating IP addresses that you are using in the network. It allows you to request for floating IP addresses from the free pool of IP addresses available to the clusters managed by the Prism Central instance. For more information, see [Floating IPs Summary View](#) on page 59.
- **Connectivity:** This page allows you to manage the following networking capabilities. For more information, see [Connectivity](#) on page 60.
 - **Gateways:** This page provides a list of network Gateways you have created and configured, and the operations you can perform on the network Gateways. For more information, see [Gateways Summary View](#) on page 60.
 - **VPN Connections:** This page provides a list of VPN connections you have created and configured, and the operations you can perform on the VPN connections. For more information, see [VPN Connections Summary View](#) on page 63.
 - **Subnet Extensions:** This page provides a list of subnets that you have extended at the Layer 2 level using VPN (point-to-point over Nutanix VPN) or VTEP (point-to-multi-point including third party). For more information, see [Subnet Extensions Summary View](#) on page 68.
 - **BGP Sessions:** This page provides a list of BGP sessions you have created and configured, and the operations you can perform on the BGP sessions. For more information, see [BGP Sessions Summary View](#) on page 72.
- **Security Policies:** This page provides a list of security policies you configured using Flow Segmentation. For more information, see [Security Policies](#) on page 75.
- **Security Dashboard:** This page provides dynamic summary of the security posture across all registered clusters. For more information, see [Security Dashboard](#) on page 75.

For information on how to configure network connections, see [Network Configuration](#) in the *Prism Central Infrastructure Guide*.

Subnets (Overlay IP subnets), Virtual private clouds, floating IPs, and Connectivity are Flow virtual networking features. These features support flexible app-driven networking that focuses on VMs and applications instead of virtual LANs and network addresses. Flow virtual networking powers network virtualization to offer a seamless network experience with enhanced security. It is disabled by default. It is a software-defined network virtualization solution providing overlay capabilities for the on-premises AHV clusters.

Security policies drives the Flow Segmentation features for secure communications. For more information, see [Flow Microsegmentation Guide](#).

Subnets

Manage a subnet from Prism Central.

- [Creating a Subnet](#)

- [Updating a Subnet](#)
- [Deleting a subnet](#)
- [Creating a subnet extension](#)
- [Assigning a Category Value to a Subnet](#)
- [Migrating VMs between VLAN and VPC networks](#)

Subnets Summary View

The **Subnets** page displays the list of subnets across all the registered clusters.

To access the **Subnets** page:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Network & Security** > **Subnets** from the **Navigation Bar**.

The **Subnets** page opens displaying the **List** tab. This tab provides information about all the subnets configured for the registered clusters.

The following table describes the fields that appear in the **Subnets** page.

Note: The fields vary based on the **View by** and **Group by** options. A dash (-) is displayed in a field when a value is not available or applicable.

Table 2: Subnets – Field Description

Field	Description	Values
Name	Displays the subnet name.	(subnet name)
External Connectivity	Displays whether or not the subnet has external connectivity configured.	(Yes/No)
Type	Displays the subnet type.	VLAN or VLAN Basic or Overlay
VLAN ID	Displays the VLAN identification number.	(ID number)
VPC	Displays the name of the VPC in which the subnet is used.	(Name of VPC)
Virtual Switch	Displays the virtual switch that is configured for the VLAN you selected. The default value is the default virtual switch <code>vs0</code> .	(virtual switch name)
Note: The virtual switch name is displayed only if you add a VLAN ID in the VLAN ID field.		
IP Prefix	Displays the IPv4 address of the network with the prefix.	(IPv4 Address/Prefix)
Cluster	Displays the name of the cluster for which this subnet is configured.	(cluster name)

Field	Description	Values
Hypervisor	Displays the hypervisor that the subnet is hosted on.	(Hypervisor)

You can perform the following actions from the **Subnets** page:

- Click the name of a subnet to open the subnet details page, which displays the detailed information about the subnet. For more information, see [Subnet Details View](#) on page 51.
- Create a subnet by clicking **Create Subnet**. For more information, see [Creating a Subnet](#) on page 93 .
- Migrate VMs between VLAN network and VPC network by clicking **Migrate**. For more information, see [Migration of VMs between VLAN Basic Subnet and VPC Subnets](#) on page 80.
- Configure network connections for a cluster by clicking **Network Config**. For more information, see [Network Configuration](#) in the *Prism Central Infrastructure Guide*.
- Filter the subnets list based on a variety of parameter values using the **Filters** pane. For more information, see [Filters Pane - Subnets page](#).
- Perform the following subnet-specific actions on a single or multiple subnets using the **Actions** dropdown menu. The **Actions** dropdown appears when one or more subnets are selected.

Table 3: Subnet Actions

Action	Description
Update	Click this action to update the subnet. For more information, see Updating a Subnet on page 105.
Extend	Click this action to create a subnet extension. For more information, see Layer 2 Network Extension Over VPN on page 127.
Manage Categories	Click this action to associate the subnet with a category or change the categories that the subnet is associated with. For more information, see Assigning a Category in the <i>Prism Central Infrastructure Guide</i> .
Delete	Click this action to delete the subnet. For more information, see Deleting Subnets, Policies or Routes on page 107.

Filters Pane - Subnets page

You can filter the information in the **Subnets** page based on the following fields that are available in the **Filters** pane.

Table 4: Filter Pane Field Description - Subnets page

Field	Description	Values
Name	Filters based on the subnet name. It returns a list of subnets that satisfy the name condition/string.	(Subnet name string)
External Connectivity	Filters based on whether the subnet has external connectivity configured or not.	(Yes/No)

Field	Description	Values
Type	Filters based on the subnet type.	(VLAN/VLAN (External)/Overlay)
VLAN ID	Filters based on VLAN identification number.	(ID number)
VPC	Filters based on the name of the VPC in which the subnet is used.	(Name of VPC)
Cluster	Filters based on the name of the cluster for which this subnet is configured.	(cluster name)
Hypervisor	Filters based on the hypervisor that the subnet is hosted on.	ESXi/AHV/Hyper-V/XenServer/Mixed Hypervisor/Null Hypervisor

Subnet Details View

The Subnet details page consists of a dashboard that provides the detailed information about the subnet.

The details page has the **Summary**, and **Throughput** tabs.

To access the details page of an individual subnet:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Network & Security > Subnets** from the **Navigation Bar**.

Prism Central displays the **Subnets** page that contains information about all the subnets configured for the registered clusters.

3. Click a subnet to open the details page of the subnet.

The **Summary** tab opens displaying the detailed information about the subnet in widgets.

Summary Tab

The **Summary** tab provides detailed information about the subnet in widgets. A dash (-) is displayed in a field when a value is not available or applicable.

The **Summary** tab has the following widgets:

Widget Name	Information provided
Properties	Provides the following: • Type — Displays the type of network like VLAN or Overlay. • VLAN ID — Displays the VLAN ID. This parameter is displayed only for VLAN networks. • VPC — Displays the VPC name. This parameter is displayed only for Overlay networks. • Cluster — Displays the cluster that the VLAN network is configured on. This parameter is displayed only for VLAN networks. • IP Address Prefix — Displays the IP address prefix configured for the network. This parameter is displayed for both VLAN and Overlay networks.

Widget Name	Information provided
IP Address Pools	Provides the following: • The IP address Pool Range assigned to the network. • The total number of used and available IPs in the cluster. • Used IPs in Subnet — Displays the number of used IPs in the subnet. • Used IPs in Pools — Displays the number of used IPs in the pool. • Free IPs in Pools — Displays the number of free IPs in the pool. • Free IPs in Subnet — Displays the number of free IPs in the subnet.
Domain Settings	Provides the following DHCP settings configured for a VM in a subnet: • Domain Name Servers — Displays the total number of DNS IP addresses. • Domain Search — Displays the VLAN domain name. • Domain Name — Displays the domain name. • TFTP Server Name — Displays the name of the TFTP server where you host the host boot file. • Boot File Name — Displays the name of the boot file that the VMs need to download from the TFTP host server.

The **Summary** tab provides the following options, at the top of the page. For more information, see the *Subnet Actions* table in [Subnets Summary View](#) on page 49.

- **Update**
- **Extend**
- **Manage Categories**
- **Delete**

Throughput Tab

The **Throughput** tab provides a graphical representation of the throughput of the subnet.

Virtual Private Clouds

You can create, view, configure and manage the virtual private clouds (VPCs) in the **Virtual Private Clouds** page.

You can deploy and manage virtual private clouds (VPCs) on a Nutanix cluster infrastructure using the Network Controller.

You can view the list of the VPCs in the **Summary** view. For more information, see [Viewing the Virtual Private Cloud Summary](#) on page 52.

You can view the details of the VPCs in the VPC details view. For more information, see [Viewing the Virtual Private Cloud Details](#) on page 54.

Viewing the Virtual Private Cloud Summary

Access the **Virtual Private Clouds** page in **Network & Security**.

About this task

To access the **Virtual Private Clouds** page, follow these steps.

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Network & Security > Virtual Private Clouds** from the **Navigation Bar**.

The **Virtual Private Clouds** page opens displaying the **List** tab. The tab provides a list of virtual private clouds you have created and configured, and the operations you can perform on them.

Virtual Private Clouds Summary View

The **Virtual Private Clouds** page displays the list of virtual private clouds (VPCs) across all the registered clusters.

The following table describes the fields that appear in the **Virtual Private Clouds** page.

Note: The fields vary based on the **View by** and **Group by** options. A dash (-) is displayed in a field when a value is not available or applicable.

Table 5: Virtual Private Clouds – Field Description

Field	Description
Name	Displays the name of the VPC. The Name of a VPC is suffixed with Transit VPC when you configure the VPC as a transit VPC.
Associated External Subnets	Displays the external subnet that the VPC is assigned to.
Categories	Displays the number of categories associated with the VPC.
Externally Routable IP Addresses	Displays the externally routable IP address.
Hypervisor	Displays the hypervisor that the VPC is hosted on.
Inter VN Traffic	Displays the traffic flowing between the virtual networks or VPCs.
Internet Traffic	Displays the traffic flowing to and from the Internet.
IPv4 Gateway	Displays the IPv4 gateway IP address.
IPv4/Subnet	Displays the IPv4 network IP with subnet prefix. For example, 10.20.30.0/24.
On-Prem Traffic	Displays the traffic flowing in the on-premises network.
VLAN ID	Displays the VLAN identification number. VLAN ID is a parameter used for Transit VPC networking in Nutanix Cloud Cluster with Microsoft Azure.

You can perform the following actions for the VPCs from the **Virtual Private Clouds** page:

- Click the name of a VPC to open the VPC details page, which displays the detailed information about the VPC. For more information, see [Virtual Private Cloud Details View](#) on page 54.

- Create a VPC by clicking **Create VPC**. For more information, see [Creating a Virtual Private Cloud](#) on page 89.
- Update or delete an existing VPC using the **Actions** dropdown menu. The **Actions** dropdown appears when one or more VPCs are selected. For more information, see [Updating a Virtual Private Cloud](#) on page 103 or [Deleting a Virtual Private Cloud](#) on page 106.
- Filter the VPC list based on a variety of parameter values using **Filters** pane. For more information, see [Filters Pane - Virtual Private Clouds Page](#).

Filters Pane - Virtual Private Clouds Page

You can filter the information in the **Virtual Private Clouds** page based on the following fields that are available in the **Filters** pane.

Table 6: Filter Pane Field Description - Virtual Private Clouds page

Field	Description	Values
Name	Filters based on the VPC name. It returns a list of IP addresses that satisfy the name condition/string.	(Virtual private cloud name string)
Associated External Subnets	Filters based on the external subnet that the VPC is assigned to.	(External Subnet)

Viewing the Virtual Private Cloud Details

You can view the details of each VPC listed in **Virtual Private Clouds**.

About this task

To access the details page of an individual VPC, follow these steps.

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Network & Security > Virtual Private Clouds** from the **Navigation Bar**.

Prism Central displays the **Virtual Private Clouds** page that contains information about all the VPCs configured for the registered clusters.

3. Click a VPC to open the details page of the VPC.

The **Summary** tab opens displaying the detailed information about the VPC in widgets.

Virtual Private Cloud Details View

The Virtual Private Cloud (VPC) details page consists of a dashboard that provides the detailed information about the VPC.

The details page has the **Summary**, **Subnets**, **Policies**, **Routes**, and **Metrics** tabs.

Summary Tab

The **Summary** tab provides detailed information about the VPC in widgets.

The **Summary** tab has the following widgets:

Widget Name	Information provided
External Connectivity	Provides the following: - Associated External Subnets — Displays the number of external subnets associated with the VPC. - Externally Routable IP Addresses — Displays the external routable IP addresses associated with the VPC.
Transit VPC	Displays Yes if the VPC is a Transit VPC. Displays No if the VPC is not a Transit VPC.
Domain Name Servers (DNS)	Displays the IP address or the FQDN of the DNS servers used by the VPC.
Associations	Provides the following: - Subnets (Overlay) — Displays the number of subnets associated with the VPC. - Policies — Displays the number of policies associated with the VPC. - Routes — Displays the number of routes associated with the VPC.
Floating IP Addresses	Provides the following: - Assigned Floating IPs — Displays the floating IP addresses assigned to the VPC. - Available Floating IPs — Displays the available floating IP addresses that can be assigned to the VPC.

Subnets Tab

The **Subnets** tab displays the list of subnets added to the VPC.

The following table describes the fields that appear in the **Subnets** tab.

Table 7: Subnets Tab – Field Description

Field	Description
Name	Displays the subnet name.
IP Range	Displays the IP address range configured for the subnet.
DHCP IP Pool	Displays the IP address pool range assigned to the subnet.
Default Gateway IP	Displays the IP address used as the default gateway by the entities in the subnet.
Actions	Action link for editing or deleting the subnet.

You can perform the following actions for a subnet from the **Subnets** tab:

- Click the name of the subnet to open the subnet details page, which displays the detailed information about the subnet. For more information, see [Subnet Details View](#) on page 51.
- Create a subnet by clicking **Create Subnet**. For more information, see [Creating a Subnet](#) on page 93.
- Update an existing subnet using the **Delete** option associated with the subnet. For more information, see [Updating a Subnet](#) on page 105.

- Delete an existing subnet using the **Delete** option associated with the subnet. For more information, see [Deleting Subnets, Policies or Routes](#) on page 107.

Policies Tab

The **Policies** tab displays information about the security-based traffic shaping policies you configured.

The following table describes the fields that appear in the **Policies** tab.

Note: The fields vary based on the **View by** option. A dash (-) is displayed in a field when a value is not available or applicable.

Table 8: Policies Tab – Field Description

Field	Description
Description	Displays the user-provided description of the policy.
Action	Displays the appropriate action for the implementation of the policy. • Permit: Permits traffic and services based on the parameters set. • Deny: Denies traffic and service based on the parameters set. • Re-route: Sends matching traffic to the next-hop IP address specified by the Reroute IP.
Priority	Displays the traffic priority.
Rule	Displays the Permit or Deny rule set for the priority.
Rule Type	Displays whether the rule is system generated or user defined.
Traffic	Displays the traffic type that the priority and rule must be applied to.
Virtual Network	Displays the ID of the subnet.
Source	Displays the source IP or subnet for which you need to manage traffic.
Destination	Displays the destination IP or subnet for which you need to set the priority.
Source Subnet	Displays the subnet IP and prefix designated as the source for the policy.
Destination Subnet	Displays the subnet IP and prefix designated as the destination for the policy.
Reroute Address	Displays the IP address to which the traffic is rerouted.
Bidirectional Policy	Displays whether the policy is bidirectional or not.
Protocol	Displays the type of protocol for which the policy is configured.
Protocol Number	Displays the protocol number for which the policy is configured.
ICMP Type	Displays the type of ICMP message associated with the policy.
ICMP Code	Displays the ICMP code of the policy.

Field	Description
Byte Count	Displays the total number of traffic bytes that matches the policy. The count is updated periodically.
Packet Count	Displays the total number of traffic packets that matches the policy. The count is updated periodically.

You can perform the following actions for a policy from the **Policies** tab:

- Create a policy by clicking **Create Policy**. For more information, see [Creating a Policy](#) on page 97.
- Perform the following actions using the **Actions** dropdown menu. The **Actions** dropdown appears when one or more policies are selected.
 - **Update:** Update the policy. For more information, see [Updating a Subnet](#) on page 105.
 - **Delete:** Delete the policy. For more information, see [Deleting Subnets, Policies or Routes](#) on page 107.
 - **Clear Counters:** Reset the counters for the selected policy.
 - **Clear All Counters:** Reset the counters for all the policies.

Routes Tab

The **Routes** tab displays the list of static routes added to the VPC.

The following table describes the fields that appear in the **Routes** tab.

Note: The fields vary based on the **View by** option. A dash (-) is displayed in a field when a value is not available or applicable.

Table 9: Routes Tab – Field Description

Field	Description
Destination Prefix	Displays the IP address and prefix of the destination.
Next Hop	Displays the next hop network or subnet for the traffic exiting the VPC.
Priority	Displays the traffic priority.
Type	Displays the type of route, local or static.
Status	Displays the status of the route, whether it is active or not

You can perform the following actions for a route from the **Routes** tab:

- View routes based on predefined criteria or create a custom view.
- Perform the following actions using the **Manage Static Routes** option:
 - **Add Static Route:** Create a static route. For more information, see [Creating Static Routes](#) on page 102.
 - Update an existing static route. For more information, see [Updating Static Routes](#) on page 106.
 - Delete a static route. For more information, see [Deleting Subnets, Policies or Routes](#) on page 107.

Metrics Tab

The **Metrics** tab displays detailed information about the VPC metrics.




Figure 11: Metrics Tab

The following table describes the fields that appear in the **Metrics** tab.

Table 10: Metrics Tab – Field Description

Field	Description
External Traffic Through	Select All External Networks (default) or (name_of_external_network_associated_with_the_VPC) from the dropdown menu. The page displays the metrics based on your selection.
Last (time_period)	Select the period for which you need to display the metrics. The dropdown menu provides the following options: • Last 24 Hours (default) • Last One Hour • Last Week
Direction of traffic	Select the direction of traffic for which you need to display the metrics. The dropdown menu provides the following options: • Both directions (default) — Includes both directions, Ingress and Egress. • Ingress — Traffic entering the externally connected subnet. • Egress — Traffic leaving the externally connected subnet.
Bandwidth	Displays graphically the bandwidth utilization of the VPC on a timeline as set in the Last (time_period) parameter.
Rx and Tx Packets	Displays graphically the received and transmitted packet volume on a timeline as set in the Last (time_period) parameter.

Floating IPs

You can access the floating IP addresses you have created and configured, from the **Floating IPs** page.

For information on floating IP addresses and their role in flow virtual networking, see the *SNAT and Floating IP Address* section in [Essential Concepts](#) on page 12.

Note: Floating IP addresses are not reachable (Pings fail) unless you associate them to primary or secondary IP addresses of VMs. For more information, see [Assigning Secondary IP Addresses to Floating IPs](#) on page 79.

Floating IPs Summary View

The **Floating IPs** page displays the list of floating IP addresses across all the registered clusters.

To access the **Floating IPs** page:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Network & Security > Floating IPs** from the **Navigation Bar**.

The **Floating IPs** page opens displaying the **List** tab. This tab provides a list of floating IPs you have created and configured, and the operations you can perform on the IPs.

The following table describes the fields that appear in the **Floating IPs** page.

Note: The fields vary based on the **View by** option. A dash (-) is displayed in a field when a value is not available or applicable.

Table 11: Floating IPs – Field Description

Parameter	Description	Values
Floating IP Address	Displays the floating IP address assigned.	(IP address)
External Subnet	Displays the name of the external subnet that the IP address is assigned to.	(Name of the assigned subnet)
Association Status	Displays the status of association between the IP address and the external subnet and VPC.	Associated
VPC	Displays the name of the VPC associated with the IP address.	(Name of the associated VPC)
VM Name	Displays the name of the VM associated with the IP address.	(Name of the assigned VM)
Private IP	Displays the private IP address assigned to the same VM. This private IP address is assigned from the internal private subnet that the network controller creates when you create a network gateway.	(IP address)

You can perform the following actions for the floating IP addresses from the **Floating IPs** page:

- Request a floating IP address by clicking **Request Floating IP**. For more information, see [Requesting Floating IPs](#) on page 92.

- Update or delete an existing floating IP address using the **Actions** dropdown menu. The **Actions** dropdown appears when one or more addresses are selected.
 - **Update:** Assign or change the assignment of the floating IP address. You can assign the floating IP address to a IP address such as a private IP address in a VPC or the primary IP address of a VM or a secondary IP address created on a VM.
 - **Delete:** Delete the floating IP address. The deleted IP address returns to the IP address pool as unused. Before you delete a floating IP address, ensure that it is not assigned to a private IP address or a VM. Change the assignment to None if it is already assigned, using the **Update** option.
- Filter the floating IP addresses list based on a variety of parameter values using **Filters** pane. For more information, see [Filters Pane - Floating IPs Page](#).

Filters Pane - Floating IPs Page

You can filter the information in the **Floating IPs** page based on the following fields that are available in the **Filters** pane.

Table 12: Filter Pane Field Description - Floating IPs page

Field	Description	Values
Floating IP Address	Filters based on the floating IP address assigned. It returns a list of IP addresses that satisfy the string.	(Floating IP address)
External Subnet	Filters based on the external subnet that the IP address is assigned to.	(External Subnet)

Connectivity

You can access network gateways, VPN connections, subnet extensions, and BGP sessions from the **Connectivity** page.

To access the **Connectivity** page:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Network & Security > Connectivity** from the **Navigation Bar**.

The **Connectivity** page opens displaying the **Gateways** tab. This tab provides a list of network Gateways you have created and configured, and the operations you can perform on the network Gateways.

- To view the VPN connections, click the **VPN Connections** tab.
- To view the subnets extended across the clusters, click the **Subnet Extensions** tab.
- To view the BGP sessions created for the clusters, click the **BGP Sessions** tab.

Gateways Summary View

The **Gateways** page displays a list of gateways created for the clusters managed by Prism Central.

To access the **Gateways** page:

1. Log in to Prism Central.

2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Network & Security > Connectivity** from the **Navigation Bar**.

The **Gateways** page opens displaying the list of network gateways you have created and configured, and the operations you can perform on the network gateways.

The following table describes the fields that appear in the **Gateways** page.

Table 13: Field Descriptions for the Gateway Page

Parameter	Description	Values
Name	Displays the name of the gateway.	(Name of gateway)
Type	Displays the gateway type.	(Local or Remote)
Service	Displays the service that the gateway uses.	(VPN or VTEP)
Service IP	Displays the IP address used by the service.	(IP address)
Status	Displays the operational status of the gateway.	(Up or Down)
Attachment Type/Vendor	Displays the type of subnet associated with the gateway.	(VLAN or Overlay-VPC name)
Connections	Displays the number of service connections (such as VPN connections) configured and operational on the gateway.	(Number)

You can perform the following actions for a gateway from the **Gateways** page:

- Click the name of a gateway to open the gateway details page, which displays the detailed information about the gateway. For more information, see [Gateway Details View](#) on page 62.
- Create a local or remote gateway with VPN or VTEP service by clicking the **Create Gateway** dropdown menu. For more information, see [Creating a Network Gateway](#) on page 108.
- Update or delete an existing gateway using the **Actions** dropdown menu. The **Actions** dropdown menu appears when one or more gateways are selected. For more information, see [Updating a Network Gateway](#) on page 115 or [Deleting a Network Gateway](#) on page 116.
- Filter the gateway list based on various parameter values using the **Filters** pane. For more information, see [Filters Pane - Gateways Page](#).

Filters Pane on the Gateways Page

You can filter the information in the **Gateways** page based on the following fields that are available in the **Filters** pane.

Table 14: Filter Pane Field Descriptions for the Gateways page

Field	Description	Values
Name	Filters based on the gateway name. It returns a list of gateways that satisfy the name condition/string.	(Gateway name string)
Service IP	Filters based on IP address used by the service.	(IP address)

Field	Description	Values
Status	Filters based on the operational status of the gateway.	(Up or Down)

Gateway Details View

The **Summary** page of an individual gateway consists of a dashboard that provides the detailed information about the gateway.

To access the **Summary** page of an individual gateway:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Network & Security** > **Connectivity** from the **Navigation Bar**.

The **Gateways** page opens displaying the list of network gateways that you have created and configured.

3. Click a gateway to view the **Summary** page of the gateway.

The gateway **Summary** page has the following widgets:

Table 15: Field Descriptions for the Gateway Widgets

Parameter	Description	Values
Properties widget		
Type	Displays the gateway type.	(Local or Remote)
Attachment Type	Displays the network entity like VLAN or VPC that the gateway is attached to.	(VLAN or VPC)
VPC or Subnet (VLAN)	Displays the name of the attached VPC or VLAN subnet.	(Name of VLAN or VPC)
Vendor (Applicable only if you select remote gateway)	Displays the name of the vendor of the gateway appliance at the remote site.	(Name of Vendor)
Floating or Private IP Address	Displays the Floating (for VPC) or Private (for VLAN) IP address assigned to the gateway.	(IP Address)
External IP (Applicable only if you select remote gateway)	Displays the IP address assigned to the remote gateway.	(IP Address that you assigned to the remote gateway.)
Status	Displays the operational status of the gateway.	(Up or Down)
Gateway Version	Displays the version of the Nutanix gateway appliance deployed.	(Version)
Cluster	Displays the name of the cluster on which the gateway is created.	(Cluster name)
Gateway VM	Displays the name of the VM on which the gateway is created.	(Name of VM - actionable link. Click the name-link to open the VM details page of the gateway VM.)
Service Configuration widget		
Service	Displays the service used by the gateway.	(VPN or VTEP or BGP)

Parameter	Description	Values
VPN Service Configuration		
External Routing	Displays the type of routing associated with the gateway for external traffic routing.	(Static or eBGP with ASN)
Internal Routing	Displays the type of routing associated with the gateway for internal traffic routing.	(Static or eBGP with ASN)
VPN Connections	Displays the total number of VPN connections associated with the gateway.	(Number - actionable link. Click the link to open the VPN connection details page for the associated VPN connection.)
View VPN Connections	Click this link to open the VPN Connections tab.	-
VTEP Service Configuration		
VXLAN (UDP) Port	Displays the VXLAN (UDP) Port for the gateway.	(Number)
Subnet Extensions	Displays the total number of subnet extensions associated with the gateway.	(Number - actionable link. Click the link to open the subnet extensions details page for the associated subnet extension.)
View Subnet Extensions	Click this link to open the Subnet Extensions tab.	-
BGP Service Configuration		
ASN	Displays the ASN of the EBGp route.	(Number)
BGP Sessions	Displays the total number of BGP sessions associated with the gateway.	(Number - actionable link. Click the link to open the BGP sessions details page for the associated BGP session.)
Serviced VPC	Displays VPC service used by the gateway.	(Name of VPC)
View BGP Sessions	Click this link to open the BGP Sessions tab.	-

You can perform the following actions for a gateway from the **Summary** tab:

- Update an existing gateway by clicking **Update**. For more information, see [Updating a Network Gateway](#) on page 115.
- Delete the gateway by clicking **Delete**. For more information, see [Deleting a Network Gateway](#) on page 116.

VPN Connections Summary View

The **VPN Connections** page displays a list of VPN connections created for the clusters managed by Prism Central.

A VPN connection represents the VPN IPsec tunnel established between local gateway and remote gateway. When you create a VPN connection, you must select two gateways between which you want to create the VPN connection.

To access the **VPN Connections** page:

1. Log in to Prism Central.

2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Network & Security > Connectivity** from the **Navigation Bar**.

The **Gateways** page opens displaying the list of network gateways.

3. Click the **VPN Connections** tab.

The **VPN Connections** page opens displaying the list of VPN connections created for the clusters.

The following table describes the fields that appear in the **VPN Connections** page.

Table 16: Field Descriptions for the VPN Connections Page

Parameter	Description	Values
Name	Displays the name of the connection.	(gateway name)
IPSec Status	Displays the connection status of IPSec tunnel.	(Connected or Not Connected)
EBGP Status	Displays the status of the EBGP gateway connection.	(Established or Not Established)
Local Gateway	Displays the name of the local gateway used for the connection.	(Name of local gateway)
Remote Gateway	Displays the name of the remote gateway used for the connection.	(Name of remote gateway)
Dynamic Routing Priority	Displays the dynamic routing priority assigned to the connection for throughput management. You can assign any value in the range of 100-1000. Nutanix Flow Virtual Networking assigns the first VPN connection the value 500 by default. Thereafter, subsequent VPN connections are assigned values decremented by 50. For example, the first connections is assigned 500, then the second connection is assigned 450, the third one 400 and so on.	(Number in the range of 100-1000. User assigned.)

You can perform the following actions for a VPN connection from the **VPN Connections** page:

- Click the name of a VPN connection to open the VPN connection details page, which displays the detailed information about the connection. For more information, see [VPN Connection Details View](#) on page 65.
- Create a VPN connection by clicking **Create VPN Connection**. For more information, see [Creating a VPN Connection](#) on page 120.
- Update or delete an existing VPN connection using the **Actions** dropdown menu. The **Actions** dropdown appears when one or more VPN connections are selected. For more information, see [Updating VPN Connection](#) on page 122 or [Deleting a VPN Connection](#) on page 122.
- Filter the VPN connection list based on various parameter values using the **Filters** pane. For more information, see [Filters Pane - VPN Connections Page](#).

Filters Pane on the VPN Connections Page

You can filter the information in the **VPN Connections** page based on the following fields that are available in the **Filters** pane.

Table 17: Filter Pane Field Descriptions for the VPN Connections page

Field	Description	Values
Name	Filters based on the VPN connection name. It returns a list of VPN connections that satisfy the name condition/string.	(VPN connection name string)
EBGP Status	Filters based on the status of the EBGP gateway connection.	(Established or Not Established)
IPSEC Status	Filters based on the connection status of IPsec tunnel.	(Connected or Disconnected)

VPN Connection Details View

The VPN Connection details page provides detailed information about a VPN connection.

The details page has the **Summary**, **Throughput**, **IPSec Logging**, and **Routing Protocol Logging** tabs.

To access the details page of an individual VPN connection:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Network & Security > Connectivity** from the **Navigation Bar**.

The **Gateways** page opens displaying the list of network gateways you have created and configured.

3. Click the **VPN Connections** tab.

The **VPN Connections** page opens displaying the list of VPN connections created for the clusters.

4. Click the name of a VPN connection to open the details page of the connection.

The **Summary** tab opens displaying the detailed information about the VPN connection in widgets.




Figure 12: VPN Connection Details

Summary Tab

The **Summary** tab provides detailed information about a VPN connection in widgets.

The following table describes the fields that appear in the **Summary** tab.

Table 18: Field Descriptions for the Summary Tab

Parameter	Description	Values
VPN Connection widget		
IPSec Status	Displays the connection status of IPSec tunnel.	(Connected or Not Connected)
EBGP Status	Displays the status of the EBGP gateway connection.	(Established or Not Established)
Dynamic Routing Priority	Displays the dynamic routing priority assigned to the connection for throughput management. You can assign any value in the range of 100-1000. Flow Virtual Networking assigns the first VPN connection the value 500 by default. Thereafter, subsequent VPN connections are assigned values decremented by 50. For example, the first connections is assigned 500, then the second connection is assigned 450, the third one 400 and so on.	(Number in the range of 100-1000. User assigned.)
Local Gateway Properties widget		
Gateway Name	Displays the name of the local gateway used for the connection.	(Name of local gateway)
Type	Displays the type of gateway.	(Local)
Attachment Type	Displays the network entity like VLAN or VPC that the gateway is attached to.	(VLAN or VPC)
VPC or Subnet (VLAN)	Displays the name of the attached VPC or VLAN subnet.	(Name of VLAN or VPC)
Tunnel IP	Displays the Tunnel IP address of the local gateway.	(IP Address)
Connection Type	Displays the connection type you selected while creating the VPN connection. The connection type may be Initiator or Acceptor of a VPN connection between the local and remote gateways. T	(Initiator or Acceptor)
External Routing	Displays the type of routing associated with the gateway for external traffic routing.	(Static or eBGP with ASN)
Internal Routing	Displays the type of routing associated with the gateway for internal traffic routing.	(Static or eBGP with ASN)
Floating or Private IP Address	Displays the Floating (for VPC) or Private (for VLAN) IP address assigned to the gateway.	(IP Address that you assigned to the local gateway with /30 prefix when you configured the VPN connection.)
Status	Displays the operational status of the gateway.	(Up or Down)
Cluster	Displays the name of the cluster on which the gateway is created.	(Cluster name)

Parameter	Description	Values
Gateway VM	Displays the name of the VM on which the gateway is created.	(Name of VM - actionable link. Click the name-link to open the VM details page of the gateway VM.)
Remote Gateway Properties widget		
Gateway Name	Displays the name of the remote gateway used for the connection.	(Name of remote gateway)
Type	Displays the type of gateway.	(Remote)
Tunnel IP	Displays the Tunnel IP address of the remote gateway.	(IP Address)
Connection Type	Displays the connection type you selected while creating the VPN connection. The connection type may be Initiator or Acceptor of a VPN connection between the local and remote gateways. T	(Initiator or Acceptor)
External Routing	Displays the type of routing associated with the gateway for external traffic routing.	(Static or eBGP with ASN)
ASN	Displays the ASN of the EBGP route. This information is only displayed if you configured EBGP as the External Routing protocol.	(Number)
Vendor	Displays the name of the vendor of the gateway appliance at the remote site.	(Name of vendor of gateway appliance)
External IP	Displays the IP address assigned to remote the gateway.	(IP Address that you assigned to the remote gateway with /30 prefix when you configured the VPN connection.)
Status	Displays the operational status of the gateway.	-
Protocol Details widget		
Service	Displays the service used by the gateway.	(VPN or VTEP)
Gateway Routes widget	Displays the status of the routes used by the gateways.	(Sent)

You can perform the following actions from the **Summary** tab:

- View the detailed information of a VPN connection. For the list of available parameters, see the *VPN Connection Summary Tab* table above.
- Update an existing VPN connection by clicking **Update**. For more information, see [Updating VPN Connection](#) on page 122.
- Delete an existing VPN connection by clicking **Delete**. For more information, see [Deleting a VPN Connection](#) on page 122.

Throughput Tab

The **Throughput** tab provides a graphical representation of the throughput of the VPN connection.

IPSec Logging

The **IPSec Logging** tab provides running logs for the IPSec tunnel of the VPN connection.

Routing Protocol Logging

The **Routing Protocol Logging** tab provides logs for the routing protocol used in the VPN connection.

Subnet Extensions Summary View

The **Subnet Extensions** page displays a list of subnet extensions created for the clusters managed by Prism Central.

To access the **Subnet Extensions** page:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Network & Security > Connectivity** from the **Navigation Bar**.

The **Gateways** page opens displaying the list of network gateways.

3. Click the **Subnet Extensions** tab.

The **Subnet Extensions** page opens displaying the list of subnet extensions created for the clusters.

The following table describes the fields that appear in the **Subnet Extensions** page.

Table 19: Field Description for the Subnet Extensions Page

Parameter	Description	Values
Name	Displays the name of the subnet extension.	(Name of subnet extension)
Type	Displays the subnet extension type.	(Across Availability Zones or To a Third Party Data Center)
Extension Over	Displays the service that the subnet extension uses.	(VPN or VTEP)
Extension Uses	Displays the name of the local network gateway that the subnet extension uses.	(Name of local network gateway)
Local Subnet	Displays the name of the local subnet that the subnet extension uses.	(Name of local subnet)
Remote Site	Displays the name of the remote network gateway that the subnet extension uses.	(Name of remote network gateway)
Connection Status	Displays the status of the connection that is created by the subnet extension.	(Not Available, Connected, or Disconnected)
	Note: Not Available status indicates that Prism Central is unable to ascertain the status.	
Interface Status	Displays the status of the interface that is used by the subnet extension.	(Connected or Down)

You can perform the following actions for a subnet extension from the **Subnet Extensions** page:

- Click the name of a subnet extension to open the subnet extension details page, which displays the detailed information about the extension. For more information, see [Subnet Extension Details View](#) on page 69.
- Extend a subnet **Across Availability Zones** or **To a Third Party Data Center** by clicking the **Create Subnet Extension** dropdown menu. You can extend a subnet using **VPN** or **VTEP** service. For more information, see [Layer 2 Network Extension](#) on page 124.
- Update or delete existing subnet extension using the **Actions** dropdown menu. The **Actions** dropdown appears when one or more subnet extensions are selected. For more information, see [Updating an Extended Subnet](#) on page 142 or [Removing an Extended Subnet](#) on page 142.
- Filter the subnet extension list based on various parameter values using the **Filters** pane. For more information, see [Filters Pane - Subnet Extensions Page](#).

Filters Pane on the Subnet Extensions Page

You can filter the information in the **Subnet Extensions** page based on the following fields that are available in the **Filters** pane.

Table 20: Filter Pane Field Descriptions for the Subnet Extensions page

Field	Description	Values
Name	Filters based on the subnet extension name. It returns a list of subnet extensions that satisfy the name condition/string.	(Subnet extension name string)
Connection Status	Filters based on the status of the connection that is created by the subnet extension.	(Connected or Disconnected)
Interface Status	Filters based on the status of the interface that is used by the subnet extension.	(Connected or Not Available)

Subnet Extension Details View

The Subnet Extension details page provides detailed information about a subnet extension.

The details page has the **Summary**, **Address Table**, and **Throughput** tabs.

To access the details page of an individual subnet extension:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Network & Security > Connectivity** from the **Navigation Bar**.

The **Gateways** page opens displaying the list of network gateways you have created and configured.

3. Click the **Subnet Extensions** tab.

The **Subnet Extensions** page opens displaying the list of subnet extensions created for the clusters.

4. Click a subnet extension to open the details page of the extension.

The **Summary** tab opens displaying the detailed information about the extension in widgets.

Summary Tab

The **Summary** tab provides detailed information about the subnet extension in widgets.

The subnet extension **Summary** tab has the following widgets:

Table 21: Subnet Extension Summary Tab Widgets

Parameter	Description	Values
Properties widget		
Type	Displays the subnet type.	(VLAN or Overlay)
VLAN ID	(For VLAN subnets only) Displays the VLAN ID of the VLAN subnet that is extended.	(VLAN ID number)
VPC	(For Overlay subnets only) Displays the name of the VPC subnet that is extended.	(Name of VPC)
Cluster	(For VLAN subnets only) Displays the cluster that the VLAN subnet belongs to.	(Name of cluster)
IP Address Prefix	Displays the network IP address with prefix, of the VLAN subnet that is extended.	(IP Address with prefix)
Virtual Switch	(For VLAN subnets only) Displays the virtual switch on which the VLAN subnet is configured.	(Virtual Switch name such as vs0 or vs1)
IP Address Pools widget		
Pool Range	Displays the range of IP addresses in the pool configured in the subnet that is extended.	(IP address range)
(Interactive Graphic Pie Chart)	Displays a dynamic pie chart that displays the statistic you hover on. Displays the following IP address statistics outside the pie chart, that you can hover on: • Total number of IP addresses available. • Used IP addresses in the subnets • Used IP addresses in the IP address pools • Free IP addresses in the subnets • Free IP addresses in the IP address pools	(IP Address statistics)
Subnet Extension widget		
Subnet Extension (properties) - Common		
Type	Displays the subnet extension type.	(Across Availability Zones or To a Third Party Data Center)
Interface Status	Displays the status of the interface that is used by the subnet extension.	(Connected or Down)
Connection Status	Displays the status of the connection that is created by the subnet extension. Not Available status indicates that Prism Central is unable to ascertain the status.	(Not Available, Connected, or Disconnected)
Local IP Address	Displays the IP address that you entered in the Local IP Address field while creating the subnet extension.	(IP Address)

Parameter	Description	Values
Local Subnet	Displays the name of the local subnet that the subnet extension uses.	(Name of local subnet)
Subnet Extension (properties) - (Only for Across Availability Zones type)		
Local Availability Zone	(Only for Across Availability Zones type) Displays the name of the local AZ that is hosting the subnet that is extended.	(Name of the local Availability Zone)
Remote Availability Zone	(Only for Across Availability Zones type) Displays the name of the remote AZ that the subnet is extended to.	(Name of the remote Availability Zone)
Remote Subnet	(Only for Across Availability Zones type) Displays the name of the remote subnet that the subnet extension connects to.	(Name of remote subnet)
Remote IP Address	(Only for Across Availability Zones type) Displays the IP address that you entered in the Remote IP Address field while creating the subnet extension.	(IP Address)
Subnet Extension (properties) - (Only for To a Third Party Data Center type)		
Local Gateway	(Only for To a Third Party Data Center type) Displays the name of the local gateway used for the subnet extension.	(Name of local gateway)
Remote Gateway	(Only for To a Third Party Data Center type) Displays the name of the remote gateway used for the subnet extension.	(Name of remote gateway)

You can perform the following actions from the **Summary** tab:

- View the detailed information of a subnet extension. For the list of available parameters, see the *Subnet Extension Details - Summary Tab Fields* table above.
- Update an existing subnet extension by clicking **Update**. For more information, see [Updating an Extended Subnet](#) on page 142.
- Delete an existing subnet extension by clicking **Delete**. For more information, see [Removing an Extended Subnet](#) on page 142.

Address Table Tab

The **Address Table** tab provides MAC Address information only when the subnet extension uses VTEP service. The tab provides the following information:

- **MAC Address:** This provides the MAC addresses of devices connected to the remote VTEP endpoint in the subnet extension.
- **Remote VTEP Endpoint:** This provides the IP address of the remote VTEP endpoint in the subnet extension.

Throughput Tab

The **Throughput** tab provides a graphical representation of the throughput of the subnet extension.

BGP Sessions Summary View

The **BGP Sessions** page displays a list of BGP sessions created for the clusters managed by Prism Central.

To access the **BGP Sessions** page:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Network & Security** > **Connectivity** from the **Navigation Bar**.

The **Gateways** page opens displaying the list of network gateways.

3. Click the **BGP Sessions** tab.

The **BGP Sessions** page opens displaying the list of BGP sessions created for the clusters.

The following table describes the fields that appear in the **BGP Sessions** page.

Table 22: BGP Sessions – Field Description

Parameter	Description	Values
Name	Displays the name of the BGP session.	(Name of BGP session)
Serviced VPC	Displays the name of the VPC that the BGP session services.	(Name of VPC)
Local Gateway	Displays the name of the local BGP gateway that the BGP session uses.	(Name of local BGP gateway)
Remote Gateway	Displays the name of the remote BGP gateway that the BGP session uses.	(Name of remote BGP gateway)
Session Status	Displays the status of the eBGP session. • Displays Established if the session is Up. • Displays Active when the network controller is attempting to establish the session.	Established or Active
Route Priority	Displays an integer number that denotes the route priority. When the route priority is assigned dynamically, then the network controller assigns integer numbers (usually between 600 and 800 starting with 700) in descending order with steps of 5. For example, the first session is assigned 700 as route priority and then when you create the second session, the controller assigns it a route priority of 695 and a third session is assigned 690. Greater the number, greater is the route priority. With dynamically assigned priority, the priority is assigned in the order of reducing priority to the order of BGP sessions created. The BGP session created first gets the highest priority 700, the second session get the second highest priority 695 and so on. You can manually assign a route priority as well by assigning any number between 300 and 900.	(Integer Number)

You can perform the following actions for a gateway from the **BGP Sessions** page:

- Click the name of a BGP session to open the details page, which displays the detailed information about the BGP session. For more information, see [BGP Session Details View](#) on page 73.
- Create a BGP session by clicking **Create BGP Session**. For more information, see [Creating a BGP session](#) on page 144.
- Update or delete an existing BGP session using the **Actions** dropdown menu. The **Actions** dropdown menu appears when one or more BGP sessions are selected. For more information, see [Updating a BGP session](#) on page 146 or [Deleting a BGP session](#) on page 147.
- Filter the gateway list based on various parameter values using the **Filters** pane. For more information, see [Filters Pane - BGP Sessions Page](#).

Filters Pane on the BGP Sessions Page

You can filter the information in the **BGP Sessions** page based on the following fields that are available in the **Filters** pane.

Table 23: Filter Pane Field Description - BGP Sessions page

Field	Description	Values
Name	Filters based on the BGP session name. It returns a list of BGP sessions that satisfy the name condition/ string.	(BGP session name string)
Session Status	Filters based on the status of the eBGP session.	(Established or Down)

BGP Session Details View

The BGP Session details page provides detailed information about a BGP session.

The details page has the **Summary**, **Routes**, and **BGP Logs** tabs.

To access the details page of an individual BGP session:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Network & Security** > **Connectivity** from the **Navigation Bar**.

The **Gateways** page opens displaying the list of network gateways you have created and configured.

3. Click the **BGP Sessions** tab.

The **BGP Sessions** page opens displaying the list of BGP sessions created for the clusters.

4. Click the name of a BGP session to open the details page of the session.

The **Summary** tab opens displaying the detailed information about the BGP session in widgets.

Summary Tab

The **Summary** tab provides detailed information about the BGP session in widgets.

The BGP session **Summary** tab has the following widgets:

Table 24: BGP Session Summary Tab Widgets

Parameter	Description	Values
Properties widget		
Session Status	Displays the overall status of the BGP session.	(Up or Down)
eBGP Status	Displays the eBGP status of the BGP session.	Established or Active
Route Priority	Displays an integer number that denotes the route priority. For more information about Route Priority, see BGP Sessions Summary View on page 72.	(Integer Number)
Local Gateway widget		
Local Gateway	Displays the name of the local BGP gateway.	(Name)
eBGP ASN	Displays the Autonomous System Number (ASN) of the local BGP gateway used by the session. It would be an integer number in the 1-65534 range (per 32-bit ASN.1 standard).	(Number)
Note: Make sure that this ASN does not conflict with any of the other on-premises BGP ASNs.		
Remote Gateway widget		
Remote Gateway	Displays the name of the remote BGP gateway.	(Name)
eBGP ASN	Displays the ASN of the remote BGP gateway used by the session. It would be an integer number in the 1-65534 range (per 32-bit ASN.1 standard).	(Number)
Note: Make sure that this ASN does not conflict with any of the other on-premises BGP ASNs.		

You can perform the following actions from the **Summary** tab:

- View the detailed information of a BGP session. For the list of available parameters, see the *BGP Session Details - Summary Tab Fields* table above.
- Update an existing BGP session by clicking **Update**. For more information, see [Updating a BGP session](#) on page 146.
- Delete an existing BGP session by clicking **Delete**. For more information, see [Deleting a BGP session](#) on page 147.

Routes Tab

The **Routes** tab provides a list of the routes used by the BGP session with the corresponding **Next Hop** details. It has the following lists:

- **Advertised** (default): The **Routes** tab opens in the **Advertised** list. The **Advertised** list provides a list of the advertised routes with the corresponding **Next Hop** details.
- **Received**: This list provides list of the routes received from remote with the corresponding **Next Hop** details.

BGP Logs Tab

The **BGP Logs** tab provides detailed live logs for the BGP session. This information can be very useful in monitoring and debugging a BGP session.




Figure 13: BGP Session Details View - BGP Logs tab sample for a BGP session

Security Policies

Security policies are defined using Nutanix Flow that provides a policy-driven security framework to inspect traffic within the data center.

For information on how to create and apply security policies on Basic VLAN Subnets, see [Flow Network Security \(formerly Flow Microsegmentation\) Guide](#).

For information on how to create and apply security policies on (advanced) VLAN Subnets and Overlay Subnets, see [Flow Network Security Next-Gen Guide](#).

For information on how to view security policies in Prism Central, see [Security Policies Summary View](#) or [Security Policy Details View](#) in the *Prism Central Infrastructure Guide*.

Security Dashboard

The Security Dashboard provides a dynamic summary of the security posture across all registered clusters.

The Security Dashboard provides a view of the most critical security parameters like cluster-based issue summary, STIG policy compliance, security hardening, and identified vulnerabilities. For more information, see [Security Dashboard](#) in the *Nutanix Security Guide*.

VIRTUAL PRIVATE CLOUD MANAGEMENT

You can deploy and manage virtual private clouds (VPCs) on a Nutanix cluster infrastructure using the Network Controller.

A Virtual Private Cloud (VPC) is an independent and isolated IP address space that functions as a logically isolated virtual network. A VPC is made up of one or more subnets that are connected through a logical or virtual router. The IP addresses within a VPC must be unique. However, IP addresses might overlap across VPCs. As VPCs are provisioned on top of another IP-based infrastructure (connecting AHV nodes), they are often referred to as the overlay networks. Tenants might spin up VMs and connect them to one or more subnets within a VPC.

Virtual Private Cloud (VPC) is a virtualized network of resources that are specifically isolated from the rest of the resource pool. VPCs manage the isolated and secure virtual network with enhanced automation and scaling. The isolation uses network namespace techniques like IP-based subnets or VLAN based networking.

AHV provides the framework to deploy VPC on on-premises clusters using the following.

- Advanced Networking subnets and DHCP management
- Multiple uplink and bridge management via virtual switch (VS)
- Virtual Private Network (VPN) gateways and connections

The Network Controller simplifies the deployment and configuration of overlay-based VPCs. The Network Controller provides the following actions for VPCs:

- Create, update and delete VPCs.
- Create, update and delete subnets within VPCs.

Note: Create subnets as necessary when you create VPCs.

- Add network security policies and services.
- Configure hybrid cloud connectivity with VPNs.

VPC Workflow

You can deploy the following types virtual private clouds (VPCs) on a Nutanix cluster infrastructure to manage the internal and external networking requirements using Flow Virtual Networking.

- *VPCs*: These are the VPCs that you create to isolate the groups of entities using overlay networks or subnets. This is the default VPC type. For more information, see *VPC* in [Essential Concepts](#) on page 12.
- *Transit VPC*: This is a *hub* VPC that VPCs connect to using one or two (NAT or No-NAT) external networks as spokes, in a a hub-and-spoke architecture to simplify the North-South connectivity. For more information, see *Transit VPC* in [Essential Concepts](#) on page 12.

The workflow to create a complete network based on VPC is described as follows:

1. Create a VPC or a transit VPC: For more information, see [Creating a Virtual Private Cloud](#) on page 89.
2. Update an existing VPC or transit VPC: For more information, see [Updating a Virtual Private Cloud](#) on page 103.
3. Add subnets to the VPC: For more information, see [Creating a Subnet](#) on page 93 to create a Subnet.
4. Update an existing subnet: For more information, see [Updating a Subnet](#) on page 105 to update a subnet.
5. Attach the subnet to VMs to VPCs: For more information, see [Attaching a Subnet to a Virtual Machine](#) on page 96.

VM IP Address Management

Primary Address

The primary IP address is assigned to a VM during initialization when the cluster provides any virtual NIC (NIC) to a VM.

- Select Assign Static IP as the **Assignment Type** to add a static IP address as primary IP address of the VM, when you attach a subnet to a VM.
- Select Assign with DHCP as the **Assignment Type** to allow DHCP to dynamically assign an IP address to the VM.
- Select No Private IP as the **Assignment Type** if you do not want to assign an IP address to the vNIC of the VM.

For more information on attaching a subnet to a VM, see [Creating a VM through Prism Central \(AHV\)](#) in the *Prism Central Infrastructure Guide*.

Secondary IP Addresses (Overlay Networks only)

For your deployment, you may need to configure multiple (static) IP addresses to a single NIC. These IP addresses (other than the primary IP address) are secondary IP addresses. A secondary IP address can be permanently associated with a specific NIC or be changed to any other NIC. The NIC ownership of a secondary IP address is important for security routing policies.

Note: You can configure secondary IP addresses only for VMs in an Overlay network.

Possible applications for secondary IP addresses include the following scenarios when you want to:

- Associate multiple floating IP addresses with one VM without creating multiple NICs (each with one primary IP address) for the VM. You can assign one floating IP address to one secondary IP address that you create for the single NIC. For information, see [Requesting Floating IPs](#) on page 92.
- Run appliances, such as load balancers, that have multiple IP addresses on each interface.
- Host applications in a High Availability (HA) configuration where the ownership of IP address moves from the active entity to the standby entity when the active entity goes down.
- Host applications in a clustered configuration where the ownership of IP address follows the leader.
- Host Nutanix Files service in a VPC as a case of clustered application.

Note: In applications that use secondary IP addresses as virtual IP addresses and the NIC ownership of the secondary IP address changes dynamically from one NIC to another, you must ensure that the ownership change is incorporated in the applications' settings or configuration. A secondary IP address can only be assigned to one VM at a time. To move the secondary IP address from the assigned VM to the another, first delete it from the assigned VM, then assign it to another VM. If the applications do not incorporate these ownership changes, incorporate the changes manually to ensure that the VPCs configured for such applications do not fail.

For information on configuring secondary IP addresses, see [Creating Secondary IP Addresses](#) on page 78.

IP Address Information

Click the **See More** link in the IP Address column in the VM details view to open the **IP Address Information** dialog box. The **IP Address Information** dialog box displays the IP addresses configured on a VM

Note: The **See More** link in the IP Address column in the VM details view and the **IP Address Information** box are available only if the VM has any secondary IP addresses configured.

Creating Secondary IP Addresses

You can assign multiple secondary IP addresses to a single vNIC.

About this task

You can add multiple secondary IP addresses to the vNIC configured on a VM. Add the secondary IP addresses to the vNIC in the **Create VM** or the **Update VM** page.

Perform the following steps to assign a secondary IP address to a vNIC configured on a VM.

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Compute & Storage > VMs** from the **Navigation Bar**.
The **VMs** page opens displaying the **List** tab.
3. Select the checkbox associated with the VM that contains the vNIC for which you want to add a secondary IP address.
4. Click **Update** from the **Actions** dropdown menu.
The **Update VM** page opens displaying the **Configuration** tab.
5. Click **Next**.
The **Resources** tab opens.
6. Go to the **Networks** section.
7. Click the **Edit** icon for the subnet that you want to add the secondary IP addresses from.
The **Update NIC** window opens.
8. Check the **Add Secondary IPs** checkbox in the **Update NIC** window.
9. Add a comma-separated list of the secondary IP addresses that you want to add to the vNIC of the VM.

Note:

Ensure that the secondary IP addresses are within the same subnet that the primary IP address of the NIC is from. The subnets are displayed in the **Private IP Assignment** section in the **Update NIC** window.

Ensure that the secondary IP address is not the same as the IP address provided in the **Private IP Assignment** field.

10. Click **Save**.
11. Click **Next** on the **Resources** and the **Management** tabs of the **Update VM** page.
If you need to make any other changes on the **Resources** and the **Management** tabs for any configurations other than adding secondary IP addresses, make the changes and then click **Next** on these tabs.
12. Click **Launch VM** on the **Review** tab after you review.

What to do next

You can view the secondary IP addresses configured on the VM in the **IP Address Information** box.

Assigning Secondary IP Addresses to Interfaces

About this task

Perform the following steps to assign the secondary IP addresses to virtual interfaces on the VM.

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Compute & Storage > VMs** from the **Navigation Bar**.
The **VMs** page opens displaying the **List** tab.
3. Click the target VM for which you want to assign a secondary IP address.
The VM details page opens displaying the **Summary** tab.
4. Click the **Console** tab.
5. Log in as a root user.
6. Run the ifconfig command as follows:

```
root@host$ ifconfig <interface> <secondary ip address> <network mask>
```

Provide the following values in the command:

Parameter	Description
<interface>	The interface of the VM such as eth0. You can provide subinterfaces such as eth0:1 and eth0:2.
<secondary IP address>	The secondary IP address that you created and want to associate with the interface.
<network mask>	The network mask that is an expansion of the network prefix of the network that the secondary IP address belongs to. For example, if the secondary IP address belongs to 10.0.0.0/24 then the network mask is 255.255.255.0.

7. Repeat the aforementioned steps for all the secondary IP addresses you want to associate with interfaces on the VM.
8. Exit from the Console.

Assigning Secondary IP Addresses to Floating IPs

About this task

After you assign secondary IP addresses to interfaces or subinterfaces on the VM, you can assign the secondary IP addresses to floating IP addresses that may be used for external connectivity.

Perform the following steps to assign a secondary IP address to floating IPs.

Procedure

- Log in to Prism Central.

- Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Network & Security** > **Floating IPs** from the **Navigation Bar**.
The **Floating IPs** page opens displaying the **List** tab.
- Perform either of the following:
 - » Click **Request Floating IP**. In the **Assign Floating IPs** section of the **Request Floating IP** window, assign floating IP addresses.

To assign floating IP addresses while requesting for them, you must have the secondary IP addresses configured and ready when you are requesting the floating IP addresses.
 - » In the **Floating IPs** page, select the checkbox associated with the floating IP address you want to assign. Click the **Update** option in the **Actions** dropdown menu.

Assign the secondary IP addresses you configured to the floating IP addresses you have.

VM and Network Migration

Flow Virtual Networking supports the following types of migrations:

- [Migration of VMs between VLAN Basic Subnet and VPC Subnets](#) on page 80
- [Migration of VLAN Basic Subnets](#) on page 85

Migration of VMs between VLAN Basic Subnet and VPC Subnets

You can migrate VMs networked in VLAN Basic Subnets to Flow Virtual Networking VPCs. The VMs networked using VLAN Basic Subnets are associated with categories. When you migrate the VLAN Basic Subnets to VPC subnets, the category associations are preserved.

Note: Flow Virtual Networking supports migration of VMs protected by protection policies from VLAN Basic networks to VPC subnets.

Migration Types

There are two types of migrations that you can select in the migration workflow.

- **Cold Migration.** For this type of migration, the incoming and outgoing connection configurations are not preserved. External connectivity for the subnet is irrelevant since the connections are not preserved.

If the source subnet is a managed subnet, the network ID and gateway is automatically populated based on the cluster and subnet selection. If the source subnet is not a managed subnet, specify the network ID and the gateway.

In both the above cases, the network ID and gateway of both the source and target networks must be the same. For example, if the network ID and gateway of the source are `10.10.10.0/32` and `10.10.10.1/32` then the target subnet must have `10.10.10.0/32` and `10.10.10.1/32` as the network ID and gateway. If the network ID and gateway are not the same then Prism central displays an error.

- **Live Migration without incoming connections.** For this type of migration, only outgoing connection configurations for the migrating VMs are preserved. Other considerations for this type of migration are:
 - During and after migration, you need to establish a subnet extension with Layer 2 connectivity between the two migrating subnets.

For more information on virtually extending a subnet at layer 2, see [Layer 2 Network Extension](#) on page 124.

- The external connection for the VPC must have NAT.
- The network ID and gateway of both the source and target networks must be the same. For example, if the network ID and gateway of the source are `10.10.10.0/32` and `10.10.10.1/32` then the target subnet must have `10.10.10.0/32` and `10.10.10.1/32` as the network ID and gateway. If the network ID and gateway are not the same then Prism central displays an error.

Conditions for Migration

You are unable to select some VMs for migration in the migration workflow because the selection button for those VMs are unavailable. When you hover on the selection button of such a VM, the pop-up message provides the reason for the unavailability of the VM for migration.

- You cannot migrate a VM with multiple vNIC. This is because a VM with vNICs in Acropolis and the Network Controller at the same time are not supported for migration. Therefore, ensure that the VM you want to migrate between VLANs and VPCs do not have multiple vNICs.
- You cannot migrate a VM which has a single vNIC with multiple IP addresses. Therefore, ensure that the VM you want to migrate between VLANs and VPCs has a single vNIC with a single IP address.
- You cannot perform cross-cluster live migration of VMs which are attached to Flow Network Security policies.
- Ensure that the IP addresses of the migrating VMs does not conflict with the IP addresses used by the VMs existing in the destination subnet. If you migrate a VM with conflicting IP address (in other words, an IP address that already belongs to another VM in the destination subnet) then an error is displayed and the migration fails for that VM.

Migrating VMs from VLAN Basic Subnets

About this task

You must have a **Super Admin** or **Prism Admin** access to migrate VMs from VLAN backed subnets to VPCs. If you are a user without **Super Admin** or **Prism Admin** level permissions, the **Migrate** button on the **Subnets** is unavailable.

You can migrate VLAN backed subnets on the **Subnets** page. Go to the **Subnets** page by clicking **Network & Security > Subnets**.

To migrate VLAN backed subnets to Flow Virtual Networking, on the **Subnets** dashboard, do the following.

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from Application Switcher Function, and navigate to **Network & Security > Subnets** from the **Navigation Bar**.
The **Subnets** page opens displaying the *List* tab.
3. Select the VLAN subnet that you want to migrate. Click **Migrate**.

4. On the **Migrate** page, do the following.
 - a. Select VLAN Basic Subnet in the **Migrate From** field.
 - b. Select Overlay Subnet in the **Migrate To** field.
 - c. Click **Proceed**.
5. On the **Migrate VMs between VLAN Network and VPC Network** page, select the **Migration Type** from the drop-down list.
You can select one of the following migration types:

- **Cold Migration.** For this type of migration, the incoming and outgoing connection configurations are not preserved.
- **Live Migration without incoming connections.** For this type of migration, only outgoing connection configurations for the migrating VMs are preserved.

For more information, see [Migration of VMs between VLAN Basic Subnet and VPC Subnets](#) on page 80.

6. To migrate VMs from VLAN to VPC, complete the configurations provided in the table and click **Next**.
You can also migrate VMs from VPC to VLAN.

Note: Click **Swap Source and Destination** link to toggle between **Migrate VMs from VLAN to VPC** and **Migrate VMs from VPC to VLAN**.

Figure 14: Migrate VMs from VLAN Network to VPC Network - Configuration

Table 25: Field Descriptions for the Configuration Tab

Field	Action	Description and Value
Source Subnet		
Cluster	Select the source cluster where the VM is located.	Name of the source cluster. (String)
Subnet (VLAN)	Select the source VLAN that networks the VM to be migrated.	Name of the source VLAN subnet. (String)

Field	Action	Description and Value
Network Address/Prefix	Enter (for unmanaged networks) the network IP address with prefix in CIDR notation. When you select a managed subnet in Subnet (VLAN) the Network Address/ Prefix	IP address of the source subnet. Use CIDR notation. For example, 10.10.10.0/32.
Gateway IP	Enter (for unmanaged networks) the Gateway IP address with prefix in CIDR notation. When you select a managed subnet in Subnet (VLAN) the Gateway IP value is automatically populated.	Gateway IP address of the source subnet. Use CIDR notation. For example, 10.10.10.1/32.
IPAM (Display only)	The IPAM status is displayed when you select a managed subnet in the Subnet (VLAN) parameter.	Displays Managed
Destination Subnet		
VPC	Select the VPC that you want to migrate the VM to.	VPC name (String)
Subnet (Overlay)	Select the Overlay subnet in the selected VPC that you want to migrate the VM to.	Name of the Overlay subnet
Network Address/Prefix	Enter (for unmanaged networks) the network IP address with prefix in CIDR notation. When you select a managed subnet in Subnet (VLAN) the Network Address/ Prefix	IP address of the source subnet. Use CIDR notation. For example, 10.10.10.0/32.
Gateway IP	Enter (for unmanaged networks) the Gateway IP address with prefix in CIDR notation. When you select a managed subnet in Subnet (VLAN) the Gateway IP value is automatically populated.	Gateway IP address of the source subnet. Use CIDR notation. For example, 10.10.10.1/32.
IPAM (Display only)	The IPAM status is displayed when you select a managed subnet in the Subnet (VLAN) parameter.	Displays Managed

Important:

For **Migrate VMs from VPC to VLAN** provide the configurations provided in the table, with the following differences.

- For **Source Subnet**, provide the VPC parameters which is the source.
- For **Destination Subnet**, provide the VLAN parameters which is the source.

7. In the **Virtual Machines** tab, select the VMs you want to migrate in the **Source** side of the tab and click **Add**.

Note: Select as many subnets as required before you click **Add**.




Figure 15: Migrate VMs from VLAN to VPC - Virtual Machines

For **Migrate VMs from VPC to VLAN**, the **Source** and destination subnets are reversed.

The selected VM or VMs are displayed on the **Destination** side of the tab.

You can select the migrating VM (**Added on the Destination side**) and click **Assign New IP** to assign a new IP address to the migrating VM after migration.

Note: The IP address of the migrating VM is persisted after migration if the existing IP address is available in the destination subnet. If you migrate a VM with conflicting IP address (in other words, an IP address that already belongs to another VM in the destination subnet) then an error is displayed on the **Migrate** tab.

- Click **Back** on the **Migrate** tab.
- On the **Virtual Machines** tab:
 - Select the migrating VM with conflicting IP address.
 - Click **Assign New IP**.

This ensures that a new IP address is assigned to the migrating VM after migration.

Click **Begin Migration** to start the migration process.

The **Migration** tab displays the progress of the migration.




Figure 16: Migrate VMs from VLAN Network to VPC - Migration

When the migration process is complete, the **Migrate** tab displays the status of the migration. It displays any errors that may have occurred during migration, the reason for failure of any VM migration.

When the status changes to **Migration Completed Successfully** with date and time stamp, the **Migration Summary of VM** table is displayed. You can filter the **Migration Summary of VM** by status using the status drop-down. There are three states: Completed, Failed and Pending. Usually, the **Migration Summary of VM** does not appear when the migration state of any migrating VM is **Pending**. Therefore, you may not find any VM listed with **Pending** state in the summary. A VM migration with pending state is displayed in **Tasks**.




Figure 17: Migrate VMs from VLAN Network to VPC - Migration status

8. Click **Close** to close the **Migrate VMs between VLAN Network and VPC Network** window after migration is complete and successful.

What to do next

You can view the migration history on the **Subnets** dashboard by clicking **Migrate > View Migration History**. The migration history table displays several attributes of the migration tasks including **Status** of the migration tasks and the **Duration** taken by the migration task to complete.

Migration of VLAN Basic Subnets

For information on VLAN Basic Subnet, and VLAN Subnets, also see [Network Types](#) on page 40, [Changing the Default VLAN Type](#) on page 42, and [Creating a Subnet](#) on page 93.

With a minimum Prism Central version of PC.2023.3 that deploys Network Controller 3.0.0 on a minimum AOS version of 6.7:

- Flow Virtual Networking supports migration of VLAN Basic Subnets (see [Essential Concepts](#) on page 12 and [Network Types](#) on page 40) to VLAN Subnets (managed by the Flow Virtual Networking controller, see [Essential Concepts](#) on page 12 and [Network Types](#) on page 40).

In the **Subnets** page that has the list of subnets, the Network Controller VLAN (VLAN or VLAN Subnet) does not have a suffix in the **Type** field. The VLAN Basic Subnet is suffixed with the word **Basic**.

- When you create VLANs using the [Creating a Subnet](#) workflow, Prism Central creates a VLAN of the type set as default.

You can change the default VLAN type that is created while you are [Creating a Subnet](#), from VLAN Subnet type to VLAN Basic Subnet type and vice versa. For information on changing the default VLAN type, see [Changing the Default VLAN Type](#) on page 42.

Migration Process

The migration process involves migrating one subnet at a time. It locks the VLAN Basic subnet on AHV and creates a corresponding VLAN Subnet on Prism Central with the same UUID and properties or attributes as the VLAN Basic Subnet. Next, the migration process updates and migrates all the vNICs to the VLAN Subnet. After all the vNICs are migrated to the VLAN Subnet, the VLAN Basic Subnet on AHV is deleted.

Since the migration process retains the UUID of the VLAN Basic Subnet, any automations that use UUIDs are protected from impact. The MAC addresses of the vNICs are also preserved after migration, thus reducing any impact to configurations and automations that use these MAC addresses.

The Prism Central VM vNICs must always remain on a VLAN Basic Subnet. Therefore, when you migrate a VLAN Basic Subnet that hosts Guests VMs and Prism Central VMs, the Prism Central VMs vNICs are migrated to a newly created VLAN Basic Subnet on AHV host. The Prism Central VM is not migrated even if it is configured in the VLAN Basic Subnet that is marked for migration.

Note: Migration is irreversible. You must create a VLAN Basic subnet and move the vNICs to that Subnet if you want to use a VLAN Basic Subnet for vNICs that were previously migrated to a VLAN Subnet.

The migration process includes a pre-check that ascertains if all the necessary conditions for migration are met.

Migration Pre-check Conditions

You must have a **Super Admin** or **Prism Admin** access to migrate a VLAN Basic Subnet to VLAN subnet. If you are a user without **Super Admin** or **Prism Admin** level permissions, the **Migrate** button on the **Subnets** is unavailable.

The migration process includes a pre-check that ascertains if the necessary conditions for migration are met as follows:

- The migration pre-check determines whether the Network Controller is enabled.
- The migration pre-check determines whether Flow Network Security is enabled. If enabled, initiate the migration process from the Policy page.

Note: Before you migrate a VLAN Basic Subnet to VLAN Subnet, migrate the attached FNS policy to FNS Next-Gen on the **Security Policies** page.

- The migration pre-check determines whether the number of vNICs or subnets included in the migration is within the scale numbers specified in *vNIC Scale* in [Network Types](#) on page 40.
- The migration pre-check determines whether the VLAN Basic Subnet to be migrated is associated with a Virtual Switch. VLAN Basic Subnets that do not have a Virtual Switch reference cannot be migrated.

For more information on virtual switches and how to change the virtual switch that the VLAN Basic Subnet is attached to, see [Virtual Switch Management](#) in the *AHV Administration Guide*.

- The migration pre-check determines whether the number of VLAN Basic Subnets included in a single migration request is equal to or less than 100. You can only migrate a maximum of 100 VLAN Basic Subnets in a single migration request.
- The migration pre-check determines whether any of the VLAN Basic Subnets have kDirect vNICs or vNICs in Trunk mode. Migration of VLAN Basic Subnets with kDirect vNICs or vNICs in Trunk mode is not supported.
- The migration pre-check determines whether the VLAN Basic Subnets are associated with any Nutanix Files VMs. Migration of VLAN Basic Subnets associated with Nutanix Files VMs is **not supported**.
- The migration pre-check determines whether the VLAN Basic Subnets are associated with any Protection Domains for disaster recovery (see [Data Protection and Recovery with Prism Element](#)). Migration of VLAN Basic Subnets associated with Protection Domains for disaster recovery is **not supported**.
- The migration pre-check determines whether any of the VLAN Basic Subnets are not managed subnets hosting the vNICs of Prism Central VMs. Migration of (managed) VLAN Basic Subnets that host Prism Central VM vNICs is **not supported**.
- The migration pre-check determines whether Microservices Infrastructure uses any of the VLAN Basic Subnets. Migration of VLAN Basic Subnets that Microservices Infrastructure uses is **not supported**.
- The migration pre-check determines whether any of the VMs in any of the VLAN Basic subnets to be migrated have vNICs in multiple VLAN Basic Subnets but not all those subnets are being migrated. Such VLAN Basic Subnets that have VMs that do not have all the vNICs in the migrating VLAN basic Subnets, **cannot be migrated**.

Migrating a VLAN Basic Subnet to VLAN Subnet

Before you begin

Ensure that Flow Virtual Networking is enabled.

About this task

Perform this task to migrate a VLAN Basic Subnet (AHV Networking based VLAN) and all the virtual NICs configured in that subnet to a newly create VLAN Subnet (Network Controller based VLAN). The target VLAN Subnet is automatically created during the procedure of migration. The outcome of this task is that the VLAN Basic Subnet configuration is fully migrated to a VLAN Subnet.

You can migrate a VLAN Basic Subnet to VLAN subnet on the **Subnets** page. Go to the **Subnets** page by clicking **Network & Security > Subnets**.

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from Application Switcher Function, and navigate to **Network & Security > Subnets** from the **Navigation Bar**.
The **Subnets** page opens displaying the *List* tab.
3. Click **Migrate**.
4. To migrate VLAN Basic Subnet to VLAN Subnet, do the following.
 - a. Select VLAN Basic Subnet in the **Migrate From** field.
You can select a maximum of 100 VLAN Basic Subnets for migration in a single request.
 - b. Select VLAN Subnet in the **Migrate To** field.
 - c. Click **Next**.

5. On the **VLAN Basic to VLAN Migration** page, click **Add** to add the VLAN Basic subnets that you need to migrate.
6. On the **Add Subnets to migration** page, from the list of subnets, select the check boxes for the VLAN Basic Subnets that you need to migrate.
Click **Add**.
7. On the **VLAN Basic to VLAN Migration** page, click **Begin Migration** to start the migration process.
The **Migrate (x) VLAN Subnet(s)?** tab displays a message about the temporary downtime, that migration cannot be aborted or paused and that migration of other subnets cannot begin until this migration process is completed.
Click **Migrate**.
8. On the **Migrating (x) Subnets** tab displays the progress of the migration.

When the migration process is complete, the **Migrating (x) Subnets** tab displays the status of the migration. It displays any errors that may have occurred during migration, the reason for failure of any VM migration.

You can filter the list of migrated VLANs using the drop-downs and filters in the table of migrated VLANs.
9. Click **Close** to close the **VLAN Basic to VLAN Migration** window after migration is complete and successful.

What to do next

You can view the migration history on the **Subnets** dashboard by clicking **Migrate > View Migration History**.

Note: If you see that the migration of any of the VLAN Basic Subnets has failed, initiate the migration for those VLAN Basic Subnets again by following all the above steps in this procedure.

VPC Management

This section provides information and procedures that you need to manage virtual private clouds (VPCs), subnets, routing policies, and static routes.

A Virtual Private Cloud (VPC) is an independent and isolated IP address space that functions as a logically isolated virtual network. A VPC is made up of one or more subnets that are connected through a logical or virtual router. VPCs allow you to manage the isolated and secure virtual network with enhanced automation and scaling. The isolation is done using network namespace techniques like IP-based subnets or VLAN based networking.

Flow Virtual Networking supports the following two types of VPCs.

VPC

The default VPC type that is referred to as *VPC* in this documentation is the one you create to isolate selected subnets of connected VMs. This is also called as User VPC or Guest VPC, specifically referred to as *VPC*.

The other VPC type is transit VPC, specifically referred to as *transit VPC* in this documentation.

Transit VPC

Overlay External Subnet for Transit VPCs

- You can only use a VLAN based network for the uplink (external connectivity) for a transit VPC. In other words, a transit VPC cannot be connected to another transit VPC.
- You can configure an Overlay subnet with external connectivity in transit VPC. When you create an Overlay external subnet, the workflow provides only transit VPCs in the **VPC** dropdown menu.
- You can configure an Overlay subnet with external connectivity (Overlay external subnet) with options such as NAT or NONAT (NAT being default) and necessary gateways for the NAT or No-NAT option.

- You can connect only regular VPCs to transit VPC using an Overlay external subnet.
You cannot attach any VMs to an Overlay external subnet. You cannot connect, for example, two regular VPCs to each other using an Overlay external subnet.
- This is in line with the behavior for the VLAN backed external subnets. The external overlay subnets are how a regular VPC will connect to a transit VPC. Two transit VPCs will not be allowed to be connected using this.

For information on VPCs, see [Essential Concepts](#) on page 12.

Creating a Virtual Private Cloud

You can create VPCs and transit VPCs on the **Virtual Private Clouds** page.

About this task

Perform the following steps to create a VPC.

Procedure

1. Log in to Prism Central.
2. From [Application Switcher Function](#), select the **Infrastructure** application and navigate to **Network & Security > Virtual Private Clouds** from the **Navigation Bar**.
The **Virtual Private Clouds** page opens displaying the **List** tab.
3. Click **Create VPC**.
The **Create VPC** window opens.

Figure 18: Create VPC

4. Provide the necessary values in respective fields.

Parameters	Description and Values
Name	Provide a name for the VPC.

Parameters	Description and Values
Transit VPC toggle switch	Toggle the Transit VPC toggle switch to Yes if you want to create a transit VPC instead of a regular VPC. For more information, see Virtual Private Cloud Management on page 76 for information about the difference between a VPC and a transit VPC.
External Connectivity	This section lets you Associate External Subnets for the VPC. A subnet with external connectivity (External Subnet) is required if the VPC needs to send traffic to a destination outside of the VPC. Note: You can add a maximum of two external subnets - one external subnet with NAT and one external subnet without NAT to a VPC. Both external subnets cannot be of the same type. For example, you cannot add two external subnets, both with NAT. You can update an existing VPC similarly. Network address translation (NAT) Gateways perform the required IP address translations required for external routing. You can also have external connectivity without NAT (No-NAT).
External Connectivity > External Subnet	Displays the name of the external subnet that is associated with the VPC.
External Connectivity > Destination Prefixes	Displays the prefixes for which this external subnet is used as the next hop. The selection is based on the longest prefix match.
External Connectivity > SNAT IP / Router IP	Displays the SNAT IP or Router IP addresses that the IPs assigned to the VPC router in the external subnet. It is used as SNAT IP in case of NAT external subnet. These addresses would be used by the physical network router as the next hop for all the networks reachable inside the VPC using a No-NAT external subnet. Note: You can specify a custom SNAT or Router IP selected from the IP address pool of the
External Connectivity > Actions	Displays the actions (as icons) that you can perform on the external subnet. The actions listed are: Edit and Delete
Associate External Subnet button	Click the Associate External Subnet button to display the Associate External Subnet window which allows you to configure the external subnet parameters.
Associate External Subnet window details	
Associate External Subnet > Subnet Type	Select the type of subnet that you have configured as the external subnet. The types you can select from are VLAN being the VLAN subnet and Overlay subnets
Associate External Subnet > External Subnet	Select an external subnet from the drop down list. By associating the VPC with the external subnet you can provide external connectivity to the VPC. When you select the external subnet, the details of the subnet like Network Address/Prefix, NAT-ed (which displays the NAT status of the subnet as Yes or No, and (for only VLAN type subnet) VLAN ID of the VLAN External Subnet are displayed in a table below the External Subnet dropdown list field.

Parameters	Description and Values
Static Routes	Configure the static routes that specify the list of prefixes for which the selected external subnet is the next hop. Also, configure the routes on the router for the return traffic to reach the VPC. For more information, see <i>External Connectivity</i> in Virtual Private Cloud Details View on page 54.
SNAT IP/Router IP	Select the appropriate option from Auto Assigned or Custom Defined. The SNAT or Router IP address is the next hop for the physical routing infrastructure. If you select the Auto Assigned, the Network Controller assigns an IP address from the IP address pool of the external subnet as SNAT or Router IP address. When you select the Custom Defined option, a table with details of available IP address pool is displayed. This table displays IP Pool Range, Used IPs in Pool, and Free IPs in Pool information for the pool. Enter an IP address selected from the Free IPs in Pool that you want to be assigned as SNAT or Router IP address, in the Custom SNAT IP / Router IP field.
External Gateway Configuration > Number of Active Hosts	Displayed only when you select a No-NAT VLAN external subnet from the Associate External Subnet > External Subnet dropdown menu) Select the Number of Active Hosts, in other words, the number of No-NAT gateways you need. You can select up to 4 gateways. If you do not select the number of gateways or active hosts, the number is pre-selected as 2, the default number of gateways or active hosts. No-NAT gateways are deployed as AHV hosts. For more information, see <i>No-NAT Gateway</i> in Essential Concepts on page 12.
Other details on the Create VPC page	
Externally Routable IP Addresses	(Optional) Add externally routable IP addresses or subnets with external connectivity without NAT. These are used by BGP Gateways to advertise routes. Ensure that the externally routable prefix (ERP) IP addresses that you provide, do not overlap with those provided for other VPCs. These prefixes are reachable from outside the VPC. When you add these prefixes to the VPC and use BGP sessions, these ERPs are advertised to the peers with the next hop as the router IP address(es) of the VPC in the No-NAT external subnet. Thus, the ERPs become reachable from the peer since the peer knows that the route passes through the VPC router. For conditions applicable to configuring ERPs in a transit VPC, see the Note in the <i>Transit VPC</i> section of Essential Concepts on page 12.
Domain Name Servers (DNS)	(Optional) DNS is advertised to Guest VMs via DHCP. This can be overridden in the subnet configuration. Click + Server IP to add DNS server IPs under IP Address and click the check mark. You can Edit or Delete an IP address you added using the options under Actions.

5. Click **Create**.

Requesting Floating IPs

About this task

User VMs or VPN gateways or many such entities require Floating IP addresses. To provide floating IP to an entity, you can request Floating IP addresses and assign them to VMs.

Perform the following steps to request a floating IP.

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Network & Security > Floating IPs** from the **Navigation Bar**.
The **Floating IPs** page opens displaying the **List** tab.
3. Click **Request Floating IP**.
The **Request Floating IP(s)** window opens.
4. Enter the information in the respective fields.

Note: Clear the **Assign Floating IPs** checkbox if you want to assign the requested IP addresses after you receive it. For more information, see [Floating IPs Summary View](#) on page 59.

Fields	Description and Values
External Subnet	Select a subnet that you configured with external connectivity. When you select an external subnet, a box displays the IP pool information for the selected external subnet. The following IP pool details are displayed. • IP Pool Ranges: Displays the range of IP addresses with the starting IP address and ending IP address of the range. • Used IPs in the Pool: Displays the number of IP addresses already used from the pool. • Free IPs in the Pool: Displays the number of unused IP addresses in the pool.
Number of Floating IPs	Enter the number of Floating IPs you want to request. You can request a maximum of 50 floating IP addresses.
Define Custom Floating IPs	Select this check box if you want to select specific IP addresses from the IP address pool range of the external subnet. When you select the check box, the Enter Floating IPs to be requested field is displayed below the check box text. Enter the specific IP addresses that you want to request as Floating IPs in this field.

Fields	Description and Values
Assign Floating IPs	Select this check box if you want to assign the Floating IPs to specific VMs in the table. Based on the number you entered in the Number of Floating IPs field, the system provides an equivalent number of rows of Search VMs and IP Address in the table. Under Search VMs, select the VM to which you want to assign a floating IP address. Under IP Address, select the IP address on the VM (primary or secondary IP address) to which you want to assign the floating IP. You can assign multiple floating IP addresses to multiple secondary IP addresses that you can create on the NIC of the VM. For information on configuring secondary IP addresses, see Creating Secondary IP Addresses on page 78.

5. Click **Save**.

What to do next

When you receive the floating IP address you requested, you can see it, assign it (if not already assigned while requesting) or delete it in the **Floating IPs** view.

Creating a Subnet

About this task

Perform the following steps to create a subnet.

Procedure

1. Log in to Prism Central.
2. From [Application Switcher Function](#), select the **Infrastructure** application and navigate to **Network & Security** > **Subnets** from the **Navigation Bar**.
The **Subnets** page opens displaying the **List** tab.
3. Click **Create Subnet**.
The **Create Subnet** window opens. The following figure displays the Create Subnet window with all the options. These options are displayed based on the values you select in the **Type** field.

Fields	Description and Values
Name	Provide a name for the subnet.
Type	Select the type of subnet you want to create. You can create a VLAN subnet or an Overlay subnet.
VLAN ID	(VLAN subnet only) Enter the number of the VLAN. Enter just the number in this field, for example 1 or 27. Enter 0 for the native VLAN. The value is displayed as vlan.1 or vlan.27 in the View pages.

Note: Provision any single VLAN ID either in the AHV network stack or in the Flow Virtual Networking (brAtlas) networking stack. Do not use the same VLAN ID in both the stacks.

Fields	Description and Values
IP Address management	(Mandatory for Overlay type subnets) This section provides the <i>Network IP Prefix</i> and <i>Gateway IP</i> fields for the subnet. (Optional for VLAN type subnet) Select this checkbox to display the <i>Network IP Prefix</i> and <i>Gateway IP</i> fields and configure the IP address details. Clearing this checkbox hides these fields. In this case, it is assumed that this virtual LAN is managed outside the cluster. Note: The DHCP Settings option is only available for VLAN subnets if you select this option.
DHCP Settings	(Optional for both VLAN and Overlay subnets) Select this checkbox to display fields for defining a domain. Selecting this checkbox displays fields to specify DNS servers and domains. Clearing this checkbox hides those fields. For more information, see Setting the DHCP Options on page 95.
Cluster (VLAN subnet only)	(VLAN subnet only) This option is available only for VLAN subnet configuration. Select the cluster that you want to assign to the subnet.
External Connectivity	Turn on this toggle switch if you want use this (VLAN or Overlay) subnet for external connectivity. The External Connectivity toggle switch is displayed as an option for an Overlay subnet, only if you associate the subnet with a transit VPC (selected in the VPC drop down menu that is displayed only when you select Overlay in the Type drop down menu). Note: • Ensure that the externally routable IP addresses (subnets with external connectivity without NAT) for different VPCs do not overlap. • Configure the routes for the external connectivity subnets with next hop as the Router or SNAT IP address. Also configure the routes on the router for the return traffic to reach the VPC. For more information, see <i>External Connectivity</i> in Virtual Private Cloud Details View on page 54.
NAT	(Option under External Connectivity) If you turn on the External Connectivity toggle switch, you can choose whether to connect to external networks with or without enabling NAT. Select the NAT checkbox to enable NAT for external connectivity for VPCs.
Virtual Switch	(VLAN subnet only) Select the virtual switch that is configured for the VLAN you selected. The default value is the default virtual switch vs0. This option is displayed only if you add a VLAN ID in the VLAN ID field.
VPC	(Overlay subnet only) Select the Virtual Private Cloud (VPC) that you want to assign to the subnet from the drop down list. You can create VPCs and assign them to Overlay subnets.

Fields	Description and Values
IP Address Pool	Defines a range of addresses for automatic assignment to virtual NICs. This field is optional for both VLAN and Overlay. For VLAN, this field is displayed only if you select the IP Address Management option. Note: Configure this field for VLAN or Overlay to complete the creation of the VPC, if you do not need external connectivity for this subnet. You must configure this field only if you need external connectivity for this subnet. Click the Create Pool button and enter the following in the Add IP Pool page: - Enter the starting IP address of the range in the Start Address field. - Enter the ending IP address of the range in the End Address field. - Under Actions, click the check mark to submit the starting and ending IP addresses you entered. Click the X mark to remove the entries.
Override DHCP Server	(VLAN subnet only) To configure a DHCP server, select the Override DHCP Server checkbox and enter an IP address in the DHCP Server IP Address field. For more information, see <i>Override DHCP Server (VLAN Only)</i> in Setting the DHCP Options on page 95.
Advanced Configuration —VLAN Basic Networking	(VLAN subnet only) Select the VLAN Basic Networking checkbox to create the Basic VLAN on AHV networking (see Basic VLANs or VLAN Basic Subnet in Essential Concepts on page 12 and Network Types on page 40).

4. Click **Create**.

Setting the DHCP Options

About this task

Selecting the **DHCP Settings** checkbox in **Create Subnet** or **Update Subnet** allows you to configure the DHCP options for the VMs within the subnet. When DHCP settings are configured for a VM in a subnet and the VM is powered on, Flow Virtual Networking configures these options on the VM automatically. If you do not configure the DHCP settings, then these options are not available on the VM automatically when you power it on.

You can enable **DHCP Settings** when you create a subnet and configure the **DHCP Settings** for the new subnet. You could also update the DHCP Settings for an existing subnet.

DHCP Settings is common to and is available on both the **Create Subnet** and the **Update Subnet** dialog boxes. To configure the **DHCP Settings**, do the following:

Procedure

- Provide the information in the **DHCP Settings** fields.

Fields	Description and Values
Domain Name Servers	Provide a comma-separated list of DNS IP addresses. Example: 8.8.8.8, 9.9.9.9

Fields	Description and Values
Domain Search	Enter the VLAN domain name. Use only the domain name format. Example: nutanix.com
TFTP Server Name	Enter a valid TFTP host server name of the TFTP server where you host the host boot file. The IP address of the TFTP server must be accessible to the virtual machines to download a boot file. Example: tftp_vlan103
Boot File Name	The name of the boot file that the VMs need to download from the TFTP host server. Example: boot_ahv2020xx

- (Optional and for VLAN networks only) Check the **Override DHCP Server** dialog box and enter an IP address in the **DHCP Server IP Address** field.

You can configure a DHCP server using the **Override DHCP Server** option only in case of VLAN networks.

The DHCP Server IP address (reserved IP address for the Acropolis DHCP server) is visible only to VMs on this network and responds only to DHCP requests. If this box is not checked, the DHCP Server IP Address field is not displayed and the DHCP server IP address is generated automatically. The automatically generated address is `network_IP_address_subnet.254`, or if the default gateway is using that address, `network_IP_address_subnet.253`.

Usually the default DHCP server IP is configured as the last usable IP in the subnet (For eg., its 10.0.0.254 for 10.0.0.0/24 subnet). If you want to use a different IP address in the subnet as the DHCP server IP, use the override option.

Attaching a Subnet to a Virtual Machine

About this task

Perform the following steps to attach a subnet to a VM.

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Compute & Storage > VMs** from the **Navigation Bar**.
The **VMs** page opens displaying the **List** tab.
3. Select the VM you want to attach a subnet to, and click **Update** from the **Actions** dropdown menu.
The **Update VM** page opens displaying the **Configuration** tab.
4. Click **Next**.
The **Resources** tab opens.
5. Click **Attach to Subnet**.
The **Attach to Subnet** window opens.

6. Provide the necessary information in the indicated fields.

- a. Select the **Subnet Name** from the dropdown menu.
- b. Select the **Network Connection State** as Connected or Disconnected.

The **Network Connection State** selection defines the state of the connection after the NIC configuration is implemented.

- c. Select the **Assignment Type**.

You can select Assign with DHCP to assign a DHCP based IP address to the VM.

You can select Assign Static IP to assign a static IP address to the VM to reach the VM quickly from any endpoint in the network such as a laptop.

7. Click **Save**.

Creating a Policy

For Policy-based routing you need to create policies that route the traffic in the network.

About this task

When you create a VPC, the Network Controller creates a default policy for the VPC. This policy is preconfigured with the Priority 1 and other default values to Deny traffic flow and service (see the table of field descriptions and values for this dialog box).

Note: You cannot update or delete the default policy.

- Policies control the traffic flowing between subnets (inter-subnet traffic).
- Policies control the traffic flowing in and out of the VPC.
- Policies do not control the traffic within a subnet (intra-subnet traffic).

Perform the following steps to create a policy.

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Network & Security > Virtual Private Clouds** from the **Navigation Bar**.
The **Virtual Private Clouds** page opens displaying the **List** tab.
3. Click the name of the VPC for which you need to create a policy.
The **Summary** tab opens displaying the detailed information about the VPC in widgets.
4. Click the **Policies** tab.
5. Click **Create Policy**.
The **Create Policy** window opens.

6. Provide the necessary values in the respective fields.

Create Policy [X]

Priority [Understand Priorities](#)
Enter priorities between 10 and 1000

Source
Any

Destination
Any

Protocol
Any

Actions
Permit

☐ Additionally Create Policy In reverse direction

If the Permit rule is overriding a Drop, then the policy rule must be set in both directions to allow bidirectional communication between the Source and Destination.

Cancel Create

Figure 19: Create Policy Page

The following table describes the fields that appear in the **Create Policy** window.

Fields	Description and Values	Value in Default Policy
Priority	The priority of the access list (ACL) determines which ACL is processed first. Priority is indicated by an integer number. A higher priority number indicates a higher priority. For example, if two ACLs have priority numbers 100 and 70 respectively, the ACL with priority 100 takes precedence over the ACL with priority 70. Note: Click the Understand Priorities link to see the Understand Priorities information box (see the illustration of this box after this table).	1

Fields	Description and Values	Value in Default Policy
Source	The source indicates the source IP or subnet for which you need to manage traffic. Source can be: - Any: Indicates any IP address. - External: Indicates an IP address that is outside the subnets configured for the VPC. - Custom: You can provide a specific Source Subnet IP with prefix.	Any
Source Subnet IP	Only required if you selected the Source as Custom. Provide the subnet IP and prefix that you need to designate as the source for the policy. Use the CIDR notation format to provide the subnet IP. For example, 10.10.10.0/24.	None
Destination	The destination is the destination IP or subnet for which you need to set the priority. Destination can be: - Any: Indicates any IP address. - External: Indicates an IP address that is outside the subnets configured for the VPC. - Custom: You can provide a specific Destination Subnet IP with prefix.	Any
Destination Subnet IP	Only required if you selected the Destination as Custom.	None
Protocol	You can also set the priority of the policy for certain protocols. Select one of the following options: - Any: Indicates any protocol. - Protocol Number: Provide an integer number that indicates the protocol to prioritize. Provide the appropriate value in the Protocol Number field. - TCP - UDP - ICMP	Any
Protocol Number	This field is displayed only if you select Protocol Number as the value in the Protocol field. The number you provide must be the IANA designated number that indicates respective protocol. For more information, see IANA Protocol Numbers.	None
Actions that you can assign to the traffic.		

Fields	Description and Values	Value in Default Policy
Permit	The Permit action permits traffic and services based on the parameters set. If the Permit rule is set to override a Drop rule (see Drop under Fallback Action in the Reroute row), then for a bidirectional Drop rule, the Permit rule must be set in both the directions to allow bidirectional communication between the Source and Destination, by selecting Additionally Create Policy in reverse direction.	Permit
Deny	The Deny action denies traffic and service based on the parameters set.	
Reroute	The Reroute action sends matching traffic to the next-hop IP address specified by the Reroute IP. Use the Reroute action to reroute internal traffic to the VPC. Select the Configure separate reroute IP for incoming and outgoing traffic check box to route the incoming traffic on an IP address to another IP address. If you do not select Configure separate reroute IP for incoming and outgoing traffic, configure a single IP address for bot incoming and outgoing traffic in Reroute IP Address (Incoming and Outgoing traffic). If you select Configure separate reroute IP for incoming and outgoing traffic, configure the IP address for the incoming traffic in Reroute IP Address (Incoming Traffic) and the IP address for the outgoing traffic a IP address configured in Reroute IP Address (Outgoing Traffic). Select a Fallback Action for the fallback traffic routing, from the following: • Pass-through that processes the traffic based on the next highest priority rule. • Drop that drops the traffic. • Allow that forwards the traffic to the destination. • No Action that reroutes the traffic to persist to the incoming and outgoing Re-route IP addresses provided. You can also persist traffic to a specified IP address that might be assigned or re-assigned to any entity by selecting this action. See <i>Reroute Policy</i> in Essential Concepts on page 12.	

Fields	Description and Values	Value in Default Policy
Forward	Forward: Forwards matching traffic to the external next-hop IP address specified by the Forward IP. Provide an IP address that the traffic needs to be forwarded to, in the Forward IP field. Note: You can apply the Forward action only if you have installed or upgraded the Network Controller to version 3.0.0 or later for the respective Prism Central version pc.2023.3 or later. You can forward the traffic using the Forward action works only if the nexthop IP address is directly connected to the logical router of the VPC. Therefore, the nexthop IP address must belong to an external subnet or the IP address range of the another subnet in a Layer 2 extended subnet. The Layer 2 subnet extension is the only usecase for the Forward action.	

Fields	Description and Values	Value in Default Policy
Additionally Create Policy in reverse direction	Select this checkbox to set a policy rule in the reverse direction if you need to setup bidirectional communication between the Source and Destination .	Select or clear

Understand Priorities

?

×

The priority of the ACL determines which ACL will be processed first. The higher the priority number higher the priority.

Example:

Priority	Source	Destination	Protocol
100	10.1.20.0/15	Any	ICMP
70	External	Any	TCP

Close

Figure 20: Understanding Priorities

- Click **Create**.

Creating Static Routes

About this task

Perform the following steps to create static routes.

To create static route, do the following in the **Create Static Routes** dialog box:

Procedure

- Log in to Prism Central.

2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Network & Security > Virtual Private Clouds** from the **Navigation Bar**.
The **Virtual Private Clouds** page opens displaying the **List** tab.
3. Click the name of the VPC for which you want to create a static route.
The **Summary** tab opens displaying the detailed information about the VPC in widgets.
4. Click the **Routes** tab.
5. Click **Manage Static Routes**.
The **Manage Static Routes** window opens.
6. Click **Add Static Route**.
7. Provide the necessary values in the respective fields.
The following table describes the fields that appear in the **Manage Static Routes** window.

Fields	Description and Values
Destination Prefix	Provide the IP address with prefix of the destination subnet.
Next Hop Link	Select the next hop link from the drop down list. The next hop link is the IP address that the traffic must be sent for the static route you are configuring.
Add Static Route	You can create multiple static routes using this option. Click this link to add another set of Destination Prefix and Next Hop Link to configure another static route.

8. Click **Save**.

Updating a Virtual Private Cloud

About this task

Perform the following steps to update a VPC or a transit VPC.

Procedure

1. Log in to Prism Central.
2. From [Application Switcher Function](#), select the **Infrastructure** application and navigate to **Network & Security > Virtual Private Clouds** from the **Navigation Bar**.
The **Virtual Private Clouds** page opens displaying the **List** tab.
3. Select the checkbox associated with the VPC you want to update, and click **Update** from the **Actions** dropdown menu.
The **Update VPC** window opens.

4. Update the necessary values in the respective fields.

The fields in the **Update VPC** window is identical to the fields in the **Create VPC** window. For more information, see [Creating a Virtual Private Cloud](#) on page 89.

Note: You cannot update the **Associate External Subnet > Number of Active Hosts**, in other words, the number of No-NAT gateways selected for the VPC.

To update the **Number of Active Hosts** for the already selected external, No-NAT VLAN network, do the following.

- Delete the associated external, No-NAT VLAN network.
- Select **Update** to save the deletion.
- Select **Associate External Subnet** and add the previously associated external, No-NAT VLAN network.
- Select the necessary number of **Number of Active Hosts** after appropriately configuring other parameters in the **Associate External Subnet** window.
- Select **Update** to save the association.

Name

test-techpubs-nonat

Transit VPC ☐ No

☒ External Connectivity

Subnets with external connectivity (External Subnets) are required to be associated with this VPC to send traffic to a destination outside of it.

[Associate External Subnet](#)

External Subnet	Destination Prefixes	SNAT IP / Router IP	Actions
vlan-nonat-techpubs-2	-	10.10.20.11 , ...	

Externally Routable IP Addresses

Must be non-overlapping (with other VPCs) addresses, between /16 and /28 netmasks

Domain Name Servers (DNS)

Cancel Update

Figure 21: Update VPC

5. Select **Update** on the **Update VPC** page.

Updating a Subnet

About this task

Perform the following steps to update a subnet.

Important: You cannot edit or update the subnet type. For example, if the subnet type is already configured as *VLAN*, you cannot modify it to an *Overlay* type subnet.

Procedure

1. Log in to Prism Central.
2. From [Application Switcher Function](#), select the **Infrastructure** application and navigate to **Network & Security > Subnets** from the **Navigation Bar**.
The **Subnets** page opens displaying the **List** tab.
3. Select the checkbox associated with the subnet you want to update, and click **Update** from the **Actions** dropdown menu.
The **Update Subnet** window opens.
4. Update the necessary values in the respective fields.

The fields in the **Update Subnet** window is identical to the fields in the **Create Subnet** window. For more information, see [Creating a Subnet](#) on page 93.

5. Click **Update** to ensure that the updates are saved in the configuration.

Category Management

A category is a key-value pair that groups similar entities. Associating a policy with a category ensures that the policy applies to all the entities in the group regardless of how the group scales with time. For example, you can associate a group of VMs with the Department: Marketing category, where Department is a category that includes a value Marketing along with other values such as Engineering and Sales.

Currently, you can associate only VMs with a category. Categories are implemented in the same way on on-premises Prism Central instances. For information on configuring categories, see the [Prism Central Infrastructure Guide](#).

Updating a Policy

About this task

Perform the following steps to update a policy.

Note: You cannot update or delete the default policy.

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Network & Security > Virtual Private Clouds** from the **Navigation Bar**.
The **Virtual Private Clouds** page opens displaying the **List** tab.
3. Click the name of the VPC for which you want to update the policy.
The **Summary** tab opens displaying the detailed information about the VPC in widgets.
4. Click the **Policies** tab.

5. Select the checkbox associated with the policy you want to update, and click **Update** from the **Actions** dropdown menu.
The **Update Policy** window opens.
6. Update the necessary values in the respective fields.
The fields in the **Update Policy** window is identical to the fields in the **Create Policy** window. For more information, see [Creating a Policy](#) on page 97.
7. Click **Update**.

Updating Static Routes

About this task

Perform the following steps to update a static route.

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Network & Security > Virtual Private Clouds** from the **Navigation Bar**.
The **Virtual Private Clouds** page opens displaying the **List** tab.
3. Click the name of the VPC for which you want to update a static route.
The **Summary** tab opens displaying the detailed information about the VPC in widgets.
4. Click the **Routes** tab.
5. Click **Manage Static Routes**.
The **Manage Static Routes** window opens.
6. Update the necessary values in the respective fields.

Note: You must configure the default route (0.0.0.0/0) to the external subnet as the next hop for connectivity outside the cluster (north-south connectivity).

For details about the fields that you can update, see [Creating Static Routes](#) on page 102.

7. Click **Save**.

Deleting a Virtual Private Cloud

About this task

Perform the following steps to delete a VPC.

Important: Prism Central does not allow you to delete a VPC if the VPC is associated with any subnets and/or VPNs. You can delete the VPC after you remove all the subnets or VPN associations from the VPC.

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Network & Security > Virtual Private Clouds** from the **Navigation Bar**.
The **Virtual Private Clouds** page opens displaying the **List** tab.

3. Select the checkbox associated with the VPC you want to delete, and click **Delete** from the **Actions** dropdown menu.
4. In the confirmation dialog box, click **Delete** to delete the VPC.
Click **Cancel** to exit without deleting the VPC.

Deleting Subnets, Policies or Routes

You can delete VPC entities such as subnets, policies or routes from the VPC details page.

About this task

Perform the following steps to delete VPC entities such as subnets, policies or routes.

Note: You cannot update or delete the default policy.

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Network & Security > Virtual Private Clouds** from the **Navigation Bar**.
The **Virtual Private Clouds** page opens displaying the **List** tab.
3. Click the name of the VPC for which you want to delete an entity.
The **Summary** tab opens displaying the detailed information about the VPC in widgets.
4. Navigate to the respective tab like **Subnets**, **Policies** or **Routes**.
5. Select the checkbox associated with the entity you want to delete, and click **Delete** from the **Actions** dropdown menu.
6. In the confirmation dialog box, click **Delete** to delete the entity.
Click **Cancel** to exit without deleting the entity.

CONNECTIONS MANAGEMENT

This section covers the management of network gateways, VPN connections and subnet extensions including operations like create, update and delete network gateways and VPN connections, and extending subnets.

Note:

You can enable network segmentation on a Layer 2 Network Extension (or extended subnet) that does not have a gateway. For more information on Layer 2 Network Extensions, see [Layer 2 Network Extension](#) on page 124. For more information, see *Segmenting a Stretched L2 Network for Disaster Recovery* in the [Securing Traffic through Network Segmentation](#) topic of the *Nutanix Security Guide*.

The Layer 2 Network Extension is also known as Layer 2 Stretch.

- For information on network gateways and their management, see [Network Gateway Management](#) on page 108.
- For information on virtual private network connections, see [Virtual Private Network Connections](#) on page 116.
- For information on Layer 2 Network Extensions, see [Layer 2 Network Extension](#) on page 124.
- For information on Border Gateway Protocol (BGP) sessions, see [Border Gateway Protocol Sessions](#) on page 143.

Network Gateway Management

You can create, update or delete network gateways that use VPN, VTEP or BGP service for connections.

Warning:

- Ensure that you add static routes to the NAT network for Prism Central, Network Time Protocol (NTP), Domain Name System (DNS), and other peer VPN, VTEP, or BGP IP addresses, when you deploy network gateways for VPN, VTEP, or BGP in a VPC with the following conditions:
 - The network gateways are connected to both NAT and no-NAT external networks.
 - The no-NAT network is the default next hop.

Without these static IP configurations, the peer gateways on the assigned floating IP addresses cannot reach the network gateways and their status is displayed as Down in Prism Central.

- Connectivity to NTP servers at time.google.com and DNS at 8.8.8.8 is mandatory for the network gateway VM to become active. If you do not have access to these resources, the status of the network gateway is displayed as Down. If you cannot open access to these services on the Internet, contact Nutanix support to change the DNS and NTP server configuration of the network gateway VM.

Creating a Network Gateway

About this task

A network gateway connects two networks together, and can be used in both VLAN and VPC networks on AHV. In other words, you can extend the routing domain of a VLAN network or that of a VPC using a connection between two gateways, one local and one remote. A network gateway pair (local and remote) may host one service such as VPN, VXLAN or BGP service that provides connectivity between the local and remote networks.

Note: You can create one network gateway with only one service such as VPN, VXLAN or BGP. The same network gateway cannot host two services at the same time. Once you create a network gateway with one service, you cannot change the service. For example, if you create a network gateway with BGP service, you cannot change it to VPN service after the network gateway is created.

You can create multiple network gateways for a VPC. Since a VPC is configured only on a Prism Central, the VPC is available to all the clusters registered to that Prism Central.

Note:

A best practice is to configure the remote gateway before you configure the local gateway especially when the gateway configuration involves entering unique parameters like eBGP ASNs in the local and remote gateways.

There are two parts in configuring a gateway (local or remote):

- (Local only) Gateway VM which the network gateway appliance deploys when you create the local network gateway.
- Service Configuration where you configure the service that you want the (local and remote) gateway to use, like **VPN** service, **VTEP** (VXLAN) service or **BGP** service.

Perform the following steps to create a VPN, VTEP or BGP service gateway.

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Network & Security > Connectivity** from the **Navigation Bar**.
The **Gateways** page opens displaying the list of network gateways you have created and configured, and the operations you can perform on the network gateways.
3. Select **Local** or **Remote** in the **Create Gateway** dropdown menu.
If you select **Local** in the dropdown menu, the **Create Local Gateway** window opens. If you select **Remote** in the dropdown menu, the **Create Remote Gateway** window opens.
4. Provide the necessary values in the respective fields as described in the table.
For example, if you select **Local** in the dropdown menu, then the **Create Local Gateway** page displays the **VM Deployment** tab. Provide the necessary values in the respective fields as described in the table, in the the **VM Deployment** and **Service Configuration** tabs.

Table 26: Local Gateway Configuration

Fields	Description	Values
VM Deployment		
Name	Enter a name for the network gateway.	(Name)

Fields	Description	Values
Gateway Attachments	(for Local gateway type only) Select the gateway attachment as VPC or VLAN. The VPN VM is deployed on a VPC VM or a cluster that has the selected VLAN respectively. 1. If you select VPC, then <i>VPC Attachment</i> is displayed. VPC is the default value for the <i>Gateway Attachments</i> field. The Gateway VM is deployed on the cluster and associated with the VPC selected in the <i>VPC Attachment</i> section. <i>VPC attachment</i> mode provides the options of eBGP and Static routing methods for external routing (configured in the External Routing Configuration section). 2. If you select VLAN, then the <i>VLAN Attachment</i> is displayed. The Gateway VM is deployed on the cluster that has the VLAN and the subnet specified in the <i>VLAN Attachment</i> section. <i>VLAN attachment</i> mode provides only the eBGP routing method for external routing.	(VLAN or VPC)
Gateway VM Deployment - VPC Attachment		
Cluster	Select the cluster on which you want to deploy the Gateway VM on.	(Name of the cluster)
VPC (If Gateway Attachment type is VPC)	Select the VPC configured on the selected cluster that you want to use for the Gateway VM deployment.	(Name of the VPC selected)
Floating IP (Optional)	Select a floating IP for the network gateway configuration. If you do not select a floating IP address then Prism Central allocates a floating IP automatically. This allocated floating IP is deleted when you delete the gateway. To request floating IPs and allocate them to subnets, see Requesting Floating IPs on page 92	(IP address)
Gateway VM Deployment - VLAN Attachment		
Cluster	Select the Cluster, from the drop down list, on which you want to deploy the Gateway VM on.	(Name of the cluster)
	Note: Only clusters with VLANs are available in the list.	

Fields	Description	Values
Subnet	Select the subnet you want to attach the Gateway VM to, from the drop down list. Note: The list includes all the subnets you created on the selected cluster. After you select the subnet, the details of the subnet are displayed in a box below the Subnet field. The details include: VLAN ID, IPAM type being Managed or Unmanaged, and Network Address with Prefix.	(Name of the VLAN subnet)
Static IP Address for VPN Gateway VM	Enter the static IP address that the Gateway VM needs to use.	(IP Address with Prefix)
Default Gateway IP	Enter the default gateway IP of the subnet for the Gateway VM.	(IP Address)
Service Configuration		
Gateway Service	Select the gateway service you want to use for the gateway.	(VPN or VTEP)
VPN Service Configuration - External Routing Configuration (This section is available for VLAN and VPC attachment types)		
Routing Protocol	1. For VPC gateway attachments: Select Static for static routing. Note: You need to create static routes for external routing and attach the route to the VPC selected in this configuration. For more information, see Creating Static Routes on page 102. 2. Select eBGP for eBGP based external routing. 3. For VLAN gateway attachments: External routing protocol is pre-set to eBGP. You cannot change the routing protocol.	(Static or eBGP)
Redistribute Connected Routes (Applicable only if VLAN type gateway attachment is selected)	(VLAN only) Select this checkbox to enable the redistribution of connected routes into the eBGP.	(Check mark or blank)
ASN (Only available if eBGP routing protocol is selected)	(For eBGP only) Enter the ASN for your on-premises gateway. If you do not have a BGP environment in your on-premises site, you can choose any number. For example, you can choose a number in the 65000 range. Note: Make sure that this ASN does not conflict with any of the other on-premises BGP ASNs. ASN must be distinct in case of eBGP.	(Number)

Fields	Description	Values
eBGP Password	(For eBGP in Local gateway type only) Enter the eBGP password for the eBGP route.	Password: The password must be between 1 and 80 characters. - Characters allowed for Pre-Shared Key for IPSec - a-z - A-Z 0-9 ~ ! @ # % ^ & * () _ - + = : ; { } [] < > , . / ? \$ - Password length: Minimum 1 and maximum 64 characters. - Characters allowed for BGP passwords - a-z - A-Z 0-9 ~ ! @ # % ^ & * () _ - + = : ; { } [] < > , . / ? \$ - Password length: Minimum 1 and maximum 80 characters.
VPN Service Configuration - Internal Routing Configuration (This section is available for VLAN attachment type only.)		

Fields	Description	Values
Routing Protocol (Between On-prem Gateway and On-prem Router)	Select the Routing Protocol to be used between on-premises Nutanix gateway and on-premises router. You can select: • Static: Select this protocol to provide a static route configuration for the VLAN gateway. • OSPF: Select this protocol to provide an OSPF routing configuration for the VLAN gateway. • iBGP: Select this protocol to provide a iBGP route configuration for the VLAN gateway. Note: For iBGP, the ASN must be the same between the Gateway appliance and the peer iBGP, when iBGP is selected as the internal routing protocol.	(Static or OSPF or iBGP)
+Add Prefix (Applicable to Static routing)	(For Static routing selected in Routing Protocol) Click this to enter a Local Prefix and click the check mark under Actions to add the prefix. If you click the X mark under Actions, the local prefix you entered is not added. The prefixes you add are advertised to all the connected peers via eBGP. The prefix must be a valid IP address with the host bits not set. You can add multiple local prefix IP addresses.	(prefix like /24)
Area ID (Applicable to OSPF protocol)	(OSPF only) Enter the OSPF area ID in the IPv4 address format.	(IPv4 address format)
Password Type	(OSPF only) Select the password type you want to set for the OSPF route. The options are: 1. MD5: Select this option to encrypt the packets with MD5 hash that can be decrypted with the MD5 password at the destination. 2. Plain Text: Select this option to set a clear-text password. 3. None: Select this if you do to set an open route without password protection	(Password)

Fields	Description	Values
Password	(OSPF only) Enter a password for the MD5 or Plain Text password type you select in the <i>Password Type</i> field. - For MD5: The password must be 1-16 characters long. Characters allowed for OSPF passwords (MD5) - a-z - A-Z 0-9 - For Plain Text: The password must be 1-8 characters long. Characters allowed for OSPF passwords (Plain text): a-z.	
Peer IP (for iBGP)	Enter the IP Address of the On-prem router used to exchange routes with the network gateway.	(IP Address)
Password	Enter a password with 1-80 characters.	(Password)
VTEP Service Configurations		
VxLAN (UDP) Port	The default value provided is 4789. Do not change this.	(Number. Default value is 4789)
BGP Service Configurations		
Serviced VPC	Select the VPC that you want to connect using the local BGP gateway.	VPC
eBGP ASN	Enter the ASN for your local gateway. You can choose any number. For example, you can choose a number in the 1-65000 range. Note: Make sure that this ASN does not conflict with any of the other local or remote BGP ASNs. Once you enter the ASN, you cannot change the ASN using the Update Gateway page.	(Number)

Table 27: Remote Gateway Configuration

Fields	Description	Values
Name	Enter a name for the network gateway.	(Name)
Gateway Service	Select the gateway service you want to use for the gateway.	(VPN or VTEP)
VPN Service Configurations		
Public IP Address	Enter the public IP address of the remote endpoint.	(IP Address)
Vendor	Select the vendor of the third party gateway appliance.	(Name of Vendor)

Fields	Description	Values
External Routing		
Protocol	1. Select Static for static routing. Note: You need to create static routes for external routing and attach the route to the VPC selected in this configuration. For more information, see Creating Static Routes on page 102. 2. Select eBGP for eBGP based external routing.	(Static or eBGP)
eBGP ASN (Only available if eBGP routing protocol is selected)	(For eBGP only) Enter the ASN for your on-premises gateway. If you do not have a BGP environment in your on-premises site, you can choose any number. For example, you can choose a number in the 1-65000 range. Note: Make sure that this ASN does not conflict with any of the other on-premises BGP ASNs. ASN must be distinct in case of eBGP.	(Number)
VTEP Service Configurations		
VTEP IP Address	Enter VTEP IP Addresses of the remote endpoints that you want to create the gateway for. You can add IP addresses of multiple endpoints in one remote gateway.	(Comma separated list of IP Addresses)
VxLAN (UDP) Port	The default value provided is 4789. Do not change this.	(Number. Default value is 4789)
BGP Service Configurations		
Service IP Address	Enter the IP Address of the remote endpoints that you want to create the gateway for.	(IP address)
eBGP ASN	Enter the ASN for the remote gateway. You can choose any number. For example, you can choose a number in the 1-65000 range. Note: Make sure that this ASN does not conflict with any of the other on-premises BGP ASNs. You can modify the ASN using the Update Gateway page.	(Number)

5. Click **Create**.

The gateways you create are displayed in the **Gateways** page.

Updating a Network Gateway

You can update a network gateway using the **Update Gateway** window.

About this task

Perform the following steps to update a gateway.

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Network & Security > Connectivity** from the **Navigation Bar**.
The **Gateways** page opens displaying the list of network gateways you have created and configured, and the operations you can perform on the network gateways.
3. Select the checkbox associated with the gateway you want to update, and click **Update** from the **Actions** dropdown menu.
The **Update Gateway** window opens.
4. Update the necessary values in the respective fields.
The fields in the **Update Gateway** window is identical to the fields in the **Create Gateway** window. For more information, see [Creating a Network Gateway](#) on page 108.

Note: You cannot modify some parameters. Such parameters are greyed and in-actionable. If you need to modify such parameters, consider creating a new gateway with the appropriate parameters and deleting the current gateway.

5. Click **Save**.

Deleting a Network Gateway

About this task

Perform the following steps to delete a gateway.

Important: You must first delete all the VPN or VTEP connections, BGP sessions or subnet extensions associated with the gateway to be able to delete a network gateway.

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Network & Security > Connectivity** from the **Navigation Bar**.
The **Gateways** page opens displaying the list of network gateways you have created and configured, and the operations you can perform on the network gateways.
3. Select the checkbox associated with the gateway you want to delete, and click **Delete** from the **Actions** dropdown menu.
4. In the confirmation dialog box, click **Delete** to delete the gateway.
Click **Cancel** to exit without deleting the gateway.

Virtual Private Network Connections

Virtual Private Network

You can use the Nutanix VPN solution to set up VPN between your on-premises clusters, which exist in distinct routing domains that are not directly connected. These distinct routing domains could either be VPCs within the same cluster or remote clusters or sites.

If you need to connect one Nutanix deployment in one site to another deployment in a different site, you can create a VPN endpoint in each of the sites. A VPN endpoint consists of a local VPN gateway, remote VPN gateway and VPN connection. Local VPN gateway can be instantiated in a VPC context or a legacy VLAN context. Launching the VPN

gateway within a VPC allows stretching of the VPC. For example, in the figure, the Blue VPC is stretched between two sites with a VPN.

VPN connections are useful in connecting two points. You can connect two VPCs in the same cluster using a VPN or VPCs in different clusters in the same site. However, VPN connection can connect only one endpoint to another endpoint. Flow virtual networking based VPN service allows you to only connect two endpoints that use Nutanix VPN based gateway service.

Virtual Tunnel End Points Based Network Extensions

To connect one endpoint to multiple endpoints or third party (non Nutanix) networks, use VXLAN (Virtual Extensible LAN) based Virtual Tunnel End Point (VTEP) service based subnet extensions. For more information, see [Layer 2 Network Extension Over VTEP](#) on page 132.

Virtual Network Connection Using BGP

Border Gateway Protocol (BGP) works in Layer 4 (application layer). It works on top of TCP at layer 2. Flow Virtual Networking allows you to create and use BGP gateways and connections at layer 3 to connect two clusters for purposes including disaster recovery.

VPN Workflow

If you need to connect one Nutanix deployment in one site to another deployment in a different site, you can create a VPN endpoint in each of the sites. A VPN endpoint consists of a local VPN gateway, remote VPN gateway and VPN connection. You can configure multiple VPN endpoints for a site.

Each endpoint must have configurations for a local VPN gateway, remote VPN gateway (pointer information for the peer local VPN in the remote site endpoint) and a VPN connection (connecting the two endpoints). Then, based on the VPN connection configuration as initiator or acceptor, one endpoint initiates a tunnel and the endpoint at the other end accepts the tunnel connection and, thus, establishes the VPN tunnel.

1. Gateways: Every VPN endpoint for each site consists of two VPN gateway configurations - Local and Remote.

Local gateway is a VM that runs the VPN protocols (IKEv2, IPSec) and routing (BGP and OSPF). Remote gateway is a pointer - database entry - that provides information about the peer remote VPN endpoint. One of the key information contained in the remote gateway is the source IP of the remote VPN endpoint. For security reasons, the local VPN gateway will accept IKEv2 packets originating only from this Source IP.

VPN gateways are of the following types:

- On premises Nutanix VPN Gateway: Represents the VPN gateway appliance at your on-premises local or remote site if you are using the Nutanix VPN solution.
- On premises Third Party Gateway: Represents the VPN gateway appliance at your on-premises site if you are using your own VPN solution (provided by a third-party vendor).

To configure third party VPN Gateways, see the relevant third party documentation.

2. VPN Connection: Represents the VPN IPSec tunnel established between local gateway and remote gateway. When you create a VPN connection, you need to select two gateways between which you want to create the VPN connection.

VPN appliances perform the following:

1. Implementation of IKEv2 and IPSec protocols.
2. Routing: Between remote sites, Flow virtual networking advertises prefixes using eBGP. Optionally it uses Static routing. Within a site, Flow virtual networking uses iBGP or OSPF to share prefixes between the Nutanix VPN appliance and the edge router.

IPSec Configuration Parameters

Nutanix supports standard encryption, authentication algorithms and DH groups.

Encryption Algorithms

- AES128
- AES256
- 3DES
- AES256GCM128

Authentication Algorithms

- MD5
- SHA1
- SHA256
- SHA384
- SHA512

DH Groups

- 14 = 2048-bit MODP group
- 19 = 256-bit random ECP group
- 20 = 384-bit random ECP group

Prerequisites for VPN Configurations

General Requirements

- Ensure that you have enabled Flow virtual networking with microservices Infrastructure.
- Ensure that you have floating IP addresses when you create VPN gateways.

Flow virtual networking automatically allocates a floating IP to a VPN gateway if you do not provide one during the VPN gateway creation. To provide floating IP during the VPN gateway creation, you can request floating IPs. For more information, see [Requesting Floating IPs](#) on page 92.

- Ensure that you have one of the following, depending on whether you are using iBGP or OSPF:
 - Peer IP (for iBGP): The IP address of the router to exchange routes with the VPN gateway VM.
 - Area ID (for OSPF): The OSPF area ID for the VPN gateway in the IP address format.
- Nutanix recommends setting the guest VM NIC MTU to 1,356 bytes for all VMs inside a VPC that send traffic over Nutanix VPN connections. This prevents fragmentation and accounts for the encapsulation overhead for VPN connections in a VPC. For more information, see the [Flow Virtual Networking MTUs](#) table.

Accounting for the 1356 byte MTU: Assuming a 1,500 byte network MTU, subtract 58 bytes for Geneve VPC encapsulation and 86 bytes for IPsec encapsulation, leaving 1,356 bytes for guest VM frames.

- Ensure that you have the following details for the deployment of the VPN gateway VM:
 - Public IP address of the VPN Gateway Device: A public WAN IP address that you want the on-premises gateway to use to communicate with the Xi VPN gateway appliance.
 - Static IP Address: A static IP address that you want to allocate to the VPN gateway VM. Use a floating IP address requested as the static IP address.
 - IP Prefix Length: The subnet mask in CIDR format of the subnet on which you want to install the VPN gateway VM. You can use an overlay subnet used for a VPC and assigned to the VM that you are using for the VPN gateway.
 - Default Gateway IP: The gateway IP address for the on-premise VPN gateway appliance.
 - Gateway ASN: ASN must not be the same as any of your on-premises BGP ASNs. If you already have a BGP environment in your on-premises site, the customer gateway is the ASN for your organization. If you do not have a BGP environment in your on-premises site, you can choose any number. For example, you can choose a number in the 0-65000 range.

Ports and Protocols

Nutanix deploys a number of ports and protocols in its software. These ports must be open in the firewalls to enable Flow Virtual Networking to function. For information on the ports and protocols used for Flow Virtual Networking, see [Ports and Protocols](#).

Endpoints and Terminations

The following endpoints and terminations occur in the course of Flow virtual networking based connections. For information on creating, updating or deleting VPN connections, see [Connections Management](#) on page 108.

Note: In a VPN connection do not configure both the gateways (local gateway and remote gateway) in an endpoint as Initiators or as Acceptors. If you configure the local gateway as Initiator then configure the remote gateway as Acceptor in one endpoint and vice-versa in the (other) remote endpoint.

VPN Endpoint Behind a Network Address Translation or Firewall Device

In this scenario, the IPSec tunnel terminates behind a network address translation (NAT) or firewall device. For NAT to work, open UDP ports 500 and 4500 in both directions.




Figure 22: VPN Endpoint Behind NAT or Firewall

Things to do in NAT

Open UDP ports 500 and 4500 on both directions

Things to do in on-premises VPN GW

Enable the business application policies to Allow the commonly-used business application ports.

IPSec Terminates on the Firewall Device

In this scenario, you do not need to open the ports for NAT (500 and 4500).

However, enable the on-premises VPN gateway to allow the traffic from the PC subnet to the advertised load balancer route where the Source port is any and the Destination port may be in the range of 1024-1034.

The PC subnet refers to the subnet where your Prism Central is running.




Figure 23: Tunnel Terminates on NAT or Firewall

Creating a VPN Connection

About this task

Create a VPN connection to establish a VPN IPSec tunnel between VPN gateways in your on-premises site. Select the gateways between which you want to create the VPN connection.

Perform the following steps to create a VPN connection.

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Network & Security > Connectivity** from the **Navigation Bar**.
The **Gateways** page opens displaying the list of network gateways you have created and configured, and the operations you can perform on the network gateways.
3. Click the **VPN Connections** tab.
The **VPN Connections** page opens displaying the list of VPN connections created for the clusters.
4. Click the **Create VPN Connection**.
5. In the **Create VPN Connection** window, provide the values in the respective fields.

Fields	Description and Values
Name	Enter a name for the connection.
VPN Connection	
IPSec Secret	Enter a secret password for the IPSec connection. To see the password, click Show . To hide the password, click Hide .
Local Gateway	Select the connection parameters on the local gateway as Initiator or Acceptor of VPN Tunnel connections.
VPN Gateway	Select the appropriate VPN Gateway as the local gateway for the VPN connection
VTI Prefix - Local Gateway	Enter a IPv4 Address with /<prefix>. Example: 10.25.25.2/30. This is the <i>VPN Tunnel Interface</i> IP address with prefix for the local gateway. The subnet for this IP address must be a /30 subnet with two usable IP addresses. One of the IP addresses is used for Local Gateway. Use the other IP address for the Remote Gateway.
Connection Handshake	This defines the type of handshake that the connection must use. There are two types of connection handshakes: 1. Initiator: The local VPN gateway acts as the initiator of the connection and thus initializes the VPN tunnel. 2. Acceptor: The local VPN gateway accepts or rejects incoming connection requests from other gateways. Note: In a VPN connection do not configure both the gateways (local gateway and remote gateway) in an endpoint as Initiators or as Acceptors. If you configure the local gateway as Initiator then configure the remote gateway as Acceptor in one endpoint and vice-versa in the (other) remote endpoint.
Remote Gateway	For a specific VPN connection, set the remote gateway as Initiator or Acceptor when you configure the VPN connection on the Remote Gateway.
VPN Gateway	Select the appropriate VPN Gateway as the remote gateway for the VPN connection.
VTI Prefix - Remote Gateway	The VPN Tunnel Interface IP address with prefix for the local gateway. Provide a IPv4 Address with /<prefix>. Example: 10.25.25.2/30. This is the <i>VPN Tunnel Interface</i> IP address with prefix for the local gateway. The subnet for this IP address must be a /30 subnet with two usable IP addresses. One of the IP addresses is used for Local Gateway. Use the other IP address for the Remote Gateway.
Advanced Settings	Set the traffic route priority for the VPN connection. The route priority uses Dynamic route priority because the priority is dependent on the routing protocol configured in the VPN gateway.
Route Priority - Dynamic Route Priority	Set the route priority as an integer number. The greater the number, higher is the priority.

6. Click **Save**.

The VPN connection you create is displayed in the **VPN Connections** page.

What to do next

The VPN connection you create is displayed in the **VPN Connections** page. Optionally, create static routes from the VPCs to the VPN connection. For information on static routes, see *What to do next* section in [VPN Connection within Same Prism Central](#) on page 122 for information.

Updating VPN Connection

About this task

Perform the following steps to update a VPN connection.

You can open the **Update VPN Connection** dialog box. The parameters in the **Update VPN Connection** dialog box are the same as those in the **Create VPN Connection** dialog box.

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Network & Security** > **Connectivity** from the **Navigation Bar**.
The **Gateways** page opens displaying the list of network gateways you have created and configured, and the operations you can perform on the network gateways.
3. Click the **VPN Connections** tab.
The **VPN Connections** page opens displaying the list of VPN connections created for the clusters.
4. Select the VPN Connection you want to update, and click **Update** from the **Actions** dropdown menu.
The **Update VPN Connection** window opens.
5. Update the necessary values in the respective fields.
The fields in the **Update VPN Connection** window is identical to the fields in the **Create VPN Connection** window. For more information, see [Creating a VPN Connection](#) on page 120.
6. Click **Save**.

Deleting a VPN Connection

About this task

Perform the following steps to delete a VPN connection.

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Network & Security** > **Connectivity** from the **Navigation Bar**.
The **Gateways** page opens displaying the list of network gateways you have created and configured, and the operations you can perform on the network gateways.
3. Click the **VPN Connections** tab.
The **VPN Connections** page opens displaying the list of VPN connections created for the clusters.
4. Select the VPN Connection you want to delete, and click **Delete** from the **Actions** dropdown menu.
5. In the confirmation dialog box, click **Delete** to delete the VPN connection.
Click **Cancel** to exit without deleting the connection.

VPN Connection within Same Prism Central

You can connect two VPCs within the same Prism Central availability zone using a VPN connection.

About this task

Assume that you have created two VPCs named **vpc-a** and **vpc-b** with overlay subnets named **subnet-a** and **subnet-b**.

To connect the two VPCs within the same Prism Central using a VPN connection, do the following.

Procedure

1. Do the following for local gateways:

- a. Create a local VPN gateway with dynamically assigned address for **vpc-a**, for example, named **local-vpn-a**. Note or write down the assigned IP address.
- b. Create a local VPN gateway with dynamically assigned address for **vpc-b**, for example, named **local-vpn-b**. Note or write down the assigned IP address.

For more information on creating a VPN gateway, see [Creating a Network Gateway](#) on page 108.

2. Do the following for remote gateways:

- a. Create a remote VPN gateway with the IP address noted in [1.a](#) on page 123 for **vpc-a**, for example, named **remote-vpn-a**.
- b. create a local VPN gateway with the IP address noted in [1.b](#) on page 123 for **vpc-b**, for example, named **remote-vpn-b**.

For more information on creating a VPN gateway, see [Creating a Network Gateway](#) on page 108.

3. Create a VPN connection between **vpc-a** and **vpc-b** named, for example, **vpn-conn-a-to-b**.

Ensure that the VTI IP addresses for the local and remote gateways is unique with /30 prefix.

Note: The *VPN Tunnel Interface* IP address with prefix for the local gateway. The subnet for this IP address must be a /30 subnet with two usable IP addresses. One of the IP addresses is used for Local Gateway. Use the other IP address for the Remote Gateway.

Ensure that you select **local-vpn-a** as the local gateway with **Connection Handshake** set as **Acceptor**.

Ensure that you select **remote-vpn-b** as the remote gateway.

4. Create a VPN connection between **vpc-b** and **vpc-a** named, for example, **vpn-conn-b-to-a**.

Ensure that the VTI IP addresses with /30 prefix for local and remote gateways are the reverse (vice versa) of what you configured for the VPN connection in previous step. For example, if in previous step you configured the VTI IP addresses as 10.20.20.5/30 for local and 10.20.20.6/30 for remote then for VPN connection in this step, configure 10.20.20.6/30 for local gateway and 10.20.20.5/30 for remote gateway respectively. These IP addresses do not need to be reachable anywhere else in the network. However, ensure that these IP addresses do not overlap with any other IP addresses assigned in the network.

Ensure that you select **local-vpn-b** as the local gateway with **Connection Handshake** set as **Initiator**.

Ensure that you select **remote-vpn-a** as the remote gateway.

What to do next

Optionally, create static routes for the subnets in the two VPCs to the VPN connections. The static routes ensure that the subnets communicate with the VPN connection.

For example,

- Create static routes in **vpc-a** with **Destination Prefix:** **subnet-b** (in **vpc-b**), **Next Hop:** **vpn-conn-a-to-b**
- Create static routes in **vpc-b** with **Destination Prefix:** **subnet-a** (in **vpc-a**), **Next Hop:** **vpn-conn-a-to-b**

For information on creating or updating static routes, see [Updating Static Routes](#) on page 106.

Layer 2 Network Extension

You can extend a subnet between on-premises local and remote clusters or sites (Availability Zones or AZs).

Layer 2 Network Extensions support seamless application migration between clusters or sites.

Note: One or more on-premises cluster or sites managed by one Prism Central instance is defined as an Availability Zone or AZ. In this section, Availability Zone or AZ refers to and must be understood as one or more on-premises clusters or sites managed by one Prism Central. Local AZ refers to local on-premises clusters or sites managed by a Prism Central instance and remote AZ refers to another on-premises cluster or site managed by another Prism Central instance.

With Layer 2 Network Extension, you can migrate a set of applications to the remote AZ while retaining their network bindings such as IP address, MAC address, and default gateway. With the subnet extension mechanism, VMs can communicate over the same broadcast domain eliminating the need to architect the network topology again, which can otherwise result in downtime.

Layer 2 Network Extension assumes that there are underlying existing layer 3 connectivity already available between the Availability Zones. You can extend a subnet from a remote AZ to the primary (Local) AZ (and other remote AZs in case of VTEP-based subnet extensions)

- You can extend a Layer 2 subnet across two Nutanix AZs over either VPN or Virtual tunnel End Point (VTEP). For more information, see [Layer 2 Network Extension Over VPN](#) on page 127.
- You can extend a Layer 2 subnet between a Nutanix AZ and one or more non-Nutanix datacenters only over VTEP. For more information, see [Layer 2 Network Extension Over VTEP](#) on page 132.

You can extend subnets for the following configurations.

- **IPAM Type.** Managed and unmanaged networks.
- **Subnet Type.** On-prem VLAN subnets and VPC subnets.
- **Traffic Type.** IPv4 unicast traffic and ARP.

- **On-prem Hypervisor.** AHV and ESXi

Note: If your cluster is ESXi, use vCenter Server to manually configure the port group attached to the subnet you need to extend. Set the security settings, **Promiscuous mode** and **Forged transmits** to **Accept** on the vSwitch.

Figure 24: ESXi Host Port Group Configuration

Prerequisites for Setting Up Subnet Extension

Ensure the following before you configure Layer 2 Network Extension between your on-premises AZs.

- Ensure that the Prism Central version supports Layer 2 Network Extension. For more information, see *Features in Flow Virtual Networking* in [Release Notes | Flow Virtual Networking](#) as applicable.

For instructions on how to upgrade a Prism Central instance through the Prism Central web console, see [Prism Central Upgrade and Installation](#) in *Prism Central Infrastructure Guide*.

- Ensure that you pair the Prism Central at the local AZ with the Prism Central at the remote AZ to use **Create Subnet Extension** wizard to extend a subnet across the AZs and facilitate bidirectional communication between these clusters or sites. Using paired availability zones it is possible to configure both VXLAN over VPN and VTEP based subnet extension. You can also extend subnets using the manual gateway and connection workflows instead of pairing the AZs.

For instructions about how to pair the local and remote AZs, see [Pairing Availability Zones](#) on page 128.

- Ensure that you set up a default static route with `0.0.0.0/0` prefix and the external network next hop for the VPC to use for any subnet extension. The static route provides NTP and DNS access for the Network Gateway appliance.

Best Practices for Subnet Extension

Nutanix recommends the following configurations to allow IP address retention for VMs on extended subnets.

- When using Nutanix IPAM ensure the address ranges in the paired subnets are unique to avoid conflict between VM IP addresses across extended subnets.
- Consider an L2 extension from a Nutanix environment to a non-Nutanix environment with the IPAM-managed VMs in the Nutanix environment. Suppose you move VMs from the Nutanix environment to the non-Nutanix environment. In that case, the Nutanix VMs are powered off, and when you try to use the VMs in the non-Nutanix environment, connectivity issues occur. The connectivity issues occur because the Network Controller continues to respond to the Address Resolution Protocol (ARP) messages of the powered-off VM. This is the default

behavior of the Network Controller. When you move the VMs from the Nutanix environment, delete the VM NICs in the Nutanix environment if you plan to retain the VMs, to prevent the ARP responses from occurring.

- If the source and target sites use third-party IPAM, ensure that there are no conflicting IP address assignments across the two sites.

Note: If the source and target sites use Nutanix IPAM, the Prism Central web console displays a message that indicates an IP address conflict if one exists.

- If connectivity between sites already provides encryption, consider using VTEP only subnet extension to reduce encryption overhead.
- Use the **Subnet Extension to a Third Party Data-Center** workflow in the following scenarios
 - To extend a subnet to more than one other AZ. This is also known as point to multipoint.
 - To extend subnets between clusters managed by the same Prism Central.
- To avoid tromboning or hair-pinning of traffic, provide valid gateway IP address for the local and remote sides of the subnet extension. If you need to route the traffic only from one side (local or remote, thus causing traffic tromboning or hair-pinning to that side) of the subnet extension, then provide a valid gateway IP address only on that side. See for more information.

Subnet Extension Workflow

You can manage Layer 2 Network Extension on the **Subnet Extensions** tab of the **Connectivity** page.

About this task

To manage Layer 2 Network Extensions, follow these steps:

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Network & Security** > **Connectivity** from the **Navigation Bar**.

The **Gateways** page opens displaying the list of network gateways.

3. Click the **Subnet Extensions** tab.

The **Subnet Extensions** page opens displaying the list of subnet extensions created for the clusters.

- » Create point-to-point Layer 2 Network Extensions between two AZs over VPN or VTEP by opening the **Create Subnet Extension Across Availability Zones** window. For more information, see [Extending a Layer 2 Subnet Over VPN](#) on page 129 for VPN-based extensions.

For more information on VTEP-based extensions, see [Extending a Layer 2 Subnet Across Availability Zones Over VTEP](#) on page 133.

- » Create point-to-point or point-to-multipoint Layer 2 Network Extensions to third party datacenters over VTEP by opening the **Create Subnet Extension To A Third Party Data-Center** window.

For more information, see [Extending a Subnet to Third Party Datacenters Over VTEP](#) on page 137.

- » Update a subnet extension that extends across AZs using the **Update Subnet Extension Across Availability Zones** window. The **Update Subnet Extension Across Availability Zones** has the same parameters and fields as the **Create Subnet Extension Across Availability Zones** window.

You can open the **Update Subnet Extension Across Availability Zones** window by following one of these steps:

- » Selecting the subnet extended across AZs in the **Subnet Extensions** and clicking the **Update** button.
- » Clicking the subnet extended across AZs in the **Subnet Extensions** and clicking the **Update** button on the **Summary** tab.

For more information, see [Updating an Extended Subnet](#) on page 142.

Layer 2 Network Extension Over VPN

Subnet extension using VPN allows seamless, secure migration to a new datacenter or for disaster recovery. VPN based Layer 2 Network Extension provides secure point to point connection to migrate workloads between Availability Zones. Consider VTEP-only subnet extension without VPN when encryption is not required.

Layer 2 Network Extension using VPN is useful:

- When the two Availability Zones (where the subnets to be extended belong) do not have any underlying secure connectivity. For example, when connecting over the Internet, VPN (IPSec) provides the necessary connectivity and encryption (security).
- Sometimes when you need to move (lift-and-shift) workloads from a VLAN subnet to a VPC subnet retaining the same VM IP addresses . You need connectivity from other subnets to workloads that have already migrated to VPC. In such cases, VPN provides the Layer 3 connectivity and encryption between the VPC segment of extended subnet to other VLAN subnets.

Prerequisites for Setting Up Subnet Extension Over VPN

- For general prerequisites to extend subnets, see [Layer 2 Network Extension](#) on page 124.
- Set up VPN gateway services and a VPN connection between local AZ and the remote AZ. The subnet extension feature supports only the Nutanix VPN solution (not a third-party VPN solution) at the both the local and remote AZs. For instructions about how to upgrade the VPN gateway VM at the local and remote clusters or sites, see [Virtual Private Network Connections](#) on page 116.

Note: Ensure that the VPN gateway version is 5.0 or higher. For instructions about how to upgrade the network gateway at the local and remote sites, see [Updating a Network Gateway](#).

- Configure subnets with the same IP CIDR prefix at the source and target sites. For example, if the IP prefix at one site is 30.0.0.0/24, the IP prefix at the other site must also be 30.0.0.0/24. The network and mask must match at both AZs.
- Configure distinct DHCP pools for the source and target sites with no IP address overlap. Separate DHCP pools ensure no IP address conflicts occur for dynamically assigned IP addresses between the two AZs.
- Procure two free IP addresses, one from each subnet, for the Network Gateway in the subnets to be extended. These IP addresses are configured as local IP address and remote IP address for the subnet extension in the **Subnet Extension** wizard. These two free IP addresses are the externally accessible IP addresses for the local gateway, and the remote gateway. Those two usable IP addresses are already contained inside the VPN connection and must not conflict with the following:
 - DHCP pools on any of the Availability Zones.
 - Gateway IP address on any of the Availability Zones.
 - IP addresses allocated to existing user VMs on any of the Availability Zones.
 - IP addresses used by Network Gateway Management NIC subnet (IP pool 100.64.1.0/24)

Limitation

To use subnet extension over a VPN, both sites must use the VPN service of the Nutanix Network Gateway. Consider VTEP-only subnet extension to connect to non-Nutanix third party sites.

Pairing Availability Zones

About this task

Note: For DRaaS, pair the on-premises AZ (Prism Central instance) only to Nutanix Cloud AZ. For reverse synchronization, you need not pair again from Nutanix Cloud AZ; Nutanix Cloud AZ captures the pairing configuration from the on-premises AZ that pairs Nutanix Cloud AZ.

To pair an AZ with another AZ or Nutanix Cloud AZ, perform the following procedure at:

- Either of the on-premises AZs for DR solution between on-premises AZs
- Either the on-premises AZ or Nutanix Cloud availability zone AZ for DR solution between on-premises AZ and Nutanix Cloud availability zone.
- Either the on-premises AZ or NC2 AZ for DR solution between on-premises AZ and NC2 AZ.
- Either of the NC2 AZs for DR solution between NC2 AZs

For more information on Prism Central-based DR solution types, see [Prism Central-Based Disaster Recovery Solution](#).

Procedure

1. Perform one of the following:

- » *On-prem to on-prem AZ, on-prem AZ to NC2 AZ, or NC2 AZ to NC2 AZ* – Select the **Infrastructure** application from [Application Switcher Function](#), and go to **Administration > Availability Zones** from the **Navigation Bar**.
- » *Nutanix Cloud AZ to on-prem AZ (DRaaS)* – Click the [Navigation icon](#) to access the **Navigation Bar**, and go to **Administration > Availability Zones**.

The **Availability Zones** page opens, displaying the paired AZs.

2. Click **Connect to Availability Zone**.

Specify the following information in the **Connect to Availability Zone** window.

a. Perform one of the following:

- » *(On-prem to on-prem AZ, on-prem AZ to NC2 AZ, or NC2 AZ to NC2 AZ)* **Availability Zone Type:** Select **Physical Location** from the dropdown menu.
- » *(Nutanix Cloud AZ to on-prem AZ (DRaaS))* **Availability Zone Type:** Select **XI** from the dropdown menu.

b. *(On-prem to on-prem AZ, on-prem AZ to NC2 AZ, or NC2 AZ to NC2 AZ)* **IP Address for Remote PC:** Enter the IP address of Prism Central running on the recovery AZ.

c. Perform one of the following:

- » *(On-prem to on-prem AZ, on-prem AZ to NC2 AZ, or NC2 AZ to NC2 AZ)* **Username:** Enter the username of Prism Central running on the recovery AZ.
- » *(Nutanix Cloud AZ to on-prem AZ (DRaaS))* **Username:** Enter the username of your Nutanix Cloud Services account.

d. Perform one of the following:

- » *(On-prem to on-prem AZ, on-prem AZ to NC2 AZ, or NC2 AZ to NC2 AZ)* **Password:** Enter the password of Prism Central running on the recovery AZ.
- » *(Nutanix Cloud AZ to on-prem AZ (DRaaS))* **Password:** Enter the password of your Nutanix Cloud Services account.

3. Click **Connect**.

Both AZs are paired with each other.

When a paired AZ is unreachable due to service interruption, missing connection, or the expired access tokens on that AZ, the **Connectivity Status** of that AZ shows Not Reachable (see [Availability Zones View](#) in *Nutanix Disaster Recovery Guide*) and the following alert is generated in **Alerts**.

Availability Zone Connection Failure: The remote availability zone AZ_URL is unreachable.

The disaster recovery operations might fail due to the unreachability. To make the paired AZ reachable, unpair the primary AZ with the recovery AZ and then pair it with the recovery AZ again.

Extending a Layer 2 Subnet Over VPN

The Layer 2 Network Extension allows VMs to communicate over the same broadcast domain to a remote site or Availability Zone (AZ).

Before you begin

For information on prerequisites and best practices for extending a Layer 2 subnet, see [Layer 2 Network Extension](#) on page 124 and [Layer 2 Network Extension Over VPN](#) on page 127.

About this task

Perform the following steps to extend a subnet from the on-premises site.

Procedure

1. Log in to Prism Central.

2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Network & Security > Connectivity** from the **Navigation Bar**.
The **Gateways** page opens displaying the list of network gateways you have created and configured, and the operations you can perform on the network gateways.
3. Click the **Subnet Extensions** tab.
The **Subnet Extensions** page opens displaying the list of subnet extensions created for the clusters.
4. Select **Create Subnet Extension > Across Availability Zones** .

5. In the **Create Subnet Extension Across Availability Zones** window, enter the necessary details as described in the following table.

Create Subnet Extension Across Availability Zones

A Subnet Extension allows VMs to communicate securely over the same broadcast domain to a remote Availability Zone. This subnet extension can be done over a VPN or a VTEP.

Extend Subnet over a ☒ VPN ☐ VTEP

• Prism Central on the remote side must be pc.10.10.10.10

• On-Prem VPN should be Nutanix VPN and must be version 5.0 or later

Local

Availability Zone

Local AZ

Subnet Type

☒ VLAN ☐ Overlay

Cluster

s12

Subnet

stvlanvlan0

VLAN ID	IPAM	Network Address/Prefix
25	Managed	10.10.10.10/24

Gateway IP Address/Prefix

10.10.10.1

Local IP Address ?

0.0.0.0

VPN Connection

vpn-connection-site1-vpc-vlan

Remote

Availability Zone

PC_10

Subnet Type

☒ VLAN ☐ Overlay

Cluster

auto_cluster_10.10.10.10

Subnet

stvlanvlan0

VLAN ID	IPAM	Network Address/Prefix
25	Managed	10.10.10.10/24

Gateway IP Address/Prefix

10.10.10.1

Remote IP Address ?

0.0.0.0

VPN Connection

vpn-connection-site2-vlan-vlan

Cancel

Save

Figure 25: Create Subnet Extension Across Availability Zones

NUTANIX

Flow Virtual Networking | Connections Management | 131

Fields	Description	Values
Extend Subnet over a	Select the gateway service you want to use for the subnet extension.	(VPN or VTEP)
Note: Configure the following fields for the Local and the Remote sides of the dialog box.		
Availability Zone	(For Local) Local AZ is pre-selected default. (For Remote) Select the appropriate AZ from the drop-down list of AZs.	(Local: Local AZ) (Remote: Dropdown list of AZs.)
Subnet Type	Select the type of subnet that you want to extend.	(VLAN or Overlay)
Cluster	Displayed if your selected VLAN subnet. Select the cluster from the dropdown list of clusters.	(Name of cluster selected from dropdown list)
VPC	Displayed if your selected Overlay subnet. Select the appropriate VPC from the dropdown list of VPCs.	(Name of VPC selected from dropdown list)
Subnet	Select the subnet that needs to be extended.	(Name of subnet selected from dropdown list)
(Network Information frame)	Displays the details of the VLAN or Overlay network that you selected in the preceding fields.	(Network information)
Gateway IP Address/Prefix	Displays the gateway IP address for the subnet. This field is already populated based on the subnet selected. For more information, see PBR-based Tromboning in L2 Extended Subnet on page 140.	(IP Address)
(Local or Remote) IP Address	Enter a unique and available IP address that are externally accessible IP addresses in Local IP Address and Remote IP Address .	(IP Address)
VPN Connection	Select the appropriate VPN Connection from the dropdown list that Flow virtual networking must use for the subnet extension. For instructions to create VPN connection, see Creating a VPN Connection on page 120.	(Name of VPN connection selected from the dropdown list)

6. Click **Save**.

A successful subnet extension is listed on the **Subnet Extension** page.

Layer 2 Network Extension Over VTEP

Layer 2 Network Extension using Virtual tunnel End Point (VTEP) allows seamless migration to new datacenters or for disaster recovery. VTEP based Layer 2 Network Extension provides point-to-multipoint connections to migrate workloads from one Availability Zone (AZ) to multiple Availability Zones without encryption. If you need security and encryption, consider using Subnet Extension over VPN.

Subnet extension using VTEP is useful:

- When both subnets that need to be stretched are Nutanix subnets (managed or unmanaged). VTEP provides an optimized workflow to stretch the two subnets.
- When both subnets are connected over an existing private and secure link that does not need additional encryption.
- When one Nutanix subnet needs to be stretched across one or more non-Nutanix networks, sites, or datacenters. Subnet Extension with third-party VTEPs provides point-to-multipoint connectivity to third party datacenters assuming that there is underlying layer 3 connectivity between these VTEPs.

VTEP-based Layer 2 Network Extension provides the following advantages:

- Layer 2 Network Extension from one AZ to multiple AZs.
- Layer 2 Network Extension between Nutanix AZs and non-Nutanix third party VTEP-based AZs.
- The Remote VTEP Gateway is a set of endpoint IP addresses. You can add endpoint IP addresses to an existing operational Remote VTEP Gateway without stopping the subnet extension services. This on-the-fly addition enables you to extend the subnets to more AZs than originally planned, or perform maintenance, without disrupting the running services or configuring new remote VTEP gateways.

Prerequisite for Setting Up Subnet Extension Over VTEP

- For general prerequisites to extend subnets, see [Layer 2 Network Extension](#) on page 124 .
- Set up VTEP local and remote gateway services on local and remote AZs. In case of point-to-multipoint extension, ensure that you create local and remote VTEP gateways on all the remote AZs that the subnet needs to be extended to.
- For each extended subnet within the same Network Gateway appliance ensure that you have unique VxLAN Network Identifiers (VNIs) that you can use for the VTEP subnet extensions. VNI may be any number between 0 and 16777215.

Extending a Layer 2 Subnet Across Availability Zones Over VTEP

The Layer 2 Network Extension over VTEP allows VMs to communicate two Availability Zones (AZ) without a VPN connection.

Before you begin

For information on prerequisites and best practices for extending a Layer 2 subnet, see [Layer 2 Network Extension](#) on page 124 and [Layer 2 Network Extension Over VPN](#) on page 127.

About this task

Perform the following steps to extend a subnet over VTEP across two availability zones (AZs).

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Network & Security > Connectivity** from the **Navigation Bar**.
The **Gateways** page opens displaying the list of network gateways you have created and configured, and the operations you can perform on the network gateways.
3. Click the **Subnet Extensions** tab.
The **Subnet Extensions** page opens displaying the list of subnet extensions created for the clusters.

4. Select **Create Subnet Extension** > **Across Availability Zones**.

Create Subnet Extension Across Availability Zones

? X

A Subnet Extension allows VMs to communicate securely over the same broadcast domain to a remote Availability Zone. This subnet extension can be done over a VPN or a VTEP.

Extend Subnet over a ☐ VPN ☒ VTEP

- Prism Central on the remote side must be pc.2021.11 or later
- On-Prem VPN should be Nutanix VPN and must be version 5.0 or later

Local

Availability Zone

Local AZ

Subnet Type ☒ VLAN ☐ Overlay

Cluster

auto_cluster_

Subnet

net-

VLAN ID	IPAM
0	Unmanaged

Gateway IP Address/Prefix

0.0.0.0/0

Local IP Address ?

0.0.0.0

Local VTEP Gateway

Remote

Availability Zone

PC_10.19.209.66

Subnet Type ☒ VLAN ☐ Overlay

Cluster

auto_cluster_

Subnet

net-

VLAN ID	IPAM
0	Unmanaged

Gateway IP Address/Prefix

0.0.0.0/0

Remote IP Address ?

0.0.0.0

Remote VTEP Gateway

Connection Properties

VxLAN Network Identifier (VNI)

MTU

1392

Cancel

Save

Figure 26: Example of Create VTEP Extension Across AZs with VLAN Subnet

5. For **Extend Subnet over a**, select VTEP.
6. Enter or select the necessary values for the parameters in the **Local** and **Remote (AZ)** sections as described in the table.

Parameters	Description and Value
Availability Zone	Displays the name of the paired availability zone at the local AZ.
Subnet Type	Select the type of the subnet - VLAN or Overlay that you are extending.
Cluster	Select the name of the cluster in the local AZ that the subnet is configured for.
Subnet	Select the name of the subnet at the local AZ for network. The VLAN ID and the IPAM - managed or unmanaged are displayed in the box below the Subnet field.
Gateway IP Address.	Enter the gateway IP address of the subnet you want to extend. Ensure that you provide the IP address in <IP-address/network-prefix> format. for example the gateway IP is 10.20.20.1 in a /24 subnet then provide the gateway IP address as 10.20.20.1/24. Note: For an unmanaged network, enter the gateway IP address of the created subnet. For more information, see PBR-based Tromboning in L2 Extended Subnet on page 140.
Local IP Address	Enter a unique and available (unused) IP address from the subnet provided in Subnet for the Network Gateway appliance.
Remote IP Address	Enter a unique and available (unused) IP address from the subnet provided in Subnet for the remote Network Gateway appliance.
Local VTEP Gateway	Select the local VTEP gateway you created on the local AZ. For more information on creating VTEP gateways, see Creating a Network Gateway on page 108.
Remote VTEP Gateway	Select the VTEP gateway you created on the remote AZ. For more information about creating VTEP gateways, see Creating a Network Gateway on page 108.
Connection Properties	
VxLAN Network Identifier (VNI)	Enter a unique number from the range 0-16777215 as VNI. Ensure that this number is not reused anywhere in the local or remote VTEP Gateways.
MTU	The default MTU is 1392 to account for 108 bytes of overhead and the standard physical MTU of 1500 bytes. VPC Geneve encapsulation requires 58 bytes and VXLAN encapsulation requires 50. However, you can enter any valid MTU value for the network, taking this overhead into account. For example, if the physical network MTU and vs0 MTU are 1600 bytes, the Network Gateway MTU can be set to 1492 to account for 108 bytes of overhead. Ensure that the MTU value does not exceed the MTU of the AHV Host interface and all the network interfaces between the local and remote AZs.

7. Click **Save**.
After the subnet is extended, the extension appears in the **Subnet Extensions** page.

Extending a Subnet to Third Party Datacenters Over VTEP

The Layer 2 Network Extension over VTEP allows VMs to communicate with multiple remote sites or Availability Zones (AZ) that may be third party (non-Nutanix) networks, or datacenters. It also provides the flexibility of adding more remote AZs to the same VTEP-based extended Layer 2 subnet. Examples of compatible VTEP gateways are switches from Cisco, Juniper, Arista, and others that support plain VXLAN VTEP termination.

About this task

Perform the following steps to extend a subnet over VTEP across multiple availability zones (AZs) or third party datacenters.

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Network & Security > Connectivity** from the **Navigation Bar**.
The **Gateways** page opens displaying the list of network gateways you have created and configured, and the operations you can perform on the network gateways.
3. Click the **Subnet Extensions** tab.
The **Subnet Extensions** page opens displaying the list of subnet extensions created for the clusters.

4. Select **Create Subnet Extension > To A Third Party Data-Center**.

Create Subnet Extension To A Third Party Data-Center

A Subnet Extension over VTEP allows VMs to communicate over the same broadcast domain to a remote site.

Local

Availability Zone

Local AZ

Subnet Type

☒ VLAN ☐ Overlay

Cluster

auto_cluster_1

Subnet

net-vlan0

VLAN ID	IPAM
0	

Gateway IP Address/Prefix

0.0.0.0/0

Local IP Address ?

0.0.0.0

Local VTEP Gateway

Remote

Remote VTEP Gateway

vtep-test-remote

Connection Properties

VxLAN Network Identifier (VNI)

MTU

1392

Cancel

Save

Figure 27: Example of Create VTEP Extension To A Third Party Data-Center with VLAN Subnet

5. Enter or select the necessary values for the parameters in the **Local**, **Remote** (AZ), and **Connection Properties** sections as described in the table.

Parameters	Description and Value
Local	
Availability Zone	Displays the name of the paired availability zone at the local AZ.
Subnet Type	Select the type of the subnet - VLAN or Overlay that you are extending.
Cluster	Select the name of the cluster in the local AZ that the subnet is configured for.
Subnet	Select the name of the subnet at the local AZ for network. The VLAN ID and the IPAM - managed or unmanaged are displayed in the box below the Subnet field.
Gateway IP Address	Enter the gateway IP address of the subnet you want to extend. Ensure that you provide the IP address in <IP-address/network-prefix> format. For example the gateway IP is 10.20.20.1 in a /24 subnet then provide the gateway IP address as 10.20.20.1/24. Note: For unmanaged network, enter the gateway IP address of the created subnet. For more information, see PBR-based Tromboning in L2 Extended Subnet on page 140.
Local IP Address	Enter a unique and available (unused) IP address from the subnet provided in Subnet .
Local VTEP Gateway	Select the local VTEP gateway you created on the local AZ. For more information on creating a remote VTEP gateway, see Creating a Network Gateway on page 108.
Remote	
Remote VTEP Gateway	Select the remote VTEP gateway you created on the local AZ. For more information on creating a remote VTEP gateway, see Creating a Network Gateway on page 108.
Connection Properties	
VxLAN Network Identifier (VNI)	Enter a unique number from the range 0-16777215 as VNI. Ensure that this number is not reused anywhere in the networks that the Prism Central and Cluster are a part of.
MTU	The default MTU is 1392 to account for 108 bytes of overhead and the standard physical MTU of 1500 bytes. VPC GENEVE encapsulation requires 58 bytes and VXLAN encapsulation requires 50. However, you can enter any valid MTU value for the network, taking this overhead into account. For example, if the physical network MTU and vs0 MTU are 1600 bytes, the Network Gateway MTU can be set to 1492 to account for 108 bytes of overhead. Ensure that the MTU value does not exceed the MTU of the AHV Host interface and all the network interfaces between the local and remote AZs.

6. Click **Save**.

After the subnet is extended, the extension appears in the **Subnet Extensions** page.

PBR-based Tromboning in L2 Extended Subnet

This topic provides information on using policy based routing for traffic tromboning in an extended subnet.

Flow Virtual Networking provides policy based routing. You can create policies to route traffic through specific routes in the network. For more information on network policy, see [Creating a Policy](#) on page 97.

When two VPCs are connected by a Layer 2 stretched or extended subnet, traffic from each VPC egress the VPC from the respective gateway of each VPC. This traffic egress route is the optimal, default traffic egress route.

When the traffic from both the VPCs egress from the gateway of one of the VPCs, the traffic route is called a tromboning traffic route. In a Layer 2 subnet extension, configure a routing policy with **Forward** action with a next hop IP address added in the Forward IP field to create a tromboning traffic route. For more information on creating a routing policy, see [Creating a Policy](#) on page 97.

Note: The **Forward** action ensures that traffic from both VPCs on either side of a Layer 2 extended subnet exits through a single specified subnet gateway (referred to as the egress gateway). To achieve this, you must configure the Forward IP field with the appropriate next hop IP address that routes traffic to the egress gateway.

This section does not cover all the possible scenarios for determining the correct next hop IP address you can provide in the Forward IP field. You must identify the appropriate IP address based on your specific network configuration. This could include the subnet gateway, the VTEP Local Gateway IP address of the other endpoint in the Layer 2 extended subnet, or the IP address of any intervening firewall VM that routes traffic to the VTEP Local Gateway at the other endpoint, as applicable to your cluster networks.

Example: Layer 2 Subnet Extension

As an example, consider that you configured a Layer 2 subnet extension across Availability Zone AZ1 and Availability Zone AZ2 as follows:

- **AZ1: On-premises**
 - Network ID: 10.1.0.0/16
 - Gateway: 10.1.0.1
 - VLAN AZ1:
 - VLAN Network ID: 10.1.100.0/24
 - VLAN Gateway: 10.1.100.11
- **VPC Prod-AZ1**
 - Subnet ID: 10.1.1.0/24
 - Subnet Gateway: 10.1.1.1
 - VTEP Local Gateway: 10.1.1.91
- VTEP L2subnet Extension IP address: 100.64.1.10 (This IP address is mapped using NAT to a Floating IP for VxLAN tunnel)

- **AZ2: On-premises or Cloud**
 - Network ID: 10.2.0.0/16
 - Gateway: 10.2.0.1
 - VLAN AZ1:
 - VLAN Network ID: 10.2.200.0/24
 - VLAN Gateway: 10.2.200.22
 - **VPC Prod-AZ2** (Identical to VPC Prod-AZ1 for L2 Subnet Extension)
 - Subnet ID: 10.1.1.0/24
 - Subnet Gateway: 10.1.1.1
 - VTEP Local Gateway: 10.1.1.92
 - VTEP L2subnet Extension IP address: 100.64.1.11 (This IP address is mapped using NAT to a Floating IP for VxLAN tunnel)

The performance on Layer 2 Subnet Extension created on VPN connection, as underlay network using non-Nutanix appliances, between on-premises subnets and AWS or Azure VPC subnets might be poor (transfer rates in KBps instead of Mbps).

PBR-based Tromboning in the Example

Create a **Forward** action routing policy for VPC Prod-AZ2 in AZ2, adding the VTEP Local Gateway IP address of VPC Prod-AZ1 (10.1.1.91) in the Forward IP field. The **Forward** action in the policy routes the traffic to gateway for the VTEP Local Gateway IP address of VPC Prod-AZ1.

Note: When you configure a routing policy with the **Forward** action, ensure that you add the appropriate next hop IP address in the Forward IP field.




Figure 28: Example of Forward action-based Tromboning

This policy trombones the traffic in the following path:

From	To
VPC Prod-AZ2	VPC Prod-AZ1 Subnet gateway 10.1.1.1—the subnet gateway for the next hop added in the Forward IP field, in this case, the VTEP Local Gateway.)
VPC Prod-AZ1 Subnet gateway 10.1.1.1	External VLAN AZ1 gateway 10.1.0.1 (VLAN AZ1 being the underlay network with external connectivity for VPC Prod AZ1)
External VLAN AZ1 gateway 10.1.0.1	External VLAN AZ2 gateway 10.2.0.1 (VLAN AZ2 being the underlay network with external connectivity for VPC Prod AZ2)

Updating an Extended Subnet

The *Update Subnet Extension Across Availability Zones* window has the same parameters and fields as the *Create Subnet Extension Across Availability Zones* window.

About this task

You can update a subnet extension that extends across AZs using the **Update Subnet Extension Across Availability Zones** or the **Update Subnet Extension To A Third Party data center** window. The **Update Subnet Extension Across Availability Zones** or the **Update Subnet Extension To A Third Party data center** window has the same parameters and fields as the **Create Subnet Extension Across Availability Zones** or the **Create Subnet Extension To A Third Party data center** window, respectively.

Based on the type of the subnet extension that you want to modify, refer to the following:

Procedure

- For information on extending a subnet over a VPN, see [Extending a Layer 2 Subnet Over VPN](#) on page 129
- For information on extending a subnet over VTEP, see [Extending a Layer 2 Subnet Across Availability Zones Over VTEP](#) on page 133
- For information on extending a subnet across third party datacenters over VTEP, see [Extending a Subnet to Third Party Datacenters Over VTEP](#) on page 137

Removing an Extended Subnet

Perform this procedure to remove the subnet extension.

About this task

This procedure deletes the extended subnet between the two Availability Zones (AZs) or between one Nutanix AZ and one or more third party subnets. Deleting the subnet extension does not automatically remove the network gateways or VPN connections that may have automatically been created by the Subnet Extension wizard. You need to separately delete these entities created automatically when the subnet was extended.

Note: Removing an extended subnet from a cluster or AZ (either source or target AZs) automatically deletes the extended subnet from the corresponding source or target AZs.

Procedure

1. Log in to Prism Central.

2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Network & Security > Connectivity** from the **Navigation Bar**.
The **Gateways** page opens displaying the list of network gateways you have created and configured, and the operations you can perform on the network gateways.
3. Click the **Subnet Extensions** tab.
The **Subnet Extensions** page opens displaying the list of subnet extensions created for the clusters.
4. Select the checkbox associated with the subnet extension you want to remove, and click **Delete** from the **Actions** dropdown menu.
5. In the confirmation dialog box, click **Remove** to remove the extension.
Click **Cancel** to exit without removing the subnet extension.

What to do next

Check the list in the **Subnet Extensions** tab to confirm that the subnet extension is removed.

Border Gateway Protocol Sessions

You can create and manage Border Gateway Protocol sessions.

VPC networking supports No-NAT connectivity. For more information on NAT and No-NAT external connectivity configurations, see [Creating a Subnet](#) on page 93.

You can configure No-NAT external connectivity for VPC subnets using IP addresses with externally routable IP address/prefix that are reachable directly (without SNAT) from the underlying infrastructure. Underlay networks can directly communicate with endpoints in VPCs using such externally routable IP address/prefix. You need to configure routes in the underlay routers to route traffic to externally routable IP address/prefix via the virtual router of the VPC. In the reverse direction, you need to configure the virtual router of the VPC to route traffic to specific infrastructure subnets via an infrastructure router. Manually configuring these routes to and from externally routable IP address/prefix in infrastructure routers is a labor-intensive and error-prone process.

Border Gateway Protocol (BGP) gateways automate the exchange of externally routable IP address/prefix (ERP), routes, and IP address/prefix sets of infrastructure routers. **BGP Sessions** configurable in **Connectivity** supports eBGP and peering with up to 5 infrastructure routers.

Other conditions applicable to BGP sessions are:

1. You can create only one BGP session for one local and remote network gateway pair.
In other words, a local BGP gateway and a remote BGP gateway can only host a maximum of one BGP session.
2. You need to have access permissions of the *VPC Admin* or *Nutanix Infra Admin* roles to create, update or delete BGP sessions.
For more information, see [Control User Access in Flow Virtual Networking \(RBAC\)](#) on page 36.
3. Advertising all the externally routable IP address/prefix (ERP) of the VPC.
4. Without an externally routable IP address/prefix, BGP session creation fails.
5. The BGP appliance can learn and install up to 250 routes.
6. The BGP session advertises a single next hop for each Externally-routable prefix (ERP) of a VPC.

A Network Gateway with a BGP Service is always associated with (servicing) exactly one VPC. A BGP session created on such gateway automatically advertises all the ERPs of the VPC.

Note: The BGP session ignores (or does not advertise) the received routes if the VPC is not associated with a routable (i.e. no-NAT) external subnet.

7. All received routes are added to the VPC routing table on FIFO (First In First Out) basis. Route installation priority is not dependent on destination IP address prefix length.

8. All received routes are added to the VPC routing table only if:

- The routes added are less than or equal to 250 routes.
- The routes to remote subnets use an IP address that is configured on a no-NAT network.

If the VPC is not associated with a no-NAT network, the BG session ignores the received routes and does not add them to the routing table.

9. You can assign a route priority between 300 and 900. If you do not assign a route priority to a BGP session, the BGP session assigns the route priority dynamically between 600 and 800 with reducing steps of 5 starting with 700.

For example, if you added one route without a priority, the BGP session assigns the route a priority of 700. When you add another route later without a priority, the BGP session assigns the new route a priority of 695.

10. Latest log messages are provided on Prism Central in the **BGP Logs** tab for easy troubleshooting.
11. The BGP session details include lists of all advertised and received routes.
12. The BGP session has a minimum 10-minute graceful restart period. If a BGP session fails for any reason, it attempts to restart over a period of 10 minutes or more. The routes of the session are preserved if the BGP session successfully restarts.

The BGP session fails when it is unable to restart after the graceful restart attempt. The route of the failed session is removed from the routing table of the VPC after session failure.

Creating a BGP session

You can create a BGP session between a local BGP gateway and a remote BGP gateway also known as BGP peer.

Before you begin

You must create a local BGP gateway for the local VLAN or VPC network to connect *from*. The BGP gateway may be created in the same VLAN or VPCC or on a different VLAN.

Similarly for the remote network, ensure that you have a remote BGP gateway configured for the VLAN or VPC to connect *to*.

For information about creating, updating or deleting network gateways with BGP service, see [Network Gateway Management](#) on page 108.

About this task

Perform the following steps to create a BGP session.

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Network & Security > Connectivity** from the **Navigation Bar**.
The **Gateways** page opens displaying the list of network gateways you have created and configured, and the operations you can perform on the network gateways.
3. Click the **BGP Sessions** tab.
The **BGP Sessions** page opens displaying the list of BGP sessions created for the clusters.
4. Click **Create BGP Session**.

5. In the **Create BGP Session** window that opens, provide the necessary values in the respective fields as described in the table.

Create BGP Session

X

BGP Sessions creates a path for exchanging routing information between a Local Gateway and a Remote Gateway.

Name

Local BGP Gateway

local-bgp

Remote BGP Gateway

remote-bgp

Dynamic Route Priority

(Autoset Dynamic Route Priority)

i

Dynamic Route Priority must be between 300 and 900.

Password

(Optional)Show

Cancel

Save

Figure 29: Create BGP Session

Parameters	Description
Name	Enter a name for the BGP session.
Local BGP Gateway	Select a local BGP gateway that you want to use for the BGP session.

Parameters	Description
Remote BGP Gateway	Select a remote BGP gateway that you want to use for the BGP session.
Dynamic Route Priority	(Optional) Enter a number between 300 and 900 as priority for the route. If you do not enter a number, Flow Virtual Networking assigns a number between 600-800. The greater the number, higher is the route priority of the session. The first automatically assigned number is 700. After that, subsequent routes requiring automatic or dynamic assignment are assigned numbers that reduce by five (5) from the previously assigned number.
Password	(Optional) Enter a password for the session. Characters allowed for BGP passwords • a-z • A-Z • 0-9 • ~ ! @ # % ^ & * () _ - + = : ; { } [] < > , . / ? \$ • Password length: Minimum 1 and maximum 80 characters. Click Show to make the password visible.

6. Click **Save**.

Updating a BGP session

You can update an existing BGP session. You cannot modify some parameters of the BGP. Such parameters are greyed and in-actionable. If you need to modify such information, consider creating a new gateway with the updated parameters and deleting the current gateway.

About this task

Perform the following steps to update a BGP session.

Note: You can only update the **Name**, **Dynamic Route Priority**, and **Password**. **Local BGP Gateway**, and **Remote BGP Gateway** are unavailable for update. If you need to modify such information, consider creating a new BGP session with the appropriate parameters and deleting the current BGP session.

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Network & Security > Connectivity** from the **Navigation Bar**.
The **Gateways** page opens displaying the list of network gateways you have created and configured, and the operations you can perform on the network gateways.
3. Click the **BGP Sessions** tab.
The **BGP Sessions** page opens displaying the list of BGP sessions created for the clusters.
4. Select the checkbox associated with the BGP session that you want to update, and click **Update** from the **Actions** dropdown menu.
The **Update BGP Session** window opens.

5. Update the necessary values in the respective fields.

The fields in the **Update BGP Session** window is identical to the fields in the **Create BGP Session** window. For more information, see [Creating a BGP session](#) on page 144.

Deleting a BGP session

You can delete an existing BGP session. If you delete a BGP session, all the routes associated with the BGP session are irretrievably deleted.

About this task

Perform the following steps to delete a BGP session.

Procedure

1. Log in to Prism Central.
2. Select the **Infrastructure** application from [Application Switcher Function](#), and navigate to **Network & Security** > **Connectivity** from the **Navigation Bar**.
The **Gateways** page opens displaying the list of network gateways you have created and configured, and the operations you can perform on the network gateways.
3. Click the **BGP Sessions** tab.
The **BGP Sessions** page opens displaying the list of BGP sessions created for the clusters.
4. Select the checkbox associated with the BGP session that you want to delete, and click **Delete** from the **Actions** dropdown menu.
Prism Central displays the **Delete BGP Session <bgp_session_name>** window with a checkbox for the message that warns you that all the active routes associated with the BGP session to be removed, causing a drop in traffic. Further, it asks you to confirm if you want to continue to delete the BGP session.
5. Select the checkbox in the warning message to make the **Delete** button available.
6. Click **Delete** to delete the BGP session.
Click **Cancel** to cancel the deletion.
The status of the **Delete** operation is displayed as a task on the **Tasks** page.

COPYRIGHT

Copyright 2025 Nutanix, Inc.

Nutanix, Inc.

1740 Technology Drive, Suite 150

San Jose, CA 95110

All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. Nutanix and the Nutanix logo are registered trademarks of Nutanix, Inc. in the United States and/or other jurisdictions. All other brand and product names mentioned herein are for identification purposes only and may be trademarks of their respective holders.

Files Manager User Guide

Files 5.2

November 4, 2025

Contents

Files Manager Overview.....	5
Deploying Files in Prism Central.....	6
Upgrading the Files Manager.....	6
Files Manager Layout.....	7
File Server Management in Prism Central.....	9
Creating a File Server in Prism Central.....	9
Configuring the File Server Basics.....	10
Configuring the Client Network.....	11
Configuring the Storage Network.....	11
Configuring the File Server Settings.....	12
Deleting a File Server.....	12
File Server Updates.....	13
Updating File Server Basics.....	14
Updating File Server Sizing.....	14
Updating the FSVM Configuration.....	15
Updating the Network Configuration in Prism Central.....	15
Updating File Server Capacity.....	17
Changing Prism Central Registration.....	18
Starting a Files Cluster in Prism Central.....	18
Stopping a Files Cluster in Prism Central.....	18
Feature Interoperability.....	19
Share and Export Management.....	20
Files on a Public Cloud.....	21
Getting Started with Files on the AWS Public Cloud.....	21
Files on AWS Requirements and Limitations.....	22
Creating an AWS Cloud Account.....	23
Files on AWS Cloud IAM Requirements.....	23
Creating a File Server on the AWS Cloud.....	26
Getting Started with Files on the Azure Public Cloud.....	28
Files on Azure Requirements and Limitations.....	29
Files on Azure Permission Requirements.....	30
Adding an Azure Cloud Account.....	34
Creating a File Server on the Azure Cloud.....	34
Files in a Virtual Private Cloud (VPC).....	37
Deploying Nutanix Files on Azure BareMetal Infrastructure.....	37
Assigning Floating IPs.....	38
RBAC for Nutanix Files.....	39

RBAC Requirements and Limitations for Nutanix Files.....	39
Managing RBAC for Nutanix Files.....	40
Built-In Role Permissions for Nutanix Files.....	40
Custom Role Permissions for Nutanix Files.....	59
Data Replication.....	62
Data Replication Navigation.....	62
Asynchronous Data Protection with Smart DR.....	64
Concepts and Terminology.....	66
Asynchronous Replication Policies.....	67
Creating an Asynchronous Protection Policy.....	67
Editing an Asynchronous Protection Policy.....	68
Switching the Direction of Replication.....	70
Updating AD and DNS Entries.....	70
Smart DR Failover.....	71
Failing Over to a Remote Site.....	72
Manually Failing Over to a Remote Site.....	74
Creating a Reverse Replication Policy.....	75
Smart DR Failback.....	76
Failing Back to the Primary Site.....	76
Synchronous Data Protection with Metro Availability on AHV.....	78
Creating a Synchronous Replication Policy.....	79
Failing Over to a Remote Site.....	80
Resuming Synchronous Replication.....	80
Editing Synchronous Replication Policies.....	80
Disabling a Synchronous Replication Policy.....	81
Deleting a Synchronous Replication Policy.....	81
Enabling a Synchronous Replication Policy.....	82
Data Synchronization with Smart Sync.....	83
Creating a Smart Sync Policy.....	84
Editing a Smart Sync Policy.....	86
Deleting a Smart Sync Policy.....	88
Enabling or Disabling a Smart Sync Policy.....	88
VDI Sync.....	89
Creating a VDI Sync Policy.....	90
Editing a VDI Sync Policy.....	91
Resolving a VDI User Profile Connection Failure.....	92
Resolving Multiple VDI User Connection Failures.....	93
Replication Jobs.....	94
Unified Namespace.....	95
Creating a Unified Namespace.....	95
Adding Nutanix File Servers to a Unified Namespace.....	96
Deleting a Unified Namespace.....	96

Adding a Third-Party File Server to a Unified Namespace.....	97
Managing Third-Party Shares of a Unified Namespace.....	98
Security Hardening.....	99
Copyright.....	100

FILES MANAGER OVERVIEW

The Files Manager is a Nutanix Files service in Prism Central.

Tip: For information on how Nutanix Files documentation is organized, see [Nutanix Files Documentation Organization](#) in the *Nutanix Files User Guide*.

The Files Manager lets you view and control all of your file servers from a single control plane. You can deploy, manage, and update file servers through the Files Manager in Prism Central. File server alerts for all registered file servers appear in a single pane for consolidated viewing, as do file server events.

The Files Manager provides multiple features for Nutanix Files, including Smart DR for share-level data protection, Smart Sync for data replication, and VDI sync to replicate VDI user data across file servers in different locations.

Prerequisites

Register your Prism Element clusters to Prism Central before enabling the Files Manager.

DEPLOYING FILES IN PRISM CENTRAL

Deploy file servers in Prism Central after upgrading the required software and enabling the Files Manager, also known as Files on Prism Central.

Note: Files Manager deployment is a one-time process. For subsequent file server creation, see [Creating a File Server in Prism Central](#) on page 9.

To deploy Files in Prism Central, do the following:

1. Meet the Files prerequisites and port requirements. For more information, see [Prerequisites](#) in the *Nutanix Files User Guide* and [Ports and Protocols](#).
2. Enable Files in Prism Central. For more information, see [Deploying Files](#) in the *Prism Admin Center Guide*.
3. Upgrade the Life Cycle Manager (LCM) in Prism Central. For more information, see the *Life Cycle Manager User Guide*.
4. Upgrade the Files Manager (FM) in LCM. For more information, see [Upgrading the Files Manager](#) on page 6.
5. To discover Files software, re-run the inventory in LCM.
6. In Prism Element, upgrade the Files Server Module (FSM). For more information, [Upgrading File Server Module \(FSM\)](#) in the *Nutanix Files User Guide*.
7. Create a file server instance using the Files Manager. For more information, see [Creating a File Server in Prism Central](#) on page 9.

Upgrading the Files Manager

Upgrade the Files Manager in LCM.

About this task

To upgrade the Files Manager, follow these steps:

Procedure

1. Log on to Prism Central and go to **Admin Center > LCM**.
2. Select **Perform Inventory** and view the latest available version in the **Updates** tab.
3. Select the **Updates** tab, then upgrade FM to a generally available (GA) version.
4. Re-run the inventory in LCM, which automatically discovers new Files software.
For more information about performing inventory and upgrades in LCM, see the [Life Cycle Manager Guide](#).
For compatibility requirements, see the *Nutanix Files Release Notes*.

FILES MANAGER LAYOUT

The Files Manager UI includes several views for file server management and insights.

The Files Manager is organized into the **Summary**, **File Servers**, **Unified Namespace**, **Data Replication**, **Alerts**, and **Events** tabs.

Summary View

The **Summary** view lets you observe highlights of file server activity from a single pane of glass, including an overview of alerts, a summary of replication policies, and a summary of file servers with replication policy errors.

File Servers View

The **File Servers** view lists all file servers registered to the Prism Central instance you are logged on to and provides high-level information about each file server registered to the instance. Clicking the name of a file server takes you to the Files Console, see [Files Console](#) in the *Nutanix Files User Guide*.

Alerts View

The **Alerts** view displays a table that lists the alerts across all registered file servers.

Alerts have the following severity levels:

Critical

The Files Manager detected a critical situation requiring immediate action. The file server has the potential to stop running, or it could run into irreparable issues.

Warning

The Files Manager detected actionable issue requiring intervention. Resolve the issue soon or a more serious issue might develop.

Info

The Files Manager detected a minor actionable problem. Do not ignore the problem and resolve it relatively soon.

Clicking the title of an alert in the **Alerts** tab displays the **Alert Details** window, which provides additional details including, but not limited to, the following:

Impact Type

Includes availability, capacity, configuration, performance, and system indicator.

Description

Provides details on the operation causing the alert and the impacted entities.

Possible Cause

A description of common causes that trigger the alert.

Recommendations

Recommendation on further actions based on the possible cause of the alert.

The **Alert Details** window also provides the following alert triaging options:

Resolve

Clicking the **Resolve** button marks the issue causing the alert as resolved. Resolving an alert indicate that the issue is fixed. After resolving the alert, the status field changes to **Resolved**.

Acknowledge

Clicking the **Acknowledge** button marks the issue causing the alert as acknowledged. Selecting **Acknowledge** cancels reminders for the condition but the alert status remains. After acknowledging the alert, the status field changes to **Acknowledged**.

Events View

The **Events** view displays a table that lists the events across all registered file servers. Events are clarified into the following **Event Type** categories: storage, user action, system action, behavioral anomaly, DR, and hardware.

Clicking the title of an event displays the **Event Details** view, which provides additional details about the event.

Data Replication Tab

The **Data Replication** tab includes several tab for replication policy management. For more information on the different types of data replication, see [Data Replication](#) on page 62. For more information on the **Data Replication** tab navigation, see [Data Replication Navigation](#) on page 62.

FILE SERVER MANAGEMENT IN PRISM CENTRAL

This topic provides an overview of file server creation and management using the Files Manager.

You can create and manage file servers through the Files Manager in Prism Central (PC). File server creation and management requires administrator privileges.

Note: For information on file server creation and management in Prism Element, see "[File Server Management](#)" in the *Nutanix Files User Guide*.

A file server usually consists of three or more file server VMs (FSVMs) combined into a logical group. You cannot reduce file servers of three FSVMs, and you cannot expand single-FSVM file servers. Additionally, single-FSVM deployments have the following limitations:

- No high availability (HA) support. Single-FSVM deployments do not include other FSVMs that an out-of-service FSVM can fail over to.
- No distributed share support. Because distributed shares and exports spread data across multiple FSVMs, single-FSVM deployments only support standard shares.

Important: When the file server capacity or the container capacity reaches 100 percent, all shares or exports within the file server become read-only and Files blocks write privileges. Prism Central displays alerts when the file server capacity or the container capacity reaches 90 percent and when either reaches 100 percent capacity.

Requirements

Meet the compatibility requirements as described in the *Nutanix Files Release Notes*. For connected site or dark site deployments, manually rerun the inventory in LCM after an upgrade.

Limitations

File servers created with the Files Manager do not support the following features:

- Prism-Element-based disaster recovery (DR), including clone and remote site activation
- Network segmentation
- Multi-VLAN

Creating a File Server in Prism Central

Follow the procedure to create a file server using the Files Manager.

About this task

Files Manager–created file servers do not support the following:

- Protection domain (PD)–based disaster recovery (DR), including clone and remote site activation
- Multi-VLAN
- Network segmentation

To create a file server, follow these steps:

Procedure

1. From the Files home page in Prism Central, select **Files Servers > + New File Server**.
The **Create a File Server** window appears.
2. Under **Provider**, select **Nutanix** from the dropdown menu.
For information on creating a file server with a cloud-based provider, see [Creating a File Server on the AWS Cloud](#) on page 26 and [Creating a File Server on the Azure Cloud](#) on page 34. For information on creating file servers in a virtual private cloud, see [Files in a Virtual Private Cloud \(VPC\)](#) on page 37.
3. Under the **Select cluster you want to create a file server on** field, select a target AOS cluster from the drop-down menu.

Note: The Files Manager only lists clusters within your administrative scope. For more information on assigning clusters to users, see [Cluster Role-Based Access Control \(RBAC\)](#) in the *Security Guide*.

4. Under **Select File Server Version**, next to the Nutanix Files version for the new file server, select the checkbox.
5. Click **Proceed**.
The **Create a New File Server** view appears.

Configuring the File Server Basics

About this task

Complete the previous task before configuring file server basics.

To configure file server basics, in the **Basics** tab of the **Create a File Server** window in Prism Central, follow these steps:

Procedure

1. In the **Name** field, enter a name for the file server.
Clients use the name to access the file server.
2. In the **DNS Domain** field, enter a unique fully qualified domain in the `file-server-name@domain` format.
3. In the **File Server Size (Logical)** field, enter the file servers's total storage size with a minimum of 1 TiB.
4. In the **Sizing Configuration** section, configure the number of FSVMs:
 - » To configure the recommended number of FSVMs, do not change the recommended value that is prepopulated.
 - » To manually specify the number of FSVMs, in the **Number of FSVMs** field, enter a value.
5. In the **FSVM Size** field, configure the FSVM size:
 - » To configure the recommended sizing configuration, do not change the recommended values that are prepopulated.
 - » To manually specify the configuration, change the values in the **VCPUs** and **Memory** fields.

Note: For single-FSVM deployments, the number of suggested FSVMs is always one.

6. Click **Next**.

Configuring the Client Network

About this task

Complete the previous task before configuring the client network.

Provide two static IP addresses for each new FSVM, one for the client-side network and one for the storage network. If the clients or the AD domain controllers are in the same subnet as the controller VM (CVM) or the storage network, configure the same client-side and storage-side networks. Otherwise, use separate client and storage networks.

If you use the same network for both clients and storage, use unique IP addresses; otherwise, clients on the same subnet as the storage network cannot access the shares.

To configure the client network, in the **Client Network** tab of the **Create a File Server** window in Prism Central, follow these steps:

Procedure

1. In the **Network** field, select a VLAN for an AHV hypervisor or a port group for ESXi.
Configured network parameters display if the target is a managed network. Unmanaged networks require additional details.
2. (Optional) Configure an unmanaged client network:
 - **Subnet Mask:** Enter the subnet mask value.
 - **Gateway:** Enter the IPv4 gateway IP address.
 - Configure IP addresses:
 - » To specify an IP range, click **Use continuous IP address** and enter the starting IP address in the **Start IP** field.
 - » To specify a non-continuous IP addresses, use the **IP Addresses (Comma Separated)** field.
Use a non-continuous IP address if the start and end IP addresses are not in a range.
The client network typically requires three IP addresses, with the exception of single-FSVM deployments, which require one IP address.
3. Click **Next**.

Configuring the Storage Network

About this task

Complete the previous task before configuring the client network.

To configure the storage network, in the **Storage Network** tab of the **Create a File Server** window in Prism Central, follow these steps:

Procedure

1. In the **Network** field, select a VLAN for an AHV hypervisor or a port group for ESXi.
Configured network parameters display if the target is a managed network. Unmanaged networks require additional details.

2. (Optional) Configure an unmanaged client network:

- **Subnet Mask:** Enter the subnet mask value.
 - **Gateway:** Enter the IPv4 gateway IP address.
 - Configure IP addresses:
 - » To specify an IP range, click **Use continuous IP address** and enter the starting IP address in the **Start IP** field.
 - » To specify a non-continuous IP addresses, use the **IP Addresses (Comma Separated)** field.
- Use a non-continuous IP address if the start and end IP addresses are not in a range.
- The client network typically requires three IP addresses, with the exception of single-FSVM deployments, which require one IP address.

3. Click **Next**.

Configuring the File Server Settings

About this task

Complete the previous task before configuring the file server settings.

To configure file server settings, in the **Settings** tab of the **Create a File Server** window in Prism Central, follow these steps:

Procedure

1. Configure DNS and NTP servers:

- a. In the **DNS Servers** field, enter the IP address of one or more DNS servers.
- b. In the **NTP Servers** field, enter the IP address or hostname of an NTP server.

2. (Optional) To use hypervisor attached disks (HAD) instead of iSCSI connections, select the **Enable direct attach mode for file server VMs** option in the **Advanced Storage Configuration** field.

For more information, see [Hypervisor Attached Disks for AHV](#).

3. Click **Next**.

4. In the **Summary** tab, review the displayed information.

5. Click **Create**.

The file server creation task begins. You can monitor the progress on the **Tasks** page.

What to do next

Nutanix recommends setting up data protection for your file server. For more information, see [Asynchronous Data Protection with Smart DR](#) on page 64.

Deleting a File Server

Delete a file server using the Files Manager.

About this task

You must leave the domain before deleting a file server.

Nutanix recommends deleting file server DNS entries before deleting the file server. Refer to *Directory Service and Domain Management* in the *Nutanix Files User Guide*.

Warning: Do not delete FSVMs or power off a file server before deleting it. If you delete FSVMs, you must force delete the file server.

Deleting a file server is permanent, and none of the deleted files can be recovered. With Prism Element–deployed file servers, if the file server protection domain snapshots are available and the Files container is present, you can recover the file server from the snapshot. Files Manager–created file servers do not have protection domain snapshots, so deleting a file server is an irrecoverable operation.

To delete a file server, follow these steps:

Procedure

1. On the Files Manager home page, select **File Servers**.
2. Select a file server to update, and then select **Actions > Delete File Server**.
The **Delete File Server** dialog box appears.
3. In the dialog box, choose one of the following options:

- » Retain all related entities.

Note: These entities can be deleted manually.

- » Delete related entities.

If you decide to delete related entities, select the **Container** and/or **Snapshots, schedules, and protection domain** checkbox.

Deleting a file server deletes all of the underlying user data related to that file server. Depending on your AOS version, you might not see the option to delete the storage container and protection domain created for that file server. You can delete that storage container and protection domain manually.

4. Select **Delete**.
The file server will be deleted.

File Server Updates

Update or expand the configuration of an existing file server deployed using Prism Central or Prism Element.

Nutanix Files lets you perform the following file server updates using the Files Manager:

- Update the name and the domain of the file server. For more information, see [Updating File Server Basics](#) on page 14.
- Update the number of file server VMs (FSVMs). For more information, see [Updating File Server Sizing](#) on page 14.
- Update the vCPU count and memory size of each FSVM. For more information, see [Updating the FSVM Configuration](#) on page 15.
- Update the network configuration of the client network and the storage network. For more information, see [Updating the Network Configuration in Prism Central](#) on page 15.
- Update the file server size. For more information, see [Updating File Server Capacity](#) on page 17.

Updating File Server Basics

Update the name and the domain of a file server using Prism Central.

About this task

Before changing the file server or domain name, leave the current domain.

For information on leaving a domain, see [Directory Service and Domain Management](#) in the *Files User Guide*.

To update file server basics, follow these steps:

Procedure

1. From the Files Manager home page, navigate to the **File Servers** tab.
2. Select the checkbox in the row of the target file server.
3. Select **Actions > Update File Server Basics**.
The **Update Name and Domain** window displays.
4. (Optional) To change the file server name, enter the new name in the **Name** field.
5. (Optional) To change the DNS domain name, enter the new name in the **DNS Domain** field.
6. Click **Update**.

What to do next

After changing the file server or domain names, delete any old DNS entries from the DNS servers, add the new DNS entries, and update the directory services configuration as needed. For more information, see [Directory Service and Domain Management](#) in the *Nutanix Files User Guide*.

Updating File Server Sizing

Update file server sizing by adding or removing file server VMs (FSVMs) on a file server using Prism Central.

About this task

Adding a new FSVM scales out a file server cluster. Removing an FSVM scales in a cluster.

Important: Scale-out and scale-in operations can cause a momentary file server connection loss. Nutanix recommends performing scale-out and scale-in operations during non-business hours.

The following limitations apply to modifying the number of FSVMS:

- You cannot scale in file servers of three FSVMS.
- You cannot scale out single-FSVM deployments.
- Single or dual-node AOS clusters only support single-FSVM deployments.
- The number of CVMs must be equal to or greater than the number of FSVMS.

To update the number of FSVMS, follow these steps:

Procedure

1. From the Files Manager home page, navigate to the **File Servers** tab.
2. Select the checkbox in the row of the target file server.

3. Select **Actions > Update File Server Sizing**.
The **Update File Server Sizing** window displays.
4. Under **Set the number of File Server VMs**, specify the number of FSVMs by entering an integer or dragging the slider.
5. Review the configuration in the **Client Network** and **Storage Network** sections.

Note: For information on updating the file server network, see [Updating the Network Configuration in Prism Central](#) on page 15 .

6. Click **Update**.

Updating the FSVM Configuration

Update the FSVM configuration by changing the vCPUs and memory resources of each FSVM using Prism Central.

Before you begin

Keep your CD-ROM with ISO attached to each FSVM before updating the resources.

About this task

Adding memory or CPU resources to FSVMs scales it out. Removing memory or CPU scales in the FSVM configuration.

Note: Scale-up is a non-disruptive operation, reducing the configuration of a file server, also known as scale-down, causes system downtime.

To update the FSVM configuration, follow these steps:

Procedure

1. From the Files Manager home page, navigate to the **File Servers** tab.
2. Select the checkbox in the row of the target file server.
3. Select **Actions > Update FSVM Configuration**.
The **Update FSVM Configuration** window displays.
4. (Optional) In the **vCPUs** field, select the total number of vCPUs for each FSVM in the cluster.
5. (Optional) In the **Memory Per File Server VM**, select the total number of memory (GiB) for each FSVM in the cluster.
6. Click **Update**.
Nutanix Files updates the vCPU and memory resources of each FSVM as specified.

Updating the Network Configuration in Prism Central

Change or re-IP the client or storage network configuration of a file server using Prism Central.

Before you begin

Stop the file server before updating the network to reduce traffic during the update process. For more information, see [Stopping a Files Cluster in Prism Central](#) on page 18.

About this task

Changing the client network and storage network of a file server VM lets you move a file server from one data center to another. If you intend to change the domain, you must leave the domain first.

Updating a network configuration on ESXi clusters requires vCenter connectivity.

For information on configuring a virtual network for guest VMs, see [Configuring a Virtual Network for Guest VM Interfaces](#) in the *Prism Web Console Guide*. For information on modifying network connections, see [Modifying Network Connections](#) in the *Prism Web Console Guide*.

Important: Updating the network configuration requires a temporary downtime. Nutanix recommends performing network updates during off-peak hours.

To update the file server network configuration, follow these steps:

Procedure

1. From the Files Manager home page, go to the **File Servers** tab.
2. In the row of the target file server, select the checkbox.
3. Go to **Actions > Update Network**.
The **Update File Server Network** window displays.
4. Follow these steps to update the network configuration:
You can change the client network, storage network, DNS, and NTP IP addresses for a file server VM.
 - a. In the **Client Network** tab, verify or update the client network details.
For more information on the fields in the **Client Network** tab, see [Configuring the Client Network](#) on page 11.
 - b. Click **Next**.
 - c. In the **Storage Network** tab, verify or update the storage network details.
For more information on the fields in the **Storage Network** tab, see [Configuring the Storage Network](#) on page 11.
 - d. Click **Update**.
If the network update operation is unsuccessful, use the original IP address details and retry the update operation.

What to do next

Modifying the client network configuration requires modifying the DNS entries. For more information, see [Updating Domain Name System \(DNS\) Entries](#).

Updating the Network Configuration to SR-IOV

Configure single root I/O virtualization (SR-IOV) on an AHV file server to automatically enable remote direct memory access (RDMA) for the NFS protocol.

Before you begin

Create a NIC profile. For more information, see [Creating a NIC Profile](#) in the *Prism Central Infrastructure Guide*.

About this task

Existing IPv6 configurations on the client network are lost after configuring the network to SR-IOV.

For information on the NFS over RDMA performance optimization feature, including minimum requirements, see [NFS Over RDMA](#) in the *Nutanix Files User Guide*. For information on enabling NFS over RDMA on ESXi, contact Nutanix Support.

To update the file server network configuration to SR-IOV, follow these steps:

Procedure

1. From the Files Manager home page, go to the **File Servers** tab.
2. In the row of the target file server, select the checkbox.
3. Go to **Actions > Update Network**.
The **Update File Server Network** window displays.
4. In the **Client Network** tab, configure the following:
 - a. In the **Network Configuration Type** section, select **SR-IOV NIC**.
 - b. From the **NIC Profile** dropdown list, select the NIC profile.
 - c. In the **Subnet Mask** field, enter the subnet mask value.
 - d. In the **Gateway** field, enter the IPv4 gateway IP address.
The client network typically requires three IP addresses, with the exception of single-FSVM deployments, which require one IP address.
5. Click **Next**.
6. (Optional) If you switch to an unmanaged network, in the **Storage Network** tab, configure the following:
 - a. In the **Network** field, select a network from the dropdown list.
 - b. In the **Subnet Mask** field, enter the subnet mask value.
 - c. In the **Gateway** field, enter the IPv4 gateway IP address.
 - d. Configure IP addresses:
 1. To specify an IP range, click **Use continuous IP address**, and in the **Start IP** field, enter the starting IP address .
 2. To specify a non-continuous IP addresses, use the **IP Addresses (Comma Separated)** field.
Use a non-continuous IP address if the start and end IP addresses are not in a range.
7. Click **Update**.
The network update operation begins, which you can monitor in the **Tasks** view.

Updating File Server Capacity

Update the size, also known as the logical capacity, of a file server using Prism Central.

Procedure

1. From the Files Manager home page, navigate to the **File Servers** tab.
2. Select the checkbox in the row of the target file server.
3. Select **Actions > Update Capacity**.
The **Capacity** window displays.

4. To change the size of the file server, enter a value in the **File Server Size (Logical)** field.

Note: You cannot reduce the size of the file server. This operation does not impact current client connections.

5. Click **Update**.
Nutanix Files increases the file server size.

Changing Prism Central Registration

Un-register a file server from an existing Prism Central instance and register it to a different one.

Procedure

1. On the old PC, do the following to remove the file server:

```
nutanix@pcvm$ afs infra.fix_pc_config pe_uuid=PE-cluster-uuid get_fs_uuids
```

Replace *PE-cluster-uuid*, with the UUID of the Prism Element cluster the file server belongs to.

2. On the new PC, do the following to add the file server:

```
nutanix@pcvm$ afs infra.fix_pc_config pe_uuid=PE-cluster-uuid fs_uuids=fs_uuid
```

Replace *PE-cluster-uuid*, with the UUID of the Prism Element cluster the file server belongs to.

Replace *fs_uuid* with the UUID of the file server.

What to do next

If you do not decommission the old PC, contact Nutanix Support.

Starting a Files Cluster in Prism Central

Start a file server deployed in Prism Central (PC).

Procedure

1. Using SSH, log on to the PC VM of the target file server.
2. List file servers on the PC VM:

```
nutanix@PCVM$ afs info.fileservers
```

Note: The `Deployment Platform` parameter in the output specifies whether the file server is deployed in Prism Element (PE) or PC.

3. Start a file server:

```
nutanix@PCVM$ afs infra.power_on fs_names=file-server
```

Replace *file-server* with the name of a file server or a comma separated list of names.

Note: For information on starting a file server in Prism Element, see [Starting a Files Cluster](#).

Stopping a Files Cluster in Prism Central

Gracefully stop a file server deployed in Prism Central (PC).

Procedure

1. Using SSH, log on to the PC VM of the target file server.

2. List file servers on the PC VM:

```
nutanix@PCVM$ afs info.fileservers
```

Note: The `Deployment Platform` parameter in the output specifies whether the file server is deployed in Prism Element (PE) or PC.

3. Stop a file server:

```
nutanix@PCVM$ <afs> infra.power_off fs_names=file-server
```

Replace *file-server* with the name of a file server or a comma separated list of names.

Note: For information on stopping a file server in Prism Element, see [Stopping a Files Cluster](#).

Feature Interoperability

The sections in this topic describe the interoperability of Files Manager features.

VDI Sync

VDI sync has the following interoperability with other Files Manager features:

- **Smart Sync**
 - Unsupported on the same share
 - Supported on other shares of the same file server
- **Smart DR**
 - Unsupported on the same share
 - Supported on other shares of the same file server
- **Public Cloud**
 - Supported on the same share

For more information, see [VDI Sync](#) on page 89.

Smart Sync

Smart Sync has the following interoperability with other Files Manager features:

- **VDI Sync**
 - See "Smart Sync" in the VDI Sync section
- **Smart DR**
 - Not interoperable for the same target share
 - Interoperable for the same source share
- **Public Cloud**
 - Interoperable on the same share

For more information, see [Data Synchronization with Smart Sync](#) on page 83.

Smart DR

The Smart DR has the following interoperability with other Files Manager features:

- **VDI Sync**
 - See "Smart DR" in the VDI Sync section
- **Smart Sync**
 - See "Smart DR" in the Smart Sync section
- **Public Cloud**
 - Partially supported with cloud-based disaster recovery

For more information, see [Asynchronous Data Protection with Smart DR](#) on page 64.

Public Cloud

Files on Public Cloud has the following interoperability with other Files Manager features:

- **VDI Sync**
 - See "Public Cloud" in the VDI Sync section
- **Smart Sync**
 - See "Public Cloud" in the Smart Sync section
- **Public DR**
 - Partially supported with cloud-based disaster recovery

For more information, see [Files on a Public Cloud](#) on page 21.

Share and Export Management

You can manage Prism Central and Prism Element deployed shares using the Files Console.

Access the Files Console from the **File Servers** view in the Files Manager by clicking the name of a file server.

For information on share and export management and the Files Console, see the following chapters in the *Nutanix Files User Guide*:

- [Files Console](#)
- [Share and Export Management](#)
- [Files Options](#)

FILES ON A PUBLIC CLOUD

Deploy Nutanix Files on a public cloud using AWS or Microsoft Azure.

The software-defined architecture of Nutanix Files lets you run it on compatible third-party infrastructure or Nutanix infrastructure with Nutanix Cloud Clusters (NC2), Nutanix Cloud Infrastructure (NCI), or as a standalone Nutanix Unified Storage cluster. AWS and Microsoft Azure provide scalable infrastructure for your public cloud. Meanwhile, Prism Central, running in your Nutanix cluster, provides the Nutanix Files control plane.

Nutanix Files on the public cloud uses cloud-based storage rather than Nutanix-backed volume groups (VGs) to back file server VMs (FSVMs).

Deploying Files on the public cloud provides the following notable features:

- Automatic storage expansion
- Automatic tiering to object storage
- Write support for files in the object store
- Management of Prism Central-based on-premises file servers from the same control plane as those on the cloud

Cloud Accounts

Deploying Files on a public cloud requires setting up a cloud account in Prism Central with AWS or Microsoft Azure account credentials. Unlike on-premises Nutanix Files deployments with file servers deployed on CVMs, public cloud file servers do not rely on CVMs and are deployed in availability zones (AZs) that are tied to specific geographic regions. You can have multiple regions associated with each cloud account.

After setting up a cloud account and meeting the prerequisites, you can create a file server on AWS or Azure.

For more information on cloud accounts, see [Creating an AWS Cloud Account](#) on page 23 and [Adding an Azure Cloud Account](#) on page 34.

Tiering

During file-server creation, you have the option to enable automatic tiering. When enabled, automatic tiering minimizes storage on the block storage tier by moving cold data, which receives fewer reads and writes, to AWS S3 or Azure Blob object storage. Automatic tiering has three tiering levels:

- *Lenient* tiering keeps the majority of data in the block storage tier rather than the object tier.
- *Balanced* tiering keeps a near-equal quantity of data in the block storage tier and the object storage tier.
- *Aggressive* tiering keeps the majority of data in the block storage tier rather than the object tier.

Getting Started with Files on the AWS Public Cloud

Provisioning and deploying Files on AWS requires following these high-level steps:

Procedure

1. Review the requirements and limitations in [Getting Started with Files on the AWS Public Cloud](#) on page 21.
2. Create a cloud account as described in [Creating an AWS Cloud Account](#) on page 23.
3. Deploy a Files file server in AWS as described in [Creating a File Server on the AWS Cloud](#) on page 26.
4. (Optional) Set up a data protection policy for your cloud-based file server as described [Asynchronous Data Protection with Smart DR](#) on page 64.

Files on AWS Requirements and Limitations

Files on AWS deployment requirements and limitations.

Deploying Files on the AWS public cloud requires setting up a cloud account in Prism Central. An AWS cloud account pulls the following details from your public cloud account to deploy a file server:

- Regions
- AZs
- Virtual private clouds (VPCs)
- Subnets

For more information on cloud accounts, see [Creating an AWS Cloud Account](#) on page 23.

Requirements

Make sure to meet the following requirements in your AWS environment:

- If Prism Central is running on-premises, not in Nutanix Cloud Clusters (NC2), set up a VPN for your AWS VPC. Files uses the VPN to reach AWS storage during file server deployment.
- Open TCP port 22 from PC to the file servers on AWS for SSH communication.
- Add the Prism Admin to the Cluster Admin role to create cloud-based file servers.
- Meet the user and file server VM (FSVM) identity and access management (IAM) requirements in AWS. For more information, see [Files on AWS Cloud IAM Requirements](#) on page 23.
- Configure an AWS VPC with one of the following subnet configurations:
 - A private subnet for file server IP address assignment and a public subnet for network address translation (NAT) internet access.
 - Or a private subnet with the following AWS endpoints:
 - An EC2 endpoint with the **Interface** type and the private subnet dedicated to file server deployment
 - An S3 endpoint with the **Gateway** type that references a route table containing the following:
 - The default managed prefix list for the S3 endpoint in the target region for file server deployment
 - The classless inter-domain route (CIDR) of the VPC dedicated to file server deployment
 - (Optional) CIDRs of on-premises file servers or CIDRs of NFS and SMB clients

For information on configuring the AWS environment, see [AWS documentation](#).

Limitations

Running Nutanix Files on AWS has the following limitations:

- Adding a node does not move storage from existing nodes to the new node like it does for on-prem deployments with volume group (VG) movement.
- You cannot delete nodes due to the AWS limitation for the number of disks per VM.
- You cannot load balance VGs; every FSVM has only one VG.
- You cannot perform hot scale-up; scaling-up requires a rolling restart of all FSVMs.
- You cannot scale up CPUs independent of memory; AWS has specific configuration constraints for FSVMs and memory.

- Deleting files and shares does not reduce underlying elastic block storage (EBS); however, deleting files and shares might reduce the S3 storage footprint.
- You can have a maximum of 60 AWS security group rules.

For information on the interoperability of Files Manager features, see [Feature Interoperability](#) on page 19.

Creating an AWS Cloud Account

Create a cloud account to deploy Nutanix Files in the AWS public cloud.

Before you begin

Do the following before you create a cloud account:

- Create a cloud user with the required permissions. For more information, see the "Requirements" section in [Files on a Public Cloud](#) on page 21.
- Generate an access key and a secret key for the cloud user.

Important: You cannot upgrade file servers with technical preview cloud file server deployments to general availability (GA) versions. Delete file servers created during the technical preview.

About this task

To create a cloud account, follow these steps:

Procedure

1. From the side pane in the Files Manager, select **Cloud Accounts > Add Cloud Account**. The **Add Cloud Account** configuration window appears.
2. In the **Cloud Account Name** field, enter a name for the new account.
3. In the **Cloud Provider** dropdown menu, select **Amazon Web Services (AWS)**.
4. In the **Connect to Amazon Web Services** section, enter the cloud credentials:
 - a. In the **Access Key** field, enter your AWS cloud user access key.
 - b. In the **Secret Key** field, enter the secret key for the cloud user's access key.
5. Under **Regions**, choose one of the following:
 - » To use the account across all regions, select **All Regions** and click **Add**.
 - » To manually specify regions, select **Preferred Regions**.
 Selecting **Preferred Regions** displays a dropdown menu that lists all possible regions.
6. If you chose the **Preferred Regions** options, select the checkbox next to each region for the cloud account and click **Add**.

Files on AWS Cloud IAM Requirements

Meet the identity and access management (IAM) role requirements for Files on the AWS Cloud.

AWS IAM User Permission Requirements

In AWS, create and assign an IAM user role to the deployment user with the required permissions listed in the following table.

Table 1: AWS IAM User Permissions

AWS Service	Required Permissions
EC2	• ec2:DescribeImages • ec2:DescribeInstances • ec2:TerminateInstances • ec2:RunInstances • ec2:StopInstances • ec2:StartInstances • ec2:ModifyInstanceAttribute • ec2:DescribePlacementGroups • ec2:DescribeSubnets • ec2:DescribeAvailabilityZones • ec2:DescribeSecurityGroups • ec2>DeleteSecurityGroup • ec2:DescribeNetworkInterfaces • ec2>DeleteNetworkInterface • ec2>CreateNetworkInterface • ec2>CreateSecurityGroup • ec2:AssignPrivateIpAddresses • ec2:AuthorizeSecurityGroupIngress • ec2:DescribeKeyPairs • ec2:RevokeSecurityGroupIngress • ec2:CreateTags • ec2:CreateVolume • ec2>DeleteVolume • ec2:AttachVolume • ec2:DescribeVolumes • ec2:DetachVolume • ec2:ModifyVolume • ec2:DescribeVpcs

AWS Service	Required Permissions
IAM	• iam:GetInstanceProfile • iam:SimulatePrincipalPolicy • iam:ListRoles • iam:PassRole
KMS	• kms:DescribeKey
S3	• s3:GetBucketLocation • s3:CreateBucket • s3:PutBucketVersioning • s3:ListBucket • s3>DeleteObject • s3:ListBucketVersion • s3>DeleteBucket • s3:PutObject • s3:ListBucketMultipartUploads • s3:AbortMultipartUpload • s3:GetBucketObjectLockConfiguration • s3:GetObject • s3:GetLifecycleConfiguration • s3:PutLifecycleConfiguration • s3:GetObjectTagging • s3:PutObjectTagging • s3:GetBucketVersioning • s3:GetObjectVersion

FSVM IAM Role Permission Requirements

In AWS, create and assign an FSVM IAM role to each file server with the required permissions listed in the following table.

Table 2: FSVM IAM Role Permissions

AWS Service	Required Permissions
EC2	• ec2:DescribeInstances • ec2:DescribeInstanceStatus • ec2:DescribeNetworkInterfaces • ec2:DescribeSubnets • ec2:DescribeVpcs • ec2:DescribeDhcpOptions • ec2:DescribeKeyPairs • ec2:DescribeRegions • ec2:DescribeVolumes • ec2:DescribeVolumesModifications • ec2:DescribeVolumeStatus • ec2:CreateTags • ec2:DescribeImages • ec2:AssignPrivateIpAddresses • ec2:AttachVolume • ec2:DetachVolume
KMS	Note: The KMS role is only required for file servers created with encryption enabled. • kms:DescribeKey • kms:Encrypt • kms:Decrypt • kms:GenerateDataKey* • kms>CreateGrant
S3	• s3:GetObject • s3:GetBucketLocation • s3:ListBucket • s3:ListBucketVersions

Creating a File Server on the AWS Cloud

Follow the procedure to create a cloud-based file server.

About this task

Cloud-based file servers do not support the following:

- Disaster recovery (DR) through PD-based-DR, including cloning and remote site activation

- Multi-VLAN
- Network segmentation

Note: For information on configuring cloud-based tiering, see the [Nutanix Files 5.1.1 Supplement](#).

To create a file server on the cloud, follow these steps:

Procedure

1. From the **Files** home page in Prism Central, click **Files Servers > New File Server**.
2. In the **Create a File Server** window, configure the following settings:
 - **Provider:** Select **AWS**.
 - **Cloud Account:** Select a cloud account. For information on cloud account creation, see [Creating an AWS Cloud Account](#) on page 23.
 - **Region:** Select one of the available regions.
 - **File Server Version:** Select the Nutanix Files version for the new file server.
3. Click **Proceed**.
4. In the **File Server Basics** tab, configure the file server basics:
 - a. In the **File Server Details** section, configure the following:
 - **Name:** Enter a name for the file server.
Clients use this name to access the file server.
 - **DNS Domain:** Enter a unique fully qualified domain.
Use the `file_server_name.domain` format.
 - **File Server Size (Logical):** Enter the total file server total storage size, which must be a minimum of 1 TB.
 - b. In the **Data Tiering and Cost Savings** section, configure data tiering:
 - (Optional) **Enable Automatic Data Tiering:** Clear the checkbox to disable automatic tiering to the AWS object store.
 - **Select the Level of Data Tiering:** From the dropdown menu, select **Lenient**, **Balanced**, or **Aggressive** tiering.
For more information on each tiering level, see [Files on a Public Cloud](#) on page 21.
 - (Optional) **Use existing bucket for data tiering:** Select the checkbox to specify an existing bucket for tiering operations.
 - (Optional) **Bucket Name:** Specify the name of an existing bucket as a tiering target for cold data.
 - c. In the **Sizing Configuration** section, configure the following:
 - **Number of File Server VMs (FSVMs):** Specify the number of FSVMs by clicking - or +, or by entering an integer value.
 - **FSVM Size:** Select a combination of vCPUs and memory resources per FSVM.

5. In the **Client Network** tab, configure the following settings:
 - **Availability Zone:** Select one of the availability zones for the specified region.
 - **Virtual Private Cloud (VPC):** Select an AWS VPC.
 - **Subnet:** Select a subnet from the subnets configured for the specified AZ
6. Click **Next**.
7. In the **Settings** tab, configure the following settings:
 - **DNS Servers:** Enter the domain name system server IP address or a comma separated list of IP addresses.
 - **NTP Servers:** Enter the network time protocol server IP address or a comma separated list of IP addresses.
 - **Key Pair:** Select a key pair from the dropdown list.
 - **Cloud IAM Role for File Server VM (FSVM):** Select a role from the dropdown list.
For more information, see [Files on AWS Cloud IAM Requirements](#) on page 23.
 - **Encryption:** Select either **none** or **encrypt data**.
 -
8. Click **Next**
9. In the **Summary** tab, review the file server information and select **Deploy**

Note: You cannot configure the name service with a Files Manager deployment. Configure the name service in the Files Console separately after you complete the deployment.

File server creation begins. You can monitor progress through the **Tasks** view.

Getting Started with Files on the Azure Public Cloud

Provisioning and deploying Files on Microsoft Azure requires following these high-level steps:

Procedure

1. Review the requirements and limitations.
For more information, see [Files on a Public Cloud](#) on page 21.
2. Create a cloud account.
For more information, see [Adding an Azure Cloud Account](#) on page 34.
3. Meet the Azure deployment requirements and review the limitations.
For more information, see [Files on Azure Requirements and Limitations](#) on page 29.
4. Meet the Azure permission requirements for Nutanix Files file server deployment.
For more information, see [Files on Azure Permission Requirements](#) on page 30,
5. Deploy a file server in Azure.
For more information, see [Creating a File Server on the Azure Cloud](#) on page 34.
6. (Optional) Set up a data protection policy for your cloud-based file server.
For more information, see [Asynchronous Data Protection with Smart DR](#) on page 64.

Files on Azure Requirements and Limitations

Deployment requirements and limitations for Files on the Microsoft Azure cloud.

Requirements

Make sure to meet the following requirements in your Azure environment:

- If Prism Central is running on-premises, not in Nutanix Cloud Clusters (NC2), set up a VPN for your Azure VPC. Files uses the VPN to reach Azure storage during file server deployment.
- Open the following ports: TCP port 22 from PC to the file servers on Azure for SSH communication.
 - Port 22, TCP and UDP, for bi-directional SSH communication between Prism Central and file servers on Azure
 - Port 80, 442, and 9440, TCP and UDP, for bi-directional communication between Prism Central and Azure
- Create an Azure subscription-level custom role with the required permissions and assign it to a managed identity to use for file server deployment. For more information, see [Files on Azure Permission Requirements](#) on page 30.
- Create a resource group for file server storage.
- (Optional) Create a key vault with the required managed identity access permissions. For more information, see [Files on Azure Permission Requirements](#) on page 30.
- (Optional) Use the key vault to create a disk encryption set and to generate an SSH key.

Networking Requirements

Configure the following networking requirements:

- In each region for file server deployment, create a virtual network with a subnet dedicated to file server deployment.
- Create an additional subnet named GatewaySubnet. For more information, see [Azure Virtual Networking](#).
- Provide each file server one IP address plus two IP addresses per node.

Limitations

Running Nutanix Files on Azure has the following limitations:

- Microsoft Azure cloud limitations apply.
- Azure Managed Disk Premium SSDs and Azure Blob storage have a quota limit. To exceed the capacity, submit an Azure quota increase request or contact Azure Support.
- Adding a node does not move storage from existing nodes to the new node like it does for on-premises deployments with volume group (VG) movement.
- You cannot delete nodes due to the Azure limitations for the number of disks per VM.
- You cannot load balance VGs; every FSVM has only one VG.
- You cannot perform hot scale-up; scaling-up requires a rolling restart of all FSVMs.
- You cannot scale up CPUs independent of memory; Azure has specific configuration constraints for FSVMs and memory.
- Deleting files and shares does not reduce underlying object storage; however, deleting files and shares might reduce the Azure Blob storage footprint.

- High availability mode can take slightly longer in Azure environments.

For information on the interoperability of Files Manager features, see [Feature Interoperability](#) on page 19.

Files on Azure Permission Requirements

Meet the identity and access requirements for Files on the Microsoft Azure cloud.

User Role IAM Requirements

Create a role for the managed identity with the required permissions listed in the following table.

Table 3: User Assigned Role Permissions

Permission Category	Required Permission
Actions	• <code>Microsoft.Network/networkInterfaces/write</code> • <code>Microsoft.Compute/disks/write</code> • <code>Microsoft.Compute/virtualMachines/write</code> • <code>Microsoft.Resources/tags/write</code> • <code>Microsoft.Compute/images/read</code> • <code>Microsoft.Compute/virtualMachines/read</code> • <code>Microsoft.Compute/virtualMachines/instanceView/read</code> • <code>Microsoft.Network/networkInterfaces/read</code> • <code>Microsoft.Resources/subscriptions/locations/read</code> • <code>Microsoft.Network/virtualNetworks/subnets/read</code> • <code>Microsoft.Compute/disks/read</code> • <code>Microsoft.Storage/storageAccounts/read</code> • <code>Microsoft.Storage/storageAccounts/blobServices/containers/read</code> • <code>Microsoft.Network/virtualNetworks/read</code> • <code>Microsoft.KeyVault/vaults/keys/write</code> • <code>Microsoft.Network/virtualNetworks/subnets/join/action</code> • <code>Microsoft.Storage/storageAccounts/managementPolicies/read</code> • <code>Microsoft.Storage/storageAccounts/managementPolicies/write</code> • <code>Microsoft.Compute/diskEncryptionSets/read</code> • <code>Microsoft.Compute/galleries/images/versions/read</code> • <code>Microsoft.Network/networkInterfaces/join/action</code>

Permission Category	Required Permission
Data Actions	• <code>Microsoft.KeyVault/vaults/keys/decrypt/action</code> • <code>Microsoft.KeyVault/vaults/keys/encrypt/action</code> • <code>Microsoft.KeyVault/vaults/keys/wrap/action</code> • <code>Microsoft.KeyVault/vaults/keys/unwrap/action</code> • <code>Microsoft.Storage/storageAccounts/blobServices/containers/blobs/read</code> • <code>Microsoft.KeyVault/vaults/keys/read</code>

Key Vault Access Policy

Create a key vault access policy for the managed role, logged-in user, and the application ID for file server deployment with the following permissions for the Key Permissions category:

- `Get`
- `List`
- `Update`
- `Create`
- `Import`
- `Delete`
- `Recover`
- `Backup`
- `Restore`
- `GetRotationPolicy`
- `SetRotationPolicy`
- `Rotate`
- `WrapKey`
- `UnwrapKey`

Application Permission Requirements

Configure app registration in Microsoft Entra to allow scripts and applications to securely authenticate and interact with Azure services using identity and access management (IAM).

When you register an app, Azure creates the following:

- An *application object* to distinguish the app
- A *service principal* to serve as the identity used within your tenant

An app registered with Entra ID acts like an identity, known as a service principal, with its own credentials (client ID, secret, or certificate). The service principal is used by automation scripts, third-party apps, and services needing access to Azure resources. For more information on app registration with Microsoft Entra, see [Microsoft Entra documentation](#).

To give your app access, in Microsoft Entra, assign a role to the service principal with the permissions listed in the following table:

Table 4: Application Permissions

Permission Category	Required Permission
Actions	• Microsoft.Resources/subscriptions/resourceGroups/read • Microsoft.Resources/subscriptions/resourceGroups/write • Microsoft.Resources/tags/read • Microsoft.Resources/tags/write • Microsoft.Resources/tags/delete • Microsoft.Network/virtualNetworks/read • Microsoft.Network/virtualNetworks/write • Microsoft.Network/virtualNetworks/listDnsResolvers/action • Microsoft.Network/virtualNetworks/subnets/read • Microsoft.Network/virtualNetworks/subnets/write • Microsoft.Network/virtualNetworks/subnets/join/action • Microsoft.Network/networkInterfaces/read • Microsoft.Network/networkInterfaces/write • Microsoft.Network/networkInterfaces/join/action • Microsoft.Network/networkInterfaces/delete • Microsoft.Network/networkInterfaces/ipconfigurations/read • Microsoft.Compute/disks/read • Microsoft.Compute/disks/write • Microsoft.Compute/disks/delete • Microsoft.Compute/virtualMachines/vmSizes/read • Microsoft.Compute/virtualMachines/read • Microsoft.Compute/virtualMachines/write • Microsoft.Compute/virtualMachines/delete • Microsoft.Compute/virtualMachines/start/action • Microsoft.Compute/virtualMachines/powerOff/action • Microsoft.Compute/virtualMachines/restart/action • Microsoft.Compute/virtualMachines/deallocate/action • Microsoft.Compute/virtualMachines/upgradeVMAgent/action • Microsoft.Compute/virtualMachines/attachDetachDataDisks/action • Microsoft.Compute/images/read • Microsoft.Compute/images/write • Microsoft.Compute/diskEncryptionSets/read • Microsoft.Compute/diskEncryptionSets/write • Microsoft.Compute/galleries/images/versions/read • Microsoft.ManagedIdentity/userAssignedIdentities/assign/action

Permission Category	Required Permission
Data Actions	• <code>Microsoft.Storage/storageAccounts/blobServices/containers/blobs/read</code> • <code>Microsoft.Storage/storageAccounts/blobServices/containers/blobs/write</code> • <code>Microsoft.Storage/storageAccounts/blobServices/containers/blobs/delete</code>
Assignable Scopes	• <code>/subscriptions/{subscription-id}</code>

Adding an Azure Cloud Account

Add a cloud account to deploy Nutanix Files in the Azure public cloud.

Before you begin

Do the following before you create a cloud account:

- Create a cloud user with the required permissions. For more information, see the "Requirements" and "Limitations" sections in [Files on a Public Cloud](#) on page 21.
- Generate an access key and a secret key for the cloud user.

About this task

To create a cloud account, follow these steps:

Procedure

1. From the side pane in the Files Manager, select **Cloud Accounts > Add Cloud Account**. The **Add Cloud Account** configuration window appears.
2. In the **Cloud Account Name** field, enter a name for the new account.
3. In the **Cloud Provider** dropdown menu, select **Microsoft Azure**.
4. In the **Connect to Microsoft Azure** section, enter the cloud credentials:
 - a. In the **Directory ID** field, enter the Azure cloud tenant ID.
 - b. In the **Subscription ID** field, enter the Azure subscription ID.
 - c. In the **Application ID** field, enter the application ID credentials in the universally unique identifier (UUID) format.
 - d. In the **Secret Value** field, enter the client secret value of the application owner.
5. Click **Add**.
The new cloud account appears in the Cloud Accounts list.

Creating a File Server on the Azure Cloud

Follow the procedure to create a cloud-based file server.

About this task

Cloud-based file servers do not support the following:

- Disaster recovery (DR) through protection domain (PD)-based-DR, including cloning and remote site activation
- Multi-VLAN
- Network segmentation

Note: You cannot configure the name service with a Files Manager deployment. Configure the name service in the Files Console separately after you complete the deployment.

To create a file server on the cloud, follow these steps:

Procedure

1. From the **Files** home page in Prism Central, click **Files Servers > New File Server**.
2. In the **Create a File Server** window, configure the following settings:
 - **Provider:** Select **Microsoft Azure**.
 - **Cloud Account:** Select a cloud account.
For information on cloud account creation, see [Adding an Azure Cloud Account](#) on page 34.
 - **Region:** Select one of the available regions.
 - **Availability Zone:** Select one of the availability zones for the specified region.
 - **File Server Version:** Select the Nutanix Files version for the new file server.
3. Click **Proceed**.
The **File Server Basics** tab appears.
4. Configure the following in the **File Server Details** section:
 - **Name:** Enter a name for the file server.
Clients use this name to access the file server.
 - **DNS Domain:** Enter a unique fully qualified domain.
 - **File Server Size (Logical):** Enter the total file server total storage size, which must be a minimum of 1 TB.

5. In the **Data Tiering and Cost Savings** section, configure data tiering:
 - (Optional) **Enable Automatic Data Tiering:** Clear the checkbox to disable automatic tiering to the Azure Blob storage.
 - **Select the Level of Data Tiering:** From the dropdown menu, select **Lenient**, **Balanced**, or **Aggressive** tiering.

For more information on each tiering level, see [Files on a Public Cloud](#) on page 21.
 - (Optional) **Use existing bucket for data tiering:** Select the checkbox to specify an existing storage account and container for tiering operations and follow these steps:
 1. In the **Storage Account ID** field, specify the name of an existing storage account to use as a tiering target for cold data.
 2. In the **Container Name** field, specify the name of the container associated with the storage account.If you do not specify an existing bucket, Azure automatically creates a new one.
6. In the **Sizing Configuration** section, configure the following:
 - **Number of File Server VMs (FSVMs):** Specify the number of FSVMs by clicking - or +, or by entering an integer value.
 - **FSVM Size:** Select a combination of vCPUs and memory resources per FSVM.
7. Click **Next**.

The **Client Network** tab appears.
8. In the **Client Network** tab, configure the following settings:
 - **Virtual Network:** Select the Azure virtual network in which to deploy the file server.
 - **Subnet:** Select a subnet from the virtual network.
9. Click **Next**.

The **Settings** tab appears.
10. In the **Settings** tab, configure the following settings:
 - **DNS Servers:** Enter the domain name system server IP address or a comma separated list of IP addresses.
 - **NTP Servers:** Enter the network time protocol server IP address or a comma separated list of IP addresses.
 - (Optional) **SSH Key Pair:** Select a key pair from the dropdown list.
 - (Optional) **Key Vault ID:** Specify the name of the key vault ID.
 - **Azure Managed Identity:** Specify the full ID of the managed identity that has been assigned the user-assigned role.
 - (Optional) **Disk Encryption Set ID:** Specify disk encryption set ID to use for disk encryption.
 - **Resource Group:** Specify the name of the Azure resource group for the deployment of all resources.
11. Click **Next**
12. In the **Summary** tab, review the file server information and select **Deploy**

File server creation begins. You can monitor the progress through the **Tasks** view.

FILES IN A VIRTUAL PRIVATE CLOUD (VPC)

Nutanix Files in a VPC runs on Microsoft Azure BareMetal Infrastructure or in your on-premises environment.

A VPC lets you have an isolated networking stack within a larger network. Configure Files in a VPC using Flow Virtual Networking.

For more information on VPC concepts, see [Essential Concepts](#) in the *Flow Virtual Networking Guide*.

VPC Requirements

Set up the VPC using Flow Virtual Networking. For VPC setup instructions, see the [Flow Networking User Guide](#).

Ensure that the VPC for Files meets the following requirements:

- For Files on Azure, have at least one associated external NAT subnet with at least one available floating IP address.
- For Files on-premises, have a NAT or no-NAT subnet.
External NAT subnets require at least one floating IP address.
- Have at least one overlay subnet on the VPC.

Note: To view overlay network details, you must have the VPC admin role.

- Have a default route for floating IP address to Controller VM (CVM) reachability.

File Server Creation in a VPC

After setting up a VPC using Flow Virtual Networking, create a file server in Prism Central using the subnets of the dedicated VPC. For information on creating a file server in Prism Central, see [Creating a File Server in Prism Central](#) on page 9. After file server deployment, initial share creation takes longer than usual (about eight minutes with a three-FSVM file server).

To access file server VMs (FSVMs) in a NAT environment for nCLI use or debugging purposes, use floating IP addresses instead of SSH.

Deploying Nutanix Files on Azure BareMetal Infrastructure

Deploy Nutanix Files in a virtual private cloud (VPC) with Nutanix Cloud Clusters (NC2) on Microsoft Azure BareMetal Infrastructure.

To deploy Nutanix Files in a VPC on Azure BareMetal Infrastructure, follow these steps:

1. Set up Nutanix Cloud Clusters (NC2) on Azure.

For more information, see the [Nutanix Cloud Clusters on Azure Deployment and User Guide](#).

Note: Ensure that you perform the required VNet peering when using the vWAN hub.

2. Use Nutanix Flow to configure a VPC that meets all of the requirements for supporting Files.

For more information, see [Files in a Virtual Private Cloud \(VPC\)](#) on page 37. For VPC setup instructions, see the [Flow Networking User Guide](#).

3. Create a file server in Prism Central using the overlay subnets of the dedicated VPC.

The limitations of Prism Central–created file servers apply to Files on Azure.

For information on file server creation steps, see [Creating a File Server in Prism Central](#) on page 9 in the [Files Manager User Guide](#).

Note: After file server deployment, initial share creation takes longer than usual (about eight minutes with a three-FSVM file server).

Assigning Floating IPs

Manage the floating IPs of a file server in a virtual private cloud (VPC).

About this task

If the client network of your VPC has a network address translation (NAT) configuration, add floating IPs to the file server before setting up features that rely on external clients reaching your VPC, such as Smart Sync and Smart DR. Similarly, after updating a file server with Smart Sync or Smart DR policies from a no-NAT to a NAT configuration, assign floating IPs to resume operations.

Procedure

1. From the Files home page in Prism Central, go to the **Files Servers** view.
2. In the row of the target file server, select the checkbox.
3. Click **Actions > Manage Floating IPs**.
4. Select the **Assign Floating IPs** checkbox.
5. Click **Save**.

RBAC FOR NUTANIX FILES

Role-based access control (RBAC) lets a Prism Central local user configure different levels of access to non-admin users based on their role within an organization.

By default, Files denies access to non-admin users unless they are explicitly given access through RBAC.

Prism Central users have the following roles:

- An *administrator* is a super admin, also known as a local user or a Prism admin in Prism Central. The super admin can perform all Files operations on all Files instances.
- A *non-admin user* is a Prism Central user without any administrator privileges.

Tip: For more information on RBAC in Prism Central, see the [Roles Management](#) chapter in the *Security Guide*.

Super admin users can define and assign the following roles to users or user groups:

- Predefined *built-in roles*:
 - The *Files Admin* role grants the permissions required to enable, create, and manage file servers.
 - The *Files Viewer* role grants the permissions required to perform read-only operations on file servers.
 - The *File Server Security Admin* role grants the permissions required for security operations like antivirus settings, ransomware detection, and virus-scan policy configuration.
 - The *File Server Share Admin* role grants the permissions required for performing share related operations, such as create, read, update, and delete (CRUD).
- File server *custom roles*

For more information on operations available for each role, see [Built-In Role Permissions for Nutanix Files](#) on page 40 and [Custom Role Permissions for Nutanix Files](#) on page 59.

Nutanix Files supports chaining permissions. For example, if you assign an *update* permission to a user, the system also assigns the *view* permission. If you have security permissions, the system assigns the *view share* permission by default.

When you use the Files Console from Prism Central, the permissions also apply. Only a super admin, a cluster admin, or users that have Prism admin or Files admin roles can launch the Files Console from Prism Central.

RBAC Requirements and Limitations for Nutanix Files

Before you configure role-based access control (RBAC) for Nutanix Files, review the requirements and limitations.

Requirements

Only a super admin can enable Files in Prism Central and create roles and access control policies for Files. For more information on the built-in roles in Prism Central, see [Roles Management](#) in the *Security Guide*.

Before configuring RBAC, enable microservices infrastructure (IAMv2) on Prism Central. For more information, see [Enabling Microservices Infrastructure](#) in the *Prism Central infrastructure Guide*.

Limitations

The following limitations apply to RBAC for Files:

- Access control policies apply only to Prism Central users and not to Files users.

- Using RBAC features requires accessing Files from Prism Central.
- Prism Central caches access control policies with a time-to-live (TTL) of five minutes. Any modification to an existing access control policy usually applies within five minutes.

Managing RBAC for Nutanix Files

In Prism Central, manage role-based access control (RBAC) for file servers.

About this task

Only users with administrator privileges can assign roles.

To manage role assignment and authorization, follow these steps:

Procedure

1. Log on to Prism Central as an administrator.

2. Review the available built-in Files roles or create a custom file server role.

For more information, see [Built-in Roles](#) and [Creating a Custom Role](#) in the *AOS Security Guide*.

For information on operations associated with different roles, see [Built-In Role Permissions for Nutanix Files](#) on page 40 and [Custom Role Permissions for Nutanix Files](#) on page 59.

Tip: When creating a custom role in Prism Central, Nutanix Files operations are listed under the **Files** entity.

3. Create an authorization policy for the role.

The authorization policy defines the scope of entities and identities (users and groups) that pertain to the role. For more information, see [Authorization Policies](#) in the *AOS Security Guide*.

Built-In Role Permissions for Nutanix Files

Note the built-in role permission requirements to perform Files operations using the Files Manager and the Files Console.

As a non-admin user, you can perform file server-related tasks after an administrator creates an authorization policy that assigns a built-in or custom role to you or your group.

For more information on the built-in roles in Prism Central, see [Roles Management](#) in the *Security Guide*.

The following table lists all of the permissions that an administrator can grant to a non-admin user to perform various file-server tasks.

Table 5: Files Operation Permission Requirements

Feature	Files Manager Operation Permission	Required Role
Antivirus	View_File_Server_Antivirus_Server	- Files Admin - Files Viewer - File Server Share Admin - File Server Security Admin - Prism Admin - Prism Viewer - Internal Super Admin - Super Admin
	Create_File_Server_Antivirus_Server	- Files Admin - Prism Admin - File Server Security Admin - Internal Super Admin - Super Admin
	Update_File_Server_Antivirus_Server	- Files Admin - Prism Admin - File Server Security Admin - Internal Super Admin - Super Admin
	Delete_File_Server_Antivirus_Server	- Files Admin - Prism Admin - File Server Security Admin - Internal Super Admin - Super Admin
	Test_File_Server_Antivirus_Server_Connection	- Files Admin - Prism Admin - File Server Security Admin - Internal Super Admin - Super Admin

Feature	Files Manager Operation Permission	Required Role
	View_File_Server_Share_Antivirus_Stats	- Files Admin - Files Viewer - Prism Admin - Prism Viewer - File Server Share Admin - File Server Security Admin - Internal Super Admin - Super Admin
	View_File_Server_Virus_Scan_Policy	- Files Admin - Files Viewer - Prism Admin - Prism Viewer - File Server Security Admin - Internal Super Admin - Super Admin
	Create_File_Server_Virus_Scan_Policy	- Files Admin - Prism Admin - File Server Security Admin - Internal Super Admin - Super Admin
	Update_File_Server_Virus_Scan_Policy	- Files Admin - Prism Admin - File Server Security Admin - Internal Super Admin - Super Admin
	Delete_File_Server_Virus_Scan_Policy	- Files Admin - Prism Admin - File Server Share Admin - Internal Super Admin - Super Admin

Feature	Files Manager Operation Permission	Required Role
	View_File_Server_AntiVirus_Infected_Files	- Files Admin - Files Viewer - Prism Admin - Prism Viewer - File Server Security Admin - Internal Super Admin - Super Admin
	Remove_File_Server_Antivirus_Infected_Files	- Files Admin - Prism Admin - File Server Security Admin - Internal Super Admin - Super Admin
	Fix_File_Server_Antivirus_Infected_File	- Files Admin - Prism Admin - File Server Security Admin - Internal Super Admin - Super Admin
	View_File_Server_Data_Protection_Policy	- Files Admin - Files Viewer - Prism Admin - Prism Viewer - Internal Super Admin - Super Admin
	Create_File_Server_Data_Protection_Policy	- Files Admin - Prism Admin - Internal Super Admin - Super Admin
Data protection		

Feature	Files Manager Operation Permission	Required Role
	Update_File_Server_Data_Protection_Policy	- Files Admin - Prism Admin - Internal Super Admin - Super Admin
	Delete_File_Server_Data_Protection_Policy	- Files Admin - Prism Admin - Internal Super Admin - Super Admin
	Failover_File_Server_Data_Protection_Policy	- Files Admin - Prism Admin - Internal Super Admin - Super Admin
	Add_File_Server_Dns_Failover_Data_Protection_Policy	- Files Admin - Prism Admin - Internal Super Admin - Super Admin
	File_Server_Resume_Replication_Data_Replication_Policy	- Files Admin - Prism Admin - Internal Super Admin - Super Admin
	View_File_Server_Data_Protection_Replication_Jobs	- Files Admin - Files Viewer - Prism Admin - Prism Viewer - Internal Super Admin - Super Admin

Feature	Files Manager Operation Permission	Required Role
DNS configuration	View_File_Server_Dns_Entries	- Files Admin - Files Viewer - Prism Admin - Prism Viewer - Internal Super Admin - Super Admin
	Update_File_Server_Dns_Entries	- Files Admin - Prism Admin - Internal Super Admin - File Server Security Admin - Super Admin
	Verify_File_Server_Dns_Entries	- Files Admin - Prism Admin - File Server Security Admin - Internal Super Admin - Super Admin
File server management	View_File_Server_Stats	- Files Admin - Files Viewer - Prism Admin - Prism Viewer - Internal Super Admin - Super Admin
	View_File_Server	- Files Admin - Files Viewer - Prism Admin - Prism Viewer - Internal Super Admin - Super Admin - File Server Share Admin - File Server Security Admin

Feature	Files Manager Operation Permission	Required Role
Migration	View_File_Server_Migration_Job	- Files Admin - Files Viewer - Prism Admin - Prism Viewer - Internal Super Admin - Super Admin
	View_File_Server_Migration_Source	- Files Admin - Files Viewer - Prism Admin - Prism Viewer - Internal Super Admin - Super Admin
	Create_File_Server_Migration_Source	- Files Admin - Prism Admin - Internal Super Admin - Super Admin
	Update_File_Server_Migration_Source	- Files Admin - Prism Admin - Internal Super Admin - Super Admin
	Delete_File_Server_Migration_Source	- Files Admin - Prism Admin - Internal Super Admin - Super Admin

Feature	Files Manager Operation Permission	Required Role
	View_Source_File_Server_Mount_Target	- Files Admin - Files Viewer - Prism Admin - Prism Viewer - Internal Super Admin - Super Admin
	Scan_Source_File_Server_Mount_Target	- Files Admin - Prism Admin - Internal Super Admin - Super Admin
	View_File_Server_Migration_Plan	- Files Admin - Files Viewer - Prism Admin - Prism Viewer - Internal Super Admin - Super Admin
	Create_File_Server_Migration_Plan	- Files Admin - Prism Admin - Internal Super Admin - Super Admin
	Update_File_Server_Migration_Plan	- Files Admin - Prism Admin - Internal Super Admin - Super Admin
	Delete_File_Server_Migration_Plan	- Files Admin - Prism Admin - Internal Super Admin - Super Admin

Feature	Files Manager Operation Permission	Required Role
	View_File_Server_Migration_Sub_Plan	- Files Admin - Files Viewer - Prism Admin - Prism Viewer - Internal Super Admin - Super Admin
	Create_File_Server_Migration_Sub_Plan	- Files Admin - Prism Admin - Internal Super Admin - Super Admin
	Update_File_Server_Migration_Sub_Plan	- Files Admin - Prism Admin - Internal Super Admin - Super Admin
	Delete_File_Server_Migration_Sub_Plan	- Files Admin - Prism Admin - Internal Super Admin - Super Admin
Notifications	View_File_Server_Notification_Policy	- Files Admin - Files Viewer - Prism Admin - Prism Viewer - Internal Super Admin - Super Admin
	Create_File_Server_Notification_Policy	- Files Admin - Prism Admin - Internal Super Admin - Super Admin

Feature	Files Manager Operation Permission	Required Role
Partner server	Update_File_Server_Notification_Policy	- Files Admin - Prism Admin - Internal Super Admin - Super Admin
	Delete_File_Server_Notification_Policy	- Files Admin - Prism Admin - Internal Super Admin - Super Admin
	View_File_Server_Partner_Server	- Files Admin - Files Viewer - Prism Admin - Prism Viewer - Internal Super Admin - Super Admin
	Create_File_Server_Partner_Server	- Files Admin - Prism Admin - Internal Super Admin - Super Admin
	Update_File_Server_Partner_Server	- Files Admin - Prism Admin - Internal Super Admin - Super Admin
	Delete_File_Server_Partner_Server	- Files Admin - Prism Admin - Internal Super Admin - Super Admin

Feature	Files Manager Operation Permission	Required Role
Quota	View_Quota_Policy_File_Server_Share	- Files Admin - Files Viewer - Prism Admin - Prism Viewer - File Server Share Admin - Internal Super Admin - Super Admin
	Create_Quota_Policy_File_Server_Share	- Files Admin - Prism Admin - File Server Share Admin - Internal Super Admin - Super Admin
	Update_Quota_Policy_File_Server_Share	- Files Admin - Prism Admin - File Server Share Admin - Internal Super Admin - Super Admin
	Delete_Quota_Policy_File_Server_Share	- Files Admin - Prism Admin - File Server Share Admin - Internal Super Admin - Super Admin
	View_Quota_Email_config	- Files Admin - Files Viewer - Prism Admin - Prism Viewer - File Server Security Admin - Internal Super Admin - Super Admin

Feature	Files Manager Operation Permission	Required Role
	Update_Quota_Email_config	- Files Admin - Prism Admin - File Server Security Admin - Internal Super Admin - Super Admin
Ransomware protection	View_File_Server_Ransomware_Config	- Files Admin - Files Viewer - Prism Admin - Prism Viewer - File Server Security Admin - Internal Super Admin - Super Admin
	Create_File_Server_Ransomware_Config	- Files Admin - Prism Admin - File Server Security Admin - Internal Super Admin - Super Admin
	Update_File_Server_Ransomware_Config	- Files Admin - Prism Admin - File Server Security Admin - Internal Super Admin - Super Admin
Recommendations	View_File_Server_Recommendations	- Files Admin - Files Viewer - Prism Admin - Prism Viewer - Internal Super Admin - Super Admin

Feature	Files Manager Operation Permission	Required Role
Self-service restore	Clear_File_Server_Recommendations	- Files Admin - Prism Admin - Internal Super Admin - Super Admin
	View_Snapshot_Changed_Contents_File_Server_Share	- Files Admin - Files Viewer - Prism Admin - Prism Viewer - File Server Share Admin - Internal Super Admin - Super Admin
	Compute_Snapshot_Change_File_Server_Share	- Files Admin - Prism Admin - File Server Share Admin - Internal Super Admin - Super Admin
	View_File_Server_Self_Service_Restore_Snapshot_Schedules	- Files Admin - Files Viewer - Prism Admin - Prism Viewer - Internal Super Admin - Super Admin
	Create_File_Server_Self_Service_Restore_Snapshot_Schedules	- Files Admin - Prism Admin - Internal Super Admin - Super Admin

Feature	Files Manager Operation Permission	Required Role
	Update_File_Server_Self_Service_Restore_Snapshot_Schedules	- Files Admin - Prism Admin - Internal Super Admin - Super Admin
	Delete_File_Server_Self_Service_Restore_Snapshot_Schedules	- Files Admin - Prism Admin - Internal Super Admin - Super Admin
	View_Snapshot_File_server_Share	- Files Admin - Files Viewer - Prism Admin - Prism Viewer - File Server Share Admin - File Server Security Admin - Internal Super Admin - Super Admin - CSI System
	reate_File_Server_Share_Snapshot	- Files Admin - File Server Share Admin - File Server Security Admin - Prism Admin - Internal Super Admin - Super Admin - CSI System

Feature	Files Manager Operation Permission	Required Role
Share management	Delete_Snapshot_File_server_Share	- Files Admin - File Server Share Admin - File Server Security Admin - Prism Admin - Internal Super Admin - Super Admin - CSI System
	Restore_File_Server_Share	- Files Admin - File Server Share Admin - Prism Admin - Internal Super Admin - Super Admin - CSI System
	View_File_Server_Share	- Files Admin - Files Viewer - Prism Admin - Prism Viewer - File Server Share Admin - File Server Security Admin - Internal Super Admin - Super Admin
	Create_File_Server_Share	- Files Admin - Prism Admin - File Server Share Admin - Internal Super Admin - Super Admin - CSI System

Feature	Files Manager Operation Permission	Required Role
	Update_File_Server_Share	- Files Admin - Prism Admin - File Server Share Admin - Internal Super Admin - Super Admin - CSI System
	Delete_File_Server_Share	- Files Admin - Prism Admin - Internal Super Admin - Super Admin - CSI System
	Enable_File_Server_Share_Worm_Compliance	- Files Admin - File Server Share Admin - File Server Security Admin - Prism Admin - Internal Super Admin - Super Admin
	Enable_File_Server_Share_Worm_Legal_Hold	- Files Admin - Prism Admin - Internal Super Admin - Super Admin - File Server Share Admin - File Server Security Admin
	Disable_File_Server_Share_Worm_Legal_Hold	- Files Admin - Prism Admin - Internal Super Admin - Super Admin - File Server Share Admin - File Server Security Admin

Feature	Files Manager Operation Permission	Required Role
Tiering	Clone_File_Server_Share	- Files Admin - File Server Share Admin - Prism Admin - Internal Super Admin - Super Admin - CSI System
	View_File_Server_Object_Store_Profile	- Files Admin - Files Viewer - Prism Admin - Prism Viewer - Internal Super Admin - Super Admin
	Create_File_Server_Object_Store_Profile	- Files Admin - Prism Admin - Internal Super Admin - Super Admin
	Update_File_Server_Object_Store_Profile	- Files Admin - Prism Admin - Internal Super Admin - Super Admin
	View_File_Server_Tiering_Config	- Files Admin - Files Viewer - Prism Admin Files Admin - Prism Admin - Internal Super Admin - Super Admin - Prism Viewer - Internal Super Admin - Super Admin

Feature	Files Manager Operation Permission	Required Role
	Create_File_Server_Tiering_Config	- Files Admin - Prism Admin - Internal Super Admin - Super Admin
	Update_File_Server_Tiering_Config	- Files Admin - Prism Admin - Internal Super Admin - Super Admin
	Delete_File_Server_Tiering_Config	- Files Admin - Prism Admin - Internal Super Admin - Super Admin
	Tier_File_Server_Data	- Files Admin - Prism Admin - Internal Super Admin - Super Admin
	Refresh_File_Server_Cold_Data	- Files Admin - Prism Admin - Internal Super Admin - Super Admin
	Recall_File_Server_Policy	- Files Admin - Prism Admin - Internal Super Admin - Super Admin

Feature	Files Manager Operation Permission	Required Role
Unified namespace	View_File_Server_Unified_Namespace	- Files Admin - Files Viewer - Prism Admin - Prism Viewer - Internal Super Admin - Super Admin
	Create_File_Server_Unified_Namespace	- Files Admin - Prism Admin - Internal Super Admin - Super Admin
	Update_File_Server_Unified_Namespace	- Files Admin - Prism Admin - Internal Super Admin - Super Admin
	Delete_File_Server_Unified_Namespace	- Files Admin - Prism Admin - Internal Super Admin - Super Admin
User mapping	Update_File_Server_User_Mapping	- Files Admin - Prism Admin - Internal Super Admin - Super Admin
VDI synchronization	View_File_Server_Home_Dir_Sync_Policy_Impacted_Users	- Files Admin - Files Viewer - Prism Admin - Prism Viewer - Internal Super Admin - Super Admin

Feature	Files Manager Operation Permission	Required Role
	Update_Primary_File_Server_Home_Dir_Sync_User	- Files Admin - Prism Admin - Internal Super Admin - Super Admin

Custom Role Permissions for Nutanix Files

Note the available permissions for file server custom roles configured in Prism Central.

With custom role creation, you can assign permissions to a role for the following operations:

- Activate File Server Data Protection Policy
- Add File Server DNS Entries
- Add File Server DNS Failover Data Protection Policy
- Clear File Server Recommendations
- Configure File Server Name Services
- Create File Server
- Create File Server Admin User
- Create File Server Antivirus Server
- Create File Server Data Consolidation Policy
- Create File Server Data Protection Policy
- Create File Server Local User
- Create File Server Object Store Profile
- Create File Server Self Service Restore Policy
- Create File Server Share
- Create File Server Share Virus Scan Policy
- Create File Server User Mapping
- Create File Server Virus Scan Policy
- Create Quota Policy File Server Share
- Deactivate File Server Data Protection Policy
- Delete File Server
- Delete File Server Admin User
- Delete File Server Antivirus Server
- Delete File Server Data Consolidation Policy
- Delete File Server Data Protection Policy
- Delete File Server Local User

- Delete File Server Object Store Profile
- Delete File Server Self Service Restore Policy
- Delete File Server Share
- Delete Quota Policy File Server Share
- Delete Snapshot File Server Share
- Disable File Server Data Consolidation Policy
- Disable File Server Data Protection Policy
- Enable File Server Data Consolidation Policy
- Enable File Server Data Protection Policy
- Failover File Server Data Protection Policy
- Perform File Server Recommendations
- Quarantine File Server Antivirus Infected Files
- Recall File Server Policy
- Remove File Server Antivirus Infected Files
- Remove File Server DNS Entries
- Rescan File Server Antivirus Infected Files
- Reset File Server Antivirus Infected Files
- Search File Server User Mapping
- Show File Analytics
- Tier File Server Policy
- UnQuarantine File Server Antivirus Infected Files
- Update Alert
- Update Alerts Policy
- Update File Server
- Update File Server Admin User
- Update File Server Antivirus Server
- Update File Server Data Consolidation Policy
- Update File Server Data Protection Policy
- Update File Server Local User
- Update File Server Object Store Profile
- Update File Server Ransomware Config
- Update File Server Self Service Restore Policy
- Update File Server Share
- Update File Server Share Virus Scan Policy

- Update File Server User Mapping
- Update File Server Virus Scan Policy
- Update Quota Policy File Server Share
- Verify File Server DNS Entries
- View Alert
- View Alerts Policy
- View Cluster
- View File Server
- View File Server Admin User
- View File Server AntiVirus Infected Files
- View File Server Antivirus Server
- View File Server Data Consolidation Policy
- View File Server Data Protected Pairs
- View File Server Data Protection Policy
- View File Server Data Protection Replication Jobs
- View File Server DNS Entries
- View File Server Local User
- View File Server Name Services
- View File Server Object Store Profile
- View File Server Ransomware Config
- View File Server Recommendations
- View File Server Self Service Restore Policy
- View File Server Share
- View File Server Share Virus Scan Policy
- View File Server User Mapping
- View File Server Virus Scan Policy
- View Quota Policy File Server Share
- View Snapshot File Server Share
- View Subnet

DATA REPLICATION

Nutanix Files offers data replication features that let you replicate data from multiple file servers and sync distributed-share data across locations.

Smart DR facilitates share-level data replication and file server-level disaster recovery, Smart Sync provides a cost-effective file-level data replication and consolidation solution, and VDI Sync lets you sync user profiles across file servers in different locations.

For more information on each feature, see the following topics:

- [Asynchronous Data Protection with Smart DR](#) on page 64
- [Data Synchronization with Smart Sync](#) on page 83
- [VDI Sync](#) on page 89

Data Replication Navigation

The **Data Replication** tab in the Files Manager provides an interface for replication policy creation, management, and monitoring.

The **Data Replication** tab includes the following subtabs: **Monitor**, **Replication Jobs**, and **Policies**.

Monitor

The monitor tab includes the **Overview** and **Disaster Recovery** subtabs.

The **Overview** provides a visual diagram of policies that involve the file servers registered to the local Prism Central cluster. The **Remote** tag appears next to file servers from remote Prism Central clusters.

The **Disaster Recovery** subtab lists existing Smart DR data protection policies and provides policy management options. Clicking a file server name in the **Primary Location** column displays a summary of the protection policy, including the protected shares.




Figure 1: Protected File Server Details

The **Actions** column displays disaster recovery options, which vary depending on the file server and replication job states. When the remote file server is on a remote Prism Central cluster, the **Actions** column displays an action to recovery AZ to perform the failover and failback DR operations.

For more information on Smart DR, see the [Asynchronous Data Protection with Smart DR](#) on page 64 chapter.

Policies

The **Policies** view displays a table that lists policies created on or interacting with file servers on the local Prism Central cluster. Clicking a policy name in the **Policies** view provides additional policy details organized into **Summary**, **Shares**, and **Replication Jobs** subtabs; VDI synchronization policies include an additional **User Connection Failures** subtab. The subtabs provide additional details, such as share replication history and user replication failures for VDI synchronization.

Clicking a share in the **Shares** tab displays the **Share Details** view. The **Share Details** view displays a **Replication Job** table filtered for jobs of the specified share. The **Share Details** view also includes additional details such as schedule compliance, which indicates replication frequency compliance, and the last successful recovery point.

Replication Jobs

The **Replications Jobs** view includes a table that lists replication jobs for policies that interact with file servers on the local Prism Central cluster. The total estimated files in the **Replication Jobs** details is an approximation and might not reflect the total number of synced files.

Clicking one of the start times displays the **Job Details** view, which displays further details. *Data synced* refers to the quantity of data synced during the replication job. *Network throughput (avg)* refers to the quantity of data synced during the replication job.

ASYNCHRONOUS DATA PROTECTION WITH SMART DR

Smart DR is a Nutanix Files disaster recovery and data protection service.

Smart DR facilitates share-level data replication and file server-level disaster recovery. In the event of a planned or unplanned loss of service, you can restore write access to protected shares by failing over to a recovery site file server. Protection policies indicate failover details, including the primary location, recovery location, and replication schedule. Rather than having a single protection policy for an entire file server, Smart DR lets you configure unique policies for different shares.

A short recovery time objective (RTO) helps ensure continuous availability of data after you fail over to a recovery site. Configuring Active Directory (AD) and domain name system (DNS) entries facilitates seamless client access redirection.

Smart DR Process Overview

After policy creation, file servers are in the *Active-Standby* state.

Planned failovers cause the recovery site to become active, so the file server state changes to *Standby-Active*. The planned failover workflow includes the following operations:

- If AD and DNS are configured on the recovery site, Files redirects clients to the recovery file server. If AD and DNS *are not* configured on the recovery site, Files does not perform client redirection.
- If a reverse replication policy is configured, you can failback to the primary site. Otherwise, you must configure a reverse replication policy before you can perform a failback. After the failback is complete, the file servers are in the *Active-Standby* state again.

The unplanned failover workflow consists of the following operations:

- Unplanned failovers cause the recovery site to become active, so the file server state changes to *Active-Active*.
 - If AD and DNS are configured on the recovery site, Files redirects clients to the recovery file server. If AD and DNS are not configured on the recovery site, Files does not perform client redirection.
- After replication can resume, the workflow varies depending on the direction of replication.
 - Resuming replication from the recovery site changes the file server state to *Standby-Active*. If a reverse replication policy is configured, you can failback to the primary site; otherwise, you have to configure the reverse replication policy before you can perform a failback. After the failback is complete, the file servers are once again in the *Active-Standby* state.
 - Resuming replication from the primary site causes the file servers to be in the *Active-Standby* state again.

For more information on file server Smart DR states, see [Concepts and Terminology](#) on page 66.

To manage the Smart DR workflow, review how to do the following:

- Configure a protection policy to replicate share data to a recovery site. For more information, see [Asynchronous Replication Policies](#) on page 67.
- Perform a planned or unplanned failover. For more information, see [Smart DR Failover](#) on page 71.
- Configure AD and DNS access for seamless client failover. For more information, see [Updating AD and DNS Entries](#) on page 70.
- Configure a reverse replication policy. For more information, see [Creating a Reverse Replication Policy](#) on page 75.

- Fail back to the primary site. For more information, see [Smart DR Failback](#) on page 76.

Multi-PC Smart DR

You can configure Smart DR policies between file servers registered to different Prism Central (PC) clusters. However, both Prism Central clusters must be in an availability zone (AZ) pair. For more information, see "[Pairing Availability Zones](#)" in the *Nutanix Disaster Recovery Guide*.

The following workflow differences apply to multi-PC Smart DR:

- You can only perform policy-related operations and monitoring on the Prism Central cluster where you created the policy.
- You can only perform Smart DR operations on the recovery site Prism Central cluster.
- In a Metro Availability environment, multi-PC Smart DR requires Metro Availability-enabled file servers to fail back to the original Prism Central instance to resume normal functionality.

Requirements

Meet the following requirements to deploy Smart DR:

- The Files Manager must have at least two file servers.

Note: Nutanix recommends using file servers on different Prism Element clusters to ensure DR after a cluster-wide outage.

- The number of FSVMs between the source and recovery sites must match.

Note: After protection policy creation, perform scale-out on the recovery file server before performing it on the primary file server.

- The primary and recovery file servers must support the same protocols.
- For AD, the primary and recovery file servers must have the same domain name.
- Meet the Files port requirements. For more information, see [Port Reference](#).
- Meet the compatibility requirements. For more information, see [Nutanix Files Release Notes](#).

Limitations

Smart DR has the following limitations:

- You cannot replicate the same share under different replication policies.
- You cannot replicate share properties modified after policy creation.
- Smart DR supports only one-to-one replication for shares.
- Smart DR does not provide automatic replication of connected shares. Manually add connected shares to the replication policy and submount the shares on the remote site. For more information, see [Connected Shares](#) in the *Nutanix Files Guide*.
- Smart DR does not support nested share replication. You can configure nested shares on the recovery site. For more information, see [Nested Shares](#) in the *Nutanix Files Guide*.
- Smart DR does not support quota policy replication.
- Even though you can replicate snapshots from the primary site, Smart DR does not support self-service restore (SSR) snapshot creation on the recovery site.

- Adding or removing FSVMs creates volume group discrepancies between the primary and recovery sites and results in replication failure.
- SSR snapshot creation and replication might be unsuccessful when shares experience volume group (VG) issues.
- After you delete a share, Files does not retain a history of replication jobs from that share.
- The minimum replication frequency is 1 minute.

Cloud-based Smart DR has the following additional limitations:

- You cannot replicate self-service restore (SSR) snapshots.
- You cannot perform seamless client access redirection and failover operations.
- Replication skips files with carriage return characters (\r and \n) in the file name.
- Replication job speed depends on the number of files transferred and the predicted total file count.
- Cloud-based Smart DR supports share-root-to-share-root replication only.
- Cloud-based Smart DR does not support archival and exclude pattern features.
- Cloud-based Smart DR is not compatible with Data Consolidation and VDI Sync.
- Standard Smart Tiering is not supported when using asynchronous replication with Smart DR to replicate files on the public cloud.

For information on the interoperability of Files Manager features, see [Feature Interoperability](#) on page 19.

Concepts and Terminology

This topic describes Smart DR data replication concepts and terminology.

The following terms are core to Smart DR concepts:

- *Availability zone (AZ)*: A physically isolated location. An availability zone indicates a unique Prism Central instance.
- *Failback*: The process for migrating operations back to the primary file server site. For more information, see [Smart DR Failback](#) on page 76.
- *Failover (planned)*: The process of migrating operations to a recovery file server for a planned event, such as routine maintenance. For more information, see [Smart DR Failover](#) on page 71.
- *Failover (unplanned)*: The process for migrating operations to a recovery file server site after an unplanned failover if a disaster occurs. For more information, see [Smart DR Failover](#) on page 71.
- *Protection policy*: The protection policy provides data protection instructions for a failover. For more information, see [Asynchronous Replication Policies](#) on page 67.
- *Primary file server*: The original file server location of the protected shares.
- *Recovery file server*: The original remote site for the protected shares.
- *Resume replication*: After an unplanned failover, when the primary file server becomes active, the sites are in an active-active state and replication stops. Resume replication to establish the direction of replication. For more information, see [Resuming Replication](#) on page 73.
- *Replication frequency*: The interval of time between replication jobs. The minimum replication frequency is 1 minute.
- *Recovery time objective (RTO)*: The amount of time required to complete a failover.

- *Replication frequency compliance*: Indicates the completion of a replication job within the replication frequency time configured for the protection policy. Alternatively, a replication frequency non-compliant replication job indicates incomplete replication during the specified replication frequency time.

The following terms refer to file server states under Smart DR policies:

- *Active-Standby*: Refers to the replication relationship between the primary and recovery file servers. Indicates that the primary file server is in read-write mode and the recovery site is in read-only mode. Replication occurs from the active file server to the standby file server.
- *Active-Active*: Refers to the replication relationship between the primary and recovery file servers after an unplanned failover. Indicates that replication stopped and both file servers are in a read-write mode.
- *Standby-Active*: Refers to the replication relationship between the primary and recovery file servers. Indicates that the primary file server is in read-only mode and the recovery site is in read-write mode. Replication occurs from the active file server to the standby file server.

Asynchronous Replication Policies

Asynchronous replication policies use Smart DR data protection.

A protection policy provides data protection instructions for a failover. Creating a protection policy requires choosing a primary location, also known as the data origin, configuring a replication schedule, and choosing a recovery location. For more information, see [Creating an Asynchronous Protection Policy](#) on page 67.

After protection policy creation begins, the task appears in the Prism Central **Tasks** view, and the new policy appears in the Files manager **Policies** view. The **Policies** view provides details about each policy created on the Prism Central instance, such as the direction of replication and the policy state. For more information on the **Policies** view, see [Data Replication Navigation](#) on page 62.

When Files protects a share for the first time, a share is created on the recovery site with an analogous configuration. If the source share permissions include one or more local users, you can create a policy only after the same local users exist on the target. This means that an admin must manually add the same local users to the target share.

Creating an Asynchronous Protection Policy

Create a Smart disaster recovery (DR) protection policy for a share or export.

About this task

Nutanix Files automatically creates recovery site shares after you create a protection policy. Files begins space consumption on the recovery site after initiating replication.

Note: Creating and updating replication policies requires the view availability zone permission for custom roles.

Procedure

1. From the side pane in the Files Manager, go to **Data Replication > Policies**.
2. Click **+ New Replication Policy**.
3. From the data replication use-case options, select **Disaster Recovery**.
4. Under **Select Replication Type**, select **Asynchronous Replication**.
5. Click **Next**.
The **Create a Data Replication Policy** page displays.

6. In the **Replication Topology** tab, in the **Primary Location (Source File Server)** pane, configure the following settings:
 - a. In the **File Server** dropdown menu, select the source file server.
 - b. Under **Shares to be protected**, click **Edit**.
The **Select Specific Shares** dialog box displays. By default, the Files Manager adds all shares to the policy.
 - c. (Optional) Clear the checkboxes next to the names of shares to exclude from the policy.
To show all shares, select the **Show All Shares** box. To search by name, enter the name of the share in the **Search By Share Names** field.
 - d. (Optional) Select the **Add newly created shares on this file server to this policy** checkbox.
 - e. Click **Done**.
7. Do the following in the **Sync Schedule** pane:
 - a. In **Sync Frequency** field, enter an integer to indicate the replication interval.
 - b. In the **Sync Frequency** dropdown, specify the replication interval type by selecting **minutes**, **hours**, or **days**.
 - c. Select either the **Start Immediately** or **Start from a specific point in time** option.
 - d. If you choose the **Start Immediately** option, skip to Step 7.
 - e. Under **Start time**, specify when to begin replication.
 - f. Under **time zone**, specify the time zone for the start time and the replication schedule.
8. In the **Primary Location (Source File Server)** pane, click **Add**.
9. In the **Recovery Location (Target File Server)** pane, configure the following settings:
 - a. In the **Availability Zone** dropdown menu, choose the availability zone of the recovery file server.
To use the local Prism Central cluster, select Local AZ.
 - b. In the **File Server** dropdown menu, choose a recovery file server.
Incompatible file servers appear as unavailable.
10. Click **Add**.
11. Do the following in the **Settings** tab:
 - a. In the **Name** field, enter a name for the protection policy.
 - b. In the **Description** field, enter a description for the policy.
12. Click **Create**.
The task for policy creation begins, and the policy appears in the **Policies** tab.

Editing an Asynchronous Protection Policy

Edit or turn off an existing asynchronous replication disaster recovery policy.

About this task

You can remove shares from a protection policy, modify the replication policy, modify the replication schedule, and change the policy name and description.

Note: Use the `afs infra.pc_fix_all_dr_policies` command for information on policy issues and to update a policy after resolving issues.

Procedure

1. From the side pane in the Files Manager, go to **Data Replication > Policies**.
2. Select the checkbox next to the name of the target asynchronous replication policy.
3. Select **Actions > Edit Policy**.
4. In the **Recovery Topology** tab, in the **Primary Location (Source File Server)** pane, configure the following settings:
 - a. In the **File Server** dropdown menu, select the source file server.
 - b. Under **Shares to be protected**, click **Edit**.
The **Select Specific Shares** dialog box displays.
 - c. (Optional) Clear the checkboxes next to the names of shares to exclude from the policy.
To show all shares, select the **Show All Shares** checkbox. To search by name, enter the name of the share in the **Search By Share Names** field.
 - d. (Optional) Select the **Add newly created shares on this file server to this policy** checkbox.
 - e. Click **Done**.
5. In the **Sync Schedule** pane, configure the following settings:
 - a. In the **Sync Frequency** field, enter an integer to indicate the replication interval.
 - b. In the **Sync Frequency** dropdown menu, specify the replication interval type by selecting **minutes**, **hours**, or **days**.
 - c. Select either the **Start Immediately** or **Start from a specific point in time** option.
 - d. If you choose the **Start Immediately** option, skip to Step 7.
 - e. Under **Start time**, specify when to begin replication.
 - f. Under **time zone**, specify the time zone for the start time and the replication schedule.
6. In the **Primary Location (Source File Server)** pane, click **Update**.
7. In the **Target File Server (Recovery Location)** pane, configure the following settings:
 - a. In the **Availability Zone** dropdown menu, choose the availability zone of the recovery file server.
To use the local Prism Central cluster, select Local AZ.
 - b. In the **File Server** dropdown menu, choose a recovery file server from the dropdown menu.
Incompatible file servers appear as unavailable.
8. Click **Next**.
9. In the **Settings** tab, complete the following fields:
 - a. In the **Name** field, enter a name for the protection policy.
 - b. In the **Description** field, enter a description for the policy.

10. Click **Update**.

The task for policy creation begins, and the policy appears in the **Policies** tab.

Switching the Direction of Replication

Switch the replication direction of an asynchronous disaster recovery policy.

Procedure

1. From the Files Manager home page, go to **Data Replication > Policies**.
2. In the **Replication Policies** table, select the checkbox next to the name of the target asynchronous replication policy.
3. Select **Actions > Switch Direction**.
4. In the **Switch Direction** dialog box, read the notice and click **Switch Direction**.
The task for the direction switch begins. Once complete, the **Policies** table displays a switched replication direction between the primary and the recovery locations. Policies for the same Prism Central (PC) instance appear in the same table. Multi-PC policies appear in the **Replication Policies** table in the Files Manager for the target Prism Central instance.

Updating AD and DNS Entries

Update Active Directory (AD) and domain name system entries for seamless client redirection.

About this task

Clients from the primary site can only perform writes on the recovery site if they have the required access permissions. To update AD and DNS entries from the target Prism Central instance, follow these steps:

Procedure

1. From the side pane in the Files Manager, go to **Data Replication > Monitor**.
2. Navigate to the **Disaster Recovery** tab.
3. In the row for the target policy, click **three-dots icon > Configure DNS and Active Directory**.
4. In the **Active Directory Credentials** section, complete the following fields:
 - **Username:** Enter an AD username in the *domain/username* format or in the UPN format *username@ADrealm*. You must have an administrator account or the following required permissions for the relevant AD organizational unit (OU) realm:
 - Create computer objects
 - Read/write service principal names

Note: If the DNS domain name of the file server is different from the AD domain name, Files requires the *Read/Write DNS host name* permission.

 - **Password:** Enter the password for the supplied username.
 - (Optional) **Preferred Domain Controller:** Specify a domain controller for join-domain operations.
You must have the permissions to modify DNS records on the domain controller.

5. In the **File Server DNS Entries** section, select a DNS entry update option from one of the following:

- » Automatically update entries (MS-DNS only)
- » Manually update entries (all other DNS)

(Optional) If you choose to manually update entries, you can also choose to specify the preferred name server in the **Preferred Name Server** field.

6. (Optional) To use the same credentials as used for AD to update the DNS entries, select the **Use the same credentials as the Active Directory** checkbox.

7. In the **Username** field, enter the username of the DNS server administrator.

8. In the **Password** field, enter the password for the username of the DNS server administrator.

9. Click **Save**.

If you did not configure DNS credentials, the Are you sure you want to proceed without providing DNS details dialog displays.

- » Click **Proceed** to proceed without providing DNS credentials.
- » Click **Cancel** to exit the dialog box and to configure the DNS options. If you choose to cancel, select **Next > Proceed** after you updated the credentials.

Smart DR Failover

This topic describes Smart DR failover concepts.

Nutanix recommends completing at least one replication job for each share under a disaster recovery policy before failing over to a recovery site.

Note: NFS clients must remount after a failover.

Both planned and unplanned failovers are user-triggered operations. The failover process includes two primary phases: activating protected shares on the recovery site and redirecting client access to the recovery site. After a failover, clients have access to protected shares on the recovery site but lose access to unprotected shares on the primary site.

Important: The number of shares failed over and latencies in programming Active Directory (AD) and domain name system (DNS) entries can cause a temporary disruption of service, lasting about a minute.

When applicable, configuring AD and the DNS details before activating the failover ensures a smooth client failover to the recovery site. Client access redirection during the failover prevents file-access issues on the recovery site. You can configure client credentials after initiating the failover, either by configuring AD and DNS credentials or by manually mounting clients using destination site IP addresses. The clients on the source site can only perform writes to the recovery site after configuring the credentials.

After you initiate the failover, the status appears in the **Tasks** view. If the failover task fails, hovering over the task in the **Tasks** view displays a thumbnail describing the issue.

Note: Sometimes failovers cause connection redirection, which can require the remount or the restart of the client.

When the failover is complete, the Files Manager activates the recovery site and sets the primary site to standby.

If the recovery site has written data, resume replication from recovery site to the primary site, which requires the primary site to be on standby. If the primary site has important data, you can resume replication from the primary site to the recovery site.

Planned Failover

Planned failovers include an option to configure a reverse replication policy to replicate data back to the primary site. Nutanix recommends reverse replication policies to facilitate faster failback. For planned failovers, incomplete replication jobs cause lengthy replication times. Planned failover operations that do not complete within the failover timeout window of four hours result in failure.

During a planned failover, protected shares on the primary site are in read-only mode until the replication of the final delta. Then Files puts the remote site shares into read-write mode. This process helps prevent data loss during a planned failover.

After recovery site activation, you can optionally configure a reverse replication policy for moving data back to the primary site.

Unplanned Failover

When the primary site goes down during a disaster, perform an unplanned failover. For unplanned failovers, incomplete replication jobs block client access to protected shares. To access protected shares on the recovery site, the Files Manager must complete at least one replication job.

For unplanned failovers, without an AD and DNS configuration, clients continue writing to the primary file server, possibly creating deltas that can result in data loss.

Failing Over to a Remote Site

Fail over a file server for disaster recovery.

About this task

To fail over a file server with a Smart DR policy, follow these steps:

Procedure

1. From the side pane in the Files Manager, go to **Data Replication > Monitor**.
2. Go to the **Disaster Recovery** tab.
3. In the row of the target file server, in the **Actions** column, click **Failover**.
The **Failover** dialog box appears and displays the **Failover Settings** tab.
4. In the **Select the type of Failover** section, specify the failover type:
 - » **Unplanned Failover**
 - » **Planned Failover**
5. Click **Next**.

6. (Optional) In the **Active Directory and DNS Configuration** tab, update the Active Directory (AD) and the domain name system (DNS) configuration. (You can also manually mount clients on the recovery site after completing the failback.)

Note: Clients from the primary site can only perform writes on the recovery site if they have the required access permissions.

- a. In the **Active Directory Credentials** section, complete the following fields:
- **Username:** Enter an AD username in the *domain/username* format or in the UPN format *username@ADrealm*. You must have an administrator account or the following required permissions for the relevant AD organizational unit (OU) realm:
 - Create computer objects
 - Read/write service principal names

Note: If the DNS domain name of the file server is different from the AD domain name, Files requires the *Read/Write DNS host name* permission.

- **Password:** Enter the password for the specified AD username.

- b. In the **File Server DNS Entries** section, complete the following fields:

- **Preferred Name Server:** Enter a fully qualified domain name (FQDN) to be used for join-domain operations.

You must have the permissions to modify DNS records on the domain controller.

- **Username:** Enter the username of the DNS server administrator.
- **Password:** Enter the DNS server administrator password.

7. Click **Failover**.

If you did not configure DNS credentials, the Are you sure you want to proceed without providing DNS details dialog displays.

- » Click **Proceed** to proceed without providing DNS credentials.
- » Click **Cancel** to exit the dialog box and to configure the DNS options. If you choose to cancel, select **Failover** once you've updated the credentials.

Results

The failover process begins. You can view the progress in the **Tasks** view. Shares become read-only on the recovery file server after the completion of the failover task.

Important: The number of shares failed over and latencies in programming Active Directory (AD) and domain name system (DNS) entries can cause a temporary disruption of service, lasting about a minute.

Resuming Replication

To resume replication, specify the active and standby sites.

About this task

When the primary and recovery sites are in an active-active state, both sites can receive writes and replication stops. Resuming replication changes the state to active-standby. During high-availability (HA) events, resume replication operations begin only after FSVMS become active. To resume replication, choose the direction of replication to indicate data to preserve.

Caution: Resuming replication can cause data loss on the site that is on standby at the end of the replication operation. When both sites are in an active-active state, both sites can receive unique writes. Choosing an active site preserves the data on that site. For example, if you resume replication from site A to site B, there may be data loss on site B.

Procedure

1. From the side pane in the Files Manager, go to **Data Replication > Monitor > Disaster Recovery tab**.
2. In the **Actions** column, select **Resume Replication**.
3. To choose the direction of replication, indicate the file server to remain active.
 - » Choose **Primary File Server to Recovery File Server**.
 - » Choose **Recovery File Server to Primary File Server** (continue to Step 5).
4. Configure Active Directory (AD) and domain name system (DNS) entries:
 - a. Under **Active Directory Credentials**, complete the following fields:
 - **Username:** Enter an AD username in the *domain/username* format or in the UPN format *username@ADrealm*. You must have an administrator account or the following required permissions for the relevant AD organizational unit (OU) realm:
 - Create computer objects
 - Read/Write service principal names
 - **Password:** Enter the password for the username.
 - (Optional) **Preferred Domain Controller:** Specify a domain controller for join-domain operations. You must have the permissions to modify DNS records on the domain controller.
 - b. Under **File Server DNS Credentials** supply the indicated information in the specified field.
 - (Optional) **Preferred name server:** Enter a preferred name server.
 - (Optional) Check the **Use the same credentials as Active Directory** selection box.
 - **Username:** Enter the username of the DNS server administrator.
 - **Password:** Enter the administrator password.
 - c. Continue to Step 6.
5. (Optional) Create a protection policy for the recovery site.
 - a. Check the **Create a Data Protection Policy** checkbox.
 - b. In the **Replication frequency** field, specify the frequency of replication with an integer.
 - c. In the **Replication frequency** dropdown menu, specify the replication interval type (minutes, hours, or weeks).
 - d. In the **Policy Name** field, enter a name for the policy.
6. Click **Resume Replication**.

Manually Failing Over to a Remote Site

When the Files Manager is not available, failover to a recovery site using a manual procedure.

Before you begin

Ensure that you have the following:

- Administrator Active Directory (AD) credentials:
 - Access to the organization unit (OU) where the primary and recovery file server computer accounts are registered
 - Read servicePrincipalName permissions
 - Write servicePrincipalName permissions
- Domain Name System (DNS) administrator privileges
- File server details:
 - Primary file server name
 - Primary file server UUID

Note: For multi-PC policies, perform the failover from the recovery site using the Prism Central UI. For more information, see [Failing Over to a Remote Site](#) on page 72.

About this task

For single-PC policies, in the event of a cluster-wide outage, you cannot access the Files Manager through the Prism Central web console to perform a failover. Instead, use the manual failover workflow to fail over file servers. To perform a manual failover, follow these steps:

Procedure

1. (Optional) From the recovery file server, obtain the UUID of the primary file server

```
nutanix@fsvm$ afs sync.list_fs_pairs
```

If the output includes two entries, use the file server UUID that is different than the file server UUID from where you run the command.

2. Activate the failover from the recovery site:

```
nutanix@fsvm$ afs dr.activate ad_username=AD-username ad_password=AD-password  
dns_domain=DNS.domain.name \  
dns_username=DNS-username dns_password=DNS-password name=primary-file-server-name  
uuid=primary-file-server-UUID
```

(Optional) If you are using NFS with LDAP or unmanaged directory services, remove the `ad_username=AD-username` and `ad_password=AD-password` parameters.

Activating a file server on a recovery site grants read/write access to the protected shares on the primary file server. Files deletes DNS entries on the primary site and creates DNS entries using the primary file server name, which point to the recovery file server.

Important: The number of shares failed over and latencies in programming Active Directory (AD) and domain name system (DNS) entries can cause a temporary disruption of service, lasting about a minute.

Creating a Reverse Replication Policy

After a Smart DR failover, create a reverse replication policy on the target to prevent unnecessary data loss.

About this task

Consider creating a reverse replication policy in the following scenarios:

- You did not create a reverse replication policy after resuming replication from the standby state.
- You deleted a reverse replication policy.

You can create a reverse replication policy after a failover when the **Actions** column displays the **Failback** option.

To create a reverse replication policy, follow these steps:

Procedure

1. From the side pane in the Files Manager, go to **Data Replication > Monitor**.
2. Go to the **Disaster Recovery** tab.
3. In the row of the target policy, select the [triple dot vertical](#) icon.
4. Select **Create a Reverse Replication Policy**.
The **Create a Reverse Replication Policy** dialog box appears.
5. Specify the replication frequency.
 - a. In the **Replication frequency** field, specify the replication frequency with an integer.
 - b. In the **Replication frequency** dropdown menu, specify the replication interval of **minutes**, **hours**, or **weeks**.
 - c. In the **Policy Name** field, enter a name for the reverse replication policy.
6. Click **Create Policy**.
The newly created reverse replication policy appears in the **Policy** tab.

Smart DR Failback

This topic describes Smart DR failback concepts.

Configuring a reverse replication policy during failover policy creation simplifies failback operations. The Files Manager deletes the reverse replication policies after completing failback operations. Creating a reverse replication policy after an unplanned failover can prevent data loss and accelerate the recovery process.

After the primary site becomes active after an unplanned failover, both sites are in an active state. Since active-active states can create data deltas, resume replication from the recovery site before failing back to the primary site to limit data loss.

First, the recovery file server is in read-only mode. After replicating the final delta, the primary site shares are in read-write mode.

If the Files Manager does not complete a failback operation during the timeout window of 4 hours, the operation is unsuccessful, which helps prevent data loss during a planned failback.

Failing Back to the Primary Site

Fail back shares to the primary site.

About this task

To fail back to the primary site, follow these steps:

Procedure

1. From the side pane in the Files Manager, go to **Data Replication > Monitor**.
2. Navigate to the **Disaster Recovery** tab.
3. In the **Actions** column, in the row for the target file server, click **Failback**.
The **Failback Settings** tab in the **Failback** window displays.
4. (Optional) In the **Active Directory Credentials** and **File Server DNS Entries** sections, update the Active Directory (AD) and DNS configuration.
Clients from the primary site can only perform writes on the recovery site if they have the required access permissions. You can also manually mount clients on the recovery site after completing the failback.
 - a. In the **Active Directory Credentials** section, complete the following fields:
 - **Username:** Enter an AD username in the *domain/username* format or in the UPN format *username@ADrealm*. You must have an administrator account or the following required permissions for the relevant AD organizational unit (OU) realm:
 - Create computer objects
 - Read/Write service principal names

Note: If the DNS domain name of the file server is different from the AD domain name, Files requires the *Read/Write DNS host name* permission.

 - **Password:** Enter the password for the supplied username.
 - **Preferred Domain Controller:** Specify a domain controller for join-domain operations.
You must have the permissions to modify DNS records on the domain controller.
 - b. In the **File Server DNS Credentials** section, complete the following fields:
 - (Optional) To update the DNS entries using the same credentials as AD, select the **Use the same credentials as the Active Directory** checkbox.
 - **Username:** Enter the username of the DNS server administrator.
 - **Password:** Enter the DNS server administrator password.
5. Click **Next**.
If did not configured DNS credentials, the Are you sure you want to proceed without providing DNS details dialog displays.
 - » Click **Proceed** to proceed without providing DNS credentials.
 - » Click **Cancel** to exit the dialog box and to configure the DNS options. If you choose to cancel, select **Next > Proceed** after you update the credentials.
6. In the **Failback** view, click **Failback**.

Results

The failback process begins. You can view the progress in the **Tasks** view. The **Disaster Recovery** tab displays the policy as activated after the failback is complete. The Files Manager deletes the reverse replication policy.

SYNCHRONOUS DATA PROTECTION WITH METRO AVAILABILITY ON AHV

Use Metro Availability as a synchronous disaster recovery solution for Prism Element and Prism Central–based file servers on AHV hypervisors.

Metro Availability uses a high-availability architecture to provide seamless failover and continuous access to services across geographically distributed clusters. As a synchronous disaster recovery solution, Metro Availability supports a recovery point objective (RPO) of 0. In the event of a primary site failure, the primary site becomes inactive and a Metro Availability policy initiates file server migration to a remote site, which becomes the active site. File servers are automatically promoted to an online state on the recovery site. Upon recovery of the primary site, in most scenarios, the system automatically activates the primary site and resumes replication; however, in some scenarios, such as network partitions, you must manually resume synchronous replication.

During a Metro Availability failover, clients can reconnect to a file server regardless of whether the file server is on the primary or the recovery site.

Failure Handling

Metro Availability failure handling minimizes downtime, prevents data corruption in split-brain scenarios, and automates most of the failover and recovery processes. Using Metro Availability, Nutanix Files ensures client access during site failures and service disruptions by facilitating synchronous replication of data between two clusters and by using a Witness Service hosted on a Prism Central cluster to resolve ownership issues.

During a failure or disruption, such as a complete site outage, a network partition between clusters, or the loss of access to the Witness, Nutanix Files determines which cluster should take over the file server. The active cluster continues serving clients independently while replication is paused. Once the failed site recovers, Nutanix Files automatically resumes replication unless a network partition occurs.

When the system encounters a network partition and cannot determine when to resume replication, manually resume synchronous replication. Network partitions occur when the primary and the recovery clusters lose communication to each other, but both clusters still have communication with the Witness Service. During a network partition event, replication is paused to prevent data inconsistency and the Witness Service grants ownership to one cluster (usually the primary), which continues serving clients independently until you manually resume synchronous replication.

Requirements

Meet the following requirements before creating a Metro Availability policy on AHV:

- Create a remote site container with the same name as the file server container on the source.
For more information, see [Creating a Container](#) in the *Prism Element Web Console Guide*.
- Make sure the recovery location has enough free resources to support a failover.
For example, if a file server has 24 GB memory, the recovery location must have 24 GB memory available for as many hosts as there are file server VMs.
- Enable disaster recovery from Prism Central settings.

Limitations

Metro Availability for Nutanix Files on AHV does not support the following:

- File servers in virtual private clouds (VPCs)
- Policies across different availability zones (AZs)
- Dense nodes

- Multi-VLAN file servers
- File servers with Protection Domain–based DR policies.
- Prism Central–based DR workflows beyond the scope of Files.
For example, workflows covered in the [Nutanix Disaster Recovery Guide](#) are not supported.
- Managed networks

Creating a Synchronous Replication Policy

Create a synchronous replication policy for a file server on the AHV hypervisor.

About this task

To create a synchronous replication policy, follow these steps:

Procedure

1. From the Files Manager side pane, go to **Data Replication > Policies**.
2. Click **+ New Replication Policy**.
3. From the data replication use-case options, select **Disaster Recovery**.
4. Under **Select Replication Type**, select **Synchronous Replication**.
5. Click **Next**.
The **Create a Data Replication Policy** page displays.
6. In the **Replication Topology** tab, configure location settings:
 - a. In the **Primary Location (Source File Server)** pane, select the source file server from the **File Server** dropdown menu.
 - b. Click **Add**.
 - c. In the **Target Cluster (Recovery Location)** pane, select the target file server from the **Cluster** dropdown menu.
7. Click **Next**.
8. In the **Settings** tab, configure general and recovery settings:
 - a. In the **Name** field, enter a name for the protection policy.
 - b. In the **Description** field, enter a description for the policy.
 - c. Use the **Execute failover after disconnectivity of** field and dropdown menu to specify an integer and time interval for how soon to execute a failover after a site failure or network failure.
 - d. In the **Network Mapping** section, use the **Virtual Network or Port Group** fields to specify network mappings.
The network of the primary location displays on the left under **Virtual Network or Port Group**. The **Virtual Network or Port Group** field to the right specifies the mapping to the recovery location.
Single network file servers have only one row for network mapping. Multi-network file servers use the first row for storage network and the second row for the client network.
9. Click **Create**.

Failing Over to a Remote Site

Fail over a file server for planned recovery.

About this task

Nutanix Files automatically performs unplanned failovers with Metro Availability; however, you can manually perform planned failovers as needed.

To perform a planned failover for a file server with a synchronous replication policy, follow these steps:

Procedure

1. From the side pane in the Files Manager, go to **Data Replication > Monitor**.
2. Go to the **Disaster Recovery** tab.
3. In the row of the target file server, in the **Actions** column, click **Failover**.
The **Failover** dialog box appears displaying the **Planned Failover** option and the **Replication Topology**.
4. Click **Failover**.
The file server failover tasks begin. After the tasks complete, the file server is active on the recovery site. Synchronous replication resumes automatically.

Resuming Synchronous Replication

Resume replication when an active synchronous replication policy has an "Out of Sync" replication status.

About this task

Resuming synchronous replication is necessary after resolving issues that cause the Metro Availability policy to have an "Out of Sync" status, which appears when a network partition exists between primary and recovery sites or when a recovery site fails.

In the event of a network partition, both the primary and the recovery site are operational, which creates a potential for divergence, so the primary cluster pauses replication. After resolving the network issue, resuming replication re-establishes data consistency between the clusters.

In the event of a recovery site failure, the source cluster continues operating while replication is paused; resuming replication synchronises the recovery site with the primary site after it's back online and reinstates data redundancy.

To resume synchronous replication, follow these steps:

Procedure

1. From the Files Manager side pane, go to **Data Replication > Policies**.
2. Select the target policy.
3. From the **Actions** dropdown menu, select **Resume Synchronous Replication**.

Editing Synchronous Replication Policies

Edit an existing synchronous replication disaster recovery policy.

About this task

Note: If you reconfigure a file server network to be different from the network in the synchronous replication policy, you must update the network mapping.

To modify an existing synchronous replication policy, follow these steps:

Procedure

1. From the Files Manager side pane, go to **Data Replication > Policies**.
2. Select the checkbox next to the name of the target synchronous replication policy.
3. Select **Actions > Edit Policy**.
4. Click **Next**.
You cannot modify source and recovery locations.
5. (Optional) In the **Settings** tab, configure the general and the recovery settings:
 - a. In the **Name** field, enter a name for the protection policy.
 - b. In the **Description** field, enter a description for the policy.
 - c. Use the **Execute failover after disconnectivity of** field and dropdown menu to specify an integer and time interval for how soon to start a failover after a site failure or network failure.
 - d. To modify network mappings, in the **Network Mapping** section, use the **Virtual Network or Port Group** fields.

The primary location's network displays on the left under **Virtual Network or Port Group**. Use the **Virtual Network or Port Group** field to the right to specify the mapping.

Single network file servers have only one row for network mapping. Double network file servers use the first row for storage network and the second row for the client network.

Alternatively, you can use the **Delete Network Mapping** and **Add Network Mapping** options.

6. Click **Update**.

Disabling a Synchronous Replication Policy

Disable a synchronous replication disaster recovery policy.

About this task

To disable a synchronous replication policy, follow these steps:

Procedure

1. From the Files Manager side pane, go to **Data Replication > Policies**.
2. Select the checkbox next to the name of the target synchronous replication policy.
3. Select **Actions > Disable Policy**.
A confirmation window for policy disablement displays.
4. Click **Disable**.
Disabling a policy immediately pauses replication jobs. To resume replication, enable the policy.

Deleting a Synchronous Replication Policy

Delete a synchronous replication disaster recovery policy.

About this task

To delete synchronous replication policy, follow these steps:

Procedure

1. From the Files Manager side pane, go to **Data Replication > Policies**.
2. Select the checkbox next to the name of the target synchronous replication policy.
3. Select **Actions > Delete Policy**.
A confirmation window for policy deletion displays.
4. Click **Delete Policy**.
Deleting a policy removes protection from entities covered by the policy. Nutanix Files does not generate any new recovery points, and in-progress replication continues until the replication job completes.

Enabling a Synchronous Replication Policy

Enable an inactive synchronous replication disaster recovery policy.

About this task

To enable a synchronous replication policy, follow these steps:

Procedure

1. From the Files Manager side pane, go to **Data Replication > Policies**.
2. Select the checkbox next to the name of the target synchronous replication policy.
3. Select **Actions > Enable Policy**.
A confirmation window for policy enablement displays.
4. Click **Enable**.
Enabling the policy resumes data replication from the last known sync state. The policy status appears as **Active** in the **Replication Policies** table.

DATA SYNCHRONIZATION WITH SMART SYNC

Smart Sync provides a cost-effective data synchronization solution.

Smart Sync lets you replicate and consolidate data from multiple file servers to a single target file server on the same or different Prism Central cluster within the same availability zone (AZ) pair. For information on AZ pairing, see [Availability Zone Pairing](#) in the *Flow Virtual Networking Guide*.

You can configure Smart Sync replication policies with file-level granularity and use one or more Smart Sync policies to define source paths for each target. You must configure policies from the target Prism Central instance, but you can monitor the policies from either the source or the target. After you configure a Smart Sync policy, the Files Manager performs replication jobs as specified in the policy. If you delete data from the source, Smart Sync deletes the data from the target, too.

Before creating a Smart Sync policy, review the [Restrictions](#) and [Limitations](#) sections.

[Data Replication Navigation](#) on page 62 provides an overview of the UI that you can use for policy creation and monitoring.

Important: You cannot upgrade file servers with technical preview Smart Sync policies to general availability (GA) versions. Delete policies created during the technical preview.

Feature Interoperability

For shares configured with both Smart DR and Smart Sync, Smart DR failover and failback operations lead to full replication of source data for the source shares. However, you must reconfigure Smart Sync policies after a failover on a share that is a source for both Smart DR and Smart Sync policies. Smart Sync and Smart DR policies are disabled during failovers and enabled after failback.

For more information on the interoperability of Files Manager features, see [Feature Interoperability](#) on page 19.

Restrictions

Smart Sync has the following restrictions:

- Smart Sync does not handle the following use cases:
 - Name-space collision on target
 - Modification of files on the target share

Note: Smart Sync overwrites conflicting files on the target with changes from the source.

- For nested shares, you must provide an absolute path to replicate the data.
- Factors like the number of files, TLDs, and hard links can impact replication job performance.
- Multi-protocol shares require configuring user mapping on both the source and the target server before configuring Smart Sync. See [User Mapping](#) in the *Nutanix Files User Guide*.

Limitations

Smart Sync has the following limitations:

- The target must support the same protocols as the source. For example, if the source uses SMB, the target must also have the SMB protocol configured.
- Replication excludes special files such as pipes, sockets, blocks, or character device files.

- You cannot use tier or recall operations from Data Lens on files consolidated on the target through Smart Sync.
- Files restricts multi-PC policy creation and management to the target Prism Central instance.
- For home share targets without support or files enabled at the root, Files does not mark root-level files failed and does not generate alert messages.
- Smart Sync does not replicate files to the root of target SMB shares.
- Smart Sync replication skips files with carriage return characters (\r or \n) in file names.
- You cannot use WORM-enabled shares as Smart Sync targets.
- Each policy has the following limits:
 - Maximum ten source file servers per policy
 - Maximum 20 source paths per source file server
 - Maximum of 100,000 directories at the root of the source path
 - Maximum of 100 source paths replicating to a target file server
 - Minimum replication frequency of 1 minute

For more information on Nutanix Files feature system limits, see [Configuration Maximums](#).

Creating a Smart Sync Policy

Replicate data from multiple sources to a single target file server.

About this task

For each policy, add a maximum of three file server sources with a maximum of four source paths per file server.

Creating and updating replication policies requires the view availability zone permission for custom roles.

To create a Data Synchronization policy, follow these steps:

Procedure

1. From the side pane in the Files Manager, go to **Data Replication > Policies**.
2. Click **+ Create New Replication** policy.
The **Create a Data Replication Policy** window displays.
3. Select the **Data Synchronization** option.
4. Click **Next**.
The **Replication Use Case: Data Synchronization** window displays.
5. In the **Recovery Topology** tab, in the **Sources** pane configure the following settings:
 - a. In the **Local AZ** dropdown, select an availability zone of the source file servers.
 - b. In the **File Server** dropdown, select a source file server.
 - c. In the **Path** field, enter the data source path.
To add multiple source paths, click the [plus](#) icon.
For source paths ending with a forward slash (/), the Files Manager copies the content from the source directly to the specified target path. For source paths ending without a forward slash, the Files Manager

creates a copy of the source folder on the target path. To sync data from source paths ending with a share name, such as /share1, create a new folder with the same name on the target. Files does not create the policy until the folder is created on the target.

To replicate nested shares, provide an absolute path of the nested share.

Caution: Make sure that there is no path collision.

- d. Click **Add**.
 - e. (Optional) To configure advanced source options, click **Advanced**.
In the **Add source** dialog window, source paths specified in the **Sources** tab appear in the **Source Paths** section.
 - To add directory or file exclusions to the source path, select the **Exclude Patterns** option.
 - Enter the directory path in the **Directory Path** field.
 - Click the **plus** icon to add additional exclusions.
 - Click **Add**.
 - f. Click **Add** in the **Sources** pane.
6. In the **Sync Schedule** pane , configure the sync schedule:
- a. In the **Sync Frequency** fields, enter an integer to indicate the time interval for syncing data, and select the interval of **minutes**, **hours**, or **days**.
 - b. To specify when to start replication, select one of the following options:
 - » **Start immediately**
 - » **Start from a specific point in time**
 - c. If you chose the **Start from a specific point in time** option, specify the start time in the **Start Time** field.
7. In the **Target Path** pane, configure the data synchronization target:
- a. From the **Local AZ** dropdown menu, select an availability zone of the source file servers.
 - b. From the **File Server** dropdown menu, select the target file server.
 - c. In the **Path** field, enter the target synchronization path.
 - d. Click **Add**.
 - e. (Optional) Click **View Folder Mapping on Target** to view the target folder mapping.
Files syncs the source data under the target path.
8. Click **Next**.

9. In the **Settings** tab, configure policy settings:
 - a. In the **Name** field, enter a policy name.
 - b. (Optional) In the **Description** field, enter a brief policy description.
 - c. (Optional) To exclude synchronization targets by file extension type, under **File Type Extension**, select the **Exclude files types during synchronization** option and in the **File Extensions** field, enter a comma-separated list of file extensions
 - d. (Optional) To prevent the target from deleting files that are deleted on the source, under **Delete Propagation**, select **Block Delete Propagation**.
10. Click **Create**.

You can view the policy creation progress under Prism Central's **Recent Tasks** view. The policy appears under **Data Replication > Policies**.

Editing a Smart Sync Policy

Update an existing Smart Sync policy.

About this task

After you create a policy, you cannot modify the target path. For multi-PC policies, you can add more file server sources.

Before updating a policy, review the *Limitations* section in [Data Synchronization with Smart Sync](#) on page 83.

Important: You cannot upgrade file servers with technical preview Smart Sync policies to general availability (GA) versions. Delete policies created during the technical preview.

Procedure

1. From the side pane in the Files Manager, go to **Data Replication > Policies**.
2. Select the checkbox next to the name of the target policy.
3. Select **Actions > Edit Policy**.

The **Update Data Synchronization Policy** view displays.
4. Make any needed updates in the **Sources** tab:
 - a. (Optional) To edit the file server source path, click the [pencil](#) icon.

You can also select **+ Add Source** to add paths.

For source paths ending with a forward slash (/), the Files Manager copies the content from the source directly to the specified target path. For source paths ending without a forward slash, the Files Manager creates a copy of the source folder on the target path. To sync data from source paths ending with a share name, such as /

share1, create a new folder with the same name on the target. Files does not create the policy until the folder is created on the target.

Note: In the event of a path collision, the policy update operation is unsuccessful.

- b. (Optional) To configure advanced source options, click **Advanced**.

In the **Add source** dialog window, source paths specified in the **Sources** tab appear in the **Source Paths** section.

- To add a directory or file exclusions to the source path, select the **Exclude Patterns** option.
- In the **Directory Path** field, enter the directory path.
- To add exclusions, click the **plus** icon .
- Click **Add**.

- c. Click **Update**.

5. (Optional) In the **Sync Schedule** pane, update the sync schedule:

- a. In the **Sync Frequency** fields, enter an integer to indicate the time interval for syncing data, and select the interval of **minutes**, **hours**, or **days**.
- b. To specify when to start replication, select one of the following options:
- » **Start immediately**
 - » **Start from a specific point in time**
- c. If you chose the **Start from a specific point in time** option, specify the start time in the **Start Time** field.

6. (Optional) To view the target folder mapping, select **View Folder Mapping on Target**.

After you update the policy, the source data syncs to the target path.

7. Click **Next**.

The **Settings** page displays.

8. (Optional) In the **Settings** tab configure the policy settings:

- a. In the **Name** field, enter a policy name.

Note: You cannot use the same name for a Smart Sync and a VDI sync policy.

- b. In the **Description** field, enter a policy description.
- c. To exclude synchronization targets by file extension type, under **File Type Extension**, select the **Exclude files types during synchronization** option to exclude synchronization targets by file extension type. In the **File Extensions** field, enter a space- or comma-separated list of file extensions.
- d. To prevent the target from deleting files that get deleted on the source, under **Delete Propagation**, select **Block Delete Propagation**
- e. Click **Update**.

9. Click **Update**.

You can view the policy update progress in Prism Central's **Tasks** view. After Files updates the policy, it lists it under **Policies**.

Deleting a Smart Sync Policy

Delete an existing Smart Sync policy.

About this task

Deleting a policy cancels future sync jobs. Files does not delete data already synced to the target.

Procedure

1. From the side pane in the Files Manager, go to **Data Replication > Policies**.
2. Select the checkbox next to the name of the target policy.
3. Select **Actions > Delete Policy**.
4. In the row for the policy to be deleted, select **three dots menu > Delete**.
The **Are you sure you want to delete the policy** confirmation dialog box displays.
5. Select **Delete Policy**.
The delete policy operation begins. Previously synced data remains on the target share.

Enabling or Disabling a Smart Sync Policy

Enable or disable an existing Smart Sync policy.

Procedure

1. Go to **Data Sync > Policies**.
2. In the row for the policy to be enabled or disabled, select the appropriate option:
 - » To enable the policy, select **Actions dropdown menu > Enable**
 - » To disable the policy, select **Actions dropdown menu > Disable**.

Note: After you disable a policy, the Files Manager pauses the current sync jobs and cancels any scheduled jobs for that policy. The files on the target keep changes from previous sync jobs.

The **Enable/Disable confirmation** dialog appears.

3. Do one of the following:
 - » If disabling, optionally select the **Abort in-progress jobs** checkbox to stop in-progress jobs, and select **Disable**.
 - » To enable, select **Enable**.

The **Policy State** column in the **Replication Policies** table displays the updated state of the policy.

VDI SYNC

Use the VDI Sync feature to replicate VDI user data across file servers in different locations on the same or different Prism Central cluster within the same availability zone (AZ) pair.

User profiles are stored on distributed shares. Creating a VDI Sync policy lets you replicate the distributed share data, including the user profile, between file servers at different assigned locations. When you travel between the assigned locations, you experience quality performance by accessing the data locally. Bi-directional replication ensures that you have access to the same data as your location changes between the two sites.

Files uses a wide-area network (WAN) to replicate data between sites, which ensures that the user profile is present on both sites and lets you take advantage of the performance improvements from using a local area network (LAN) connection at either site. As a result, input and output (I/O) operations, like file opens, reads, and writes, do not experience the latency issues associated with internet connections. VDI Sync also includes guardrails, including user-level locks, to protect your data, prevent data corruption, and prevent most conflicts. Conflicts can occur when there is more than one connection attempt to the same share. However, when conflicts occur, VDI Sync lets you resolve the conflicts and connection failures. For more information, see the [Conflict and Login Failure Handling](#) section.

[Data Replication Navigation](#) on page 62 provides an overview of the Data Replication UI, which includes policy monitoring tools.

If you configured File Analytics or Data Lens for file servers at either location, you do not need to make these products VDI Sync aware, because Data Lens and File Analytics work independently on each site. For more information, see the [File Analytics Guide](#) and the [Data Lens User Guide](#).

Conflict and Login Failure Handling

Although unlikely, one of the following scenarios might occur:

- A *conflict* occurs when you try to log on to both sites at almost the same time, and both sites accept you. During a conflict, the system marks you as a conflicted user and logs you out from both sites. You cannot log on to either site until an admin resolves this conflict using the **resolve failure** or **allow user connections** actions.
- A *login failure* occurs when you try to log on to a site while still logged onto another site, which causes an ACCESS_DENIED error. In this scenario, an admin uses the **resolve failure** or **allow user connection** actions to let you switch sites. The file server that you are logged on to has *ownership* of your user profile.

The **allow user connections** action lets you fail over all users to a single site. However, if you see any user conflicts, the admin must identify the conflicted users and manually fail over each user one by one using the **resolve failure** action. For more information, see [Resolving a VDI User Profile Connection Failure](#) on page 92.

Note: The VDI desktop configuration controls the routing of the VDI user session to a specific file server.

Restrictions

VDI Sync has the following restrictions:

- You must add an entire share in a VDI Sync policy. Sub-paths are not permitted.
- User profiles must be top-level directories (TLDs) of distributed shares.
- Legacy SMB shares created with versions earlier than Files 3.5 require enabling multi-protocol to perform a metadata transfer before enabling VDI Sync. For assistance with the metadata transfer, contact Nutanix Support.

Limitations

VDI Sync does not support the following:

- Common shares that are accessible by multiple users

- Standard, multi-protocol, or nested shares
- NFS exports
- Top-level directory (TLD) access by anyone except the owner or the backup admin
- Policies between file servers from different domains

For information on the interoperability of Files Manager features, see [Feature Interoperability](#) on page 19.

Creating a VDI Sync Policy

Create a policy to sync share data between two different locations.

About this task

To create a VDI data synchronization policy, follow these steps:

Procedure

1. From the side pane in the Files Manager, go to **Data Replication > Policies**.
2. Click **+ Create New Replication** policy.
The **Create a Data Replication Policy** window displays.
3. Select the **VDI Data Synchronization** option.
4. Click **Next**.
The **Replication Use Case: VDI Data Synchronization** window displays.
5. In the **Details** pane (File Server #1), configure the first file server:
 - a. (Optional) To change the availability zone from the current Prism Central instance that you are logged on to, in the **Local AZ** dropdown menu, select a different Prism Central instance.
 - b. From the **File Server** dropdown menu, select a file server.
 - c. From the **Shares** dropdown menu, select a share.
 - d. (Optional) To add more shares to the policy, click **Add Share**.
 - e. Click **Add**.

Note: Files server #1 becomes the preferred location for the policy.

6. In the **Sync Schedule** pane, configure the sync schedule:
 - a. In the **Sync Frequency** fields, enter an integer to indicate the time interval for syncing data, and select the interval of **minutes**, **hours**, or **days**.
The minimum frequency is 10 minutes.
 - b. Select one of the following options to specify when to start replication:
 - » **Start immediately**
 - » **Start from a specific point in time**
 - c. If you chose the **Start from a specific point in time** option, specify the start time in the **Start Time** field.

7. In the **File Server #2** pane, configure the second file server:
 - a. (Optional) To change the availability zone from the current Prism Central instance that you are logged on to, in the **Local AZ** dropdown menu, select a different Prism Central instance.
 - b. From the **File Server** dropdown menu, select a file server.
 - c. From the **Shares** dropdown menu, select a share.
 - d. Click **Add**.
8. Click **Next**.
9. In the **Settings** tab, configure policy settings:
 - a. In the **Name** field, enter a policy name.
 - b. (Optional) In the **Description** field, enter a policy description.
10. Click **Create**.

A policy creation task begins and the policy appears in the **Policies** table.

Editing a VDI Sync Policy

You can perform some updates to a VDI Sync policy using the **Edit** option.

About this task

You can add shares to a policy, update a policy name and description, and update the sync schedule.

Caution: To edit share sync relationships, Nutanix recommends deleting a policy and creating a new one to prevent data loss.

Delete the existing VDI sync policy and create a new one in the following scenarios:

- To modify one of the file server sites. You cannot edit the file servers or change the preferred locations for a policy.
- To modify the existing share mapping within a policy.

To edit a VDI sync policy, follow these steps:

Procedure

1. From the side pane in the Files Manager, go to **Data Replication > Policies**.
2. Select the checkbox next to the name of the target policy.
3. Select **Actions > Edit Policy**.
4. (Optional) To configure the shares on the first file server, in the **File Server #1** pane, follow these steps:
 - a. (Optional) To add more shares to the policy, click **Add Share**.
 - b. (Optional) To remove shares from the policy, click the **minus** icon.

5. (Optional) To configure the sync schedule, in the **Sync Schedule** pane, configure the following settings:
 - a. In the **Sync Frequency** fields, enter an integer to indicate the time interval for syncing data, and select the interval of **minutes**, **hours**, or **days**.
The minimum frequency is 10 minutes.
 - b. To specify when to start replication, select one of the following options:
 - » **Start immediately**
 - » **Start from a specific point in time**
 - c. If you chose the **Start from a specific point in time** option, specify the start time in the **Start Time** field.
6. Click **Next**.
7. (Optional) To configure policy settings, in the **Settings** tab, update the following fields.
 - a. (Optional) In the **Name** field, update the policy name.
 - b. In the **Description** field, update the description field.
8. Click **Update**.

Note: If updating a VDI Sync policy is unsuccessful after adding or removing a share, delete the policy and create a new one.

Resolving a VDI User Profile Connection Failure

Resolve VDI sync user profile connection failures and conflicts by failing over to a target site.

About this task

You can resolve a single user profile connection or multiple connections at a time.

For instructions on resolving multiple connections at a time, see [Resolving Multiple VDI User Connection Failures](#) on page 93.

To resolve the connection for a single user, follow these steps:

Procedure

1. From the side pane in the Files Manager, go to **Data Replication** > **Policies**.
2. Select the checkbox next to the name of the target policy.
3. From the **Users with Profile Connection Failures** table, select the checkbox for the target user.
4. Click **Resolve Failure**.
The **Resolve Failures** dialog appears.
5. From the **Select the file server to serve the selected user connections** dropdown menu, select a file server.

The file server you select determines the site the user connection fails over to.

Caution: Resolving a failure can cause data loss if the target user performed writes during the current or previous session on the former site that was not chosen for serving user connections.

6. Click **Resolve**.
Files triggers the **resolve failures** action.

Resolving Multiple VDI User Connection Failures

Resolve VDI sync user profile connection failures and conflicts for multiple users.

About this task

You can resolve a single user profile connection or multiple connections at a time.

For instructions on resolving the connection for a single user, see [Resolving a VDI User Profile Connection Failure](#) on page 92.

To resolve the connection for multiple users, follow these steps:

Procedure

1. From the side pane in the Files Manager, go to **Data Replication > Policies**.
2. Click the policy name.
3. From the **Actions** dropdown menu, select **Allows User Connections**.
4. From the **Select the file server to serve the selected user connections** dropdown menu, select a file server.

The file server you select determines the site the user connections fail over to.

Caution: Resolving a failure can cause data loss if the target user performed writes during the current or previous session on the former site that was not chosen for serving user connections.

5. Select **Failover**.
Files triggers the **Allow User Connections** action.

REPLICATION JOBS

The Replications Jobs dashboard tracks the replication progress for each Smart DR, Data Sync, and VDI Sync policy.

The **Replications Jobs** dashboard includes a table that lists replication jobs for policies that interact with file servers on the local Prism Central cluster, excluding synchronous replication jobs with Metro Availability. For information on replication jobs with Metro Availability, see [Synchronous Data Protection with Metro Availability on AHV](#) on page 78.

The total estimated files in the **Replication Jobs** details is an approximation and might not reflect the total number of synced files. Creating a policy can trigger multiple replication jobs depending on the number of shares under the protection policy.

If the replication source doesn't experience any data churn between jobs, Files skips replication, and the **Replication Jobs** table displays the **Skipped** status.

Clicking the replication job source displays the **Job Details** view, which provides further replication details. The *Network throughput (avg)* refers to the average quantity of data replicated in one second. Bandwidth restrictions and large quantities of replication data affect the replication time and the upcoming jobs. A replication job that takes longer than the specified replication frequency appears as replication frequency non-compliant in the **Replication Jobs** view.

Additionally, replication jobs have the following behaviors:

- Replication jobs that start during off-peak hours might have more available bandwidth.
- The Files Manager only maintains 50 replication jobs per share.
- Replication jobs for empty SMB shares might show some bytes replicated, which represents the metadata of alternate data streams (ADS).

Note: If replication jobs for a share are non-compliant, Nutanix recommends updating the replication frequency schedule.

UNIFIED NAMESPACE

A unified namespace lets you extend file server storage by leveraging the storage of multiple file servers.

During unified namespace creation, you must designate a local core member file server for Files to use as the unified namespace. After you create the unified namespace, you can add file servers to it while still being able to access the added file servers using their original namespaces. File servers added to a unified namespace are known as non-core member file servers, which can be third-party file servers.

Requirements

- You must create at least one file server before creating a unified namespace.
- File servers must be registered to the same Prism Central instance.
- Protection-domain-based disaster recovery in unified namespace environments requires registering recovery sites to the same Prism Central domain.

Note: Nutanix Files deactivates unified namespace policies for core member file servers that migrate outside of the Prism Central domain. Files reactivates the policies after the master file server rejoins the original Prism Central domain.

Limitations

The unified namespace feature has the following limitations:

- Each unified namespace can have a maximum of ten file servers.
- Share names must be different across all file servers registered to the unified namespace.
- Nutanix Files does not support virtual private cloud (VPC) network-based configurations for unified namespace file servers.
- Nutanix Files does not support NFSv3 exports as non-core member file servers.
- You cannot perform NFS export migration on edge file servers.
- You cannot modify or delete the core member file server in a unified namespace.
- If you delete a unified namespace, Files deletes all of the unified namespace metadata; however, Files preserves the metadata of the individual file servers, which you can delete separately.
- Deleting member file servers requires using the force delete operation.

Note: Force deleting non-core member file servers doesn't delete that file server's metadata from the unified namespace.

Creating a Unified Namespace

Create a unified namespace for your file servers.

About this task

You must create at least one file server before creating a unified namespace.

For more information, see [Creating a File Server in Prism Central](#) on page 9.

To create a unified namespace, follow these steps:

Procedure

1. In the Files Manager navigation pane, go to the **Unified Namespaces** tab.
2. Click **+ Create Unified Namespace**.
The **Create a Unified Namespace** page displays.
3. Under **Members**, click **Add File Servers**.
 - a. Select the checkbox next to each file server to add to the namespace.
 - b. Click **Add**.
4. Under **Namespace Access**, select the core member file server from the **Core Member** dropdown list.
5. Click **Create**.
Files starts a task to create the unified namespace. The unified namespace appears in the **Unified Namespaces** table.

Adding Nutanix File Servers to a Unified Namespace

Add file servers to a unified namespace.

About this task

To add Nutanix file servers to a unified namespace policy, follow these steps:

Procedure

1. In the Files Manager navigation pane, go to the **Unified Namespaces** tab.
2. Select the checkbox in the row of the target unified namespace.
3. From the **Actions** dropdown list, select **Update**.
The **Update Unified Namespace** page displays.
4. Click **Add File Servers > Nutanix File Server**.
For information on adding third-party file servers to a unified namespace, see [Adding a Third-Party File Server to a Unified Namespace](#) on page 97.
The **Add File Servers** dialog box displays.
5. In the **Add File Servers** dialog box, select the file servers to add and un-select the file servers to remove from the unified namespace.
6. Click **Update**.
7. In the **Update Unified Namespace** view, under **Namespace Access**, click **Update**.
The task to add the file servers to the namespace begins.

Deleting a Unified Namespace

Delete a unified file server namespace.

About this task

To delete a unified namespace policy, follow these steps:

Procedure

1. In the Files Manager navigation pane, go to the **Unified Namespaces** tab.
2. Select the checkbox in the row of the target unified namespace.
3. Go to the **Actions** dropdown, and select **Delete**.
The **Delete Unified Namespace** dialog box displays.
4. Read the message in the **Delete Unified Namespace** dialog box and click **Delete**.
The deletion task begins.

Adding a Third-Party File Server to a Unified Namespace

Add a third-party file server to an existing unified namespace as a non-core member file server.

About this task

To add a third-party file server, follow these steps:

Procedure

1. In the Files Manager navigation pane, go to the **Unified Namespaces** tab.
2. In the **Unified Namespaces** table, click the name of a unified namespace to add a third-party file server.
3. In the **Overview** tab of the unified namespace, click **Update**.
4. In the **Update Unified Namespace** view, click **Add File Servers > Third Party File Server**.
The **Add Third Party File Server** dialog box displays.
5. Select where to add the third-party file server from:
 - » To add a new third-party file server, select **Add a new third-party file server** and do the following:
 1. In the **Name** field, enter the file server name.
 2. In the **File Server Address** field, enter the fully qualified domain name (FQDN) or the IP address of the third-party file server.
 - » To add a third-party file server from Move, select **Select from existing third-party file servers** and do the following:
 1. From the file server table, select one of the third-party file servers in Move.

You can add the file server to the unified namespace before initiating Move migration. For more information, see [Files Migration](#) in the *Move User Guide*.
6. Click **Add**.
The third-party file server appears in the **File Servers Added** table on the **Update Unified Namespace** view.
7. In the **Update Unified Namespace** view, under **Namespace Access**, click **Update**.
The task to add the file servers to the namespace begins.

What to do next

Specify the third-party file server shares and exports to place under the unified namespace policy. For more information, see [Managing Third-Party Shares of a Unified Namespace](#) on page 98.

Managing Third-Party Shares of a Unified Namespace

Specify shares and exports of a third-party file server to include in or remove from a unified namespace.

About this task

To add or remove shares from a unified namespace, follow these steps:

Procedure

1. In the Files Manager navigation pane, go to the **Unified Namespaces** tab.
2. In the **Unified Namespaces** table, click the name of the target unified namespace.
3. In the **Overview** tab of the unified namespace, click **Manage Shares**.
4. In the **Shares** section of the **Manage Shares in Namespace** dialog box, click **Update Shares** and follow these steps:
 - a. Click **Download the CSV**, which downloads the CSV file template of shares and exports to be added to the unified namespace.
 - b. In your text editor, use the downloaded template to specify the names of shares and exports to keep under the unified namespace.
Use the *file-server-name,share-name,protocol,protocol* format. For example, use *FS,Users,NFS,-* or *FS,Users,SMB,NFS*.
Add or delete shares from the CSV file to manage the shares under the unified namespace. Changes to the CSV file overwrite the previous list of shares for the unified namespace.
 - a. Under **Upload the CSV**, click **Browse**.
 - b. From your local file manager, select the updated CSV file that contains the names of shares and exports to be added to the unified namespace.
5. Click **Upload**.
The updated list of shares appears in the **Shares** tab of the **Unified Namespace** view.

SECURITY HARDENING

Use the Nutanix command-line interface (nCLI) or Files CLI to customize your Files Manager security configuration settings.

The following table provides Prism Central–based hardening commands for Nutanix Files:

Table 6: Files on Prism Central Hardening Configurations

Description	Command
Edit the security configuration of a file server.	<pre>nutanix@FSVM\$ afs security.edit_security_config</pre>
Enable the U.S. Department of Defense knowledge consent banner.	<pre>nutanix@FSVM\$ afs security.edit_security_config enable_banner=true false</pre>
Enable the core-dump settings to let the file server VM generate stack traces for any cluster issue.	<pre>nutanix@FSVM\$ afs security.edit_security_config enable_core=true false</pre>
Enable the strong password policy.	<pre>nutanix@FSVM\$ afs security.edit_security_config enable_high_strength_password=true false</pre>
Generate the file server SSL certificate.	<pre>nutanix@FSVM\$ afs security.ssl_certificate_generate</pre>
Get the security configuration of a file server.	<pre>nutanix@FSVM\$ afs security.get_security_config</pre>
Get the file server secure sockets layer (SSL) certificate.	<pre>nutanix@FSVM\$ afs security.get_ssl_certificate</pre>
Schedule weekly execution of the advanced intrusion detection environment.	<pre>nutanix@FSVM\$ afs security.edit_security_config enable_aide=true false</pre>
Update the personal information exchange (pfx) file password for the SSL certificate.	<pre>nutanix@FSVM\$ afs security.ssl_certificate_change_pfx_file_password</pre>

COPYRIGHT

Copyright 2025 Nutanix, Inc.

Nutanix, Inc.

1740 Technology Drive, Suite 150


San Jose, CA 95110

All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. Nutanix and the Nutanix logo are registered trademarks of Nutanix, Inc. in the United States and/or other jurisdictions. All other brand and product names mentioned herein are for identification purposes only and may be trademarks of their respective holders.

Nutanix Files User Guide

Files 5.2

October 30, 2025



NUTANIX

Contents

Introduction to Nutanix Files.....	6
Nutanix Files Documentation Organization.....	8
Nutanix Files Deployment.....	10
Port Requirements.....	10
Prerequisites.....	10
The File Server Module.....	13
Installing Nutanix Files.....	13
Nutanix Files Upgrades.....	14
 Nutanix Files Overview.....	 16
File Server View in Prism Element.....	16
Nutanix Files Console.....	18
Viewing File Server Details.....	18
Viewing Share Details.....	19
File Analytics.....	20
 File Server Management.....	 21
Creating a File Server.....	21
Configuring File Server Basics.....	22
Configuring the Client Network.....	23
Configuring the Storage Network.....	24
Configuring Directory Services.....	26
Review.....	29
Deleting a File Server.....	30
File Server Updates.....	31
Updating File Server Basics.....	32
Scaling File Server VMs.....	32
Updating File Server VM Memory and vCPU Resources.....	33
Changing a File Server VM Password.....	33
Logging on to a File Server VM.....	34
Setting Timezones.....	35
Starting a Nutanix Files Cluster.....	35
Stopping a Files Cluster.....	36
Stopping a File Server VM.....	36
 File Server Network Management.....	 37
Managing the DNS Server.....	37
Network Segmentation.....	38
Configuring Network Segmentation.....	38
Multi-VLAN Network Management.....	39
Adding a Client Network.....	40
Listing Client Networks.....	41
Deleting a Client Network.....	41
Scaling Out A Multi-VLAN File Server.....	42
Configuring Client Exceptions.....	43
Configuring NFS Client Exceptions.....	43

Updating the File Server Network Configuration in Prism Element.....	44
Updating the Primary Network of a File Server.....	45
Share and Export Management.....	46
Creating an SMB Share.....	46
Creating an NFS Export.....	50
Creating a Multi-Protocol Share or Export.....	55
Multi-Protocol Support for Nutanix Files.....	61
Modifying a Share or Export.....	63
Deleting a Share or Export.....	64
Distributed Share Access.....	64
Configure Preferred Shares.....	65
Continuously Available SMB Shares.....	65
Connected Shares.....	65
Connecting Shares.....	66
Disconnecting Shares.....	66
WORM Shares.....	67
WORM Share Interoperability.....	68
Creating a WORM Share (CLI).....	68
Deleting a WORM Share.....	69
Nested Shares and Exports.....	70
Directory Service and Domain Management.....	72
Share and Export Permissions.....	72
Directory Services and Authentication.....	73
Updating File Server Directory Services.....	74
Setting Active Directory Machine Account Password Expiry.....	74
Joining a Domain.....	74
Leaving a Domain.....	75
Updating Domain Name System (DNS) Entries.....	76
Disjoint Domains.....	78
Configuring Disjoint Domains.....	78
User Management.....	80
User Mapping.....	80
Configuring User Mapping.....	81
Changing User Mapping Behavior.....	83
Managing Roles.....	84
Managing REST API Roles.....	84
Managing SMB Limited Local Users.....	85
Quotas.....	86
Creating User Quota Policies.....	89
Creating Directory-Level Quota Policies.....	90
Editing or Deleting a Quota Policy.....	91
Role-Based Access Control (RBAC) for Files.....	91
Nutanix Files Options.....	92
File Server and Share Clones.....	92
Cloning a File Server.....	93
Cloning a Share or Export.....	93
File Blocking.....	94
Blocking Files on a File Server.....	95

Enabling SMB Symlinks.....	95
Managing Access-Based Enumeration for SMB Shares.....	96
Durable SMB File Handles.....	97
Non-Disruptive Upgrades (NDU) for SMB.....	97
Antivirus (AV) Scanning for SMB Shares.....	97
Configuring Antivirus Scanning for SMB.....	98
Antivirus Tab.....	100
Performance Optimization.....	103
Hypervisor Attached Disks for AHV.....	103
Workload Optimization.....	103
Modifying the Workload Type.....	104
Removing Rebalancing Guardrails.....	104
File System Compression.....	104
Managing the Performance Profile.....	105
NFS Over RDMA.....	105
1N:2D Fault Tolerance.....	106
Disabling SmartOps Shares.....	106
Data Management.....	108
Nutanix Files Data Protection and Recovery in Prism Element.....	108
Configuring Protection Domain-Based Disaster Recovery.....	109
Activating Disaster Recovery.....	113
Metro Availability on ESXi.....	114
Smart Tiering.....	116
Standard Smart Tiering Overview.....	116
Smart Tiering Requirements and Recommendations.....	117
Configuring a Tiering Location.....	119
Setting a Capacity Threshold.....	121
Creating a Tiering Policy.....	122
Tiering Data Manually.....	123
Updating a Tiering Configuration.....	123
Recalling Tiered Data.....	123
Self-Service Restore.....	124
Enabling Self-Service Restore.....	125
Configuring Snapshot Schedule Policies.....	126
Taking On-Demand Snapshots.....	127
Restoring a Share.....	128
Manually Restoring a Share.....	128
Retrieving Files with Self-Service Restore.....	129
Deleting Self-Service Restore Snapshots.....	130
Setting Custom Hourly Snapshot Times.....	130
Share Migration.....	130
Migrating Shares.....	131
Re-attempt Share Migration.....	134
Configuring Backup for Distributed Shares.....	134
Files Data Collection.....	135
High Availability.....	135
Encryption.....	135
Nutanix Files REST APIs.....	137
Using the Files API Explorer.....	137

Security Hardening.....	138
Troubleshooting.....	143
Disable distributed file system (DFS) support on a witness share.....	143
Invalid Mounts After Authentication Change.....	143
Client Access Denial (NFS Protocol).....	143
Clients Cannot Mount Shares.....	143
Client Side Network Mapping.....	144
Connecting to Authentication Services.....	144
Constraint Violation.....	144
Disable Directory Leasing on Snapshot.....	144
DNS Missing SRV Records.....	145
Domain Controller Issues.....	145
Finding IP Addresses.....	145
Identifying the Share Owner.....	145
Invalid Credential.....	146
NLM Locks.....	146
Network Cannot Expand.....	146
NFS Cache Expiry During Share Migration.....	146
Restore Target Share Access Upon Canceling Migration Plan.....	147
NTLM Authentication Issues.....	147
Share Copying.....	147
Stale Statistics.....	148
Time Difference.....	148
Unsuccessful Authentication.....	148
Copyright.....	149

INTRODUCTION TO NUTANIX FILES

Nutanix Files is a software-defined, scale-out file storage solution that lets you share files in a centralized and protected location.

Nutanix Files eliminates the requirement for a third-party file server and uses a scale-out architecture that provides file services to clients through the Server Message Block (SMB) and Network File System (NFS) protocols. A logical file server instance, sometimes referred to as a Files cluster, consists of one or more file server VMs (FSVMs), which encompass shares and exports. You can create multiple file servers within a single Nutanix cluster.

Nutanix Files creates a volume group for every FSVM to provide stable storage for persistent states and audit events. During a service outage, the states, storage, and events of a volume group fail over to another FSVM. Nutanix Files also creates a dedicated container for every file server instance.

Tip: For performance best practices, sizing recommendations, migration guidance, and an in-depth technical overview of Files architecture, see [Solutions Documentation](#) and filter for "Unified Storage."

Nutanix Files offerings include the Files Manager for a unified control plane of all file servers and the deployment of file servers in Prism Central. File Analytics and Data Lens provide statistics and file server insights. For more information, see the [File Analytics User Guide](#), [Data Lens User Guide](#), and the [Files Manager User Guide](#).




Figure 1: File Server Components

File Shares and Exports

SMB shares and NFS exports encapsulate file directories and offer two types of shares or exports:

Distributed: A distributed share, also known as a home share or sharded export, spreads data across all of the FSVMs on the file server to improve the performance and scalability of client connections.

Note: For information on managing top-level directories (TLDs) in distributed shares, see the [MMC Snap-In for Nutanix Files](#).




Figure 2: Distributed Share or Export

Standard: A standard share, also known as a general purpose share or non-distributed export, contains all of the data on a single FSVM. A standard share or export serves data and connections from a single FSVM.




Figure 3: Standard Share or Export

For more information on standard and distributed shares, see [Share and Export Management](#) on page 46.

Features

Nutanix Files includes the following features:

- SMBv2, SMBv3, NFSv3 and NFSv4 protocol support.
- Multi-protocol support, see [Multi-Protocol Support for Nutanix Files](#) on page 61.
- AHV and ESXi hypervisor support.

- High availability for both VMs and data, see [High Availability](#) on page 135.
- Load balancing through scale-up and scale-out, see [Performance Optimization](#) on page 103.
- Data management, synchronization, and protection., see [Nutanix Files Data Protection and Recovery in Prism Element](#) on page 108.

You can also use Files as a storage class for Nutanix Kubernetes Engine (NKE) using the CSI Volume Driver. For more information, see [Nutanix Kubernetes Engine](#) and [CSI Volume Driver](#) documentation.

For a description of features added with every major Nutanix Files release, see the [Nutanix Files Release Notes](#).

Networking

Nutanix Files uses storage and client networks as follows:

- *Storage network*: Enables communication between the FSVMs and the Controller VMs.
- *Client network*: Enables communication between the clients and the FSVMs, allowing clients to access the file shares and exports, and enables client network communication with the directory services.

For more information, see [File Server Network Management](#) on page 37. To separate Files iSCSI connections to a segmented network, see [Network Segmentation](#) on page 38.




Figure 4: Files Networking

Nutanix Files Documentation Organization

The Nutanix Files documentation is organized into two user guides and the release notes.

- The [Nutanix Files User Guide](#) describes Files concepts and administrative tasks.
- The [Files Manager User Guide](#) covers file server management in Prism Central (PC).
- The [Nutanix Files Release Notes](#) cover updates, features, and issues in each Files release.

For more information on the contents of each document, see the "Documentation Organization" table below.

Table 1: Documentation Organization

Topics	Document
Feature overviews and instructions for administering and managing file servers and shares in the Prism Element web console and the Files Console, including the following: • Network management • Directory service and domain management • User management and quotas • Data management and performance optimization using the Files Console and Prism Element web console • NearSync, Async, and Metro Availability on ESXi • Smart Tiering • Self-service restore • Share migration • High availability • Controller VM (CVM) and file server VM (FSVM) based command-line interface (CLI) procedures • Troubleshooting	<i>Nutanix Files User Guide</i>
Feature overviews and instructions for administering and managing file servers in PC, including the following: • Nutanix Files on AWS and Azure clouds • Nutanix Files in a virtual private cloud (VPC) • Prism Central–based file server creation • Role-based access control (RBAC) for Files • Smart DR • Smart Sync • VDI Sync • Metro Availability on AHV • Unified namespace • Files on Azure BareMetal Infrastructure	<i>Files Manager User Guide</i>
Feature descriptions and compatibility requirements, resolved issues, and known issues.	<i>Nutanix Files Release Notes</i>

Additional Resources

For more information about Nutanix Files, see the following resources:

- For information about performance best practices, sizing recommendations, migration guidance, and an in-depth technical overview of Files architecture, see the [Solutions Documentation](#).
- For information about system limits, see the [Configuration Maximums](#).
- For information about CVE resolutions, see the [Nutanix Vulnerability Database](#).
- For information about Unified Storage solutions for your environment, see the [Validated Designs](#).
- For information on supported AOS upgrade paths, see the [Software Product Interoperability](#).
- For licensing information, see the [Nutanix Cloud Platform Software Options](#).
- For information on configuring and managing licenses for your clusters, see the [License Manager Guide](#).

Nutanix Files Deployment

An overview of Nutanix Files deployment.

To deploy Files in a Nutanix cluster, do the following:

Note: To segregate Files iSCSI connections to a segmented network and to deploy Files, see [Network Segmentation](#) on page 38.

1. Satisfy the prerequisites and port requirements, see [Prerequisites](#) on page 10 and [Port Requirements](#) on page 10.
2. Install the Files software using the Life-Cycle Manager (LCM), see [Nutanix Files Upgrades](#) on page 14.
3. Create a file server instance (Files cluster), see [Creating a File Server](#) on page 21.
4. Create one or more file shares (SMB) or exports (NFS), see [Share and Export Management](#) on page 46.

To deploy Files in a network-segmented environment, see [Network Segmentation](#) on page 38.

Port Requirements

Files has various firewall requirements depending on the protocols and services being used.

The [Port Reference](#) provides detailed port information for Nutanix products and services, including port sources and destinations, service descriptions, directionality, and protocol requirements.

Prerequisites

Satisfy the requirements, and review the recommendations and limitations before deploying Nutanix Files.

Requirements

Complete the following tasks before deploying Files:

- Install the File Server Module (FSM). For more information, see [Nutanix Files Upgrades](#) on page 14.
- Configure and define the storage and client networks.
- Configure a dedicated iSCSI data services IP address for Files.
- Set up a network time protocol (NTP) server.
- For ESXi, register all ESXi host nodes in the AOS cluster to the same vCenter.

The following requirements depend on the protocol and directory-service configuration of the file server:

- For the SMB protocol, ensure that the distributed file system (DFS) for Windows clients is turned on, which is the default configuration.
- For the NFS protocol, ensure that NFS clients have unique names.
- If using Active Directory (AD) for user authentication, have the credentials of the domain administrator or a user with delegated permissions. For more information, see [Required Active Directory Information](#) on page 12.
- If you use the Lightweight Directory Access Protocol (LDAP) for NFS with permissions required for search, have the credentials of the bind distinguished name (DN).

Minimum File Server Configuration

File servers require the following minimum compute and memory resources:

- A minimum of four vCPUs per file server VM (FSVM)
- A minimum of 12 GiB of memory per FSVM
- A number of CVMs equal to or greater than the number of FSVMs

Network Requirements

The storage network requires at least one additional IP address than the number of FSVMs; in other words, the number of IP addresses must be equal to the number of FSVMs plus one. The client network requires the same number of IP addresses as the number of FSVM nodes. Single-FSVM deployments require one IP address for the storage network and one IP address for the client network.

If the client and storage networks are different, clients on the storage network cannot access shares or exports.

If you use the same network for both client and storage, the IP addresses must be unique.

To segregate Files iSCSI connections to a segmented network, see [Network Segmentation](#) on page 38.

Required Information for File Server Configuration

Collect the following information for each required service and entity before deploying Files.

Service or Entity	Required Information	Description
Active Directory (Optional)	Windows AD domain	The directory service domain for the file server
	AD admin account	An admin with domain administrator or delegated permissions
LDAP (Optional)	LDAP URL	An LDAP server name or IP address (with optional port number)
	Base DN	The distinguished name of the entry for starting the records search
	Bind DN	The distinguished name of entry for performing search (optional when anonymous bind is disallowed)
DNS server	Bind password	The password to use for bind DN
	DNS server names	The server for FSVM name resolution and access to external services

Note: If the LDAP server mandates a bind for queries, Enter the bind DN.

Service or Entity	Required Information	Description
NTP server	NTP server name	The server for time synchronization between the file server and the AD service
iSCSI data services	iSCSI data services IP address	The iSCSI IP address for connecting FSVMs to storage For more information on iSCSI storage, see the Nutanix Volumes Guide .
Managed networks	Storage network	The VLAN that connects the CVM to the FSVM
	Client network	The VLAN that connects the FSVM to the AD and DNS
Unmanaged networks	Storage network gateway	The VLAN that connects the CVM to the FSVM
	Storage network subnet	The network subnet for storage communication
	Storage network IP address range	The network range for storage communication
	Client network gateway	The VLAN that connects the FSVM to AD and DNS
	Client network subnet	The network subnet for client communication
	Client network IP address range	The network range for client communication

Required Active Directory Information

Configuring AD is optional. Nutanix Files uses the following AD parameters to join the domain:

- Fully qualified domain name (FQDN)
- A user name that authenticates on the AD on the domain
- A password for the account used to connect to the AD server
- (Optional) Organizational unit (OU)
- (Optional) Preferred domain controller

A domain user authenticates to the file server on the domain controller and creates Files computer accounts, related SPN entries, and Files DNS entries when using Microsoft DNS. If you are in the same domain, you can use the user principal name (UPN) or SamAccountName. If you are outside of the domain but in the same forest, then use the UPN.

Note: Files only allows the "at" symbol (@) to specify the domain. Files does not allow UPNs with multiple @ symbols. For example, user@mydomain.com is valid but user@name@mydomain.com is invalid. Files does not allow the forward slash (/) and the backward slash (\) in the UPN.

Files uses the user name password to authenticate to AD and to create the Files computer account. By default, Files creates the computer account in the Computers container. The OU contains the computer account Files creates. List the AD users permissions for the machine account on the OU. In an organization with complex hierarchies, create the computer account in a container by using a forward slash (/) to denote hierarchies. For example, organizational_unit/inner_organizational_unit. Do not use the forward slash in the organizational unit name.

Recommendations

We recommend that you do the following before deploying Nutanix Files:

- Enable the distributed resource scheduler (DRS).
- Configure separate storage and client networks even though a single client and storage network is permitted.
- Set up a domain name system (DNS) server.

Limitations

The [Configuration Maximums](#) page provides details on Files system limits.

The [Nutanix Files Release Notes](#) provide release-specific details on supported configurations and software compatibility.

The File Server Module

Nutanix Files requires a compatible version of the File Server Module (FSM).

The FSM provides Files life-cycle and GUI components. Files relies on AOS for some control plane components.

Before upgrading or installing Files, upgrade or install a compatible version of the FSM using the Life Cycle Manager. For more information about steps on performing inventory and upgrades in LCM, see the [Life Cycle Manager Guide](#).

Installing Nutanix Files

Use the 1-click feature in Prism Element (PE) to install Files.

About this task

Installing Files also requires installing compatible versions of the required software.

For information on software compatibility requirements, see "Nutanix Files Software Compatibility" in the [Nutanix Files Release Notes](#).

Procedure

1. Log onto Prism Element (PE) with your credentials.
2. Click the **gear icon**.
3. In the **General Settings** menu, click **Upgrade Software**.
4. In the **Upgrade Software** pane, click the **File Server** tab.

5. Complete the installation process by doing one of the following.

- » To download the Files version you want to install, go to Step 6.

Note: Nutanix recommends using the download option in Prism Element.

- » To upload the Files binary from a different source, go to Step 7.

Note: If you do not see a target Files version, refer to the Nutanix Files [Upgrade Paths](#) page on the Support Portal to confirm that the target version is 1-click enabled.




Figure 5: Upgrade Software: File Server tab

6. (Optional) To download and install software, click **Download** next to the target Files version. The download and installation process takes a few minutes. Once the download completes, click **Continue** to deploy a file server in Prism Element. For more information, see [Creating a File Server](#) on page 21. To install the Files Manager and deploy file servers in Prism Central, see [Deploying File Servers with the Files Manager](#) in the *Files Manager User Guide*.
7. (Optional) To install uploaded software using the File Server binary, click the **upload the File Server binary** link, and do the following:

- a. **File server metadata file:** Click the **Choose File** button and select the target Files metadata .json file.

Note: The [Downloads](#) page provides Files, FSM, and Files Manager installation files.

- b. **File server binary file:** Click the **Choose File** button and select the target Files binary .qcow2 file.
- c. To upload the installation files, click **Upload Now**. The installation process takes a few minutes.

Nutanix Files Upgrades

Upgrade Nutanix Files and the required components using the Life-Cycle Manager (LCM).

Caution: Nutanix recommends performing Files upgrades during non-business hours. Upgrading Files can cause a brief FSVM downtime and a temporary disconnection of connected clients.

Before upgrading Files, check LCM to ensure that you have compatible versions of the required software. For information on the software requirements for each release, see "Nutanix Files Software Compatibility" in the [Nutanix Files Release Notes](#). Upgrading Files also requires upgrading the File Server Module (FSM) and the Files Manager using LCM.

For information on performing inventory and upgrades using LCM, see the [Life Cycle Manager Guide](#).

Consider the following before performing a Files upgrade:

- After an AOS upgrade, compression on the Files container resumes after a one-hour delay.
- (ESXi only) When performing one-click ESXi hypervisor upgrades, we recommend turning off anti-affinity rules on all FSVMs. Turning off the anti-affinity rules allows the FSVMs to move to different hosts during upgrades. After the hypervisor upgrades successfully, turn on the anti-affinity rules on the FSVMs.
- (ESXi only) ESXi hosts that belong to multiple vSphere clusters or are deployed across multiple data-centers might experience limitations. See [KB 5369](#) for more information.

NUTANIX FILES OVERVIEW

Manage file servers and shares from the Files Console and the **File Server** view in Prism Element (PE).

The **Files Server** view in Prism Element is the landing page for Files. The **File Server** view provides an overview of your file servers on a PE cluster and basic file server management options.

The Files Console is a GUI for the management of file servers and shares. To access the Files Console from the **File Server** view in PE, click **Launch Files Console** next to the target file server.

To manage file servers created in Prism Central, access the Files Console by navigating to file servers from the Files Manager, see the [Files Manager User Guide](#).

File Server View in Prism Element

The **File Server** view in Prism Element provides basic information about each file server in the cluster.

To get to the **File Server** view, select **File Server** from the pull-down menu in Prism Element.

File Server View Layout

The **File Server** view includes the following sections and options:

- Action buttons to create a file server. For more information, see [Creating a File Server](#) on page 21.
- An entities table displays information about each file server. You can filter the table contents by entering a string in the search field located above the table.

Note: The **Recommendations** column is disabled. See **Recommendations** in the Files Console dashboard.

- A **File Summary** pane displays high-level details about file servers on the cluster, and, after selecting a file server, the **File Server Details** pane displays summary information for that file server.
- A **gear icon** with options to download the table content in the CSV or JSON format.
- A list of operations that you can perform, including the following: **Launch Files Console**, **Clone**, **Update**, **Protect**, and **Delete**.

Note: The values for the parameters do not account for features applied by AOS or space used at the storage container level.




Figure 6: File Server View

Details Pane

Selecting a file server in the table presents detailed information in the **File Server Details** pane. The following tables describe the fields.

Table 2: File Server Details Fields

Parameter	Description	Values
Name	The file server name.	(name)
DNS domain name	The name of the domain that the file server is registered to. "Not Protected" indicates that the file server is not currently in a protection domain.	(DNS domain name)
Open connections	The number of open connections.	(integer)
Share/Export Count	The total number of shares or exports.	(integer)
Space used	The total amount of storage space used within the file server currently.	xx [GiB TiB]

Parameter	Description	Values
Space used by snapshots	The amount of space used within the file server to store snapshots currently.	xx [MiB GiB TiB]
Total available space	The amount of available (unused) storage space currently on the file server.	xx [MiB GiB TiB]
Size	The size of the file system of the file server.	xx [TiB]
Protection domain	The name of the protection domain that includes this file server. Clicking the name displays the Data Protection view for that protection domain. "Not Protected" indicates that the file server is not in a protection domain. See the Data Protection and Recovery with Prism Element guide for information on protection domains.	(protection domain name)
Storage container	The name of the storage container of the file server. Clicking the name displays the Storage Container view for that storage container.	(storage container name)
Protocol	The protocols used by the file server.	SMB, NFS, or both
SMB directory service	The SMB protocol always uses Active Directory as the directory service.	Active Directory
NFS directory service	The NFS protocol has multiple options for the directory service.	Unmanaged, Active Directory, or LDAP
Client-side network	The name of the primary network used by clients. Click View IPs to view all the IPs for the primary network. Click More to view the additional client networks.	(network name)
Storage network	The name of the network used for storage.	(network name)
Memory	Used memory.	GiB
CPU	Total CPUs.	Numerical
Data reduction ratio	Data reduced using file-system-level compression.	x:x (numerical)
Data savings	Amount of data saved using file-system-level compression.	MiB

Nutanix Files Console

The Files Console provides administrative tools and dynamically updated information for a file server and its shares and exports in different tabs.

Access the Files Console from the **File Server** view in Prism Element by selecting **Launch Files Console** or by clicking the name of a file server in the Files Manager.

Viewing File Server Details

View file server details in the Files Console.

About this task

To view more details for a file server, follow these steps:

Procedure

1. Navigate to the Files Console.

For more information, see [Nutanix Files Console](#) on page 18.

2. (Optional) To view a visualization of the data usage on a filer server, in the **Dashboard** view, select the **Capacity Summary** pane.

1. To view the sum of live data, snapshot data, and tiered data minus the tiered data in the snapshot, see **Total data (logical)**.
2. To view the actual disk usage of the on-premises hardware, see **Total used (physical)**.

Note: The total physical capacity is reduced by compression and does not include tiered data.

3. To view the ratio of the logical space used to the physical space used, see **Data reduction**.

Note: A 1:1 ratio indicates that compression is turned off. The calculation excludes tiered data.

4. To view the data from all snapshots, including Nutanix Self-Service Restore, third-party changed file tracking (CFT) snapshots, Data Lens, File Analytics, and Smart DR, see **Snapshot data**.

Note: Snapshot data does not include protection domain-based snapshots.

5. To view the data that exists in a snapshot but is marked for deletion, see **Tiered data in snapshot**.
 6. To view the allotted capacity for the filer server minus the total data, see **Free capacity**.
 7. To view the total data tiered to an object store not including data tiered in a snapshot, see **Total tiered data**.
3. (Optional) To view a rolling time interval monitor of the storage space used for data and snapshots on the filer server, in the **Monitoring** dropdown menu, select **Usage > Storage Used**.
 1. To display the value for the time specified on the horizontal axis, hover over the data.
 2. To isolate a data set, select or clear the **Spaced Used by Dataset** and **Space Used For Snapshots** boxes.
 4. (Optional) To view the average latency across a rolling time interval monitor, in the **Monitoring** dropdown menu, select **Performance > Latency**.
 1. To display the value for the time specified on the horizontal axis, hover over the data.
 2. To add the **Write Latency**, **Read Latency**, and **Metadata Latency** data set options, select the **Show I/O and Metadata Breakdown** option.
 5. (Optional) To add the **Write Throughput** and **Read Throughput** data options to the **Throughput** graph in the **Performance** tab, select the **Show I/O and Metadata Breakdown** option.
 6. (Optional) To add the **Write IOPS**, **Read IOPS**, and **Metadata IOPS** data options to the **IOPS** graph in the **Performance** tab, select the **Show I/O and Metadata Breakdown** option.

Viewing Share Details

View share details in the Nutanix Files Console.

About this task

To view more details for a share, follow these steps:

Procedure

1. Navigate to the Files Console.

For more information, see [Nutanix Files Console](#) on page 18.

2. (Optional) To switch from the **General** view to the **Metrics** view, in the share table in the **Shares & Exports** view, select the **View By** dropdown menu.
3. (Optional) To open the share details view, which consists of the **Summary**, **Snapshots**, **Quota Policies**, **Antivirus**, and **Metrics** tabs, select the name of the share in the Files Console.
4. (Optional) To view the unique share data that does not include shared data or snapshots, in the **Capacity Summary** pane of the **Dashboard** view, select **Live data**.

File Analytics

File Analytics provides data and statistics on the operations and contents of a file server.

Once you deploy File Analytics, Files adds a File Analytics VM (AVM) to the Files cluster. A single AVM supports all file servers in the cluster, but you must enable Analytics separately for each file server. Files protects the data on the AVM and keeps it in a separate volume group.

For deployment steps and administrative guidance, see the [File Analytics Guide](#).

FILE SERVER MANAGEMENT

Create a new file server and manage its configuration.

Creating and managing a file server requires administrator privileges. You can perform most file server management tasks in the **File Server** view in Prism Element and through the Files Console. You can start, stop, and manage file server VMs (FSVMs) through the VM tab in Prism Element, see "VM Management" in the [Prism Web Console Guide](#) for more details.

A file server usually consists of three or more FSVMs combined into a logical group. The [Scaling File Server VMs](#) on page 32 section describes restrictions for increasing and reducing the number of FSVMs.

Single-FSVM deployments have the following limitations:

- No High Availability (HA) support. Single-FSVM deployments do not include other FSVMs that an out-of-service FSVM can fall back on, see [High Availability](#) on page 135.
- No distributed share support. Since distributed shares and exports spread data across multiple FSVMs, single-FSVM deployments only support standard shares.

Important: When the file server capacity or the container capacity reaches 100 percent, all shares or exports within the file server become read-only and Files blocks write privileges. Prism Element displays alerts when the file server capacity or the container capacity reaches 90 percent and when either reaches 100 percent capacity.

Creating a File Server

Create a file server in Prism Element.

Before you begin

For deployments using the ESXi hypervisor, add all ESXi hosts under the same ESXi cluster and, during deployment, create one FSVM per ESXi host to use the high availability and the VMware Distributed Resource Scheduler (DRS) features.

About this task

The following collection of tasks describe how to configure a file server in Prism Element:

- [Configuring File Server Basics](#)
- [Configuring the Client Network](#)
- [Configuring the Storage Network](#)
- [Configuring Directory Services](#)
- [Review](#)

Complete all of the configuration tasks to create a file server.

Note: To create a Prism Central-managed file server, see [Creating a File Server](#) in the *Files Manager User Guide*.

To create a file server in Prism Element, follow these steps:

Procedure

1. In Prism Element, go to the **File Server** view .

For more information, see [File Server View in Prism Element](#) on page 16.

2. In the upper-right corner, click **+ File Server**.
3. If the **New File Server: Pre-Check** window appears, review the displayed information and address any unsatisfied prerequisites before continuing.

Files checks your current environment and either verifies that it satisfies the prerequisites or identifies where the environment does not meet the requirements. A blue check mark indicates a satisfied prerequisite. Any unchecked items need attention. Files checks for the following: the Files software version, available data services IP addresses, available networks, and, for ESXi, the vCenter registration status of the cluster.
4. Complete each of the following configuration procedures:
 - [Configuring File Server Basics](#)
 - [Configuring the Client Network](#)
 - [Configuring the Storage Network](#)
 - [Configuring Directory Services](#)
 - [Review](#)

Configuring File Server Basics

About this task

Complete the previous task before configuring file server basics.

To configure file server basics, in the **Basics** tab of the **Create a File Server** window in Prism Element, follow these steps:

Procedure

1. In the **Name** field, enter a name for the file server.

Clients use the file server name to access the file server. The fully qualified name (file server name + domain) must be unique.
2. In the **DNS Domain** field, enter a fully qualified domain.
3. In the **File Server Storage** field, enter the file server total storage size.

The minimum is 1 TiB.
4. **Capacity Configuration:** Review the recommended connections and throughput performance values.

Files automatically recommends the number of file server VMs, vCPUs per VM, and memory per VM.
5. (Optional) To keep the recommended capacity configuration, click **Next**.
6. (Optional) To change the configuration, click **Customize**.

The **File Server Capacity Configuration** window displays.

 - a. In the **SMB Concurrent Connections** field, enter the target number of connections.
 - b. In the **NFS Throughput in MBps** field, enter the throughput amount in megabits per second (MBps).
 - c. Click **Save**.

7. (Optional) To manually specify the FSVM parameters, click **Customize > Configure manually**.

The **File Server Capacity Configuration** window displays.

- a. In the **Number of File Server VMs** field, specify the number of the file server VMs (FSVMs) for the file server.
- b. In the **vCPUs Per VM** field, specify the number of vCPUs for each FSVM.
- c. In the **Memory Per VM** field, specify the memory capacity in gigabytes (GiB) for each FSVM.
- d. Click **Save**.

Configuring the Client Network

About this task

Complete the previous task before configuring the client network.

Provide two static IP addresses for each new FSVM, one for the client-side network and one for the storage network. If clients or the AD domain controllers are in the same subnet as the controller VM (CVM) or the storage network, configure the same client-side and storage-side networks. Otherwise, use separate client and storage networks.

If you use the same network for both clients and storage, use unique IP addresses; otherwise clients on the same subnet as the storage network cannot access the shares.

In the **Client Network** tab of the **Create a File Server** window in Prism Element, follow these steps:

Procedure

1. Depending on the type of hypervisor, do one of the following:

- » For AHV, in the **VLAN** dropdown list, select a VLAN.
- » For ESXi, in the **Port Group** dropdown list, select a port group.

Configured network parameters display if the target is a managed network. Unmanaged networks require additional details.

2. (Optional) Configure an unmanaged client network:
 - a. In the **Subnet Mask** field, enter the subnet mask value.
 - b. In the **Gateway** field, enter the IPv4 gateway IP address.
 - c. Click **IP Addresses**.
A table for entering IP addresses displays.
 - d. In the **From** field, enter the starting IP address.

Note: Files usually requires four IP addresses. Files requires one IP address for single-FSVM deployments.

- e. In the **To** field, enter the ending IP address.
 - f. Click **Save**.
 - g. (Optional) To enable the IPv6 protocol, select the **Add IPv6** checkbox and follow these steps:
 1. In the **Prefix Length** field, enter an integer for the prefix length.
The default IPv6 prefix is 64.
 2. In the **Gateway** field, enter the alpha-numeric, colon-separated IPv6 gateway address.
 3. (Optional) To enter IP addresses, click **+ IP Addresses**, enter the starting IP address in the **From** field, and enter the total number of IP addresses in the **No. of IPs** field.
Usually Files requires three IP addresses.
 4. Click **Save**.
 3. In the **DNS Resolver IP** field, enter an IP address or a comma-separated list for multiple addresses for one or more DNS servers.
- Note:** The DNS resolver IP and the NTP servers fields might pre-populate based on the target VLAN or port group.
4. In the **NTP Servers** field, enter an IP address, or a comma-separated list for multiple addresses, for one or more NTP servers.
 5. Click **Next**.

Configuring the Storage Network

About this task

Complete the previous task before configuring the storage network.

For ESXi deployments, do not use the Controller VM (CVM) backplane network. Nutanix Files does not support the CVM backplane network, which can cause future upgrade operations to fail.

In the **Storage Network** tab of the **Create a File Server** window in Prism Element, follow these steps:

Procedure

1. Depending on the type of hypervisor, do one of the following:

- » For AHV, in the **VLAN** dropdown list, select a VLAN.
- » For ESXi, in the **Port Group** dropdown list, select a port group.

Configured network parameters display if the target is a managed network. Unmanaged networks require additional details.

The screenshot shows a web form for configuring network details. At the top, there is a dropdown menu labeled 'VLAN' with the selected option 'vlan.3101-managed - Managed'. Below this is a table with four rows of configuration details:

VLAN ID	vlan.3101-managed
Subnet Mask	255.255.255.0
IP address range	10.10.10.1 to 10.10.10.14
No. of IP addresses	1 IP(s) will be used out of 14 available IP(s)

Figure 7: Storage Network Details (managed VLAN selected)

2. (Optional) Configure an unmanaged client network:

- a. In the **Subnet Mask** field, enter the subnet mask value.
- b. In the **Gateway** field, enter the IPv4 gateway IP address.
- c. Click **IP Addresses**.
A table for entering IP addresses displays.
- d. In the **From** field, enter the starting IP address.

Note: Files usually requires four IP addresses. Files requires one IP address for single-FSVM deployments.

- e. In the **To** field, enter the ending IP address.
- f. Click **Save**.
- g. (Optional) To enable the IPv6 protocol, select the **Add IPv6** checkbox.
 1. In the **Prefix Length** field, enter an integer for the prefix length. The default IPv6 prefix is 64.
 2. In the **Gateway** field, enter the alpha-numeric, colon-separated IPv6 gateway address.
 3. (Optional) To enter IP addresses, click **+ IP Addresses**, enter the starting IP address in the **From** field, and the total number of IP addresses in the **No. of IPs** field.

Files usually requires at least three IP addresses, but single-FSVM deployments require one IP address.
 4. Click **Save**.

3. In the **DNS Resolver IP** field, enter IPv4 addresses for one or more DNS servers.
For multiple entries, use a comma-separated list.

Note: The DNS resolver IP and the NTP servers fields might pre-populate based on the target VLAN or port group.

4. In the **NTP Servers** field, enter the NTP server names or IP addresses.
For multiple entries, use a comma-separated list.

5. Click **Next**.

Configuring Directory Services

About this task

Complete the previous task before configuring the directory services.

In the **Directory Services** tab of the **Create a File Server** window in Prism Element, follow these steps:

Procedure

1. (Optional) To select one or more protocols to use, in the Directory Services tab, check the **Use SMB Protocol** checkbox, the **Use NFS Protocol** checkbox, or both.

You can skip this step and select the protocols after file server creation, but you cannot use the file server until you select the protocols.

Note: Files supports using LDAP for Active Directory (AD).

2. (Optional) Configure directory services for the SMB protocol:

- a. In the **Active Directory Realm Name** field, enter the AD realm name for join domain operations.
- b. In the **Username** field, enter an AD username in the *domain|username* format or in the UPN format *username@ADrealm*.

You must have an administrator account or the following required permissions for the relevant AD organizational unit (OU) realm:

- **Create computer objects:** Allows the delegated admin to create a machine account in the specified OU path, which is a requirement for Files to join the domain.
- **Read and write service principal names:** Helps set the service principal names (SPNs) for the machine account. Kerberos access to Files requires SPNs.

Caution: Removing the file server computer object in AD can cause file server service disruption.

If the DNS domain name of the file server is different from the AD domain name, the following permission is also required: **Read/Write DNS host name**.

- c. In the **Password** field, enter the password for the username.
- d. (Optional) To make the user specified in the **username** field a file server administrator, select the **Make This User a File Server Admin** checkbox.
- e. (Optional) To configure the preferred domain controller and organizational unit name, select the **Show Advanced Options** checkbox and do the following:
 1. (Optional) In the **Preferred Domain Controller** field, specify a domain controller for join-domain operations in a multi-domain controller Active Directory configuration.

You must have the permissions to modify DNS records on the domain controller. Files discovers and uses the local domain controller by default.
 2. (Optional) In the **Organizational Unit** field, enter your preferred organizational unit name.

Files creates the machine account in the Computers OU container by default.
 3. (Optional) In the **Sub-OU** field, enter a sub-organizational unit name using the following format: *parentcompanyOU/sub-ou1/sub-ou2/sub-oux*.

For example, if the distinguished name is OU=Infrastructure, OU=Servers, OU=Test, OU=Services,DC=contoso, DC=com enter the OU path using the following format, Services/Test/Servers/Infrastructure.
- f. (Optional) To overwrite an existing machine account during the join-domain operation if one exists with the same name as the file server, select the **Overwrite Existing Files Machine Account (if present)** checkbox.
- g. (Optional) To use AD credentials for adding DNS entries, select the **Add Files Server DNS Entries Using the Same Username And Password** checkbox.

When using Microsoft DNS, ensure that the AD user account has DNS admin rights.

Note: The default machine account password expiry period is 0, meaning the password does not expire. To update the machine account password expiry period, see [Setting Active Directory Machine Account Password Expiry](#) on page 74.

3. (Optional) To configure NFS protocol directory services, select the authentication method from the **User Management and Authentication** dropdown list and do the following:

Note: If you selected the unmanaged method, you do not need to configure user management and authentication details.

- a. (Optional) Configure **User Management and Authentication: LDAP:**

1. In the **LDAP details** section, if the target LDAP/LDAPS server details do not appear, click **+ New LDAP Server** and, in the **Server URI** column, enter the uniform resource identifier (URI) address of the server.

Note: To add any additional LDAP servers, repeat this step.

For LDAPS servers, use the LDAPS format as follows: `ldaps://ldap-servername.company.com:port`.

2. In the **Actions** column, click **Save**.
3. In the **Base DN** field, enter the base distinguished name (search starting point).
4. (Optional) In the **Bind DN** field, enter the distinguished name at which to bind to the server.

The bind DN is required when the LDAP server mandates a bind for queries.

5. In the **Bind password** field, enter the bind password.

- b. (Optional) For **User Management and Authentication: Active Directory** configure the following:

Note: When NFS uses AD, you do not need to configure user mapping for multi-protocol access. If you previously configured AD with SMB, the following fields do not appear because they are already configured for the SMB protocol.

- **Enable Identity Management for Unix (RFC 2307):** Select this checkbox if you have configured RFC 2307 for Active Directory.
- **Active Directory Realm Name:** Displays the Active Directory realm name (read-only).
- **Username:** Enter an AD username in the `domain|username` format or in the UPN format `username@ADrealm`.

You must have an administrator account or the following required permissions for the relevant AD organizational unit (OU) realm:

- `Create computer objects`
- `Read/write service principal names`

Note: If the DNS domain name of the file server is different from the AD domain name, the `Read/write DNS host name` permission is also required.

- **Password:** Enter the password for the user.
- (Optional) **Show Active Directory Advanced Options:** To display additional configuration options, select the checkbox. Otherwise, continue to Step 3.c.
- (Optional) **Preferred Domain Controller:** Specify a specific domain controller for join-domain operations in a multi-domain controller Active Directory configuration.

You must have the permissions to modify DNS records on the domain controller. Files discovers and uses a site local domain controller by default.

- (Optional) **Organizational Unit:** Enter your preferred organizational unit name.

Files creates the machine account in the `Computers` OU container by default.

Note: The default machine account password expiry period is 0, meaning that the password does not expire. To update the machine account password expiry period, see [Setting Active Directory Machine Account Password Expiry](#) on page 74.

- (Optional) **Overwrite Existing Files Machine Account (if present):** To overwrite an existing machine account during the join-domain operation if one exists with the same name as the file server, select this checkbox
- c. (Optional) To modify the default NFS protocol version for all exports on the file server, check the **Show NFS Advanced Options** box.
By default, Files supports both NFSv3 and NFSv4 protocols for exports.
 1. (Optional) To disable NFSv3 by default for all exports, clear the **Enable NFSv3 by default for all exports** checkbox.
 2. (Optional) To disable NFSv4 by default for all exports, clear the **Enable NFSv4 by default for all exports** checkbox.
 3. In the **NFSV4 Domain** field, enter the NFSv4 domain name.
Files uses the DNS domain name to map NFSv4 names to user identifiers (UIDs) and group identifiers (GIDs); clients and servers must agree on the mapping.
- d. Click **Next**.

Review

About this task

Complete the previous task before reviewing the file server configuration.

In the **Review** tab of the **Create a File Server** window in Prism Element, follow these steps:

Procedure

1. In the **Summary** tab, review and verify the displayed configuration information.

Create File Server

Basics Client Network Storage Network Directory Services **Summary**

File Server test1

Storage Capacity: 1 TiB

Configuration: 1 File Server VMs, 4 VCPUs each, 12 GiB RAM each

Protection

☒ Create a Protection Domain and a default schedule (highly recommended)

Protection Domain will be automatically created with a default local schedule of 1 snapshot every day and retaining the 2 latest snapshots.

You can add more schedules as well as remote sites on the DR page. [Learn More](#)

Protection Domain Name

NTNX-test1

Back Cancel Create

Figure 8: Create File Server: Summary Tab

2. (Optional) To change default protection domain, enter a new name in the **Protection Domain Name** field.
The default protection domain name is `NTNX-file_server_name`.
3. Click **Create**.
Creating the file server begins. Monitor the progress on the **Tasks** page. You can either select **Tasks** from the main menu, or click the tasks icon and then **View All Tasks**.

What to do next

To add DNS entries, see [Updating Domain Name System \(DNS\) Entries](#) on page 76.

For deployments using the ESXi hypervisor, disable VM monitoring. See the VMware vSphere documentation for procedures describing how to enable or disable VM monitoring.

Deleting a File Server

Delete a file server in Prism Element.

About this task

Do not delete FSVMs or power off a file server prior to deleting it. If you delete FSVMs, you will need to force delete the file server.

To delete a file server, do the following:

Caution: Deleting a file server is permanent, and none of the deleted files can be recovered.

Procedure

1. Go to the File Server view (see [File Server View in Prism Element](#) on page 16), select the target file server, and click the **Delete** action button.
2. In the dialog box, choose one of the following options:

- » Retain all the related entities.
- » Delete specific related entities.

Deleting a file server deletes all the underlying user data related to that file server. However, depending on your AOS version, you may not see the option to delete the storage container and protection domain created for that file server. You can delete that storage container and protection domain manually.

3. Click the **Delete** button.




Figure 9: Delete Window

- a. If you decide to delete specific related entities, choose which entities to delete from the following options (otherwise continue to the next step):
 - **Snapshots, schedules, and protection domain**
 - **Container**
4. (Optional) delete file server DNS entries. Click **Open DNS settings**, see [Updating Domain Name System \(DNS\) Entries](#) on page 76.

Deleting DNS and AD records requires administrator privileges.
 5. Click **Delete**.

File Server Updates

Update or expand the configuration of an existing file server.

You can modify an existing file server in the following ways:

- To update the name, domain, or storage capacity, see [Updating File Server Basics](#) on page 32.

- To update the number of file server VMs (FSVMs), see [Scaling File Server VMs](#) on page 32.
- To update the vCPU count and memory size for each FSVM, see [Updating File Server VM Memory and vCPU Resources](#) on page 33.
- To update the network configuration, including the client network, storage network, domain name server (DNS), or NTP server, see [Updating the File Server Network Configuration in Prism Element](#) on page 44 .
- Update blocked file types, see [Blocking Files on a File Server](#) on page 95.

Updating File Server Basics

Update the name, the domain, and the size of a file server.

Procedure

1. In the [Nutanix Files Console](#) on page 18, go to **Configuration > Platform**.
2. Click **Update > File Server Basics**.
The **File Server Basics** window appears.
3. Do the following in the indicated fields:

Tip: Before changing the file server or domain name, click the **Leave Domain** link to leave the current domain (see [Leaving a Domain](#) on page 75). Changing a file server name does not remove existing DNS entries, which you must remove manually.

- a. To change or update the file server name, enter the new name in the **File Server Name** field.
- b. To change the DNS domain name, enter the new name in the **Domain** field.

Tip: After changing the file server or domain names, delete any old DNS entries from the DNS servers, add the new DNS entries (see [Updating Domain Name System \(DNS\) Entries](#) on page 76), and update the Directory Services configuration as needed (see [Updating File Server Directory Services](#) on page 74).

- c. To change the size of the file server, enter a value in the **File Server Size (Logical)** field.

Note: You cannot reduce the size of the file server. This operation does not impact current client connections.

4. When the entries are correct, click **Update**.

Scaling File Server VMs

Scale out to add new file server VMs (FSVMs) to a file server cluster or scale in to remove FSVMs.

About this task

Scale-out and scale-in operations might cause a momentary connection drop while accessing file servers. Nutanix recommends performing scale-out and scale-in operations during non-business hours.

The following limitations apply to changing the number of FSVMs:

- You cannot reduce file servers of three FSVMs.
- You cannot expand single-FSVM deployments.
- Single or dual-node AOS clusters only support single-FSVM deployments.
- You can scale in a cluster by up to three FSVMs at a time.

- The number of CVMs must be equal to or greater than the number of FSVMs.

To change the number of FSVMs, do the following:

Procedure

1. In the [Nutanix Files Console](#) on page 18, go to **Configuration > Platform**.
2. Click **Update > Scale Out / Scale In**.
The **Scale Out or Scale In** window appears.
3. Under **New Capacity**, enter an integer for the new number of FSVMs.
4. (Optional) If you scale out, under **Client Network**, confirm the new **Start IP Address** and **End IP Address**.
You can enter different IP addresses.
5. (Optional) If you scale out, under **Storage Network**, confirm the new **Start IP Address** and **End IP Address**.
You can enter different IP addresses.
6. Click **Update**.

Updating File Server VM Memory and vCPU Resources

Scale up or scale down the vCPUs and memory resources per file server VM (FSVM).

Before you begin

Make sure to keep CD-ROM with ISO attached to each FSVM before updating the resources.

About this task

Scale-up is a non-disruptive operation, but scale-down causes system downtime. To update the memory and vCPU capacity, follow the steps as indicated.

Procedure

1. In the [Nutanix Files Console](#) on page 18 go to **Configuration > Platform**.
2. Click **Update > Scale Up / Scale Down**.
The **Scale Out or Scale In** window appears.
3. Do the following in the **Updated Configuration** section.
 - a. In the **vCPUs** field, select the target (total) number of vCPUs for each FSVM in the cluster.
 - b. In the **Memory** drop-down, select the target (total) number of memory (GiB) for each FSVM in the cluster.
 - c. Click **Update**.

Changing a File Server VM Password

Update the password of a file server VM (FSVM).

About this task

To change an FSVM password, follow these steps:

Procedure

1. Log on to a FSVM with SSH, see [Logging on to a File Server VM](#) on page 34
2. Change the Nutanix password.

```
nutanix@fsvm$ sudo passwd nutanix
```

3. Respond to the prompts, providing the current and new nutanix user password.

```
Changing password for user nutanix.  
Old Password:  
New password:  
Retype new password:  
passwd: all authentication tokens updated successfully.
```

The password must meet the following complexity requirements:

- At least 8 characters long
- At least 1 lowercase letter
- At least 1 uppercase letter
- At least 1 number
- At least 1 special character
- At least 4 characters difference from the old password
- Should not be among the last 10 passwords

The operation updates the password of every FSVM on the file server.

Logging on to a File Server VM

Log on to a file server VM (FSVM) using SSH.

About this task

To SSH on to a FSVM, follow these steps:

Procedure

1. Retrieve the IP address for the FSVM.

- a. SSH onto the CVM.

```
$ ssh nutanix@CVM_IP_address
```

- b. When prompted, enter the CVM password.

- c. List the FSVM IP addresses.

```
nutanix@cvm$ ncli fs ls
```

FSVM IP addresses appear in the following format **Nvm IP Addresses: *external IP, internal IP.***

2. SSH onto the FSVM using the internal IP address.

```
nutanix@cvm$ ssh internal_ip_address
```

Setting Timezones

Set the timezone for all FSVMs on the file server.

About this task

The time zone persists across all FSVMs when Files upgrades.

To set a timezone, follow these steps:

Procedure

1. SSH onto any FSVM in your cluster, see [Logging on to a File Server VM](#) on page 34.
2. Get the timezone for your region.

```
nutanix@fsvm$ afs fs.get_timezone
```

Note: To list all available timezones, see the contents of the `zoneinfo` file for the desired region using the `ls /usr/share/zoneinfo/region` command.

The command returns your timezone into a region and city format.

3. Set the timezone according to the output.

```
nutanix@fsvm$ afs fs.set_timezone "region_name/city_name"
```

Replace *region_name* and *city_name* with the region and city name specified in the timezone output. For example, `afs fs.set_timezone Asia/Kolkata`.

Starting a Nutanix Files Cluster

Start a Files cluster on a CVM using the CLI.

About this task

A deployed filer server, also known as a Files cluster, consists of multiple file server VMs (FSVMs).

To start a file server, follow these steps:

Procedure

1. Using SSH, log on to a Controller VM of the Nutanix cluster that deployed Files.
2. Get a list of file servers.

```
nutanix@cvm$ afs info.fileservers
```

3. Choose to start a single file server or all file servers.

» Start a single file server.

```
nutanix@cvm$ afs infra.start fs_names
```

Replace *fs_names* with the name of a single file server or with a comma-separated list of multiple file servers.

» Start all file servers.

```
nutanix@cvm$ afs infra.start *
```

Note: For steps on stopping and starting a Nutanix cluster, see "Node Management" in the administrative guide for the target hypervisor.

Stopping a Files Cluster

Stop a Files cluster using the provided commands.

About this task

Before you stop a Nutanix cluster that is running Files, first stop the Files cluster (set of file server VMs running on the nodes). To stop a Files cluster, do the following:

Procedure

1. Using SSH, log on to a Controller VM of the Nutanix cluster that is running Files.
2. List file server names.

```
nutanix@cvm$ ncli file-server list
```

3. Stop a single file server.

```
nutanix@cvm$ afs infra.stop fs_name
```

Replace *fs_name* with the name of the file server.

Stopping a File Server VM

Use the nCLI or Nutanix Files CLI to gracefully turn off or restart a single FSVM.

About this task

Check the high availability state before turning off or restarting an FSVM because these operations can impact user connections and cause a temporary downtime.

To gracefully turn off or restart an FSVM, follow these steps:

Procedure

- Gracefully turn off an FSVM.

```
nutanix@fsvm$ fsm_shutdown --poweroff now
```

- Restart an FSVM.

```
nutanix@fsvm$ fsm_shutdown -r now
```

FILE SERVER NETWORK MANAGEMENT

A file server uses internal and external networks.

During file server creation, Files requires configuring an external network for client communication and an internal (backend) network for controller VM (CVM) communication. The external client network configured during file server creation becomes the primary network, which Files uses for Active Directory (AD) and domain name system (DNS) mapping. You cannot delete the primary network, but you can modify it after file server creation by adding additional client networks, see [Multi-VLAN Network Management](#) on page 39. The Files Console only displays the primary network.

When a file server has multiple networks, the first network configured remains the primary network, unless you update the network configuration. After you add an additional network using the CLI, you can update the primary network through the user interface. Changing the primary network requires leaving the domain, which causes service disruptions.

Network management includes the following operations:

- [Adding a Client Network](#) on page 40
- [Listing Client Networks](#) on page 41
- [Deleting a Client Network](#) on page 41
- [Updating the Primary Network of a File Server](#) on page 45
- [Scaling Out A Multi-VLAN File Server](#) on page 42

Managing the DNS Server

Files manages domain name system (DNS) mappings for the primary network.

Before you begin

For multi-VLAN file servers, ensure that the netmask ordering policy is enabled on the DNS server.

About this task

Files only handles DNS auto-registration and un-registration for the primary network. For a non-primary network, configure the DNS mapping manually on the DNS server. Files requires DNS registration after you perform scale out operations. Files requires you to un-register from a DNS server after you perform network delete or scale in operations.

Follow the steps as indicated to get the required DNS mapping for manual configuration:

Procedure

1. Log on to a file server VM (FSVM).
2. Get the JSON file with network information.

```
nutanix@fsvm$ afs dns.get_register_info
```

Note: This command creates the JSON file `/tmp/tmpfile_name.json`. *file_name*, which is a temporary file name. The JSON file contains a mapping of NVM node name to the IP addresses for all of client networks.

3. (Optional) if you need to include or exclude the network IP addresses for registering to the DNS server, edit the JSON file .

4. Manually register all of the networks to the DNS server.

```
nutanix@fsvm$ sudo net ads dns register_nodes --user=username jsonfilename=filename
```

Replace *username* with the file server administrator. Replace *filename* with the JSON file containing the network information.

5. (Optional) check the configured DNS mappings.

```
nutanix@fsvm$ afs dns.verify_external_networks show_reverse_ptr_records=true  
external_network_uuid=external_network_uuid
```

Replace *external_network_uuid* with the unique user identifier (UUID) of the manually configured external network.

Network Segmentation

Deploying Files in a Volumes network segmented environment isolates the file server communication between FSVMs and CVMs onto a dedicated subnet, a Volumes segment.

Important: Before using network segmentation with Files, meet the following conditions.

- Enable network segmentation for Nutanix Volumes, see "Service-Specific Traffic Isolation" and "Service-Specific Settings and Configurations" section in the [AOS Security Guide](#).

Consider using network segmentation with Files for the following use cases:

- You do not have enough IP addresses in the management network of CVM to allocate them for Files internal network. Also, you do not want to use the IP addresses from another routable subnet for Files internal IP address as you have set up firewall rules and do not want data traffic (iSCSI) between Files and CVM to go through firewalls.
- You have networks with varying capacity, and you want to use the one with maximum bandwidth for storage and the one with lesser bandwidth for management. For example, 10G network for storage and 1G network for management.

Network segmentation enhances security, resilience, and cluster performance by isolating a subset of traffic to its own network. By configuring network segmentation on a cluster, Files does the following:

- Files storage data uses its own network, segmented from other services as well as from the cluster backplane.
- Files management uses its own network, segmented from default management network on the cluster.

Note: Files only supports using network segmentation with Nutanix Volumes.

To deploy Files in a Volumes network segmented environment, see [Configuring Network Segmentation](#) on page 38.

Configuring Network Segmentation

Determine the method for configuring a Volumes network-segmented environment for your file server.

Before you begin

Make sure you have met the requirements for network segmentation, see [Network Segmentation](#) on page 38.

About this task

To configure network segmentation, follow these steps:

Procedure

1. Enable Volumes network segmentation on the AOS cluster.

For more information, see the [Service-Specific Traffic Isolation](#) and [Service-Specific Settings and Configurations](#) sections in the *Security Guide*.

2. Assess your setup to determine the method for configuring network segmentation.
 - » If you enabled network segmentation for Volumes in Prism Element and have not deployed a file server, deploy the file server using the network segmented subnet as the Files storage network. For more information, see [Creating a File Server](#) on page 21.
 - » If you enabled network segmentation for Volumes in Prism Element and have existing file servers, re-IP the storage network to a Volumes network segmented subnet. For more information, see [Updating the File Server Network Configuration in Prism Element](#) on page 44.
 - » If you deployed a file server on the same subnet where you plan to configure network segmentation, perform manual steps to allow communication between the file server and the CVMs after configuring networking segmentation. For more information, see [KB-17093](#).

Multi-VLAN Network Management

Management of multiple client networks after file server deployment.

You can configure multiple non-routable client networks on a single file server. You can also configure share access over a subset of configured client networks. Only file servers configured in a dual-VLAN network state support multi-VLAN networks.

The following non-exhaustive list describes some instances when a multi-VLAN network is an appropriate solution:

- A single client network runs out of IP addresses.
- Logical isolation on a single client network becomes complex.
- Routing traffic from multiple client networks to a single client network leads to complex routing and firewall management.
- Multi-tenancy of client networks requires a multi-VLAN configuration.

Important: Ensure that you create an unmanaged network with VLAN ID = 0 for AHV, or, for the ESXi hypervisor, create an unmanaged network with VLAN ID = 4095. For information on creating a network on the AHV hypervisor, see [Network Management](#). For information on creating VLAN 4095 on ESXi, see the [vSphere Administration Guide for Acropolis](#).

Multi-vSwitch Support (AHV Only)

For AHV, you can configure multi-VLAN networks of multiple vSwitches on the same file server. The multi-vSwitch support requires at least two virtual switches (vSwitch), and each vSwitch must have at least one unmanaged network with `VLAN ID = 0`.

For information on steps to create a vSwitch, see [Creating or Updating a Virtual Switch](#) in the *Prism Element Web Console Guide*.

Notes and Limitations

Consider the following limitations before deploying files on a multi-VLAN network:

- You can have a maximum of 10 client networks.
- File servers configured in a single-VLAN network state do not support multi-VLAN.

- Only IPv4 networks support multi-VLAN (IPv6 networks do not). If you have already configured an IPv6 network, remove the IPv6 network before using multi-VLAN.
- You cannot add storage or Controller VM (CVM) networks as client networks.
- Nutanix on AWS does not support multi-VLAN.
- An active multi-VLAN configuration has the sysctl parameter enabled (arp_ignore=2). An in-active multi-VLAN configuration has the arp_ignore parameter disabled (arp_ignore=0).
- You can have a maximum of 4 vSwitches for adding VLANs to a file server.
- Using multiple vSwitches is only supported for AHV.

Adding a Client Network

Add additional client networks to a file server.

Before you begin

Before you can add a network to a file server, you must configure the network on the CVM. For more information, see the [AHV Administration Guide](#) and the [vSphere Administration Guide for Acropolis](#).

About this task

To add client networks, do the following:

Procedure

1. Log on to the CVM with SSH:

```
$ ssh nutanix@cvm_ip_address
```

2. List existing client networks to discover existing networks on the CVM:

```
nutanix@cvm$ acli net.list
```

3. Log on to a file server VM (FSVM) with SSH.

For more information, see [Logging on to a File Server VM](#) on page 34.

4. Because the required JSON parameters vary depending on your environment, use one of the following three JSON templates:

For managed AHV networks use the following JSON template:

```
[
  {
    "virtual_network_uuid": "uuid"
  }
]
```

For unmanaged AHV networks, use the following JSON template:

```
{
  "virtual_network_uuid": "uuid",
  "gateway_ip": "gateway-ip-address",
  "netmask_ip": "netmask-ip-address",
  "ip_list": [
    "fsvm-ip-address",
    "fsvm-ip-address",
    ...
  ]
}
```

```
}
```

The following use cases require using the `trunk_virtual_network_uuid` parameter:

- AHV environments where multiple networks with VLAN ID 0 are present for the same vSwitch (in both single and multiple switch environments)
- ESXi standard vSwitch environments where multiple port-groups with VLAN ID 4095 on the same vSwitch
- ESXi distributed vSwitch (DVS) environments, which also requires a `vlan_id` parameter for each port-group or network you add to the file server

For the above use cases, use the following JSON template:

```
[
  {
    "virtual_network_uuid": "uuid",
    "vlan_id": "vlan-id",
    "trunk_virtual_network_uuid": "uuid"
  }
]
```

Replace `virtual_network_uuid` with the universal unique identifier for the CVM network you want to add from the list in Step 2.

Note: For ESXi, use `nutanix@CVM$ acli uhura.network.list` to list UUIDs.

Replace `gateway_ip` with the gateway IP address for the CVM network to add from the list in Step 2.

Replace `netmask_ip` with the netmask IP address for the CVM network to add from the list in Step 2.

Replace `fsvm-ip-address` with the IP address for the new FSVM. The number of `ip_address` entries should match the number of FSVMs on the file server. The IP addresses must be sequential.

Replace `trunk_virtual_network_uuid` with the trunk VLAN UUID.

Note: Use the `afs help net.add_external` command to view the JSON file format.

5. Add a client network using a JSON file:

```
nutanix@fsvm$ afs net.add_external input-json-file-path
```

Replace `input-json-file-path` with the file path to a local JSON that contains the input parameters file.

Listing Client Networks

List the configured client networks on the file server.

About this task

To list the client networks, do the following:

Procedure

1. Log on to a file server VM (FSVM) with SSH. See [Logging on to a File Server VM](#) on page 34.
2. List the configured external client networks.

```
nutanix@fsvm$ afs net.list_external
```

A list of configured networks appears.

Deleting a Client Network

Remove configured client networks from a file server.

About this task

To remove a client network, do the following:

Procedure

1. List existing client networks. See [Listing Client Networks](#) on page 41.
2. Log on to a file server VM (FSVM) with SSH. See [Logging on to a File Server VM](#) on page 34.
 - a. To delete networks, list the universally unique identifiers (UUIDs).

```
nutanix@fsvm$ afs net.delete_external network_uuid_list [key=val ...]
```

Replace `network_uuid_list` with a comma-separated list of network UUIDs for the networks you want to delete.

Note: To view a list of network UUIDs, use the `afs net.list_external` command.

Scaling Out A Multi-VLAN File Server

Add additional networks to scaled out FSVMs.

About this task

Scaling out FSVMs configures the new nodes with the primary network only, which creates a partial network state on the file server. In a partial network state, you cannot create shares or rebalance volume groups (VGs), but you can perform upgrade and scale-in tasks.

To add the remaining non-primary networks to the scaled-out nodes, use the network patch operation as follows:

Procedure

1. List existing client networks. See [Listing Client Networks](#) on page 41.
2. Add FSVMs to the cluster. See [Scaling File Server VMs](#) on page 32.

After adding FSVMs, the **Alerts** tab displays the following: The file server is in a partial network state.
3. Log on to a file server VM (FSVM) with SSH. See [Logging on to a File Server VM](#) on page 34.
4. Resolve the partial network state using a JSON file that includes details of the remaining networks.

```
nutanix@fsvm$ afs net.patch_external input-json-file
```

Replace `input-json-file` with the path to a local JSON file formatted as indicated in the following template:

```
[
  {
    "uuid": "uuid"
  },
  {
    "uuid": "uuid",
    "ip_list": [
      "ip_address",
      "ip_address",
      ...
    ]
  },
  {
    "uuid": "uuid",
    "trunk_virtual_network_uuid": "uuid"
  },
  ...
]
```

]

Note: You can also use the `afs help net.patch_external` command to view the JSON file format.

For managed networks, replace the `uuid` with universally unique identifier (UUID) of the network.

For unmanaged networks, replace the `uuid` with the UUID of the network and replace `ip_address` with IP addresses for the network, adding as many comma separated values as newly added FSVMs.

The `trunk_virtual_network_uuid` parameter is required for the following use cases:

- For AHV, where multiple networks with VLAN ID 0 are present for the same vSwitch, in both single and multiple switch environments
- For an ESXi standard vSwitch, where multiple port-groups with VLAN ID 4095 are present on the same vSwitch
- For an ESXi distributed vSwitch (DVS), which also requires a `vlan_id` parameter for each port-group or network you add to the file-server

Performing the procedure resolves The file server is in a partial network state alert.

Configuring Client Exceptions

Configure the client networks exception list for exports and shares.

About this task

Configuring client exceptions defines the share availability on a subset of client networks. By default, all shares are available on all client networks. After you have added multiple client networks, the shares can be made available only on a subset of configured client networks.

Note: Client exception values created using a command-line interface (CLI) do not appear in the UI.

To configure client exceptions for SMB shares, do the following:

Procedure

Log on to an FSVM with SSH.

```
nutanix@fsvm$ afs smb.client_exceptions smb-share-name whitelist  
ip_net_list=subnet1,subnet2
```

Replace `smb-share-name` with the name of the SMB share.

Replace `subnet1` and `subnet2` with the subnet IP addresses in classless inter-domain routing (CIDR) notation.

Configuring NFS Client Exceptions

About this task

To configure client network exceptions for NFS exports, follow these steps:

Procedure

1. From the **File Server** view in Prism Element (PE), click **Launch Files Console**.
2. From the File Console, select an NFS share and click **Update**.
3. In the **Basics** tab, click **Next**.
4. In the **General Settings** tab, click **Next**.

5. In the **Protocol Settings and Permissions** tab, follow these steps:
 - a. Change **Default Access** to **No Access** in the **General Settings** tab.
 - b. Click **Add Exceptions**, and then add the subnets in CIDR format in the **clients with read-only access** and **clients with read-write access** fields.
6. Click **Next**.
7. Verify the details on the **Summary** tab and click **Update**.
For more information, see [Creating an NFS Export](#) on page 50.

Updating the File Server Network Configuration in Prism Element

Change or re-IP the file server client or storage network configuration in Prism Element.

Before you begin

Stop a file server before updating the network to reduce traffic during the update process. For more information, see [Stopping a Files Cluster](#).

About this task

Changing the client network and storage network for a file server VM lets you move a file server from one datacenter to another. If you intend to change the domain, you must leave the domain first.

Updating a network configuration on ESXi clusters requires vCenter connectivity.

For information on configuring a virtual network for guest VMs, see [Configuring a Virtual Network for Guest VM Interfaces](#) in the *Prism Web Console Guide*. For information on modifying network connections, see [Modifying Network Connections](#) in the *Prism Web Console Guide*.

Important: Updating the network configuration requires a temporary downtime. Nutanix recommends performing network updates during off-peak hours.

To update the file server network configuration, follow these steps:

Procedure

1. In Prism Element, go to the **File Server** view.
For more information, see [File Server View in Prism Element](#) on page 16.
2. From the **File Servers** table, select a file server.
3. From the options under **File Servers** table, go to **Update > Network Configuration**.
4. In the **Update File Server Network** window, update the file server configuration:
You can change the client network, storage network, domain name system (DNS), and network time protocol (NTP) IP addresses for a file server VM.
 - a. In the **Client Network** tab, verify or update the client network details, including the VLAN, IPv6, DNS, or NTP entries.
Files display the following message after configuring client networks on the file server: *Number* Clients Networks are configured for the file server. Only the primary network can be updated through the UI below.

Once the primary network is updated the additional clients networks already configured will be cleared. Use the CLI to re-configure additional external networks.

For more information on the client network configuration fields, see [Configuring the Client Network](#) on page 23.

- b. (Optional) To update the DNS entries, click **file server DNS entries** and delete the existing DNS entries. Modifying the client network configuration requires modifying the DNS entries. You can modify the DNS entries after completing the network updates. For more information, see [Updating Domain Name System \(DNS\) Entries](#) on page 76.

- c. Click **Next**

- d. In the **Storage Network** tab, verify or update the storage network details.

For more information on **Storage Network** fields, see [Configuring the Storage Network](#) on page 24.

- 5. Click **Save**.

If the network update operation is unsuccessful, use the original IP address details and try the update operation again.

Updating the Primary Network of a File Server

After adding additional networks, you can update the primary client network.

Before you begin

Before updating the primary network, you need to have other networks configured and you need to leave the domain. For more information, see [Multi-VLAN Network Management](#) on page 39 and [Leaving a Domain](#) on page 75.

Procedure

1. Log on to a file server VM (FSVM) with SSH, see [Logging on to a File Server VM](#) on page 34.
2. Update the primary network.

```
nutanix@fsvm$ afs net.update_external_primary_network network_uuid
```

Replace *network_uuid* with the network unique universal identifier (UUID) for the primary network.

What to do next

After updating the primary network, join the domain, see [Joining a Domain](#) on page 74.

SHARE AND EXPORT MANAGEMENT

Before creating a share or export, review the available features and configurations.

You can create file shares (SMB), exports (NFS), or multi-protocol shares for a file server.

You have the option of creating distributed or standard shares and exports. A distributed share or export is the repository for the personal files of a user, and a standard share is the repository shared by a group. A distributed share or export distributes data across all file server VM (FSVM) in a cluster, while standard shares store all data on a single FSVM. You can store files in the root directory of a distributed NFS share but not in the root directory of a distributed SMB share.

For more information on share and export permissions, see [Share and Export Permissions](#) on page 72.

To create a share or export, see one of the following:

- [Creating an SMB Share](#) on page 46
- [Creating an NFS Export](#) on page 50
- [Creating a Multi-Protocol Share or Export](#) on page 55

Limitations

The following limitations apply to shares and exports:

- Distributed shares are only available on deployments of three or more FSVMs.
- Do not use Windows Explorer to create new top-level directories, as you then cannot rename any folders created with the default New Folder name. For more information, see [Troubleshooting](#) on page 143.
- Windows clients do not support using UTF8 encoding for naming files and directories on NFSv3 exports.
- You cannot access top-level directories (TLDs) or shares through the FSVM short name.
- You cannot modify TLDs from Mac clients.
- Files does not support custom extended attributes (EA) on SMB shares.

For a list of additional configuration limitations, see [Configuration Maximums](#).

Creating an SMB Share

About this task

Create an SMB share on a file server.

Note: Use multi-protocol shares for mounting SMB shares to Linux clients. For more information, see [Creating a Multi-Protocol Share or Export](#) on page 55.

To create an SMB share, follow these steps:

Procedure

1. Go to the **Shares & Exports** view in the Files Console.
For more information, see [Nutanix Files Console](#) on page 18.

2. Click **+New Share or Export**.




Figure 10: + New Share or Export Button

The **Create Share or Export** window appears and displays the **Basics** tab.

3. In the **Basics** tab, complete the following fields:

- **Name:** Enter a name for the share that adheres to the following naming conventions:
 - Files does not allow the following names:
 - `Global`
 - `Printers`
 - `admin$` (reserved)
 - `ipc$` (reserved)
 - `homes` (reserved)
 - Share names are not case sensitive.
 - Each share name must be unique.
 - A share name cannot start or end with a blank space.
 - A share name cannot start with a dot (.).
 - A share name cannot include the following characters: `!@#%~'\\"^&$()[]|+=*{}\\/\<>:; , ' .`
 - Names that end with the dollar symbol (\$) are for hidden shares.
 - Comply with the name length limits documented in [Configuration Maximums](#).
- (Optional) **Description:** Enter a description for the share.
- (Optional) **Share path:** To create a nested share, specify the path of the parent share. Nested shares inherit some properties from the parent directory. For more information, see [Nested Shares and Exports](#) on page 70 for details.
- (Optional) **Create this share at the specified path:** To create a connected share, select the checkbox. For more information, see [Connected Shares](#).
- (Optional): **Max Size:** To set an upper size limit, enter the maximum share size in GiB.

Leaving the field empty means no upper limit exists for the size of the share. Setting a value changes the capacity in the client view.

For information on sizing limits, see [Configuration Maximums](#).

- **Select Share or Export type :** Select **Standard** or **Distributed**.

Note: By default, shares on single-FSVM deployments are standard.

A *Distributed share*, also known as a home share, load-balances user data across multiple FSVMs by distributing top-level directories.

A *Standard share*, also known as a general-purpose share, serves data and connections from a single FSVM.

A distributed share is frequently used as the repository for the personal files of a user, while a standard share is frequently the repository shared by a group.

Tip: When using a distributed share for user profiles, facilitate better load distribution by creating a top-level directory for each user.

- **Primary Protocol Access:** Select the **SMB (Ideal for Windows Clients)** option.

- (Optional) **Enable multi-protocol access for NFS clients:** To enable the multi-protocol feature, select the checkbox.

To create and configure a multi-protocol share, follow the steps in [Creating a Multi-Protocol Share or Export](#) on page 55. Otherwise, continue to the next step.

For information on the multi-protocol feature, see [Multi-Protocol Support for Nutanix Files](#) on page 61.

4. Click **Next**

5. In the **General Settings** tab, complete the following fields:

- a. (Optional) To enable snapshots of the share contents, select the **Enable Self Service Restore** checkbox, and select a snapshot schedule policy from the **Snapshot Policy** dropdown menu.

For more information, see [Self-Service Restore](#) on page 124.

- b. (Optional) To save space and reduce data on the share through in-line compression of written data, select the **Enable Compression** checkbox.

For more information, see [File System Compression](#) on page 104.

Note: Enabling compression at the container level makes the **Enable Compression** option unavailable. The **Enable Compression** option does not appear on clusters upgraded from earlier versions that do not support file system compression.

- c. (Optional) To block files with specific character patterns in their names, select the **Blocked File Types** checkbox.

For more information, see [File Blocking](#).

A field appears for blocked file types.

Enter a comma-separated list of character patterns of directories and file names blocked from the share.

- d. (Optional) To enable write once, read many (WORM) support on the share, select the **Enable WORM** checkbox.

Important: You cannot disable WORM after creating the share.

1. Under **Cool Off Interval**, enter an integer to specify the duration that the file must be unmodified to transition to the WORM state, and use the dropdown menu to specify the interval of **minutes** or **seconds**
2. Under **Retention Period**, enter an integer to specify the duration that Files must retain the file before the file transitions into the expired state, and use the dropdown menu to specify the interval of **days** or **weeks**.

For more information on WORM, see [WORM Shares](#) on page 67.

6. Click **Next**.

7. In the **Protocol Settings and Permissions** tab, complete the following fields:

- (Optional) **Enable Access Based Enumeration (ABE)**: To restrict user TLD browsing access to only the files and folders that they have access permissions for, select the checkbox.
For more information, see [Managing Access-Based Enumeration for SMB Shares](#) on page 96.
- (Optional) **Allow Long Names for Files and Folders**: To let users create file names up to 1023 bytes (255 characters) long, select the checkbox.
- (Optional) **Encrypt SMB3 Messages**: To enable message encryption between the file server and clients, select the checkbox.

For more information, see [Encryption](#) on page 135.

- (Optional) **Enable Continuous Availability (CA)**: To enable CA on the share, select the checkbox.
For more information, see [Continuously Available SMB Shares](#) on page 65.
- (Optional) **Share Permissions**: To customize the default share permissions for all users and groups on the share, click **Add User or Group** and follow these steps:

1. In the **User or Group** column, enter the name of the user or group to apply the permissions to.

Specify a user or a group in the sAMAccountName or NETBIOS\sAMAccountName format where sAMAccountName shares the domain with the file server.

2. In the **Type** column, from the dropdown menu, choose between the **Deny** or **Allow** permission option.
3. In the **Read**, **Change**, and **Full Control** columns, select the **denied** or **allowed** file modification actions for the target user or group.
4. Click the [check mark](#) icon to add the permissions.

For more information about permission assignment, see [Share and Export Permissions](#) on page 72.

8. Click **Next**.

9. In the **Summary** tab, review the configuration information.

10. (Optional) To edit the configuration, click **Back**.

11. When the specified configuration is correct, click **Create**.

What to do next

Map the newly created share in your namespace.

Creating an NFS Export

Create an NFS export on a file server.

About this task

Note: Nutanix Files does not support mounting NFSv4 exports on Windows clients. Use NFSv3 exports or multi-protocol shares. For more information, see [Creating a Multi-Protocol Share or Export](#) on page 55.

To create an NFS export, follow these steps:

1. Go to the **Shares & Exports** view in the Files Console and click **+New Share or Export**.



NUTANIX

2. In the **Basics** tab, complete the following fields:

- **Name:** Enter a name for the share that adheres to the following naming conventions:

- Files does not allow the following names:

- `Global`
- `Printers`
- `admin$` (reserved)
- `ipc$` (reserved)
- `homes` (reserved)

- Share names are not case sensitive.

- Each share name must be unique.

- A share name cannot start or end with a blank space.

- A share name cannot start with a dot (.).

- A share name cannot include the following characters: `!@#%~'\\"^&$()[]|+=*{}\\/\\"<>:; , ' "`

- Names that end with the dollar symbol (\$) are for hidden shares.

- Comply with the name length limits documented in [Configuration Maximums](#).

- (Optional) **Description:** Enter a description for the share.

- (Optional) **Share path:** To create a nested share, specify the path of the parent share.

Nested shares inherit some properties from the parent directory. For more information, see [Nested Shares and Exports](#) on page 70 for details.

- (Optional) **Create this share at the specified path:** To create a connected share, select the checkbox.

For more information, see [Connected Shares](#).

- (Optional): **Max Size:** To set an upper size limit, enter the maximum share size in GiB.

Leaving the field empty means that no upper limit exists for the size of the share. Setting a value changes the capacity in the client view.

For information on sizing limits, see [Configuration Maximums](#).

- **Select Share or Export type :** Select **Standard** or **Distributed**.

Note: By default, shares on single-FSVM deployments are standard.

A *Distributed share*, also known as a home share, load-balances user data across multiple FSVMs by distributing top-level directories.

A *Standard share*, also known as a general-purpose share, serves data and connections from a single FSVM.

A distributed share is frequently used as the repository for the personal files of a user, while a standard share is frequently the repository shared by a group.

Tip: When using a distributed share for user profiles, facilitate better load distribution by creating a top-level directory for each user.

- **Primary Protocol Access:** Select the **NFS** option.

- (Optional) **Enable multi-protocol access for SMB clients:** To enable the multi-protocol feature, select the checkbox.

To create and configure a multi-protocol share, follow the steps in [Creating a Multi-Protocol Share or Export](#) on page 55. Otherwise, continue to the next step.

For information on the multi-protocol feature, see [Multi-Protocol Support for Nutanix Files](#) on page 61.

3. Click **Next**

4. In the **General Settings** tab, complete the following fields:

- a. (Optional) To enable snapshots of the share contents, select the **Enable Self Service Restore** checkbox, and select a snapshot schedule policy from the **Snapshot Policy** dropdown menu.

For more information, see [Self-Service Restore](#) on page 124.

- b. (Optional) To save space and reduce data on the share through in-line compression of written data, select the **Enable Compression** checkbox.

For more information, see [File System Compression](#) on page 104.

Note: Enabling compression at the container level makes the **Enable Compression** option unavailable. The **Enable Compression** option does not appear on clusters upgraded from earlier versions that do not support file system compression.

- c. (Optional) To block files with specific character patterns in their names, select the **Blocked File Types** checkbox.

For more information, see [File Blocking](#).

A field appears for blocked file types.

Enter a comma-separated list of character patterns of directories and file names blocked from the share.

- d. (Optional) To enable write once, read many (WORM) support on the share, select the **Enable WORM** checkbox.

Important: You cannot disable WORM after creating the share.

1. Under **Cool Off Interval**, enter an integer to specify the duration that the file must be unmodified to transition to the WORM state, and use the dropdown menu to specify the interval of **minutes** or **seconds**
2. Under **Retention Period**, enter an integer to specify the duration that Files must retain the file before the file transitions into the expired state, and use the dropdown menu to specify the interval of **days** or **weeks**.

For more information on WORM, see [WORM Shares](#) on page 67.

5. Click **Next**.

6. In the **Protocol Settings and Permissions** tab, complete the following fields:

- **Authentication:** Select the authentication method from the dropdown menu.

The authentication options are **None**, **System**, **Kerberos 5**, **Kerberos 5i**, and **Kerberos 5p**.

Files does not support Kerberos 5, Kerberos 5i, and Kerberos 5p with the NFSv3 protocol.

For LDAP or an unmanaged directory service, use **None** or **System** authentication; both options are client-based but have different authorization rules. For AD, use **None**, **System**, or user-based authentication with Kerberos. For more information, see [Directory Services and Authentication](#) on page 73.

You can map regular users to anonymous users to restrict access to the NFS export. Setting the NFS authentication type to **None** maps all client users to anonymous users.

Note: Changes to the authentication method can cause invalid mounts. Resolve invalid mounts by remounting the authentication type as a value for parameter sec in the mount command on the client. For example, use "-o sec=krb5" if you change the export authentication type to Kerberos 5.

- **Default Access (For all clients):** From the dropdown menu, select the default access permissions from one of the following options:

- **Read-Write:** Clients can mount the NFS export and write to it.
- **Read-Only:** Clients can mount the NFS export and read data but not write to it.
- **No Access:** Clients cannot mount the NFS export nor read or write from it.

- **Squash:** Select the **None**, **Root Squash**, or **All Squash** value from the dropdown menu.

The squash option controls the access privileges of root client users (users with UID 0). The **None** value gives root users super-user access privileges to the export, letting them create, edit, and delete files from any

user. The **Root Squash** option disables super user access and maps all root users to anonymous users. The **All Squash** value maps all users to anonymous users.

- (Optional) **Override anonymous IDs for squashing:** To enter anonymous user identifier (UID) and group identifier (GID) values other than the default, select the checkbox.

Both the anonymous UID and GID have a default value of -2.

- **Anonymous UID:** Enter the anonymous UID value for the export.
- **Anonymous GID:** Enter the anonymous GID value for the export.
- (Optional) **Add Exceptions:** To refine the read, write, and no-access permissions, click **Add Exceptions**.

The **Add Exceptions** dialogue box displays **Clients**, **Access**, and **Squash** fields.

- In the **Clients** field, enter a comma separate list of clients for the target **Access** permission level.

Exceptions take precedence over the default, so any client listed in one of the exception fields gets that level of permission instead of the default permission. For IPv4 clients, exceptions can be complete IP addresses, wildcards, or subnets. For IPv6 clients, Files only accepts complete IP addresses as exceptions.

Caution: If a client IP address matches more than one exception rule, the client may experience access issues.

Use one of the following formats:

- Absolute IP address (for example, 129.144.0.0)
- Hostname of client (for example, client.domain.com/)
- CIDR format for specifying subnets (for example, 129.144.0.0/24)
- Netgroups (for example @it_admin)

Netgroups can limit access to hosts, but Files does not support Netgroups that limit access to users.

Note: If you add or remove a client from a Netgroup, the change reflects on the client after 30 minutes.

- Wildcards (for example clients*.domain.com).

To use wildcards, configure reverse domain name system (DNS) records for clients.

After you add an exception, it appears on the **Protocol Settings and Permissions** tab under **Exceptions for Client Access**. For more information about permission assignment, see [Share and Export Permissions](#) on page 72.

7. Click **Next**.
8. In the **Summary** tab, review the configuration information.
9. (Optional) To edit the configuration, click **Back**.
10. When the specified configuration is correct, click **Create**.

Creating a Multi-Protocol Share or Export

Create a multi-protocol share or export.

About this task

Nutanix Files does not support the multi-protocol feature for shares with long name support already turned on.

For information on the multi-protocol feature, see [Multi-Protocol Support for Nutanix Files](#) on page 61.

To create a multi-protocol share or export, follow these steps:

Procedure

1. Go to the **Shares & Exports** view in the Files Console and click **+New Share or Export**.




Figure 12: Create a New Share Button

2. In the **Basics** tab, complete the following fields:

- **Name:** Enter a name for the share that adheres to the following naming conventions:

- Files does not allow the following names:

- `Global`
- `Printers`
- `admin$` (reserved)
- `ipc$` (reserved)
- `homes` (reserved)

- Share names are not case sensitive.

- Each share name must be unique.

- A share name cannot start or end with a blank space.

- A share name cannot start with a dot (.).

- A share name cannot include the following characters: `!@#%~'\\"^&$()[]|+=*{}\\/\\"<>:; , ' "`

- Names that end with the dollar symbol (\$) are for hidden shares.

- Comply with the name length limits documented in [Configuration Maximums](#).

- (Optional) **Description:** Enter a description for the share.

- (Optional) **Share path:** To create a nested share, specify the path of the parent share.

Nested shares inherit some properties from the parent directory. For more information, see [Nested Shares and Exports](#) on page 70 for details.

- (Optional) **Create this share at the specified path:** To create a connected share, select the checkbox.

For more information, see [Connected Shares](#).

- (Optional): **Max Size:** To set an upper size limit, enter the maximum share size in GiB.

Leaving the field empty means that no upper limit exists for the size of the share. Setting a value changes the capacity in the client view.

For information on sizing limits, see [Configuration Maximums](#).

- **Select Share or Export type :** Select **Standard** or **Distributed**.

Note: By default, shares on single-FSVM deployments are standard.

A *Distributed share*, also known as a home share, load-balances user data across multiple FSVMs by distributing top-level directories.

A *Standard share*, also known as a general-purpose share, serves data and connections from a single FSVM.

A distributed share is frequently used as the repository for the personal files of a user, while a standard share is frequently the repository shared by a group.

Tip: When using a distributed share for user profiles, facilitate better load distribution by creating a top-level directory for each user.

- **Primary Protocol Access:** Select the **NFS** or **SMB** option.

- **Enable multi-protocol access for SMB clients** or **Enable multi-protocol access for NFS clients**: Select the checkbox to enable the multi-protocol feature.

3. Click **Next**

4. In the **General Settings** tab, complete the following fields:

- (Optional) To enable snapshots of the share contents, select the **Enable Self Service Restore** checkbox, and select a snapshot schedule policy from the **Snapshot Policy** dropdown menu.
For more information, see [Self-Service Restore](#) on page 124.
- (Optional) To save space and reduce data on the share through in-line compression of written data, select the **Enable Compression** checkbox.
For more information, see [File System Compression](#) on page 104.

Note: Enabling compression at the container level makes the **Enable Compression** option unavailable. The **Enable Compression** option does not appear on clusters upgraded from earlier versions that do not support file system compression.

- (Optional) To block files with specific character patterns in their names, select the **Blocked File Types** checkbox.
For more information, see [File Blocking](#).
A field appears for blocked file types.
Enter a comma-separated list of character patterns of directories and file names blocked from the share.
- (Optional) To enable write once, read many (WORM) support on the share, select the **Enable WORM** checkbox.

Important: You cannot disable WORM after creating the share.

- Under **Cool Off Interval**, enter an integer to specify the duration that the file must be unmodified to transition to the WORM state, and use the dropdown menu to specify the interval of **minutes** or **seconds**
- Under **Retention Period**, enter an integer to specify the duration that Files must retain the file before the file transitions into the expired state, and use the dropdown menu to specify the interval of **days** or **weeks**.

For more information on WORM, see [WORM Shares](#) on page 67.

5. In the **Protocol Settings and Permissions** tab, complete the SMB and NFS configuration sections.

Tip: The primary protocol determines the protocol settings that appear first.

- a. Complete the following fields to configure SMB protocol settings:

- (Optional) **Enable Access Based Enumeration (ABE):** To restrict user TLD browsing access to only the files and folders that they have access permissions for, select the checkbox.

For more information, see [Managing Access-Based Enumeration for SMB Shares](#) on page 96.

- (Optional) **Allow Long Names for Files and Folders:** To let users create file names up to 1023 bytes (255 characters) long, select the checkbox.

- (Optional) **Encrypt SMB3 Messages:** To enable message encryption between the file server and clients, select the checkbox.

For more information, see [Encryption](#) on page 135.

- (Optional) **Enable Continuous Availability (CA):** To enable CA on the share, select the checkbox.

For more information, see [Continuously Available SMB Shares](#) on page 65.

- (Optional) **Share Permissions:** To customize the default share permissions for all users and groups on the share, click **Add User or Group** and follow these steps:

1. In the **User or Group** column, enter the name of the user or group to apply the permissions to.

Specify a user or a group in the sAMAccountName or NETBIOS\sAMAccountName format where sAMAccountName shares the domain with the file server.

2. In the **Type** column, from the dropdown menu, choose between the **Deny** or **Allow** permission option.

3. In the **Read**, **Change**, and **Full Control** columns, select the **denied** or **allowed** file modification actions for the target user or group.

4. Click the [check mark](#) icon to add the permissions.

For more information about permission assignment, see [Share and Export Permissions](#) on page 72.

- b. Complete the following fields to configure NFS protocol settings:

- **Authentication:** Select the authentication method from the dropdown menu.

The authentication options are **None**, **System**, **Kerberos 5**, **Kerberos 5i**, and **Kerberos 5p**.

Files does not support Kerberos 5, Kerberos 5i, and Kerberos 5p with the NFSv3 protocol.

For LDAP or an unmanaged directory service, use **None** or **System** authentication; both options are client-based but have different authorization rules. For AD, use **None**, **System**, or user-based

authentication with Kerberos. For more information, see [Directory Services and Authentication](#) on page 73.

You can map regular users to anonymous users to restrict access to the NFS export. Setting the NFS authentication type to **None** maps all client users to anonymous users.

Note: Changes to the authentication method can cause invalid mounts. Resolve invalid mounts by remounting the authentication type as a value for parameter sec in the mount command on the client. For example, use "-o sec=krb5" if you change the export authentication type to Kerberos 5.

- **Default Access (For all clients):** From the dropdown menu, select the default access permissions from one of the following options:
 - **Read-Write:** Clients can mount the NFS export and write to it.
 - **Read-Only:** Clients can mount the NFS export and read data but not write to it.
 - **No Access:** Clients cannot mount the NFS export nor read or write from it.
- **Squash:** Select the **None**, **Root Squash**, or **All Squash** value from the dropdown menu.

The squash option controls the access privileges of root client users (users with UID 0). The **None** value gives root users super-user access privileges to the export, letting them create, edit, and delete files from any user. The **Root Squash** option disables super user access and maps all root users to anonymous users. The **All Squash** value maps all users to anonymous users.

- (Optional) **Override anonymous IDs for squashing:** To enter anonymous user identifier (UID) and group identifier (GID) values other than the default, select the checkbox.

Both the anonymous UID and GID have a default value of -2.

- **Anonymous UID:** Enter the anonymous UID value for the export.
- **Anonymous GID:** Enter the anonymous GID value for the export.
- (Optional) **Add Exceptions:** To refine the read, write, and no-access permissions, click **Add Exceptions**.

The **Add Exceptions** dialogue box displays **Clients**, **Access**, and **Squash** fields.

- In the **Clients** field, enter a comma separated list of clients for the target **Access** permission level. Exceptions take precedence over the default, so any client listed in one of the exception fields has that level of permission instead of the default permission. For IPv4 clients, exceptions can be complete

IP addresses, wildcards, or subnets. For IPv6 clients, Files only accepts complete IP addresses as exceptions.

Caution: If a client IP address matches more than one exception rule, the client might experience access issues.

Use one of the following formats:

- Absolute IP address (for example, 129.144.0.0)
- Hostname of client (for example, client.domain.com/)
- CIDR format for specifying subnets (for example, 129.144.0.0/24)
- Netgroups (for example @it_admin)

Netgroups can limit access to hosts, but Files does not support Netgroups that limit access to users.

Note: If you add or remove a client from a Netgroup, the change reflects on the client after 30 minutes.

- Wildcards (for example clients*.domain.com).

To use wildcards, configure reverse domain name system (DNS) records for clients.

After you add an exception, it appears on the **Protocol Settings and Permissions** tab under **Exceptions for Client Access**. For more information about permission assignment, see [Share and Export Permissions](#) on page 72.

6. Complete the following fields in the **Multi Protocol Access** section:

- (Optional) **Allow simultaneous read access to same files:** To enable read access for different clients to the same files at the same time, select the checkbox.
- **Symlink creation from non-native NFS:** To allow symlink creation on NFS shares from SMB clients, select the checkbox.

7. Click **Next**.

8. In the **Summary** tab, review the configuration information.

9. (Optional) To edit the configuration, click **Back**.

10. When all of the information is correct, click **Create**.

What to do next

To configure user mapping between NFS and SMB users, see [User Mapping](#) on page 80.

Multi-Protocol Support for Nutanix Files

Multi-protocol support for Files provides SMB and NFS clients read-and-write access to the same share or export.

When you choose a native protocol, it determines access-control options for that share or export. You can then manage all access-control options using the native protocol. Once you create a share or export, you cannot change the native protocol.

Note: For non-native NFS access, group membership for trusted domain users cannot contain all groups when the groups are local groups from remote trusted domains.

Consider the following protocol-specific options and caveats before deciding on a primary protocol.

Authentication

SMB clients have only one authentication type, Active Directory (AD), that they can use to access exports. NFS clients can access SMB shares using AD, LDAP, system, or unmanaged protocols. Files uses the native protocol to authenticate all non-native shares.

Authorization

Files uses user mapping to map IDs and user or group names, which authorizes access to shares and exports from native and non-native protocols. Configure user mapping when the NFS authentication of a multi-protocol share is not Kerberos (when NFS uses Kerberos, the users and groups for NFS and SMB are the same).

An IT administrator must configure user mapping, see [User Mapping](#) on page 80.

Access-Based Enumeration

Files supports access based enumeration (ABE) for native and non-native SMB shares.

Access Control Lists (ACLs)

ACLs are stored in the format of the native protocol. Modify ACLs using the native protocol. NFS clients do not honor share-level ACLs. When you create an object using a non-native protocol client, the object inherits access control details from the parent folder.

Alternative Data Stream (ADS)

Files supports alternative data streams (ADS) for native and non-native SMB shares.

Antivirus

Files supports antivirus protection on native and non-native SMB shares.

If you access a file that has a virus through an NFS client-created symlink or hard-link, only the path of the first access appears in the quarantine table.

Audit

Files supports audits for native and non-native shares.

Backup and Restore

Files supports back-up and restore only using the native protocol.

Caution: Backup using a non-native client can result in metadata and ADS loss.

Change Notifications

There is no directory change notification support across protocols. Namespace changes made to a share from one protocol take up to a minute to reflect for a client of the second protocol.

Concurrent Access

Files supports concurrent read access but not concurrent write access for multi-protocol shares. Additionally, you cannot perform reads on a file that is open for writes with another protocol.

High-Availability

Files equips multi-protocol enabled shares and exports with high-availability. Files does not honor the NFS grace period on SMB clients.

Namespace Compatibility

Since all SMB lookups are case insensitive, and all NFS lookups are case-sensitive, SMB clients can encounter multiple objects with the same name string but different cases.

NFS clients can create multiple files on SMB shares that use the same name-string but different cases. This leads to issues for Windows clients that find files with the same name but in different cases. Files maps NFS objects with illegal or special characters in the namespace to a new SMB name.

Quota

Files supports quota for multi-protocol shares, see [Quotas](#) on page 86.

Stats

Files combines stats for native and non-native shares.

Symlinks

NFS clients can create symlinks on native and non-native NFS shares and follow them using a native or non-native protocol.

Limitations

- Files silently ignores permission change requests for non-native protocol clients.
- NFS file or directory names that include trailing spaces are treated as illegal characters on SMB namespaces.
- Native SMB shares cannot enable the NFS squash option.
- NFS users might not see ownership details on native SMB shares. For NFS users, default owner and group on native SMB shares appears as `root/root`. However, SMB users see the default owner and group as `BUILTIN administrators`.
- Template mapping ignores the domain of users and groups.
- Files does not support multi-protocol compatibility for legacy NFS exports.
- When SMB clients create share objects on non-native shares, the objects derive the mode information from their parent share.
- If SMB is the primary protocol, you can only apply a quota to SMB users.
- Files and folders with illegal characters display on the client, but you cannot open or browse them.

Modifying a Share or Export

Modify a share or an export.

About this task

The following limitations apply to modifying a share or export:

- Shares and exports created using earlier versions of Files might not support new features.
- After renaming a share or export, Files retains access to the original share name during the current connection. Close and reopen the browser or CLI to display the updated name.

For information and steps on creating or updating a quota policy, see [Creating User Quota Policies](#) on page 89. For a list of modifiable features on shares and exports with explicit paths, see [Explicit Paths for Shares and Exports](#).

To update a share or export, follow these steps:

Procedure

1. In the Files Console, navigate to the **Shares & Exports** tab.
2. In the row for the target share, click the **three dots menu** > **Update**.
3. (Optional) In the **Basics**, **General Settings**, or **Protocol Settings and Permissions** tabs, update the share configuration.
For information about each field, see the [Creating an SMB Share](#) on page 46, [Creating an NFS Export](#) on page 50, or [Creating a Multi-Protocol Share or Export](#) on page 55 topics.
4. Click **Update**.

What to do next

After renaming an SMB share, manually re-map the share to the client.

Deleting a Share or Export

Delete a share or an export through the Files Console.

About this task

Warning: Deleting a share or export is permanent, and the deleted share or export cannot be recovered.

To delete a share or export, follow these steps:

Procedure

1. Remove the share or export contents and disconnect all clients.
Ensure you delete all directories and files within the share or export.

Note: All shares and exports with explicit paths must be deleted before deleting the parent share.

2. In the Files Console, select the target entry in the share table of the **Share & Exports** tab.
3. In the row for the selected entry, click **three dots menu** > **Delete**.

Caution: When deleting a share or export with an explicit path, the data and storage is not removed, see [Nested Shares and Exports](#) on page 70.

Distributed Share Access

Access home directories without requiring the physical distributed share or export in the profile location.

Home directories are a type of distributed share. If the user profile exists in one of the distributed shares, access the distributed share directly using the universal naming convention (UNC) path.

If your username matches the directory name, you can reach it over the short UNC path: `\Files_server\SamAccountName` (instead of `\\Files_server\home_share\SamAccountName`).

Note: If your home directory exists in multiple distributed shares, Files matches to the first share created chronologically.

Distributed shares appear with other shares in Files when enumerating shares on the file server. Files enables distributed home share support by default. Learn more about distributed shares in the [Nutanix Files Solutions Documentation](#).

Before modifying distributed share settings, confirm the following:

- Home directories are top-level directories in the distributed share.
- You have the required permissions. See the "Permissions" section in [Share and Export Management](#) on page 46.

Files supports using the following commands for FSVMs at the distributed share location

1. To enable or disable distributed home share support:

```
nutanix@fsvm$ afs smb.set_conf "enable user homes" yes|no section=global
```

2. To enable dollar user distributed share access.

```
nutanix@fsvm$ afs smb.set_conf "template user homeshare" "%U$" section=global
```

Configure Preferred Shares

Configure access to distributed directories without the physical distributed share and the profile.

Before you begin

Ensure that you have met the requirements, see [Distributed Share Access](#) on page 64.

About this task

To configure preferred shares for direct access, follow this step:

Procedure

- Configure preferred home shares.

```
nutanix@fsvm$ afs smb.set_conf "preferred home shares" "home-share1, home-share2" section=global
```

Replace *home-share1* and *home-share2* with the names of the homes shares.

Continuously Available SMB Shares

Persistent file handles facilitate continuous availability (CA) for SMB shares.

During a disruption of service, SMB shares remain continuously available by using persistent file handles and high-availability. Persistent file handles reduce the period of data unavailability by automatically reconnecting users to the file service in use before the disruption.

Files stores open lock information on persistent volume groups, available on every FSVM in the cluster. During a fail-over, the fail-over FSVM can access the volume group. Persistent file handles let clients wait to reconnect without notifying you about the connection failure. As a result, disruptions are seamless to you and applications.

During a lock protection interval, Files does not prevent NFS clients from accessing the file. Nutanix Files only prevents SMB clients from accessing files during the disconnected state.

Note: The periodic handler cleans up stale locks for the disconnected persistent file handles every 24 hours.

Synchronous writes on CA shares may impact performance. Nutanix recommends continuous availability for shares with less metadata and more I/O intense workloads, such as VHD or VHDx based profile disks, where performance impact is minimal. As a result, CA is best suited for distributed shares.

Connected Shares

Connect standard or distributed shares in the namespace of another standard or distributed share.

Connecting shares creates a unique, continuous namespace. Use connected shares to distribute data across multiple FSVMs from a specified directory.

Distributed shares let you shard data only on the top-level directory (TLD) level. To shard data from a lower-level directory, connect shares by submounting a distributed share onto the path of another share or export. Files distributes that data from the mount point across multiple FSVMs. The distribution of data achieves load balancing.

The connected, continuous namespace lets you have different settings within the namespace, such as directory-level quotas. For more details on directory-level quotas, see [Creating Directory-Level Quota Policies](#) on page 90. Rather than connecting existing shares, you can also create a new share with an explicit path, see [Nested Shares and Exports](#) on page 70.

Note: For multi-protocol enabled shares, a submount point can apply to multiple directories due to the case-sensitive NFS behavior and case-insensitive SMB behavior. As a result, Files uses all directories with the same name but different cases as submount points.

Connected Share Limitations

The following limitations apply to custom-level sharding:

- The submount path cannot point to the root of a parent share.
- The submount path cannot contain snapshots.
- The submount path on a parent share must be empty before configuring submounts.
- Both parent and child shares must use the same type of NFS authentication.
- Nutanix does not allow nesting submounts. For example, if you have a child share `child1` submounted on the submount path `/parent/dir1/dir2/child1`, you cannot submount `child2` anywhere on the path.
- Nested shares and submounted child shares cannot have the same directory hierarchy.
- Stale locks for disconnected persistent file handles get cleaned up every 24 hours.
- With NFSv3, you cannot directly mount the relative path of a connected share. Either mount the connected share root as an independent share or mount the parent share and then access the connected share.

Connecting Shares

Shard, or connect, a standard share or export.

About this task

To shard shares, follow these steps:

Procedure

1. Sub-mount a distributed share onto a standard share to connect shares.

```
nutanix@fsvm$ afs share.edit child-share-name submount_path=submount-path
```

Replace *submount-path* with the file path to the mount point from the standard share. Replace *child-share-name* with the name of the child share.

2. List the connected shares.

```
nutanix@fsvm$ afs share.list_all_submounted_shares parent_share_name=parent-share
```

Replace *parent-share* with the parent share name.

Disconnecting Shares

Disconnect connected shares or exports.

About this task

To disconnect a child share from a parent share, update the mount point of the child share.

To disconnect shares, follow these steps:

Procedure

1. Replace *child-share-name* with the name of the child share.

```
nutanix@fsvm$ afs share.edit child-share-name submount_path=''
```

2. List all child shares connected to a parent share.

```
nutanix@fsvm$ afs share.list_submounted_shares parent-share-name
```

Note: Due to caching issues, child shares and exports can still appear as mounted.

What to do next

For NFS, remount the parent share on the client.

WORM Shares

Write-once, read-many (WORM) shares offer file retention to prevent inadvertent file deletion.

WORM shares block file modification and deletion operations. As a result, only those applications that do not rely on file modification or deletion work on WORM shares.

Note: WORM cannot be enabled on existing non-worm shares and exports. Also, you cannot apply worm on a newly (already) created share. You can only enable share-level retention (WORM) while creating a new share.

Share-level retention is also called *autocommit*. Configuring a cool-off interval for a share specifies the time for committing. When you do not change a file for the duration specified by the cool-off interval, the change time *ctime* remains unchanged. After the cool-off interval, the file transitions into the *locked* state and goes into retention, also called *commit*. The file remains in retention as specified by the `retention_time` parameter, or the default retention period specified for that share. After the retention time, the file transitions into the *expired* state where you can delete the file but not modify it.

You can edit the file during the cool-off interval. However, once the file is in the retention period, only privileged users can make edits to the file or delete it. If the retention period expires, Files does not support privileged delete operations.

WORM shares created with compliance mode block privileged users from deleting shares in the retention period. Compliance mode also includes the `legal_hold` variable, which, when toggled **on**, prevents the deletion of expired files. Once you create a WORM compliant share, you cannot disable WORM compliance mode.

The following applies to WORM enabled shares:

- If you set the DOS read-only bit, no one can delete the file. To perform the delete operations, you can turn the DOS read-only bit off.
- You cannot modify WORM cool-off interval and retention period parameters after share creation, but you can toggle the legal hold. You can modify share properties unrelated to WORM in the Files UI.
- To delete a file as a privileged user, use the following command:

```
nutanix@fsvm$ afs worm.priv_delete share_name path=file_path
```

To create and configure a WORM share or export, refer to one of the following workflows:

- [Creating an SMB Share](#) on page 46

- [Creating an NFS Export](#) on page 50
- [Creating a Multi-Protocol Share or Export](#) on page 55
- [Creating a WORM Share \(CLI\)](#) on page 68

WORM Share Interoperability

This topic describes the interoperability of write-once, read-many (WORM) enabled shares with other features.

Consider the following interoperability, before using the WORM feature:

- You cannot enable tiering and WORM on the same share.
- Nested shares created within a WORM share have the same WORM configuration as the parent share.
- WORM shares do not retain their properties after SmartSync. You can configure the WORM policy on a remote site.
- Files supports Smart DR for WORM shares. After a Smart DR failover, the destination share has the same properties as the source share.
- Restoring a backed up WORM share converts it into regular share.
- A cloned WORM share has the same WORM configuration as the source.
- Files does not support retention for directories and symbolic links.
- Files share conversion retains the WORM properties.
- Nutanix does not recommend anti-virus (AV) scan on a WORM share or export.
- Files does not support anti-virus policy configuration for WORM shares. Configured anti-virus policies skip WORM shares from scans.
- In the *locked* state, you cannot modify or delete alternate data streams (ADS), like regular files. The following applies In the *expired* state when the DOS read-only bit is off:
 - You can modify treams associated with the base file.
 - You cannot modify the base file and the associated attributes including modified time (mtime) and change time (ctime).

To create and configure a WORM share, see [Creating a WORM Share \(CLI\)](#) on page 68.

Creating a WORM Share (CLI)

You can create and configure a write-once, read-many (WORM) share using Files CLI commands on a File Server VM (FSVM).

Before you begin

Make sure you have read about WORM shares and their interoperability with other features. See [WORM Shares](#) on page 67 and [WORM Share Interoperability](#) on page 68.

About this task

You can do basic WORM share creation, deletion and querying of properties using the `afs` CLI commands.

Note: You cannot enable WORM on existing non-WORM shares and exports.

To create and configure a WORM share, do the following:

Procedure

1. Using SSH, log on to a Controller VM of the Nutanix cluster that deployed Files.
2. Start a file server.
For more information, see [Starting a Nutanix Files Cluster](#) on page 35.
3. Create a WORM share with default values.

Note: If no values are entered for the parameters, default values will be set.

```
nutanix@fsvm$ afs share.add share_name worm_enabled=true
```

4. Create and configure a WORM share with non-default values.

```
nutanix@fsvm$ afs share.add share_name worm_enabled=true worm_retention_period=time in seconds  
worm_cooloff_interval=time in seconds worm_type=ShareLevel
```

(Optional) You can use one or more of the following parameters as needed while creating a share.

Note: You cannot modify the cool-off interval or the retention period after creating the WORM share.

- `worm_compliance=`: Specify as true to block privileged users from deleting shares in the retention period (requires `worm_enabled=true`): Activates compliance mode. Compliance mode lets you set the `legal_hold`.
- `legal_hold=`: Prevents the deletion of expired files by all users.
- `description=`: Provide share description.
- `hard_quota_size_gb=`: Enter the hard quota size in GB.
- `protocol=`: Provide the value as NFS or SMB.
- `share_type=`: Provide the share type as general or distributed.
- `size_gb=`: Enter the size of the share in GB.
- `timeout=`: Enter the timeout in seconds.
- `worm_cooloff_interval=`: Specify the cool off interval for WORM in seconds. The default value for `cooloff_interval` is 600 seconds (10 minutes).
- `worm_enabled=`: Specify the value as true or false.
- `worm_type=`: Specify the type of WORM as `ShareLevel`.
- `worm_retention_period=`: Specify the value of the retention period in seconds. The default value for retention period is 31449600 seconds (52 weeks).

5. (Optional) Check the attributes of the WORM share.

```
nutanix@fsvm$ afs worm.spec share_name=share_name
```

What to do next

For information on deleting an expired WORM share, see [Deleting a WORM Share](#) on page 69.

Deleting a WORM Share

Admin can delete a write once, read many (WORM) share using the Files CLI or the Files Console.

About this task

To delete a WORM share, follow these steps:

Procedure

1. Using SSH, log on to a Controller VM of the Nutanix cluster that deployed Files.

2. Start a file server.

For more information, see [Starting a Nutanix Files Cluster](#) on page 35.

3. Delete a WORM share.

```
nutanix@fsvm$ afs share.delete share_name
```

Replace *share_name* with the name of the share.

What to do next

For deleting shares from the Files Console, see [Deleting a Share or Export](#) on page 64.

Nested Shares and Exports

Use nested shares and exports to hide share paths or to restrict access to parts of a tree.

Unless specified, Files creates shares and exports with an implicit path */share-name*. To create a share or export with a different path, use the **Share Path** field during share or export creation to specify an explicit path for the share. Specifying an explicit path creates a nested share or export. The explicit path for the nested share must begin with the parent share or export. For example, to create an explicit path for share1 you can specify the explicit path */share1/dir1/dir2*. Child shares and exports are nested under the parent shares.

To connect existing shares, see [Connected Shares](#) on page 65.

The following rules and limitations apply for creating, deleting, or modifying shares and exports with explicit paths:

- New nested shares and exports cannot start with the name of another share or export created using an explicit path.
- Explicit paths must start with the forward-slash (/) character.
- When you create a share or export, an explicit path must contain a parent share or parent export name and the directory path. Do not use the root directory.
- Files does not support changed file tracking (CFT) backup for nested shares and exports.
- When you enable or disable multi-protocol on a parent share, the action reflects for all nested shares or exports for that parent.
- Nested shares and exports inherit both the primary and secondary protocols, which you cannot modify at the nested-share level.
- The share path is case-sensitive for the NFS protocol.
- When you delete any of the directories in an explicit path, the path becomes inaccessible.
- Namespace changes under the root of the nested share do not appear immediately when accessed through the parent share path.

Consider the following recommendations:

- Delete all nested shares and exports before deleting the parent share.
- When a path becomes inaccessible after the deletion of a directory in an explicit path, you can recover the path by creating a directory with the same name.

Feature Modification

You can modify some of the properties of nested shares. Most of the time you cannot modify inherited properties. The following table lists modification options.

Table 3: Nested Share Modification Options

Feature	Default State	Modification
SMB access based enumeration (ABE)	Disabled	Enabled
NFS advanced settings	Enabled	Enabled
Add quota policy	Disabled	Disabled
AV settings	Inherited from parent	Disabled
Max size	Inherited from parent	Disabled

DIRECTORY SERVICE AND DOMAIN MANAGEMENT

Administration of directory services and file server domains.

The domain and directory service configuration is an essential step of file server management. Join a domain and update domain name system (DNS) entries during or after file server creation. This chapter describes domain and directory service management options and operations.

The directory services configuration specifies the primary and, if applicable, secondary protocol of a file server. If you intend to create multi-protocol shares and exports on file servers with secondary protocols, configure user mapping. For more information, see [User Mapping](#) on page 80 and [Multi-Protocol Support for Nutanix Files](#) on page 61.

Share and Export Permissions

The authorization setup on native shares and exports determines user and group access permissions to share data.

You can only make access permissions changes using the native protocol.

Native SMB Shares

You can set share permissions when creating or updating a primary SMB share in Files. See [Creating an SMB Share](#) on page 46.

You can add New Technology File System (NTFS) permissions for any valid user or group to the share root or to the sub-folders and files inside the share. Access to share objects is based on share permissions and NTFS permissions on the object. SMB share permissions, SMB share ACLs, and NTFS permissions on SMB shares only apply to SMB clients.

To access multi-protocol enabled shares using a non-native protocol, configure user mapping, see [User Mapping](#) on page 80.

You can also use the Windows native Microsoft Management Console (MMC) shared folders plug-in to manage share ACLs on SMB shares. Use the [MMC Snap-In for Nutanix Files](#) to manage NTFS permissions for distributed native SMB shares. On multi-protocol enabled shares, Files performs permission checks based on the native ACLs.

Default permissions for SMB shares are as follows:

- **Distributed shares:**
 - Domain administrator: Full access
 - Domain user: Read only
 - Creator owner: Full access (inherited only)
- **Standard shares:**
 - Domain administrator: Full access
 - Domain user: Full access
 - Creator owner: Full access (inherited only)

Native NFS Exports

Unix mode bits, set from NFS clients, enforce permissions on NFS exports. By default, the NFS export root directory sets sticky bits, which are permission bits that allow only object owners and root users to remove objects. As an

object owner, you can change the permissions on the files and directories that you own. By default, super users can only change export permissions after disabling root squashing.

Files does not support NFS ACLs.

Default permissions for NFS exports are as follows:

- **Distributed and standard exports:**
 - The initial permissions set on a NFS export (distributed and standard) is not configurable from Prism Element and defaults to mode 1777.
 - The choice of mode 1777 allows normal (non-root) users to create files and directories in any share once a share has been created. The user owns the files and directories. In addition, the sticky bit (setting 1 in the 1777) ensures that only the owner (and root) can delete the created files.

You can change the permissions from any client that has root privilege, but this requires disabling root squashing for the client.

Directory Services and Authentication

This topic provides an overview of directory services and authentication options for Files.

Depending on the primary and secondary protocols configured for your file server, directory service and authentication options vary. Configure directory services at the file-server level in the **Directory Services** tab during or after file server creation. Configure share-level authentication during share or export creation in the **General Settings** tab.

File servers using the SMB protocol use Active Directory (AD) as the directory service. File servers using the NFS protocol can use Lightweight Directory Access Protocol (LDAP), AD, or unmanaged directory services. For NFS and multi-protocol shares, configure authentication during share creation in the **General Settings** tab. For SMB shares, authentication is preconfigured once you set up the AD service for the file server.

The following describes each directory service and available authentication methods.

Active Directory

Files supports AD with NFS and SMB. To set up AD, specify the AD realm name, username, and password. If you have multiple domain controllers, you can configure a preferred domain controller (PDC). If you do not configure a PDC, Files uses any available domain controller.

Using the RFC 2307 standard is an optional setting in AD.

AD with SMB uses Kerberos and Windows NT LAN Manager (NTLM) authentication. SMB shares also support local user access.

For AD with NFS, you can use none, system, or user-based authentication with Kerberos. Each Kerberos type has varying levels of encryption, with Kerberos 5 being the lowest and Kerberos 5p the highest. The none authentication option does not have any security between users: each user can read and write to all other user's files. With the system authentication method, each user has a user ID that restricts other users from accessing their files.

LDAP

Files supports LDAP with NFS only. When you configure LDAP as the directory service, Files uses Netgroups to control client access. The Files LDAP setup is based on the RFC 2307 standard.

To configure LDAP, assign an LDAP server and a base distinguished name (DN). The base DN specifies where the search for users begins. The file server binds to the LDAP server. For added security, set up a bind DN for the LDAP server. Without the bind DN, Files binds to the LDAP server as anonymous.

Unmanaged

Files supports the unmanaged directory service with NFS only. The unmanaged option lets you have none or system authentication, both of which rely on client-based rather than user-based authentication. To set up client-based authentication, specify default client access and exceptions during export creation.

Updating File Server Directory Services

Update the authentication and user management configuration of a file server.

About this task

To update the directory services configuration, do the following:

Procedure

1. In the Files Console (see [Nutanix Files Console](#) on page 18), go to **Configuration > Authentication**.
2. In the **Directory Services** tab, select the protocols you want to use and complete the fields for the selected protocols.

You can select SMB, NFS, or both. Each protocol contains more fields to complete. See the **Directory Services** tab step in [Creating a File Server](#) on page 21 for detailed instructions.

3. When all of the fields are correct, click **Update**.

Setting Active Directory Machine Account Password Expiry

Set an expiry period for the Active Directory (AD) machine account password.

About this task

The Files AD machine account password expiry value is zero by default, so, by default, the password remains unchanged.

For security reasons, if you want to keep the password renewed, follow the steps as indicated to change the password expiry to 30 days.

Note: Renewing the password frequently (more often than once per week) can cause service issues due to possible AD synchronization issues in a multi-DC environment.

To set an expiry period for the AD machine account password, follow these steps:

Procedure

1. Get the AD machine account password.

```
nutanix@FSVM:~$ afs smb.get_conf "machine password timeout"  
[global]
```

2. Update the password expiry period to 30 days.

```
nutanix@FSVM:~$ afs smb.set_conf "machine password timeout" 2592000 section=global
```

The machine password timeout value is in seconds.

Joining a Domain

Join a file server to a domain.

Before you begin

Sync the file server time and Network Time Protocol (NTP) time before joining to a domain.

About this task

Join a domain either during or after file server creation. To join a domain during file server creation, see [Creating a File Server](#) on page 21.

When joining to an Active Directory (AD) domain, Files adds the following entries to AD.

- Machine account with the Files cluster name
- SPN entries

If the AD computer objects exist from previous unresponsive join attempts, the join domain process is unresponsive. Before attempting to join to the domain, remove the previous Files computer objects from AD. Otherwise, do not remove the file server computer object in AD, as that can cause disruptions to file server services.

To join to a domain after file server creation, follow the steps as indicated:

Procedure

1. In the Files Console, go to **Configuration > Authentication**.

For more information, see [Nutanix Files Console](#) on page 18.

2. In the **Directory Services** tab, select one or more protocols you want to use and complete the fields for the selected protocols. You can select SMB, NFS, or both. Each protocol contains more fields to complete, including fields about AD. See the **Directory Services** tab step in [Creating a File Server](#) on page 21 for detailed instructions.

Enter an AD username in the *domain|username* format or in the UPN format *username@ADrealm*. You must have an administrator account or the following required permissions for the relevant AD organizational unit (OU) realm:

- **Create computer objects:** Allows the delegated admin to create a machine account in the specified OU path, which is a requirement for Files to join the domain.
- **Read and write service principal names:** Helps set the service principal names (SPNs) for the machine account. Kerberos access to Files requires SPNs.

3. When all the fields are correct, click **Update**.

Leaving a Domain

Disconnect a file server from a domain.

About this task

Leaving a domain is a prerequisite for the following operations:

- Changing the file server name
- Joining the file server to a different domain
- Deleting the file server
- Changing the primary file server

Leaving a domain deletes the service principal name (SPN) entries and the Files computer accounts created on Active Directory.

Caution: Leaving a domain stops the Nutanix Files service. Without a domain connection, clients cannot use shares, exports, or permissions. Join a file server to a domain to resume services.

Follow these steps to leave a domain:

Procedure

1. In the Files Console, go to **Configuration > Authentication**.
2. In the **Directory Services** tab, click **Leave Domain**.

Directory Services User Mapping

Select the protocol you plan to use and configure directory services options.

☒ Use SMB Protocol

Active Directory is necessary for using SMB protocol

Active Directory Realm Name

child4.afs.minerva.com

Leave Domain

☒ Use NFS Protocol

You can use AD, LDAP or Unmanaged for NFS user management and authentication

User Management And Authentication

Unmanaged

☐ Show NFS Advanced Options

Update

Figure 13: File Server Basics Window

What to do next

After changing the file server name or domain, follow these steps to establish a file server connection to another domain:

1. Remove old domain name system (DNS) entries from one or more DNS servers.
2. If moving to a different domain, update the file server DNS domain configuration. For more information, see [Updating File Server Basics](#) on page 32.
3. Add the new DNS entries, see [Updating Domain Name System \(DNS\) Entries](#) on page 76.
4. Update the directory services, see [Updating File Server Directory Services](#) on page 74.

Updating Domain Name System (DNS) Entries

Modify or add to existing domain name system (DNS) entries.

About this task

Administrators can choose to update or add more DNS entries for the following scenarios:

- Creating a file server
- Changing the file server name or target domain
- Removing a node
- Adding a node

Files adds entries to the DNS server to ensure the availability of file services to clients using a Files namespace. Files uses DNS entries to resolve DFS referrals.

The total number of DNS entries added to the specified DNS server is three times the number of nodes in the Files cluster. Use the following format for DNS entries:

- **File server name for each FSVM IP address:** `file_server_name.domain_fqdn`
- **FSVM name for its IP address :** `fsvm_name.domain_fqdn`.
- **Reverse external FSVM IP address:** `fsvm_external_ip_reversed-in-addr.arpa`

To update the DNS configuration, follow these steps:

Procedure

1. In the Files Console, go to **Configuration > Update DNS Entries**.

For more information, see [Nutanix Files Console](#) on page 18.

2.

» **Automatic (MS-DNS only)**

» **Manual (SII other DNS)** continue to Step 4.

To automatically configure DNS entries, select the **Automatic (MS-DNS only)** option. The automatic option is available only when using Microsoft DNS. If you use a different DNS server or you see a message that Files did not find the entries on the DNS server, use the manual method.

3. Automatically configure DNS entries:

Note: Files can only add PTR entries for the /24 zone. Check the subnet of the external network to confirm the zone. Files does not add PTR records automatically for external networks that are part of a large DNS zone. For these networks, follow the manual process instead.

- a. **Select an action:** Select **Add/Update entries** or **Delete Entries** from the pull-down menu.

Selecting **Add/Update entries** adds or updates the required file server entries to the DNS server automatically. You do not need to enter them yourself. Selecting **Delete Entries** automatically deletes the current file server entries from the DNS server.

- b. (Optional)**Preferred Name Server:** Enter a preferred name server.
- c. **Username:** Enter the user name of the DNS server administrator.
- d. **Password:** Enter the administrator password.
- e. Click **Submit**.

4. To manually configure DNS entries, click the **Manual (All other DNS)** option and do the following:

- a. Copy the entries from the table and add them to the DNS server.

You can either copy the entries manually or click the **Copy All** link to copy the entire table. Click **Actions** to download that table in the CSV or JSON format. To delete the current file servers entries on the DNS server, use this list to determine which entries to delete.

Note: For file server name, add A records only. For FSVM names, add both the A records and PTR records.

- b. Click **Verify**.

After copying the entries to the DNS server, verify that they are correct. Clicking the **Verify** button checks each entry in the table against the DNS server entries. A check mark in the **Is Verified** column indicates that entry is present on the DNS server. There should be a check mark for every entry.

Repeat steps for any entries that do not have a check mark in the **Is Verified** column.

The most likely reason (other than a network issue) for a missing check mark is a copy-and-paste error when transferring that entry to the DNS server. Check your work and repeat this process until Files verifies all entries.

Disjoint Domains

Configure different DNS domain and active directory (AD) realm values through the disjoint domains feature.

Your deployment must have the following requirements prior to configuring disjoint domains:

- A running AOS cluster
- A DNS domain name
- An AD realm name

Configuring Disjoint Domains

To use different domain name system (DNS) domain and Active Directory (AD) realms, configure disjoint domains.

Before you begin

Create a file server.

About this task

To configure disjoint domains, follow these steps:

Procedure

1. After creating the file server configure the protocols and join the domain.

```
nutanix@fsvm$ afs fs.configure_name_services AD-realm-name username protocol
```

"SMB", "NFS", "NFS,SMB", "SMB,NFS" and "None" are valid arguments for *Protocol*.

(Optional) use one of following instead:

```
nutanix@fsvm$ afs fs.configure_name_services AD-realm-name username protocol  
organizational_unit=organizational unit
```

```
nutanix@fsvm$ afs fs.configure_name_services AD-realm-name username protocol  
overwrite=true/false
```

2. Configure Microsoft DNS. For more information, see [Updating Domain Name System \(DNS\) Entries](#) on page 76.

If the DNS domain name and AD realm names are not the same, use the full user principal name (UPN) of an administrator or user.

USER MANAGEMENT

Configure user and group options.

This chapter describes options and steps for managing users and groups on the file server, including their roles, quotas, and mapping across protocols.

The directory services configuration specifies the primary and, if applicable, secondary protocol of a file server. If you intend to create multi-protocol shares and exports on file servers with secondary protocols, configure user mapping, see [User Mapping](#) on page 80 and [Multi-Protocol Support for Nutanix Files](#) on page 61.

User Mapping

Configure user mapping on file servers that have multi-protocol shares or exports.

User mapping lets you access the same share or export using native and non-native protocols. You can retain your user-mapping configuration while configuring the directory services. User-mapping configurations are on a file-server level, which extend to all shares and exports on the file server.

When NFS is the primary protocol with unmanaged authentication, map the Active Directory (AD) username to the NFS client user ID, and map the AD primary group name to the NFS client group ID (GID).

The following restrictions apply to directory service use across protocols:

- NFS clients accessing SMB shares can use IDs from AD, LDAP, or client-generated unmanaged IDs .
- If the NFS security type is **not** Kerberos, configure user mapping for multi-protocol shares and exports
- LDAP usernames cannot be numeric.

Note: Group identifiers (GID) and user identifiers (UID) can appear mismatched because of the access point. The first part of the UID/GID is a config-based range. The last part of the UID/GID is the relative ID (RID) of the user, which is based on the SID. Clients and file servers use different config ranges, so the first part of the GID/UID can appear mismatched.

Mapping Behavior

Files uses the following mapping behaviors:

- *UserGroups*: Default mapping for new shares and exports. Files maps a non-native user to a native user. Files ignores all groups of the non-native user and uses groups of the native user for authorization. Files does not use the groups of the non-native user in the access token.
- *MappedGroups*: Files maps the non-native user and the associated groups to a native user and the respective native groups. The native groups can also be groups of the native user. The access token has a set of user and group identities. Use MappedGroups for the following use cases:
 - Deny access to a specific group.
 - Give access to users based on groups in the ACLs of a file.
- *Legacy mapping*: Mapping for legacy multi-protocol shares and exports (shares and exports on file servers created with versions before Files 3.6.1). Legacy non-native SMB shares use the MappedGroups mapping behavior, and legacy non-native NFS exports use the UserGroup mapping mechanism. You cannot modify the legacy-mapping behavior.

Note: With legacy mapping, Files requires group mapping for the primary group of SMB users to access native NFS exports.

Mapping Configurations

User mapping includes the following mapping configurations:

- **Search:** Use the **Search** tab to search for mapping rules of a user or a group.
- **Rule-Based Mapping:** Use the **Rule-Based Mapping** tab to configure a mapping rule for AD and LDAP users. The following options apply:
 - SMB name matches NFS user name.
 - No template mapping.
- **Explicit Mapping:** Explicit mapping overrides rule-based mapping. The rule entered first takes precedence for users and groups that have multiple mappings configured. Explicit mapping consists of two mapping subcategories: a **one-to-one mapping** list and **wildcard based mapping**.

You can use the one-to-one mapping list to manually enter or upload a csv file that maps users across protocols. Use wildcards for many-to-one mapping. Do not use wildcards on both ends of a user-mapping entry. You can also deny share and export access for a specific user or group.
- **Default Mapping:** Default mapping is the simplest mapping method to configure. You can map all non-native SMB users and groups to a specific native NFS user or group, and conversely. Default mapping also specifies how to handle users that have not had user mapping configured.

The **Summary** tab shows the configured mapping rules and the order in which the rules are prioritized. User-mapping rules take priority in the following order:

1. Deny access rules
2. One-to-one mapping
3. Wildcard mapping
4. Template mapping
5. Default mapping

Configuring User Mapping

Configure non-native user access to multi-protocol shares and exports.

About this task

User mapping connects users between protocols. The following restrictions apply to user mapping:

- User mapping does not support the user name format (UPN).
- Organization units (OUs) and Active Directory (AD) group names can include spaces, but the SMB protocol does not recognize trailing spaces.
- Only one-to-one mappings and template mappings are bi-directional, wildcard and default mappings are not.

For information on mapping rules and types, see [User Mapping](#) on page 80.

To configure user mapping, follow these steps:

Procedure

1. In the **Files Console**, go to **Configuration > Authentication > User Mapping**.

For more information, see [Nutanix Files Console](#) on page 18.

2. (Optional) To configure explicit mapping rules, in the **Explicit Mapping** section, click **Configure** or **Edit**.

a. (Optional) To map single users or single groups, configure one-to-one mapping:

1. (Optional) To map users or groups manually, click **+ Add manually** and follow these steps:

1. In the **SMB Name** field, enter the name of an SMB user or group.
2. In the **NFS ID**, enter the ID of an NFS user or group.
3. From the **User/Group** dropdown menu, select if the mapping is for a **User** or a **Group**.
4. Click the **check mark icon**.

2. (Optional) To upload a file that specifies one-to-one mapping rules, follow these steps:

1. Click **upload a user mapping CSV file**.

If mapping already exists, in the confirmation window, click **Delete Manual Mapping**.

2. Click **Choose File** and select a CSV file with one-to-one mappings from your browser.

Format the CSV file to include three columns with the indicated information in the following order:

- Name of an SMB user or group
- Name of an NFS user or group
- Indication of whether the mapping is for a user or group

3. Click **Upload**.

b. (Optional) To map multiple users or groups to one user or group, follow these steps:

1. Click **+ Add wildcard mapping**.

Note: Files does not support user mapping entries that have wildcards on both ends. Wildcards do not support the use of numeric NFS IDs.

2. In the **Priority** field, choose the priority for the rule.

The lower the number, the higher the priority. For example, 1 indicates the highest priority.

3. In the **SMB Name** field, enter the SMB user or group name.
4. In the **NFS ID** field, enter the numeric NFS user or group ID.
5. From the **User/Group** dropdown menu, specify if the mapping for a **User** or a **Group**.
6. Click the **check mark icon**.

c. (Optional) To deny access to users or groups, follow these steps:

1. In the **Deny Access** section, click **Add SMB or NFS users**.

Note: Deny access rules take the highest priority.

2. To deny access, add comma-separated users and groups in one or more of the following fields:

- **SMB users to be denied access to NFS exports**
- **SMB groups to be denied access to NFS exports**
- **NFS users to be denied Access to SMB shares**
- **NFS groups to be denied access to SMB shares**

d. Click **Save**.

3. (Optional) To configure rule-based mapping map for SMB and NFS users and groups, choose one of the following default rules:

- » SMB name matches NFS name.
- » No template mapping.

Click **Save**.

4. (Optional) To set up default-mapping rules when no applicable rule-based or explicit mapping rules exist for the user or group, follow these steps:

- a. Choose one of the following options from the **SMB Users With No NFS Mapping** dropdown:

- » Map to specific Unix user and group.
- » Deny access to NFS export.

For the **Map to specific Unix user and group** option, configure the following:

- Enter a value in the **Unix UID** field.

Note: If NFS has LDAP, use a user or group name. Otherwise, use a user or group ID.

- Enter a value in the **Unix GID** field.

- b. In the **NFS Users With No SMB Mapping** field choose from one of the following options:

- » Deny access to SMB share.
- » Map to specific AD user and group.

For the **Map to specific AD user and group** option, configure the following:

- Specify a user in the **SMB User** field.

Note: If NFS has LDAP, use a user or group name. Otherwise, use a user or group ID.

- Specify a user in the **SMB Group** field.

- c. Click **Save**.

5. (Optional) To delete all mapping rules for all users and groups, click **Purge All Mapping**.

Caution: Clicking **Purge All Mapping** permanently removes all existing mapping rules.

Changing User Mapping Behavior

Change the user mapping behavior for multi-protocol shares.

About this task

For more information on mapping behaviors, see [User Mapping](#) on page 80.

To update user mapping behavior, follow this step:

Procedure

- Use the Files CLI to update user mapping behavior.

```
nutanix@fsvm~$ afs idmap.set_identity_mapping_behavior options
```

Replace the *options* with one of the following mapping behaviors: *Legacy*, *MappedGroups*, *UserGroups*.

Managing Roles

Manage roles by adding, removing, or modifying administrator privileges.

About this task

You can use a user or a group to create the following two types of Files administrators:

- *File server admin*: The file server admin can manage all file server operations, modify the access permissions for all users in all the shares/exports, and back up and restore data on the file server.
- *Backup admin*: The backup admin can back up and restore data on the file server (but does not have other administrative permissions).

Note: Assign a backup service account (AD user or group) the backup admin role to prevent insufficient access issues.

Follow these steps to manage roles:

Procedure

Creating a New Role

1. In the Files Console (see [Nutanix Files Console](#) on page 18), go to **Configuration > Manage Roles**. The **Manage roles** view displays.
2. To add an administrator, click **+ New User** in the **Add Users** section. A line for new credentials appears at the bottom of the users list.
3. Do the following to configure a role for a user or group:
 - a. In the **User** field, enter the Active Directory user or group name.
Use the *samAcctName* or NETBIOS*samAcctName* format. Replace *samAcctName* with the SAM-account-name.
 - b. In the **Role** field, select **File Server Admin: Full access** or **Backup Admin: Backup access only** from the pull-down list.
 - c. To add the user, click the **check mark** icon.
 - d. To add more administrators, repeat these steps.
The target users or groups must sign out for changes to take effect.

Modifying Roles

4. (Optional) To modify an administrator, in the row for the user, click **three dot menu > edit** and update the name or role.
5. (Optional) To delete an administrator, in the row for the user, click the **three dot menu > edit** for that user.

Managing REST API Roles

Manage REST API access for all users on a file server.

About this task

To add, modify, or remove REST API access, follow these steps.

Procedure

1. In the Files Console, go to **Configuration > Manage roles**.
2. To add a REST API user, click **+ New User** in the **REST API access users** section.
A new line for the new user appears at the bottom of the list.
3. (Optional) To add a user, follow these steps:
 - a. In the **Username** field, enter the username requiring REST API access.

Note: You cannot add "admin" as a username.

- b. In the **Password** field, type in a password for that user.
- c. In the same row, click the **check mark icon** to save the configuration.
- d. To add more users, repeat this step.

Modifying a REST API User

4. (Optional) To modify a username or password, click **three dots menu > Edit** in the row for the target entry.
5. (Optional) To delete a user, click **three dots menu > Delete** in the row for the target entry.

Managing SMB Limited Local Users

Manage limited local users and groups on your file server.

About this task

To add SMB limited local users to a file server for SMB share access, use the command line or the Microsoft Management Console (MMC) Local Users and Groups snap-in. Adding limited local users using the snap-in adds them to the built-in Local Users group. Users in the Local Users group inherit the group's default permissions.

Limited local user management includes the following limitations:

- No access to non-native shares
- No user mapping of limited local users
- No quotas for limited local users
- No limited local user renaming
- A maximum of 100 limited local users per file server

The limited local user must join a domain or a standalone file server joined to an Active Directory domain.

Use the following commands to manage limited local users:

Procedure

Limited Local User Management

- List limited local users.

```
nutanix@fsvm$ afs user.list
```

- Add a limited local user.

```
nutanix@fsvm$ afs user.add limited-localuser-name
```

Replace *limited-localuser-name* with the name of the limited local user.

- Edit a limited local user.

```
nutanix@fsvm$ afs user.edit limited-localuser-name
```

Replace *limited-localuser-name* with the name of the limited local user.

- Delete a limited local user.

```
nutanix@fsvm$ afs user.delete limited-localuser-name
```

Replace *limited-localuser-name* with the name of the limited local user.

Limited Local Group Management

- List the members of a limited local group.

```
nutanix@fsvm$ sudo net sam listmem group-name
```

Replace *group-name* with the name of a limited local group.

Note: Local groups are in BUILTIN\Users and BUILTIN\Administrators directories.

- Add a limited local user to a limited local group.

```
nutanix@fsvm$ sudo net sam addmem 'group-name' 'file-server-name\user-name'
```

Replace *group-name* with the name of a limited local group.

Replace *file-server-name\user-name* with the name of a file server to join and the user name.

- Add a limited local user to the file server ADMIN group.

```
nutanix@fsvm$ afs misc.add_admin_user ADMIN file-server-name\user-name
```

Replace *file-server-name\user-name* with the name of a file server to join and the user name.

- Delete a limited local group.

```
nutanix@fsvm$ sudo net sam delmem 'group-name' 'file-server-name\user-name'
```

Replace *group-name* with the name of a limited local group.

Replace *file-server-name\user-name* with the name of a file server to join and the user name.

Quotas

Nutanix Files has multiple quota types, which include notifications and policies.

Set quotas to allot the storage space a user or group can use.

There are two quota levels:

- *User*: Sets a specific amount of storage for a single user. For example, if an administrator allots only 1 GB, then you cannot use more than 1 GB – the total storage capacity for you is limited to 1 GB.
- *Group*: Sets the amount of space for each user in that group. For example, a group with a policy of 10 GB and 10 users can potentially use 100 GB of data (10 x 10 = 100 GB) under that quota policy.

Alternatively, rather than configuring quotas for specific users, groups, or directories, you can configure a maximum share size to restrict the amount of storage space used by a share. To limit the space in a directory, see [Creating Directory-Level Quota Policies](#) on page 90.

Note: File Analytics and Data Lens provide insights on user usage. For more information see the [File Analytics Guide](#) and the [Data Lens User Guide](#).

Notifications

You can configure email alert notifications that Files sends when user or group quotas are near the maximum threshold. Files can send the alerts to you and other recipients. Emails notifications alert the recipient when the quota is near maximum and when it is near full consumption. When the quota reaches 90 percent consumption, Files sends warning emails to the recipients. When the quota reaches 100 percent consumption, Files sends alert emails to the recipients. If the quota has a soft limit, you can continue to consume over 100 percent of the storage and Files will send an email notification to the recipients every 24 hours.

Policies

A quota policy specifies the consumption limit and enforcement type for all quota levels as configured by the administrator. Enforcement types determine if a user or group can continue to use the quota once they consume their share. See the enforcement types descriptions in the following table.

Caution: Quota policy enforcement begins several minutes after policy creation. Therefore, if you reach the quota limit before the interval is complete, Files raises the alert but does not enforce the quota.

Quota Configurations	Description
User or Group	The designated name for a specific user or group.
Quota	The limit of quota space (in GB).
Enforcement Type	• Hard Limit: Prevents further writes once quota limit is reached. • Soft Limit: Does not prevent writes. Sends email notifications to email recipients. Note: You cannot set both a soft and hard limit for the same user or group.
Email Recipients	Enable the email recipients box and enter the email addresses for recipients Files should notify about hard and soft quota limits.

Applying Quota Policies

Files resolves quota policies as follows:

- If you have defined a Files user-level quota, then recipients receive the quota from this user-level policy.
- If you have not defined a user level, but you have defined multiple group-level policies, then Files applies the policy with the most space.
- If you have not defined a user or group policy for any given user, Files chooses the quota default policy.
- For distributed shares and exports, each user has one home directory. Therefore, quota applies only to the first user directory at the root of a distributed share.




Figure 14: Quota Policy Hierarchy

Note: If you add a new AD group and add a quota policy for that group, changes take 24 hours to reflect. For a faster update, contact Nutanix Support.

Multi-Protocol Limitations

The following limitations apply when you enable both SMB and NFS read-and-write access on a share or export.

- You can only apply a quota to users and groups of the primary protocol. For example, if SMB is the primary protocol, you can only apply a quota to SMB users.
- When you map multiple non-native users or groups to a single native user or group, Files only applies a quota to the first non-native user or group.

- Quota applies to non-native users mapped to native users who belong to a group quota.

Creating User Quota Policies

Create a user or group quota policy.

About this task

Files implements user and group quota types that balance storage per a user of a share or export.

About this task

To configure the quota levels for users or groups follow these:

Procedure

Creating a new quota policy

1. In the **Files Console**, go to **Data Management > User Quotas**.
For more information, see [Nutanix Files Console](#) on page 18.
2. In the **User Quotas** tab, click **Create User Quota**.
The **Add User Quota** dialog window displays.

3. Configure the user quota policy:

- a. Under **Share Selection**, select the target share or export from the dropdown menu.
- b. Under **Add Users**, select to add the quota policy for an **Individual User** or a **User Group**.
- c. In the **Username** field, enter the target user or group names.
- d. In the **Quota Limit** field, enter a number for the quota limit in GiB.
- e. (Optional) To set a soft limit that allow quota writes once a user or group reaches the limit, select the **Allow writes when quota is reached** checkbox.

The soft limit does not prevent writes, but notifies recipients after reaching 90% and 100% of the quota limit.

Note: You cannot set both a soft and hard limit for the same user or group.

- f. (Optional) To send quota notifications to admin, select the **Alert Files admins on Prism Central or Prism Element** checkbox.
- g. (Optional) To send email notifications to users, select the **Send email notifications** checkbox and follow these steps:
 1. Select **Email Files user(s) in the quota**.
 2. (Optional) To notify additional users or groups about the quota policy limit, in the **Additional Email Recipients** field, enter a comma separated list of email addresses.
 3. (Optional) To customize the alert, click **Customize email template**.
 1. Under **Email Subject**, enter a subject for the alert emails.
 2. In the email body section, enter the custom message to appear in quota alert emails.
 3. Click **Save**.

The **Add Quota Policy** window is displayed.

 1. (Optional) To customize the alert, click **Customize email template**.
 1. Under **Email Subject**, enter a subject for the alert emails.
 2. (Optional) In the email body section, enter the custom message to appear in quota alert emails.
 3. Click **Save**.
4. Click **Add**.

Creating Directory-Level Quota Policies

Set a directory-level quota by connecting shares.

About this task

Files supports size limits for shares and exports. Set a directory quota by sub-mounting a share or export onto the path of a standard share.

To set directory-level quotes, follow these steps:

Procedure

1. Create a share or export and specify the max size, see [Creating an SMB Share](#) on page 46.
Alternatively, you can apply a share size limit to an existing share or export, see [Modifying a Share or Export](#) on page 63.
2. Connect the share or export to a directory of a standard share or export, see [Connecting Shares](#) on page 66.
Files applies the quota (max size) for the connected share or export to the directory.

Editing or Deleting a Quota Policy

Update or delete a user or group quota policy.

About this task

To update or delete a quota policy, follow these steps:

Procedure

1. Click the share or export in the **Shares & Exports** tab.
2. Go to the **Quota Policies** tab
3. (Optional) To edit an existing quota policy, click the **three dot menu** > **edit**.
 - a. Edit the existing policy by updating the amount of space, the enforcement type, or the email notifications and recipients.
For information on each field, see [Creating User Quota Policies](#) on page 89.
 - b. Click **Update**.

Note: To remove a quota, change the share size to 0.

4. (Optional) To remove a quota policy, in the row for the target policy, click the **three dot menu** > **delete**.

Role-Based Access Control (RBAC) for Files

You can configure RBAC with Files on Prism Central (PC) using the Files Manager.

For information on RBAC, see the following documentation:

- [Role-Based Access Control Workflows for Files](#) in the *Files Manager Guide*.
- [Controlling User Access \(RBAC\)](#) in the *Security Guide*.
- [Operations Entities](#) in the *Prism Central Guide*.

NUTANIX FILES OPTIONS

Nutanix Files provides a number of features to customize your implementation.

File Server and Share Clones

You can clone a share or a file server from a snapshot at a local or a remote site.

File Server Cloning

Perform file server cloning operations from the **File Server** view in Prism Element. You can only clone Prism Element-based file servers.

By default, the file server clone is unprotected. Be sure to enable the protection domain if you plan to protect the file server clone. Nutanix Files cannot clone snapshots taken on earlier versions of Files, and Files cannot migrate or replicate file server clones to clusters on earlier AOS and Files versions.

File server cloning lets you accomplish the following tasks without impacting the original Files cluster:

- Creating backups at the primary and secondary sites
- Undertaking disaster recovery (DR) tests at a secondary site
- Recovering a file server or share from a specific point in time
- Spinning up a file server at the primary or remote site for testing or development purposes

Important: Renaming cloned file servers can create file server protection domain inconsistencies. To avoid issues, delete the existing protection domain, rename the file server, then create a new protection domain.

Share-Level Cloning

Share-level cloning lets you have multiple clones of the same share. You can create share-level clones with Prism Element and Prism Central–based file servers.

If you delete a parent share, the share cloned from the most recent snapshot becomes the parent share.

Share cloning includes the following limitations:

- By default, the clone depth limit is set to one, meaning that you cannot clone a share clone.
- A single share can have a maximum of ten share clones.
- Share cloning does not support the following operations:
 - Restoring the parent share if the share has clones
 - Cloning shares with tiered files
- Files does not support cloning shares with the following features enabled:
 - Nested shares
 - Write once, read many (WORM) shares

Note: You can clone WORM shares in enterprise mode.

- Smart Tiering

Cloning a File Server

Clone a file server from a snapshot.

About this task

The file server clone is a thin copy that consumes minimal storage space. You cannot clone a file server to a different storage container.

To clone a file server, follow these steps:

Procedure

1. In the **File Server** view in Prism Element, select the target file server.
2. From the operations list, click **Clone**.
The **Clone File Server** window displays.
3. In the **Snapshots** tab, configure the following:
 - **Name of Cloned File Server:** Enter a name for the new file server.
The name of the clone must be different from the original file server name.
 - **Domain:** Enter a fully qualified domain.
 - **List of Snapshots:** Select the snapshot to use for the clone.
Select an existing snapshot or select **Take a new snapshot** to take a new snapshot of the file server.
4. Click **Next**.
5. In the **Client Network** tab, provide client network details.
For information on configuring the client network, see [Configuring the Client Network](#) on page 23.
6. Click **Next**.
7. In the **Storage Network** tab, provide storage network details.
For information on configuring the storage network, see [Configuring the Storage Network](#) on page 24 .
8. Click **Next**.
9. In the **Directory Services** tab, select one or more protocols to use and provide the configuration details.
For information on configuring directory services, see [Configuring Directory Services](#) on page 26.
10. Click **Create**.

Cloning a Share or Export

Create a clone of a share or an export.

About this task

To clone a share, follow these steps:

Procedure

1. In the [Nutanix Files Console](#) on page 18, go to the **Shares & Exports** tab.
2. In the row for the target share, click the [triple dot horizontal](#) icon, then click **Clone Share**.
The **Clone Share** window displays.

3. In the **Share Name** field, enter a name for the share clone.

Note the following naming conventions:

- Share names are not case-sensitive.
- Each share name must be unique.
- A share name cannot start or end with a blank space.
- A share name cannot start with a period.
- A share name must follow a regular expression: `^[^!@#%~'\\"^&$()\\[\\]|+=*{}\\|/\\\\\\ \\<>:; , ' `] * [$] ? $`
- Names that end with the \$ are hidden shares.
- Files does not allow the following names: **Global**, **Printers**, **admin\$**, **ipc\$**, and **homes**.
- See [Configuration Maximums](#) for name length limits.

4. In the **Description** field, enter a description for the share clone.

5. In the **Select a Snapshot** section, select one of the following options:

- » New Snapshot
- » Select from existing snapshot

If you chose the **Select from existing snapshot** option, select the target snapshot from the snapshot table.

6. Click **Clone**.

File Blocking

Restrict specific files or file types from appearing on a file server or share.

Specify a character pattern of file names or extensions to block files. Use an asterisk (*) as a wildcard for multiple characters or a question mark (?) as a wildcard for a single character.

Note: The question mark character (?) only matches UTF-8 single byte ASCII characters. The question mark character (?) does not apply to multibyte unicode characters.

Files applies the file blocking policy to all levels of a share or export, which disables the ability to create files with the specified character pattern in the name. An attempt to create blocked files results in an error. Share-level file blocking overrides the files blocked on the file-server level.

Files allows a maximum of 300 file blocking patterns on a file server.

After enabling file blocking, Files does not permit the following operations:

- Creating a file with the blocked character pattern.
- Renaming an existing file to one with the blocked character pattern.
- Duplicating a file with the blocked character pattern.
- Moving a file with a blocked character pattern.

You can still perform read-and-write operations on existing blocked files.

To enable file blocking on a file server, see [File Server Updates](#) on page 31. To enable file blocking on a share or export, see [Modifying a Share or Export](#) on page 63.

Files blocked on a share level appear in the **Summary** tab of the create and update share windows. Files blocked on a file server level appear in the **Blocked File Types** view of the Files Console..

Blocking Files on a File Server

Block the creation and modification of files on a share, export, or file server.

About this task

For information about the file blocking feature, see [File Blocking](#) on page 94. To block files on the share or export level, see [Modifying a Share or Export](#) on page 63.

To block files with specific character patterns in their names, follow these steps:

Procedure

1. In the [Nutanix Files Console](#) on page 18, go to **Configuration > Blocked File Types**.
2. In the **Blocked File Types** field, enter or modify a comma-separated list of file extensions for blocked file types.
3. Click **Save**.

Enabling SMB Symlinks

Enable symlink support over the SMB protocol.

About this task

A symbolic link, also known as a symlink or a soft link, is a file type that contains a path to a target file or directory. Although symlink support is enabled by default, Files requires enabling the following link configurations on Windows clients to use symlinks on SMB shares:

- Remote-to-remote
- Remote-to-local

Note: By default, symlinks are disabled on Windows clients. To enable, symlink access on Windows clients, use the following command.

```
> fsutil behavior set SymlinkEvaluation R2R:1 R2L:1
```

SMB clients must have symlink access enabled to access symlinks and follow the targets.

Files does not support the following functionalities for symlinks:

- Directory symlinks at the root of distributed shares.=
- Alternate data streams (ADS)
- Durable or persistent file handles
- Opportunistic locks (oplocks) and leases

Additionally, there is no interoperability support between SMB and NFS on Multi-protocol shares for symlinks with absolute targets.

Note: Hycu does not support symlinks. Commvault does not support cross-share symlinks.

To turn symlink support on or off, follow these steps:

Procedure

- Enable symlink support on the file server.

```
nutanix@fsvm$ afs smb.set_conf "enable smb symlinks" "True" section=global
```

- Disable symlink support on the file server.

```
nutanix@fsvm$ afs smb.set_conf "enable smb symlinks" "False" section=global
```

Disabling SMB links only prevents the creation of new symlinks. Existing symlinks remain accessible.

Managing Access-Based Enumeration for SMB Shares

Access-based enumeration (ABE) for the SMB protocol restricts user-access by only letting you view the files and folders you have read-access to when browsing content on a file server.

About this task

ABE is a Microsoft Windows SMB protocol feature that filters the list of available files and folders on the file server to only include those you have read-access to. The filtering ensures that Files enforces read-and-write privileges for all users and that information can remain confidential. ABE controls the user visibility of shared folders on mounted file system shares based on the user permissions.

To enable ABE during share creation, see [Creating an SMB Share](#) on page 46. To activate ABE after a group membership of a user changes, remove all previous share sessions, remount the share, and reconnect existing client connections. Depending on the number of folders or files on a share, enabling ABE can slightly affect performance.

To enable ABE, follow these steps:

Procedure

1. In the Files Console, go to the **Shares & Exports** tab.
2. In the row for the target share, click the **three dots menu** > **Edit**.
3. In the **Protocol Settings and Permissions** tab, follow these steps:
 - a. (Optional) To enable ABE, select the **Enable Access Based Enumeration (ABE)** checkbox.
For standard shares, Files applies ABE to all folders and files on the share. For a distributed share, by default, Files only applies ABE to the share and top-level directories (TLDs).
 - b. (Optional) To enforce ABE on the entire subtree, use the following command:

```
nutanix@fsvm$ afs smb.set_conf "abe tld subtree" "yes" section=global
```

- c. (Optional) To disable ABE, unselect the **Enable Access Based Enumeration (ABE)** checkbox.
- d. (Optional) To list the shares that have ABE enabled, use the following command:

```
nutanix@fsvm$ afs smb.share_info display_abe_only=True
```

Output displayed is similar to the following.

```
Share Name: dep-sh1
Share Type: DEFAULT
ABE Enabled: ENABLED
NVM Host: NTN.example.com
```

```
Share Name: home-sh1
Share Type: HOMES
ABE Enabled: ENABLED
```

4. Click **Save**.

Durable SMB File Handles

Durable handles facilitate seamless SMB client reconnection.

Nutanix Files automatically enables durable handles for SMB shares. Durable handles let SMB clients survive a temporary client-side connection loss after opening a file, allowing transparent client reconnection within a timeout.

Durable handle reconnection times increase when a previous session, with multiple durable handles, must flush the data and metadata before re-establishing the connection.

Durable file handles have the following limitations:

- Files does not grant durable file handles for client negotiations with a dialect lower than SMB 3.0.
- Files does not support durable file handles for the directories or alternate data stream (ADS) files.
- Files does not support durable reconnects during server-side copy operations due to the client-side limitation of not attempting to restore the source handle first.

Non-Disruptive Upgrades (NDU) for SMB

NDU support for the SMB protocol provides a seamless reconnect experience during a planned outage for SMB connections to non-continuously available shares.

Note: NDU support is a best-effort solution without any guarantees for SMB client reconnection.

Nutanix Files upgrades one node at a time, which is known as a rolling upgrade. Services on one FSVM fail over to another FSVM before an upgrade begins and fail back to the source FSVM after completing the node upgrade. If the failover or failback succeeds within a minute, client applications can seamlessly reconnect to Files without any failures.

Applications that cannot tolerate network disruptions of more than a minute might time out during upgrades. In such cases, the client application retries to establish the connection to Files. To retry establishing a connection, services save the client connection state to the persistent store just before an FSVM upgrade begins, and then restore the state on the FSVM that takes over after services become available. This functionality makes most SMB client applications stay connected during the upgrade process.

We recommend using continuously available shares for applications that require strong reconnect guarantees, for example FSLogix and other container-based user profile applications. For more information, see [Continuously Available SMB Shares](#) on page 65.

Limitations

The non-disruptive upgrade feature does not support file copy over remote desktop protocol (RDP) onto an SMB share.

Antivirus (AV) Scanning for SMB Shares

Nutanix Files supports the Internet Content Adaptation Protocol (ICAP) to enable communication with external servers hosting third-party antivirus software.

Antivirus software scans files to help provide protection against viruses when you open, close, read from, or write to files. By default, shares have antivirus scanning disabled.

For a list of compatible security software, see the [Compatibility Matrix](#). Files requires two or more ICAP servers. Nutanix recommends having a minimum number of scanning threads that is 11 times the number of FSVM nodes or (11 * number of FSVM nodes).

Overview

Nutanix Files performs the following tasks with ICAP servers when a client requests to read, write, open, or close a file:

1. Files determines that the file requires scanning.
2. Files sends files that require scans to the ICAP server with a scan request.
3. The ICAP server scans the file and reports the scan results to Files.
4. Files quarantines and denies access to unsafe files.
5. If the file is clean or disinfected, then Files allows the client access to the file.

File States and Events

The following terms describe the status of antivirus scanning.

- *Quarantined*: A scanned file that the antivirus scan qualifies as unsafe. Files blocks access to the file until you manually change the file state.
- *Unquarantined*: The administrator moves the file from the quarantined state to allow client access. Files does not rescan unquarantined files.

The following terms describe antivirus scanning events:

- *Cleaned*: The antivirus scanner has scanned and cleaned the file. This process overwrites the original file. Using this feature requires the disinfected virus file function on the ICAP server.
- *Quarantined*: A scanned file qualified as unsafe. Files blocks all access to the file until the administrator manually changes the file state.
- *Unquarantined*: The administrator manually moves the file from the quarantined state to allow client access. Files does not rescan unquarantined files.
- *Deleted*: File removed from the file system.

Configuring Antivirus Scanning for SMB

Configure and enable antivirus scanning for SMB shares.

About this task

To view the entire antivirus scanning configuration page, you need the `View File Server Antivirus Server` and `View File Server Virus Scan Policy` permissions.

After configuring the antivirus scan, enable the scan for each share Files should scan.

To configure an Antivirus scanning, follow these steps:

Procedure

1. In the [Nutanix Files Console](#) on page 18, go to **Configuration > Antivirus**.

2. Follow these steps to connect the ICAP server:

a. Click **+ Connect ICAP Server**.

A new row appears for new ICAP server details.

b. Enter the indicated information in the corresponding fields:

- **IP address or hostname**
- **Port** (the default port number is 1344)
- **Service name**

Note: Only use this field if you have changed the default service name of the ICAP server. The default service name for most ICAP servers is `avscan`.

- **Description**

c. Click the **check mark icon**.

For a detected antivirus server, the software tests the validity of the configured server and updates the status to OK.

d. Ensure the connection status automatically updates to **OK**.

e. Click **Next**.

3. Complete the **Scan Settings**.

You can override settings through the share-level antivirus settings.

Note: Nutanix recommends two or more ICAP servers.

- a. **Scan on Write:** Scans saved and updated files (a write operation).
- b. **Scan on Read:** Scans opened files (read operation).
Nutanix recommends to always enable **Scan on Read**.
- c. **File extensions to be excluded:** Add one or more file extensions to exclude from the scan.
Ensure the settings match the file type configuration of your ICAP servers.
Nutanix recommends adding the following file extensions for user profiles when using the Files antivirus scanning:

- .dat
- .ini
- .pol

Note: When the ICAP server incorrectly quarantines files with specific extensions, adding this file type extension to the ignore list only prevents future file quarantines. To access the incorrectly quarantined files, remove them from quarantine.

- d. **File Size:** Specify a file size to limit scanned files according to size.
- e. To specify **Advanced Settings**, do the following:
 - **Scan Timeout:** Set the maximum amount of time that a scan can take before timing out.
 - **Block access to files if scan cannot be completed (recommended):** Selecting this option blocks access if the ICAP servers are unavailable or cannot scan the file for any reason.

Important: If you do not choose the block access option, repeated scan requests might cause severe latency when accessing shares if the ICAP server is unreachable.

- f. Click **Save**.

4. Enable the antivirus scanning on each share.

- a. Go to the **Shares & Exports** tab and click on the target share.
- b. In the **Share Details**, go to **Actions > Configure Antivirus**.
The **Configure Antivirus** setup window displays.

- c. **Note:** To **Disable Antivirus scan**, uncheck the **Enable antivirus scan** and click save.

Check the **Enable antivirus scan** box.

- d. Check the **Enable antivirus scan** box.
- e. (Optional) Change the settings for the share (see Step 3 for details).
- f. Click **Save**.

Antivirus Tab

The layout and elements of the **Antivirus** tab in the Files Console.

The **Antivirus** tab displays dynamic information about scanned files. The tab includes the following sub-tabs:

- ICAP servers
- Reports
- Quarantined Files
- Unquarantined Files

To view this information, in the Files Console, go to **Monitoring > Antivirus**.

ICAP Servers

The ICAP servers tab displays the scanned files information for each ICAP server.

- **ICAP Server Statistics:** The table displays information such as port number, description, files scanned, disconnect count, average latency, connection status, and actions available.
- **Average Latency:** This graph displays the latency times for the scans (in milliseconds).
- **Files Processed** or **Data Processed:** Click the **Files Processed** drop-down arrow to select the files processed or data processed graph. The processed files graph displays the number of scanned files. The data processed graphs display the amount of processed data (in GiB).
- **Queue Length:** The number of files in the scan queue.

Reports

The **Reports** tab displays the information about the scanning period and share status.

- **Scan Period:** This information displays the files scanned, threats detected, number of files cleaned, and number of files quarantined during each scanning period.
- **Share Status:** Displays the state of the scanned share. The parameters includes: file path, threat description, ICAP server, time, action taken on share.

Quarantined Files

The **Quarantined Files** tab displays the files that contain a virus. The antivirus software places virus-infected files into quarantine where clients cannot read or write the files. An administrator can perform the following actions on the quarantined files:

- **Rescan:** Rescan the files that have been quarantined.
- **Unquarantine:** Move the files out of quarantine. The selected file is then available for use.
- **Delete:** Delete the quarantined file permanently.

Unquarantined Files

The **Unquarantined Files** tab displays files manually released from quarantine. You can use unquarantined files. Files does not rescan unquarantined files again until the administrator resets the unquarantine state. Perform the following actions on unquarantined files.

- **Reset:** Move the files to a normal state that is not quarantined or unquarantined. In this state, the next access to the file triggers the scan.
- **Quarantine:** Move the files to quarantine to block read and write access.

Note: Reaching the limit of the number of files in both the **Quarantined** and **Unquarantined** tables impacts scanning through the UI. The web consoles alerts you when the number of files has reached 80 percent of the number of files supported.

PERFORMANCE OPTIMIZATION

Files performance optimization notifies you when the file server is under load and needs a change to improve performance.

Performance optimization includes scale-up and scale-out options. The **Dashboard** tab in the Files Console includes a **Recommendations** widget that indicates performance optimization options.

Scale-up recommendations occur when an FSVM reaches 95% client connection utilization over a two hour time window. If performance is slow due to storage group disruption, the recommendations for optimal performance options appear in the **Recommendations** widget. When possible, perform scale-out and rebalancing operations during maintenance windows, as scaleout and rebalancing disrupts existing connections.

Hypervisor Attached Disks for AHV

Use hypervisor-attached disks (HAD) to attach volume groups (VGs) directly to file server VMs (FSVMs).

Enabling the direct attach mode for volume groups enhances performance functionality by replacing iSCSI communication between VGs and FSVMs with direct communication. HAD reduces dependencies on network connections, isolates traffic, and eliminates unencrypted iSCSI communication between FSVMs and volume groups.

You can create file servers with direct attach mode during file server creation in Prism Central. For HAD file server creation with Prism Element, you must enable direct attach mode, which creates all new file servers with HAD. If you enable HAD in Prism Element, disable direct attach mode to create file servers without HAD. To enable HAD mode in Prism Element, contact Nutanix Support.

Limitations

Direct attach mode has the following limitations:

- HAD is only available for newly created file servers.
- Metro Availability with Files on AHV is not supported.
- Cross-hypervisor disaster recovery with ESXi is not supported.
- For Protection Domain (PD) based Disaster Recovery, both sites must meet minimum Nutanix Files and Files Manager version requirements.

Workload Optimization

Files provides multiple options for optimizing the performance of your file server.

There are three types of share workload types: default, random, and sequential. Designating a workload type determines the file system characteristics (including block size) used for the share, which optimizes the resource usage and performance of certain workloads. For example, workloads with small input and output (I/O) on large files perform more efficiently with small block sizes.

The different workload types have the following specifications.

- **Default:** Uses 64 KB per block. Does not provide specified optimization. The share can perform all workloads varying performance
- **Random:** Uses 16 KB per block. Optimized for small I/O workloads.
- **Sequential:** Uses 1 MB per block. Optimized for large I/O workloads. Requires a minimum of 24 GB memory per FSVM.

The **Shares** view in the Files Console includes a **Metrics** tab and a **Performance** sub-tab, which displays write, read, and metadata I/O per second in the **IOPS** graph. Use the data from the graph to configure the workload type for the share, see [Modifying the Workload Type](#) on page 104.

Caution: If you modify the share type without following the workload optimization guidance as specified, share performance can degrade.

- If the I/O sizes for read and write operations are less than or equal to 16 KB and the file sizes equal to 10 MB or more, use the **Random** workload type.
- If the I/O sizes for read and write operations are less than or equal to 1 MB and the file sizes equal to 10 MB or more, use the **Sequential** workload type.
- If the I/O sizes for read and write operations do not match the criteria for **Random** or **Sequential** workload types, use the **Default** workload type.

Modifying the Workload Type

To optimize performance, modify the workload type of a share or an export.

About this task

Changing the workload type of a share changes the performance characteristics only for the files created after the change.

To modify a workload type, follow this step:

Procedure

- Modify the workload size.

```
nutanix@fsvm$ afs share.edit share-name share_workload_type=workload
```

Replace *share-name* with the name of the share or export. Replace *workload* with one of the following workload types: `default`, `sequential`, or `random`.

Note: For more information on workload size, see [Workload Optimization](#) on page 103.

Removing Rebalancing Guardrails

You can unblock rebalancing guardrails.

About this task

Files guardrails prevent initiating rebalancing operations during high-volume periods.

To remove load-balancing guardrails on a file server, follow this step:

Procedure

- Unblock load-balancing guardrails on the file server.

```
nutanix@fsvm$ afs lb.unblock_rebalancing
```

File System Compression

File system compression reduces the input and output (I/O) load, iSCSI traffic, space usage, and the amount of data on a share or an export.

You can enable file system compression at the share and export level. Earlier versions of Files included compression at the container level. Files applies share-level compression during ingest operations, compressing incoming data in-line prior to writing it to storage. As a result, share-level compression reduces storage traffic between Files and AOS.

Note: Only clusters created with later versions of Files and AOS support file system compression. The option to enable file system compression does not appear on clusters created with earlier versions of Files and AOS.

To enable or disable compression on the share or export level, see the following:

- [Creating an NFS Export](#) on page 50
- [Creating an SMB Share](#) on page 46
- [Creating a Multi-Protocol Share or Export](#) on page 55
- [Modifying a Share or Export](#) on page 63

Managing the Performance Profile

Enable or disable the performance profile, which lets you create file servers that meet the performance requirements for artificial intelligence (AI) and machine learning (ML) workloads.

About this task

The performance profile does not support the following:

- Protection domain (PD) based DR operations like file server cloning and file server activation
- Deployment of PC, VPC, or cloud-based file servers
- Using iSCSI extensions for remote direct memory access (iSER) for storage traffic when a file server is in high availability mode and with more than one FSVM on the same host

Caution: Nutanix recommends enabling the performance profile on one file server in a cluster. Enabling the performance profile on more than one file server in a cluster can cause performance degradation.

To manage the performance profile, do the following:

Procedure

- Manage the performance profile for file servers.
 - » Enable the performance profile before creating file servers for AI and ML workloads.

```
nutanix@cvm$ afs infra.set_performance_profile true
```
 - » If you do not want to create file servers for AI and ML workloads, disable the performance profile.

```
nutanix@cvm$ afs infra.set_performance_profile false
```
 - » Enable the maximum performance profile.

```
nutanix@cvm$ afs infra.set_max_read_performance_profile file-server-name
```

Replace *file-server-name* with the target file server.

Note: We recommend contacting Nutanix Support to configure maximum read performance.

NFS Over RDMA

Remote direct memory access (RDMA) technology lets you optimize data transfer for Nutanix Files deployments.

The NFS over RDMA performance configuration supports latency-sensitive and high-bandwidth workloads such as machine learning and analytics. Using RDMA adapters between storage and host memory systems, Nutanix Files can

transfer data more efficiently, resulting in optimized CPU usage and reduced latency. Additionally, support for GPU direct storage (GDS) facilitates data transfer directly from file servers to GPU memory.

Performance improvements include a fivefold reduction in latency for electronic design automation (EDA) and software build workloads and a threefold improvement in single client throughput.

NFS Over RDMA with AHV

Nutanix recommends that the cluster running NFS over RDMA meet the following minimum requirements:

- The controller VM (CVM) needs a minimum of 16 vCPU and 64 GiB memory
- Each file server VM (FSVM) needs a minimum of 16 vCPU and 64 GiB memory

To enable NFS over RDMA for AHV hypervisors, update the client network to the single root I/O virtualization (SR-IOV) configuration in Files on Prism Central. For more information, see [Updating the Network Configuration to SR-IOV](#) in the *Files Manager User Guide*.

The following features are not supported when SR-IOV is enabled:

- Protection domain (PD)-based DR
- Live migration of VMs between hosts with SR-IOV NICs
- Multi-VLAN with VLAN trunking

If you plan to start a host upgrade, turn off any file servers that are already configured with SR-IOV before starting the host upgrade.

NFS Over RDMA with ESXi

As a prerequisite for ESXi hypervisors, configure Mellanox OpenFabrics Enterprise Distribution (OFED) drivers on clients before enabling NFS over RDMA.

To enable NFS over RDMA for ESXi, contact Nutanix Support.

1N:2D Fault Tolerance

Enable one node or two disk fault tolerance on a container level to enhance data resiliency.

One node or two disk (1N:2D) fault tolerance mode helps protect your data by decoupling disk and node failures. Enable this feature to provide enhanced storage efficiency and resiliency on dense node clusters.

Meet the following precheck requirements to enable 1N:2D on a newly deployed file server:

- Enable 1N:2D fault tolerance before creating a file server.
- Enable fault tolerance (FT2) with redundancy factor 3 on the AOS cluster.
- Confirm that each node in the dense node platform hosts at least 20 hard disk drives (HDDs).
- Ensure that no external FT2 containers are present.

Note: Files rejects attempts to create redundancy factor 2 file servers in 1N:2D and redundancy factor 3-enabled environments.

If a file server fails to meet the precheck requirements during deployment, Files deploys it with 1N:2D disabled.

To enable 1N:2D fault tolerance, contact Nutanix Support.

Disabling SmartOps Shares

By default, Nutanix Files adds a SmartOps share to each new file server to accelerate the speed and efficiency of recursive file deletion.

About this task

The file server offload mechanism facilitated by SmartOps shares lets you optimize large-scale file deletion operations. As an added benefit, you can delete top-level directory (TLD) content on distributed shares without relying on the MMC Snap-In for Nutanix Files.

Files enables the SmartOps share feature by default.

To disable the SmartOps share feature on a file server, follow this step:

Procedure

- Disable a SmartOps share:

```
nutanix@FSVM$ afs smartops.disable
```


DATA MANAGEMENT

Files has multiple file server data recovery and management options.

Files provides several features for managing the data on your file server. Manage data availability using data recovery features, which include self-service restore (SSR) and several types of disaster recovery (DR). Consolidate and organize your data using data management features.

Files provides the following DR features:

- **Async and NearSync DR:** Async and NearSync DR are Prism Element (PE) based DR solutions that replicate data to a protection domain on a set schedule. The *Data Protection and Recovery with Prism Element* guide provide information on both solutions.
- **Metro Availability:** Unlike Async and NearSync DR, Metro Availability provides synchronous data replication to the protection domain.
- **Smart DR:** Smart DR replicates data with share-level granularity to a recovery file server.

Files provides the following data management features:

- **Smart Tiering:** Smart Tiering maximizes space on the file server by moving stale data to an object store. Use the Files Console to configure standard Smart Tiering policies and Data Lens for advanced Smart Tiering. See the [Smart Tiering](#) on page 116 and the [Data Lens User Guide](#) for more information.
- **Data Synchronization:** Data synchronization is a Prism Central (PC-based) feature that lets you consolidate data from one or more source paths on a file server to another file server. Manage synchronization policies and locations using the Files Manager (Files on PC). For information on data synchronization, see [Data Synchronization](#) in the *Files Manager User Guide*.

Important: You cannot upgrade file servers with technical preview data synchronization policies to general availability (GA) versions. Delete policies created during the technical preview.

- **Self-Service Restore:** SSR takes read-only snapshots at the share level. Based on your snapshot retention configuration, you can recover a share, data from a deleted file, or an earlier file version. File-level SSR for SMB lets you restore and manage previous versions of Files at the system level, while file-level SSR for NFS lets you restore files by manually copying read-only snapshot versions. Both DR and SSR let you configure a snapshot interval, but with SSR, you cannot set a specific time to take the snapshot. See [Self-Service Restore](#) on page 124 for more information.

Nutanix Files Data Protection and Recovery in Prism Element

Nutanix Files provides disaster recovery (DR) in Prism Element through customizable protection domains and protection policies.

Tip: For information on Files disaster recovery in Prism Central, including with share-level replication, see [Asynchronous Data Protection with Smart DR](#) and [Synchronous Data Protection with Metro Availability](#) in the *Files Manager User Guide*.

As part of DR, Files creates a protection domain for a file server and the entities within the file server, such as VMs and volume groups, during file server creation. By default, Files adds all entities on the file server to the protection domain. You can provide custom names for the file server protection domains at the time of setting up the file server. For more information, see [Creating a File Server](#) on page 21. If a file server does not have a specified protection domain name, the default protection domain name is `NTNX-file_server_name`. If you clear the checkbox for the protection option during file server creation, Files does not create the protection domain.

For more information on managing protection domains, see [Data Protection and Recovery with Prism Element](#).

To activate Async or NearSync DR, enable and set up the schedule for snapshots and replication for the protection domain. Files Async and NearSync DR take snapshots when the preceding snapshot is complete. Async has a 60-minute recovery point objective (RPO), and NearSync has a 1-minute RPO. For information on activating synchronous replication with Metro Availability, see [Metro Availability on ESXi](#) or [Synchronous Data Protection with Metro Availability](#).

Files creates a dedicated container for each file server instance, which cannot be used by another file server, VM, or for any other purpose. This requirement also applies to remote containers used for replication. Before replicating a file server container to a remote site, make sure that the remote container, like the local container, is not used for any other purpose. The remote site must have at least the same number of nodes as the number of FSVMs in the file server. To ensure feature parity after activating the file server on the remote site, both sites must have the same AOS version.

Note the following functionalities and limitations:

- You can restore all self-service restore (SSR) and Windows Previous Version (WPV) snapshots that exist at the time of the protection domain (PD) snapshot locally or remotely.
- Files does not support DR with one-hour RPO on high-density nodes. To check the supported frequency for the snapshots across various hardware configurations, refer to the "Snapshot Frequency" table in the *Resource Requirements Supporting Snapshot Frequency (Asynchronous, NearSync)* section of the [Data Protection and Recovery with Prism Element](#) guide .
- Files does not perform frequent suspend and resume operations of the protection domain. Files performs protection domain suspend and resume operations during meta-data pinning on flash drives.

Files supports cross-hypervisor NearSync and Async DR.

Configuring Protection Domain-Based Disaster Recovery

Set up asynchronous Prism Element–based file server disaster recovery for planned or unplanned migration.

About this task

To protect the VGs associated with a file server, Nutanix recommends using this Files Console workflow and not the Prism Element workflow. You need vStore mapping to initiate the full replication of protected entities to the specified target container. You can only define the vStore mapping for the file server using Files disaster recovery workflow.

To set up protection-domain based file server asynchronous disaster recovery, follow these steps:

Procedure

1. If you have not done so already, configure a remote backup site to the local cluster.
See [Configuring a Remote Site \(Physical Cluster\)](#) in the *Data Protection and Recovery with Prism Element* guide.

Note: The remote site must have at least the same number of nodes as the number of FSVMs in the file server.

2. In the **Files Console** (see [Nutanix Files Console](#) on page 18), go to **Data Management > Protection**. The **Disaster Recovery** and **Self Service Restore** tabs appear. In the **Disaster Recovery** tab, the **Protection Domain (PD) Based** and **Smart DR** sections specify if the indicated DR type is enabled.

After creating a file server, Files automatically adds the file server to a newly created protection domain. However, when the file server does not have a protection domain, the **Configure** action link displays in the **Protection Domain (PD) Based** section with the **not enabled** status.

3. If the **Configure** action link displays in the **Protection Domain (PD) Based** section, click **Configure**. Files redirects to the **File Server** view in Prism Element.

4. In the file server table, click the row of the target file server.
5. Under the file server table, click **Protect**.
The **Protection configuration: [file-server-name]** window appears.
6. In the **Disaster Recovery** section, do the following:
 - a. (Optional) in the **Protection Domain Name** field, update the name of the protection domain.
 - b. Click **Protect File Server**.
7. (Optional) configure a schedule for disaster recovery.
Prism Element creates a default schedule for every protection domain.
 - a. In the Files Console, go to **Data Management > Protection**.
 - b. Click **Manage on Prism Element**.
Files redirects you to Prism Element.
 - c. See [Creating a Protection Domain Schedule](#) on page 110 to add or modify the protections schedule.
8. Configure the local and remote container mapping.
 - a. If you did not map the local and remote containers when configuring a remote site (vStore name mapping), create a new remote container.
In Metro Availability enabled environments, avoid using identical container names. Ensure that the name of the remote container is unique, then map the containers to each other.

Note: For more information, see the [Creating a Storage Container](#) section in the *Prism Web Console Guide*.

Changing a vStore mapping causes associated protection domains to initiate full replication of protected entities to the newly specified target container as if it were initial replication. A mapping change therefore results in the overconsumption of storage resources at the remote site. Contact Nutanix Support for help with cleaning up snapshots in the previously specified container.

- b. Ensure that the remote and the local containers have symmetric configurations and that both containers map to each other by configuring explicit mapping between the source and the destination container on the remote vStore site.

For more information, see [vStore Data Protection](#) in the *Data Protection and Recovery with Prism Element* guide.

Creating a Protection Domain Schedule

Create a snapshot schedule for Files protection domains.

About this task

Files disaster recovery uses snapshot schedules for protection domains. Make sure to also save the snapshot schedule on the remote site.

To create a protection domain schedule, follow these steps:

Procedure

1. In Prism Element, go to the **Data Protection** tab in the pull-down menu.
2. In the **Table** view and **Async DR** tab, select the protection domain from the table.

3. Click **Update**.

The **Update Protection Domain** window appears.

4. In the **Schedule** tab, click **New Schedule** (or, to update an existing schedule, click the **pencil icon**).

5. Complete the indicated fields:

The screenshot shows the 'Schedule' configuration window. On the left, under 'Configure your local schedule', there are three radio buttons: 'Repeat every' (selected), 'Repeat weekly', and 'Repeat monthly'. The 'Repeat every' option has three input fields for 'minute(s)', 'hour(s)', and 'day(s)'. The 'Repeat weekly' option has checkboxes for S, M, T, W, T, F, S. The 'Repeat monthly' option has a 'Day of month' input field with 'e.g., 1,10,20'. Below these are 'Start on' and 'End on' date and time pickers. On the right, under 'Retention policy', the 'Local' checkbox is checked, and the 'keep the last' input field is set to '1' snapshots. A message states 'Remote sites have not been defined for this cluster.' At the bottom right are 'Cancel' and 'Create Schedule' buttons.

Figure 15: Create the File Server Protection Domain Schedule

- a. **Repeat every [minutes|hours|days]:** Click the appropriate circle for minutes, hours, or days and then enter the desired number in the box for the scheduled time interval.

The interval cannot be less than 1 minute.

Note: Intervals of less than 60 minutes use NearSync disaster recovery. NearSync schedules inherit requirements and limitations of AOS NearSync, see "[Requirements of Data Protection with NearSync Replication](#)" in the *Data Protection and Recovery with Prism Element* guide.

- b. **Repeat [weekly|monthly]:** Select which days to run the schedule.
- If you select weekly, select the boxes for the days of the week the schedule should run.
 - If you select monthly, enter one or more integers (in a comma-separated list) to indicate which days in the month to run the schedule. For example, to run the schedule on the 1st, 10th, and 20th days, enter "1,10,20".
- c. **Start on:** Enter the start date and time in the indicated fields.
The default value is the current date and time. Enter a new date if you want to delay the schedule from starting immediately.
- d. **End on:** To specify an end date, check the box and then enter the end date and time in the indicated fields.
The schedule does not have an end date by default, and the schedule runs indefinitely unless you enter an end date.
- e. **Retention Policy:** Enter the number of snapshots to save locally and at the remote sites.
- Enter a number in the **Local** line "**keep the last ## snapshots**" field. The default is 1.
 - Enter the number of snapshots to save on the **Remote Site** in the "**keep the last ## snapshots**" field. This number can be different from the number that you have entered in the **Local** line. This replication is

an async replication. After the replication completes, the protection domain is going to be available in the **Async DR** tab of the remote site.

- The saved snapshots equal to the value entered in the **keep the last ## snapshots** field + 1. For example, if you entered 20 as the value for **keep the last ## snapshots** field, Files saves 21 snapshots. When Files takes the next (22nd) snapshot, Files deletes the oldest snapshot and replaces it with the new snapshot.

Note: If too many schedules have the same start time, replications can fall behind. To avoid this issue, stagger start times across schedules.

Activating Disaster Recovery

Recover a file server using protection domain (PD) based disaster recovery after a planned or unplanned migration.

About this task

To fail over and fail back a protection domain, follow these steps:

Procedure

1. Fail over the protection domain.

Note: The name for the automatically created protection domain contains NTNIX as a prefix followed by the file server name.

Unactivated protection domains display shaded indicators next to the protection domain name. Activated protection domains display green indicators next to the name. For more information, see [Failover and Failback Operations for Asynchronous and NearSync DR](#) in the *Data Protection and Recovery with Prism Element* guide.

2. Fail back a protection domain.

See the [Failing Back a Protection Domain](#) topic in the *Data Protection and Recovery with Prism Element* guide for steps.

What to do next

Activate the file server for planned or unplanned migration. For more information, see [Activating a File Server](#) on page 113.

Activating a File Server

Activate an inactive file server after Protection Domain–based Disaster Recovery.

About this task

To activate a file server, follow these steps:

Procedure

1. In the Prism Element **File Server** view, select the target file server.

When a file server is inactive, (Needs activation) appears next to the file server name and an **Activate** button appears in the action button list under the file server table.

2. Click **Activate**.

The **Activate File Server** window appears.

3. Complete the indicated fields in the **Client Network** and **Storage Network** tabs.

Some fields populate from the information provided during file server creation. See [Creating a File Server](#) on page 21 for more information about each field.

4. When all of the information is complete, click **Save**.

The file server configuration updates, and the (Needs activation) message disappears indicating the file server is now active.

Metro Availability on ESXi

Use Metro Availability as a synchronous disaster recovery solution for file servers on ESXi hypervisors.

Metro Availability is a disaster recovery solution that provides synchronous data replication to a remote site. Configuring Metro Availability for Files requires creating a Metro Availability protection domain for the file server. After configuring the protection domain, you can monitor and manage the configuration from the **Data Protection** tab in Prism Element.

For more information on Metro Availability, see the [Data Protection and Recovery with Prism Element Guide](#).

To deploy Metro Availability for Files, complete the following tasks:

- Configure the Metro Availability environment for Files. For more information, see [Configuring Metro Availability on ESXi](#) on page 115.
- Enable Metro Availability for Files. For more information, see [Enabling Metro Availability on ESXi](#) on page 116.

Requirements

Meet the following requirements before deploying Metro Availability for Files:

- Each file server must have its own protection domain.
- The minimum number of nodes on the remote site must be equal to or greater than the maximum number of FSVMs per file server on the source site.
- Meet the general ESXi Metro Availability requirements.

For more information, see [Conditions for Implementing Data Protection \(Metro Availability\)](#) in the *Data Protection and Recovery with Prism Element Guide*.

Limitations

The following limitations apply to Metro Availability with Nutanix Files:

- You cannot configure Metro Availability on a cloned file server or clone file servers in Metro Availability protection domains.
- Metro Availability uses the same set of IP addresses and networks for file server VMs on the target and the primary sites.
- Upgrade operations, adding and removing FSVMs, and renaming the file server temporarily suspend Metro Availability.
- Metro Availability only supports automatic failures and witness failure scenarios.
- If an ESXi host reboots while the protection domain is in a decoupled state, the host might be in a permanent all-paths-down (APD) state, which requires manually unmounting and remounting the datastore.

For information on general limitations of Metro Availability, see [Data Protection and Recovery with Prism Element Guide](#).

Failovers

Metro Availability for Files supports witness and automatic resume failure handling.

For unplanned failovers in the witness mode, FSVM migration and promotion of the protection domain occurs automatically. Unplanned failovers in the automatic resume mode require manually promoting the protection domain. For more information, see [Site Failure Handling of Active Metro Protection Domains \(Manual, Automatic Resume, or Witness Mode\)](#) in the *Data Protection and Recovery with Prism Element* guide.

Planned failovers in the witness mode or in the automatic resume mode require powering off FSVMs, manually migrating the FSVMs to the remote site, and promoting the protection domain on the recovery site. For more information, see [Failing Over a Protection Domain \(Planned Failover\)](#) in the *Data Protection and Recovery with Prism Element* guide.

Primary sites that are in a healthy state after a Metro Availability failover require manually re-enabling Metro Availability. For more information, see [Re-enabling Metro Availability Configuration After Failover](#).

Failbacks

Failbacks, irrespective of the failure handling mode, require manually migrating the FSVMs to the original site in vSphere and manually triggering the promote operation. For more information, see [Metro Availability on ESXi](#) in the *Data Protection and Recovery with Prism Element* guide.

Configuring Metro Availability on ESXi

Configure Metro Availability for Files on an ESXi hypervisor.

About this task

To configure Metro Availability, follow these steps:

Procedure

1. Make sure that the file server to protect is not part of a protection domain by checking the **File Server** table in Prism Element.
For more information, see [File Server View in Prism Element](#) on page 16.
2. Configure remote sites.
For more information, see [Configuring a Remote Site \(Physical Cluster\)](#) in the *Data Protection and Recovery with Prism Element* guide.
3. Configure the vSphere cluster in vCenter.
For information on navigating vCenter, see [VMware documentation](#).
 - a. Configure AOS hosts to vSphere.
For more information, see [vSphere Metro Storage Cluster for Nutanix Metro Availability](#) and [VMware vSphere Settings](#).
 - b. Ensure that the vSphere Cluster Services (vCLS) VMs are in the correct container.
For more information, see [vSphere Cluster Services on Nutanix Clusters](#).
4. Register the witness VM on the primary and the recovery sites.
For more information, see [Installing a Witness VM](#) in the *Data Protection and Recovery with Prism Element* guide.
The witness VM appears on both sites.
5. Create a remote site container with the same name as the source site container.
For more information, see [Creating a Storage Container](#) in the *Prism Web Console Guide*.

6. Enable Metro Availability on the file server.

For more information, see [Enabling Metro Availability on ESXi](#) on page 116.

What to do next

Manage your Metro Availability deployment in the **Protection** view of Prism Element.

Enabling Metro Availability on ESXi

Enable Metro Availability for Files on an ESXi hypervisor.

About this task

You must configure all of the required entities before enabling Metro Availability.

For information on required configurations, see [Metro Availability on ESXi](#) on page 114.

Procedure

1. In the **File Server** view in Prism Element, click **Protect**.
2. Under **File Server Protection Domain**, click **Create a Metro Protection Domain**.
3. Do the following in the indicated fields:
 - **Protection Domain Name:** Files automatically suggests the file server name as the name for the protection domain. Keep the name as is or provide another name.
 - **Failure Handling Mode:** Choose either **Witness** or **Automatic**. For more information on both options, see [Metro Availability Witness Option](#).
 - **Remote Site:** Select the IP address of the remote site.
4. Click **Create and Enable Protection Domain**.

What to do next

Before failback, ensure that Metro is in the standby-active mode, with the primary site in standby.

Smart Tiering

Tier data using Smart Tiering.

You can free up space on your file server by tiering data to an object store. There are two ways to perform tiering: manage and define policies through the Files Console using Standard Tiering or through [Nutanix Data Lens](#) using Advanced Tiering.

Standard Smart Tiering Overview

Standard Smart Tiering manages hot, warm, and cold file server data.

Use Smart Tiering to maximize available file server space by moving cold data from the file server to an object store. Tiered storage does not contain the full data set. The full metadata and pointers to tiered data remain on the primary storage.

Nutanix supports using Nutanix Objects, Microsoft Azure, AWS Standard, AWS IA, or Wasabi (S3-compatible storage) for object storage, which you must configure before setting up a tiering profile. You can create a tiering profile with the AWS Glacier instant retrieval storage class in addition to Object Store, AWS, and S3-compliant storage classes.

Follow these steps to enable tiering on the file server:

1. Review the [Limitations](#) section.

2. Review the [Smart Tiering Requirements and Recommendations](#) on page 117 section.
3. Configure the tiering location and profile, see [Configuring a Tiering Location](#) on page 119.
4. Configure the tiering capacity threshold, see [Setting a Capacity Threshold](#) on page 121.
5. Configure a tiering policy, see [Creating a Tiering Policy](#) on page 122.

Limitations

- You can add a migration target share to a tiering profile only after the migration plan completes.
- You can have only one tiering profile and tiering location for each file server.
- You cannot evict tiered file data from the primary storage until all snapshots containing the file expire.
- You cannot tier files that are smaller than 64 KiB or larger than 5 TiB.
- You cannot write to tiered (offline) files.
- You can only configure tiering at the file-server level.
- After re-enabling versioning, tiering operations need a minimum of 10 minutes to resume.
- Files cannot compress data while tiering data to the object store.
- Files does not tier sparse files.
- Standard Smart Tiering is not supported when using asynchronous replication with Smart DR to replicate files on the public cloud.
- None of the object stores support WORM-enabled buckets.
- As part of Wasabi, garbage collection of the objects does not work.

Terminology

- *Capacity threshold* – The percentage of file server capacity you must reach to trigger tiering operations.
- *Object store* – The secondary storage service.
- *Tiering* – Migrating cold data from primary to secondary storage.
- *Tiering profile* – The secondary storage location configuration, which includes the object store, security keys, and the storage bucket.
- *Tiering policy* – Defines shares to be included or excluded from tiering by file size, and cool off period.
- *Tiering schedule* – Specifies times for tiering data.
- *Tiered data in snapshot* – Data from a tiered file that exists in a snapshot and remains on premises until you delete the snapshot.
- *Data retention* – After file deletion from a file server, the corresponding object is deleted only after the specified retention period.

Smart Tiering Requirements and Recommendations

Meet the indicated requirements to configure and administer Smart Tiering.

General Requirements

Meet the following general requirements regardless of object store type:

- Enable versioning on the object store bucket.
- Disable object lock on the object store bucket.

Note: For Advanced Smart Tiering with Data Lens, refer to the [Data Lens User Guide](#) for requirements and administrative guidance.

AWS Simple Storage Service (S3) Requirements

Ensure that you have the following AWS IAM user permissions:

- s3:GetBucketObjectLockConfiguration
- s3:PutObject
- s3:GetObject
- s3:GetLifecycleConfiguration
- s3:AbortMultipartUpload
- s3:PutLifecycleConfiguration
- s3:CreateBucket
- s3:GetObjectTagging
- s3:ListBucket
- s3:PutObjectTagging
- s3:GetBucketVersioning
- s3:GetObjectVersion
- s3:DeleteObject

Ensure that the security keys have the following permissions on the S3 bucket:

- Set LifeCycle policy
- Read and write
- Set tags on object
- List and write on object

Nutanix Objects Object Store Requirements

The following requirements apply when using Nutanix Objects as the object store:

- Use the bucket owner secret and access key for tiering.
- Ensure that the owner has read and write permissions.

Security Recommendations

Unless otherwise noted, the following recommendations apply to all object store types. Nutanix recommends using the following security best practice:

- Implement security best practices for AWS S3 as described on the AWS documentation portal.
- Do not use the same bucket for storing non-Nutanix Files data.
- Do not manually delete or overwrite any object in the bucket as this action results in data loss.
- Enable encryption on the bucket.
- Have the permissions to query the bucket location.

- Do the following when changing security keys:
 - Ensure that both keys are valid when changing from key1 to key2.
 - Retire key1 from s3 only after updating key2 on the file server.
- On AWS, Files sets the life-cycle rule on the bucket to garbage collect deleted objects. Do not overwrite or remove the rule from S3. Do not apply another life-cycle rule to the bucket.
- For Nutanix Objects, provide security keys of the owner of the bucket.

Configuring a Tiering Location

To tier data, configure a secondary storage object store.

About this task

A tiering profile includes the tiering location and defines the secondary storage target, which is the object store.

To configure a tiering location, follow these steps:

Procedure

1. In the Files Console, go to **Data Management > Smart Tiering > Standard**.
For more information, see [Nutanix Files Console](#) on page 18.
2. Click **Enable and Configure Tiering**.
3. In the **Tiering Configuration** pane, under **Tiering Location**, click **Add Tiering Location**.

4. Update the fields as indicated:

- **Name:** Enter a name for the tiering profile.
- **Store type:** Select an object store type.
- **Service host:** Enter a URL of the tiering service host.

Supported URL formats for S3:

- `https://s3.<region>.amazonaws.com`, for example: `s3.us-east-1.amazonaws.com`
- <https://s3.amazonaws.com>, here Files tries to get the location constraint by querying the server.

Note: Do not specify the virtually hosted-style bucket URL. For **service host**, only specify the service host name or s3 endpoint (for example, `s3.<region>.amazonaws.com` or `s3.amazonaws.com`). Specify the bucket name in the **bucket name** field.

- **Access key:** Enter the S3 access key ID for the object store.
- **Secret key:** Enter the secret access key for accessing the object store.
- **Bucket configuration**

- **Bucket name:** Enter a name of the bucket in the object store.

Note: You cannot change the bucket name.

- **Retention period:** Enter a period for data retention after deletion from the file server.

Note: The default retention period is 5 years and the minimum is 6 months.

- **Certificate**

- **Certificate:** Upload a secure sockets layer (SSL) certificate file.
- **Validate certificate:** Check this box to enforce SSL certificate verification during tiering operations after profile creation.

Configuring a Tiering Location (Azure)

To tier data, configure a secondary storage object store.

About this task

A tiering profile includes the tiering location and defines the secondary storage target, which is the object store.

To configure a tiering location, follow these steps:

Procedure

1. In the Files Console, go to **Data Management > Tiering > Standard**.
For more information, see [Nutanix Files Console](#) on page 18.
2. Click **Enable and Configure Tiering**.
3. In the **Tiering Configuration** pane, under **Tiering Location**, click **Add Tiering Location**.

4. Update the fields as indicated:

- **Name:** Enter a name for the tiering profile.
- **Store type:** Select **Azure** or **Azure Cool**.

Note: Once configured, you cannot modify the store type.

- **Service host:** Enter a URL of the tiering service host.

• For example: `blob.core.windows.net`

Note: Once configured, you cannot modify the service host.

- **Storage account name:** Enter the name of the Azure storage account associated with the account key.
- **Account key:** Enter the account access key for accessing your Azure Blob Storage account.
- **Container configuration**

- **Container name:** Enter a name of the Azure Blob Storage container.

Note: You cannot change the container name.

- **Retention period:** Enter a period for data retention after deletion from the file server.

Note: The default retention period is 5 years and the minimum is 6 months.

- **Generate JSON:** Use this option to manually configure the life-cycle policy in the Azure web console.
- **Certificate**
 - **Certificate:** Upload a secure sockets layer (SSL) certificate file.
 - **Validate certificate:** Check this box to enforce SSL certificate verification during tiering operations after profile creation.

Setting a Capacity Threshold

The capacity threshold specifies when to tier data.

About this task

To set a capacity threshold, follow these steps:

Procedure

1. In the Files Console, go to **Data Management > Smart Tiering**.
2. In the **Tiering Configuration** pane, click **Set a Capacity Threshold**.
3. To tier data based on a threshold of occupied capacity, check the **Set a Capacity Threshold** box.
 - a. In the **% of Primary Storage** field, enter an integer that indicates the capacity percentage to trigger tiering (minimum 1%).

4. Under **When to tier**, select one of the following options:

- » To trigger tiering manually after crossing the capacity threshold, select the **Manual** option.
- » To configure a schedule for tiering data, select **Every Day** option.

Note: If you choose the **Manual** option, follow the steps in [Tiering Data Manually](#) on page 123 to initiate tiering after you observe that the file server reaches the configured capacity threshold.

5. If you chose the **Every Day** option, do the following in the indicated fields (otherwise, continue to the next step).

Note: The actual tiering operation on the file server can begin and end shortly after the scheduled time.

- **Frequency:** Use the **Every Day** option and do the following (or click **set a custom window for each day**):
 - **Start time:** Specify when to start tiering operations.
 - **End time:** Specify when to end tiering operations.
- (Optional) to customize the tiering schedule, click **set a custom window for each day**, and do the following:
 - Select the desired tiering days.
 - For each day, specify the schedule for tiering operations in the **start time** and **end time** fields.

6. Click **Save**.

Creating a Tiering Policy

Add a tiering policy to specify when to tier cold data to the object storage.

About this task

To create a tiering policy, follow these steps:

Procedure

1. In the Files Console, go to **Data Management > Smart Tiering**.
For more information, see [Nutanix Files Console](#) on page 18.
2. In the **Tiering Configuration** pane, under **Tiering Policy**, click **Create a Tiering Policy**.
3. Under **Filters for Data to be Tiered**, do the following:
 - **Access Time:**
 - Enter an integer to specify the data age.
 - In the dropdown, select the unit of measurement (days, weeks, months, years).
 - **File Size (greater than):**
 - Enter an integer for the minimum tiering file size (minimum 64 KB).
 - In the dropdown, select the unit of measurement (KB, MB, GB).
4. (Optional) To specify shares excluded from tiering, check the **Exclude Shares** box under **Shares**.

5. Click **Save**.

If you set a schedule for the capacity threshold, file server automatically triggers the tiering operation in Nutanix Files according to the policy and threshold configuration. If you chose to manually trigger tiering, you must trigger the operation after the file server meets the criteria in the threshold and policy configuration.

Tiering Data Manually

If you have not selected to tier data on a schedule, initiate tiering manually.

About this task

See [Setting a Capacity Threshold](#) on page 121 for details on configuring a manual tiering policy.

To tier data manually, follow these steps:

Procedure

1. In the Files Console, go to **Data Management > Smart Tiering**.

For more information, see [Nutanix Files Console](#) on page 18.

2. Click **Tier Data**.

An unavailable **Tier Data** button indicates that file server has not found enough tiering candidates on the file server or that the file server has not reached the capacity threshold.

A dialogue window asks you to confirm the action.

Updating a Tiering Configuration

Edit an existing tiering configuration.

About this task

To update a tiering configuration, follow these steps:

Procedure

1. In the Files Console, go to the **Dashboard** page.

For more information, see [Nutanix Files Console](#) on page 18.

2. In the **Data Management** pane, click **Smart Tiering**.

3. In the **Tiering Configuration** section, do one or more of the following:

- (Optional) To update the tiering profile configuration, click **Edit** under **Tiering Location**. Update one or more fields as described in [Configuring a Tiering Location](#) on page 119. You can only update the access key, secret key, profile name, or the certificate.
- (Optional) To update the capacity threshold, click **Edit** under **Capacity Threshold**. Update one or more fields as described in [Setting a Capacity Threshold](#) on page 121.
- (Optional) To update the tiering policy, click **Edit** under **Tiering Policy**. Update one or more fields as described in [Creating a Tiering Policy](#) on page 122

Recalling Tiered Data

Recall tiered data from an object store to the primary storage.

About this task

To manually recall data, follow these steps:

Procedure

Use the corresponding command for the data type you wish to recall.

- » Recall a tiered file.

```
nutanix@fsvm$ afs tiering.recall share_name file_paths=file_path
```

Replace *file_path* with a comma separated list of target file paths starting from the share root, and replace *share_name* with the name of the share.

Note: You can use a maximum of 32 file paths a time.

- » Recall files using inodes.

```
nutanix@fsvm$ afs tiering.recall share_name inodes=inodes
```

Replace *inodes* with a comma separated list of target inodes, and replace *share_name* with the name of the share.

- » Recall all files on a share.

```
nutanix@fsvm$ afs mscan.create_scan share_list=share_name receiver_type=RecallFiles
```

Replace *share_name* with a comma separated list of share names.

Self-Service Restore

Self-service restore (SSR) lets you restore a previous version of a file or share.

With SSR, Files takes snapshots of share data in intervals defined by a snapshot schedule policy or on-demand. You can enable SSR during or after share creation. The feature is disabled by default.

Files takes snapshots of the stored cluster data at the share and file levels. Files allows client-side restoration with file-level granularity and server-side restoration with share-level granularity. Server-side restoration is faster and takes approximately the same time regardless of the share size or the number of files.

Use the client to view and manage file-level snapshot data. You don't need administrator level permissions for file-level snapshot data retrieval.

Nutanix recommends scheduling snapshots in regular or frequent intervals to provide same-day protection against accidental deletions or corruption.

You can configure minute-based (every 15 minutes), hourly (every hour), daily, weekly, and monthly snapshots schedules. By default, SSR takes a snapshot every hour, retains the most recent 24 snapshots, and deletes the oldest SSR snapshot after exceeding the retention count for the snapshot type.

The snapshot retention count corresponds to the retention period. For example, for daily snapshots, Files deletes the oldest snapshot and creates a new one every day.

Files supports share updates for both standard and distributed shares. To enable SSR during share creation, see [Creating an SMB Share](#) on page 46 or [Creating an NFS Export](#) on page 50. To enable SSR after share creation, see [Enabling Self-Service Restore](#) on page 125.

Limitations

Consider the following limitations before enabling SSR:

- SSR for SMB does not restore streams or attributes in directories.
- Files does not support SSR at the root of distributed shares or exports.
- Files has a combined limit of 400 snapshots per policy for all schedule types.

- Share-level restore limitations:
 - Restoring a share cancels metadata scan jobs.
 - Restoring a share is unsuccessful when a snapshot is in the database but not in the file system, which can occur during service disruptions or deletion of snapshots.
 - Share restore rollback doesn't re-enable paused features, for example Smart Tiering.
 - An alert might not appear for surpassing the maximum number of TLDs after share-restore SSR.
 - You can have a maximum of 2,000 sub-mounted shares under the restored share.

Share-Level Restore

You can restore share-level snapshots using the Files Console, see [Restoring a Share](#) on page 128.

Restoring a share takes the same amount of time irrespective of the amount of data. The share restore feature only restores the share data and not the complete metadata, like the state of share features and share permissions. For NFS shares, existing client mounts encounter a stale mount error and need to remount. After restoring a share, Files automatically re-enables snapshots.

See information on share-level limitations below.

Table 4: Feature Parity

Feature	Parity
Share migration	You cannot restore shares on an incomplete or active migration plan.
Smart Sync	You cannot restore shares that are part of a Smart Sync policy.
Smart DR	You cannot restore shares that are part of a Smart DR policy.
Tiering	In the event of unresponsive non-critical tiering operations, you will get an alert.

File-Level Restore

You can view removed or overwritten files and choose a snapshot from the history of the target file. For file-level snapshots, the primary protocol determines the method for restoring the data. For SMB, you can restore files through the system. For NFS, Files only provides read-access and you must manually restore the files by copying and pasting the file data. For more information, see [Retrieving Files with Self-Service Restore](#) on page 129.

Enabling Self-Service Restore

Enable self-service restore (SSR) on a share or an export.

About this task

To enable SSR, follow these steps:

Procedure

1. Go to the **Shares & Exports** view in the Files Console.
For more information, see [Nutanix Files Console](#) on page 18.

2. In the row for the target share, click the **three dot menu** > **Edit**.
The **Update Share & Export** window appears.
3. Click **Next**.
4. In the **General Settings** tab, select the **Enable Self Service Restore** checkbox.
5. In the **Snapshot Policy** dropdown menu, select the SSR policy.
For information on configuring an SSR snapshot policy, see [Configuring Snapshot Schedule Policies](#) on page 126.
6. Click **Next** > **Save**.

What to do next

Add a snapshot schedule for SSR. For more information, see [Configuring Snapshot Schedule Policies](#) on page 126.

Configuring Snapshot Schedule Policies

Configure file server protection by adding a self-service restore (SSR) snapshot policy.

Before you begin

Ensure that SSR is enabled on the target shares. For more information, see [Enabling Self-Service Restore](#) on page 125.

About this task

Before you configure an SSR snapshot policy, Nutanix Files assigns the default SSR policy to shares and exports that have SSR enabled.

After you configure additional SSR policies, you can assign any SSR policy during or after share creation. For more information, see [Modifying a Share or Export](#) on page 63.

Nutanix Files takes daily, weekly, and monthly snapshots at midnight (00:00 UTC) and at zero minutes for hourly snapshots. With the exception of hourly snapshots, you cannot change the time when Files takes a snapshot. However, you can change the snapshot frequency for any snapshot type.

To configure snapshot schedules, follow these steps:

Procedure

1. In the **Files Console**, go to **Data Management** > **Protection** > **Self Service Restore**.
2. In the **Snapshot Policies** section, click **+ Add New Policy**.
The **Add New Policy** window appears.
3. Under **SSR Policy Name**, enter a name for the SSR policy.
4. Click **+ Add New Schedule**.

5. Complete the following fields to configure a schedule:

- a. For **Type**, select **minutes**, **hourly**, **daily**, **weekly**, or **monthly** snapshot interval .

Note: You can only have one schedule type per file server. For example, a single file server cannot have two hourly snapshot schedules.

- b. For **Frequency**, enter an integer to specify the frequency of snapshots for minute-based, hourly, and daily schedules. for weekly schedules, specify the day of the week, and for monthly schedules, specify the date. By default, the schedule does not have an end date and runs indefinitely.
- c. For **Snapshot Retention**, enter an integer to specify the number of snapshots to retain for the policy.

6. Click the [blue checkmark icon](#).

7. (Optional) To add additional schedules to the same policy, repeat the procedure beginning with Step 4.

8. Click **Save**.

Modifying a Snapshot Policy

Modify or delete an existing self-service restore (SSR) policy.

About this task

To modify or delete an existing SSR policy, follow these steps:

Procedure

1. In the **Files Console**, go to **Data Management > Protection > Self Service Restore**.

2. In the **Snapshot Policies**, click the **Three dot menu > Edit**.

3. Update the policy configuration:

- a. (Optional) To update the policy name, in **SSR Policy Name** field, enter a new name.
- b. (Optional) To add a schedule to the policy, click **+ Add New Schedule**.
For information on each field, see [Configuring Snapshot Schedule Policies](#) on page 126.
- c. (Optional) To delete an existing schedule, in the **Actions** column, click the [remove](#) icon next to the schedule.

Important: Deleting a schedule deletes all of the snapshots for that schedule type. For example, deleting an hourly schedule deletes all hourly snapshots.

- d. (Optional) To edit an existing snapshot schedule, in the **Actions** column, click the [edit](#) icon next to the schedule.
For information on each field, see [Configuring Snapshot Schedule Policies](#) on page 126.
- e. (Optional) To delete a policy, in the **Actions** column, click the [remove](#) icon next to each schedule.

4. Click **Update**.

The updated policy applies to all shares under the policy.

Taking On-Demand Snapshots

Take a self-service restore snapshot at any time.

About this task

Self-service restore must be enabled on a share or export before taking an on-demand snapshot.

To take an on-demand snapshot, follow these steps:

Procedure

1. Go to the **Shares & Exports** view in the Files Console.
For more information, see [Nutanix Files Console](#) on page 18.
2. In the row for the target share, click **three dot menu** > **Take Snapshot**.
The **Take Snapshot** dialog box appears.
3. In the **Snapshot Name** field, enter a name for the snapshot.
4. Select the radio button for one of the following expiry options:
 - » No Expiration
 - » Expire on a specific dateIn the **Expiry Date** field, select an expiry date.
5. Click **Take Snapshot**.

Restoring a Share

Restore a share or an export from a snapshot.

About this task

When you restore a share from a snapshot, Files deletes snapshots taken after the target snapshot. Disable the following features before you attempt to restore a share: Smart-DR, Smart Sync, share conversion, and share migration.

About this task

To restore a share, follow these steps:

Procedure

1. In the Files Console, go to the **Shares & Exports** tab.
2. In the row for the target share, click the **three dot menu** > **Restore Share**.
The **Restore Share** window appears.
3. Select the target snapshot from the table.

Important: Restoring a share is irreversible and deletes all newer snapshots.

4. Click **Restore**.
The restore task begins and Files pauses new snapshots on the target until the operation ends. If the restore task completes but the post-restore task fails, use the following command `afs share.restore_commit share-name`.

What to do next

If you disabled any of the following features, re-enable them: Smart-DR, Smart Sync, share conversion, share migration.

Manually Restoring a Share

Manually restore a share from a Nutanix Files or third-party snapshot.

About this task

To manually restore a share from a snapshot, do the following:

Procedure

- Restore the share:

```
$ afs share.restore_from_snapshot share-name snapshot-name
```

Replace *share-name* with the name of the share and replace *snapshot-name* with the name of the snapshot.

Retrieving Files with Self-Service Restore

Retrieve SMB file snapshots using self-service restore (SSR).

About this task

Accessing file snapshots requires browsing the target directory on the client.

To retrieve SMB files, follow these steps:

Procedure

1. On the client, go to the target directory and look for **Previous Versions** as described in Microsoft documentation.

Previous versions of files typically display in order of the date modified. Therefore, unmodified files do not appear. Previous versions of folders display every available snapshot.

To see deleted files or directories, access a previous version of the parent folder and restore the file or directory.

2. Open and manage the previous versions according to the Microsoft documentation.

Retrieving NFS Files with SSR

About this task

Enabling SSR on NFS exports lets you access snapshot versions of a directory using a hidden `.snapshot` sub-directory. Browse the sub-directory for the intended snapshot and restore it by manually copying the file or its content.

To retrieve an NFS file with SSR, follow these steps:

Procedure

1. Go to the snapshots of the target directory.

```
$ cd /dir1/dir2/.snapshot
```

Replace */dir1/dir2* with the directory path.

You can only access the `.snapshot` directory using the `cd` command. The `.snapshot` directory is not visible otherwise. Using the `ls -a` command does not show the `.snapshot` directory. For distributed NFS exports, the `.snapshot` sub-directory is in top-level directory (TLD).

Note: When the absolute path to the snapshot directory is longer than 3922 characters, the attempt to browse into the snapshot can fail with a file name too long error message.

2. List the snapshots for the target directory.

```
$ ls
```

3. Browse the snapshots and copy the desired data.

Deleting Self-Service Restore Snapshots

Manually delete self-service restore (SSR) snapshots to reclaim disk space.

About this task

By default, Files deletes the oldest SSR snapshot after exceeding the retention count for that snapshot type.

To delete an SSR snapshot, follow these steps:

Procedure

1. List the snapshots by creation time:

```
nutanix@fsvm$ afs snapshot.list share_name=share_name
```

2. List the space occupied by a single snapshot or range:

```
nutanix@fsvm$ afs snapshot.reclaimable_space uuid_start:uuid_end
```

3. Remove snapshots:

- » Remove a single snapshot using the universally unique identifier (UUID):

```
nutanix@fsvm$ afs snapshot.remove share_name=share_name  
snapshot_uuid_list=snapshot_uuid
```

Note: You can delete multiple snapshots by specifying a comma-separated list of UUIDs.

- » Remove multiple snapshots using labels by defining the label:

```
nutanix@fsvm$ afs snapshot.remove share_name=share_name label=label
```

Labels define if the snapshot schedule is hourly, daily, weekly, or monthly. Deleting snapshots using labels deletes all snapshots for the specified label on the share.

Setting Custom Hourly Snapshot Times

Set a custom hour for self-service restore (SSR) snapshots.

About this task

To take hourly snapshots at a custom time, indicate the time by specifying the number of hours offset from midnight UTC.

Procedure

Set snapshot times:

```
nutanix@fsvm$ afs snapshot.set_ssr_hourly_offset 1 - 23
```

Replace `1 - 23` with an integer to specify the time offset from UTC. For example, if the offset is 9 and the frequency is 12 hours, Files takes hourly snapshots at 9 AM and 9 PM UTC, compared to the default of 12 AM and 12 PM UTC.

Share Migration

Share migration lets you copy data from a specified path on one share to another, including from third-party sources to Nutanix Files.

Note: Files fully supports the migration of live shares. Nutanix recommends running the final or cut-over iteration on a static share.

Each time you initiate share migration, Files starts an iteration of the migration copy process. If some files fail to migrate, you can copy the failed files to the target using the migration retry procedure, see [Re-attempt Share Migration](#) on page 134. Files trims the forward slash (/) from the beginning and the end of the source path and migrates the content inside the path.

Note: Once the number of failed files exceeds a threshold, the migration retry procedure cannot succeed, and you must run another iteration of the migration.

If the source share is not static and the data changes either during or after the migration, run another iteration to copy the changed data from the source to the target. If you run an iteration when the source and target shares are in sync, then Files shows zero copied entities (files, folders, streams). To ensure the migration of all data, run multiple iterations.

For more information about restoring target share access upon cancellation of a migration plan, see [Troubleshooting](#) on page 143.

Note: You can also migrate shares using Move. Move provides a user interface for file migration. For more information, see [Files Migration](#) in the *Move User Guide*.

Supported Features

Share migration supports the following.

- Migrating and retaining symlinks to the target shares.
- Source shares with SMB encryption enabled.
- A mirroring option for the data between the source and the target. Files replicates any renames and deletes on the source to the target.

NFSv3 Export Migration

Meet the following requirements:

- Configure AUTH_SYS security on the source. On Linux-based servers, you can do this using the `sec=sys` export option.
- Disable root squashing on the server to allow the Files migrator to read and transfer all files. On Linux-based servers, use the `no_root_squash` export option.

Limitations

Files does not do the following for share migration:

- Migrate system access control lists (SACLs).
- Migrate directory junctions.
- Support multi-protocol on shares that already have long-name enabled.
- Support custom extended attributes (EA) on SMB shares. Therefore, all custom extended attributes present on the source files are ignored during migration.
- Support deduplication. You must disable deduplication during share migration.
- Migrate hard-link files. As a workaround, retry the iteration twice.

Migrating Shares

Migrate share data from another share.

About this task

To migrate shares, follow these steps:

Procedure

1. Provide source file server details.

```
nutanix@fsvm$ afs migration.source_add source_alias=source_server_alias /  
source_fs_fqdn="source_server_IP_address_or_FQDN" source_user='domain\\user_name'
```

Do not add a destination file size. For NFS migration, `source_user` is not required.

For example:

```
nutanix@fsvm$ afs migration.source_add source_fs_fqdn="filer1.child4.afs.minerva.com"  
source_alias=filer1 source_user='CHILD4\\backup_user'  
Password for user:  
Registering source (filer1: filer1.child4.afs.minerva.com).  
Source created.
```

Note: The file server must be able to access the source user domain.

Files asks you to provide the source user password.

2. If you have not already, create a target share using the Files user interface.
3. Check that Files registered the source.

```
nutanix@fsvm$ afs migration.source_list
```

For example:

```
nutanix@fsvm$ afs migration.source_list  
+-----+-----+-----+-----+  
+-----+  
| SOURCE ALIAS | SOURCE FQDN | SOURCE USER | START TIME |  
END TIME |  
+-----+-----+-----+-----+  
+-----+  
| filer1 | filer1.child4.afs.minerva.com | CHILD4\\backup_user | 0 | 0  
|  
+-----+-----+-----+-----+  
+-----+
```

4. Create the migration plan.

```
nutanix@fsvm$ afs  
migration.plan_create name=plan_name source_alias=source_server_alias /  
source_share=share/path target_share=target_share target_share_path=path/within/share  
/  
plan_type=NFS|SMB  
migrate=true|false
```

```
mirror=true|false
```

Setting the mirror option to *true* honors rename and delete operations. For consolidation to the same target share, set the mirror option to *false*.

Setting migrate option to *true* initiates the migration as soon as you create the plan.

Example 1: Migrate the source share to the target share.

```
nutanix@fsvm$ afs migration.plan_create name=plan1 source_alias=filer1
source_share=share1 target_share=home plan_type=NFS migrate=false mirror=true
```

Example 2: Migrate a path on the source share to a target share.

```
nutanix@fsvm$ afs migration.plan_create name=plan1 source_alias=filer1
source_share=share1/users target_share=home plan_type=NFS migrate=false mirror=true
```

Example 3: Migrate a path on the source share to a path on the target share.

```
nutanix@fsvm$ afs migration.plan_create name=plan1 source_alias=filer1
source_share=share1/users/user1 target_share=home target_share_path=user1
plan_type=SMB migrate=true mirror=true
```

Note:

- Share migration requires target_share_path when the target is a directory within the target share.

(optional) add more arguments.

```
json_output=          Render output in json format
```

```
total_bytes=          The total number of Bytes to be migrated.
```

```
total_files=          The total number of files to be migrated.
```

Tip: Setting the total_bytes or total_files arguments lets Files display the progress of the migration.

5. Start the migration.

```
nutanix@fsvm$ afs migration.plan_migrate name=plan_name
```

For example:

```
nutanix@fsvm$ afs migration.plan_migrate name=plan1
```

Note: Use the allow_unknown_sids= parameter if you have files or directories with local security identifiers (SIDs) on the source.

```
allow_unknown_sids=    Copy entities with an unknown owner or group SID.
                        When set to false, Files does not copy entities with
an                      unknown user or group SID.
```

```
                        false, Files fails the migration when the root of
the                      source share path has files.(For distributed shares
                        only.
```

```
ignore_long_tld_names= Ignore TLD names with more than 255 bytes on source.
If                      set to false, Files fails migration when the name of
a
```

255

TLD at the root of the source share is longer than bytes long. (For distributed shares only.)

Note: You can rerun the migration multiple times after completion.

6. Check the status of the migration.

Setting verbose to true provides extra migration status details for every FSVM on the file server.

```
nutanix@fsvm$ afs migration.plan_status name=plan_name verbose=true|false
```

Note: Checking the status of a migration plan returns a granular status view of all nodes on the file server. For any node excluded from the migration plan, Files reports a succeeded status.

7. (optional) Use one or more of the following commands as needed.

migration.plan_summary	Display a summary of each iteration of the plan.
migration.plan_pause	Pause the migration.
migration.plan_resume	Resume the migration.
migration.plan_cancel	Cancel the migration.
migration.source_update	Update the source.
migration.source_delete	Delete the source.
migration.source_list	List configured sources.
migration.plan_update	Update the migration plan.

Re-attempt Share Migration

Retry share migration for files that failed to migrate.

About this task

Before retrying share migration for failed files, Nutanix recommends running the share migration operation multiple times to ensure the migration of all source data.

About this task

To re-attempt share migration, follow these steps:

Procedure

1. Get a list of files that failed to migrate in the last iteration of the plan.

```
nutanix@fsvm$ afs migration.plan_get_failed_files_list name=plan_name /  
out_file=output_directory_path [iter_index=-1]
```

2. Retry the migration for the files that failed to migrate in the last iteration of the plan.

```
nutanix@fsvm$ afs migration.plan_retry_failed_files name=plan_name
```

Configuring Backup for Distributed Shares

Specify a backup server for distributed shares and exports.

About this task

Third-party solutions do not automatically back up distributed shares. Follow the steps as indicated to update the backup configuration syntax and allow the backup of distributed shares and exports.

Procedure

1. Check the global backup configuration.

```
nutanix@fsvm$ afs smb.get_conf "backup hosts"
```

2. Update the backup configuration.

- » Add a backup server by replacing *backup-server-ip* with the IP address for the backup server. For multiple backup servers, enter a comma separated list of the *backup-server-ip*.

```
nutanix@fsvm$ afs smb.set_conf "backup hosts" backup-server-ip section=global
```

- » Remove a backup server.

```
nutanix@fsvm$ afs smb.del_conf "backup hosts" section=global
```

3. Validate changes to the backup configuration.

```
nutanix@fsvm$ afs smb.get_conf "backup hosts"
```

Files Data Collection

Files supports data collection with Pulse.

The feature known as Pulse collects Files diagnostic system data and sends it to Nutanix Support. After you enable Pulse, Files synchronizes Pulse configurations from the Controller VM to the file server VMs. Synchronized data includes the Pulse enablement status and the mechanism chosen for streaming data to Nutanix Support. If Pulse data cannot reach Nutanix Support, administrators receive alerts in Prism Element. To enable or disable Pulse, see ["Configuring Pulse"](#) in the *Prism Web Console Guide*.

High Availability

High Availability (HA) is the failover mechanism for file server VMs (FSVMs).

HA for Files ensures that an FSVM on clusters of two or more FSVMs can fail over to another FSVM during a service disruption. All clusters of two or more FSVMs have HA enabled by default.

When an FSVM experiences an issue causing a failover, Files reassigns the IP address of the FSVM to another FSVM in the cluster. The IP address of the out-of-service FSVM remains available. However, the shares and exports on the impacted FSVM remain unavailable for several minutes. For NFS, it takes about 30 seconds to perform read operations again and about 120 seconds to gain write access due to the NFS grace period. For SMB, the failover process typically takes between 30 and 180 seconds.

Affinity rules do not affect HA, so multiple FSVMs can share a single host during an HA event.

Encryption

Nutanix Files supports AOS software encryption and in-flight message encryption for SMB3 shares.

You can apply AOS software encryption to Files by activating it through Prism Element. Ensure that you are running a compatible version of AOS. For more information, see [Configuring Data-At-Rest Encryption \(Software Only\)](#) in the *AOS Security Guide* and AOS feature support in the [Nutanix Files Release Notes](#).

Both NFS and SMB traffic is unencrypted by default. To encrypt communication between SMB clients and Files, enable SMB3 message encryption on all SMB shares. For more information, see [Modifying a Share or Export](#) on page 63. After enabling message encryption, Files encrypts messages on the file server side and decrypts them on the client side (only on new connections for the share). Clients that do not support encryption (Linux, Mac, Windows 7) cannot access a share with encryption enabled.

You can encrypt NFS traffic by configuring authentication. To encrypt all communication between NFS clients and Files, disable NFSv3 and only enable NFSv4, configure Active Directory for NFS, and enable Kerberos 5p authentication for all NFS exports. For more information, see [Updating File Server Directory Services](#) on page 74 and [Modifying a Share or Export](#) on page 63.

NUTANIX FILES REST APIS

Access the Files REST APIs using the Files REST API explorer.

The Files REST API Explorer offers developers tools to customize the Files experience using Files v4 REST APIs. You can access the Files REST API Explorer through the Files Console or through an FSVM. The Files service v4 APIs are independent from Prism Element (PE) and Prism Central (PC) APIs. However, the platform APIs, for operations such as create, clone, update or delete a file server, remain in Prism Element (PE).

To access the Files API Explorer in the Files Console (see [Nutanix Files Console](#) on page 18), go to **REST API Explorer**.

Using the Files API Explorer

Use the REST API Explorer for Nutanix Files.

Before you begin

Make sure you have added a REST API user and authorized them, see [Managing REST API Roles](#) on page 84.

About this task

To access the Files V4 APIs, see the [API and SDK Reference](#) on the Nutanix Developer Portal.

Note: Nutanix Files V4 APIs are only available for Prism Central.

To access the REST API explorer, follow these steps:

Procedure

1. In the [Nutanix Files Console](#) on page 18, go to **Admin > REST API explorer**.
2. From the left navigation pane, select the target API.
You can also find APIs using the search bar.
3. Click **Try it out**.
List of parameters, their descriptions, and response types are listed.
4. Enter the required parameters and response type values.
5. Click **Execute**.
The REST API Explorer lists the responses.

SECURITY HARDENING

Use the Nutanix command-line interface (nCLI) or Files CLI to customize your Files security configuration.

Table 5: General Files Hardening Configuration

Description	Command or Settings
Support file server configuration of the SCMA policy.	<pre>ncli file-server get-security-config fs-name=file-server-name</pre>
Schedule weekly execution of advanced intrusion detection environment (AIDE)	<pre>ncli file-server edit-security-params fs-name=file-server-name enable-aide=true</pre>
Enable the strong password policy.	<pre>ncli file-server edit-security-params fs-name=file-server-name enable-high-strength-password=true</pre>
Enable the Department of Defense knowledge consent banner of the US department.	<pre>ncli file-server edit-security-params fs-name=file-server-name enable-banner=true</pre>
Change the default schedule of running the SCMA. The schedule can be hourly, daily, weekly, and monthly.	<pre>ncli file-server edit-security-params fs-name=file-server-name schedule=hourly</pre>
Disable the core-dump settings to let the file server VM generate stack traces for any cluster issue.	<pre>ncli file-server edit-security-params fs-name=file-server-name enable-core=false</pre>
Note: On a file server, this parameter turns both the core and the kernel dump salt status on or off.	
When a high governance official must run the hardened configuration.	Have the following settings. <pre>Enable Aide : true Enable Core : false Enable High Strength Password : true Enable Banner : false Schedule : HOURLY</pre>
When a federal official must run the hardened configuration.	Have the following settings. <pre>Enable Aide : true Enable Core : false Enable High Strength Password : true Enable Banner : true Schedule : HOURLY</pre>
Back up the DoD banner file.	Run the following command, and repeat on all FSVMs. <pre>nutanix@FSVM\$ sudo cp -a /srv/salt/security/AFS/sshd/DODbanner /srv/salt/security/AFS/sshd/DODbannerbak</pre>

Description	Command or Settings
Modify DoD banner file.	Run the following command, and repeat on all FSVMs. <pre data-bbox="850 289 1442 346">nutanix@FSVM\$ sudo vi /srv/salt/security/AFS/sshd/DODbanner</pre>
Create a separate mount for audit logs on new file servers.	<pre data-bbox="850 386 1442 556">nutanix@cvm\$ allssh 'echo "-- audit_logs_separate_fs=true" >> ~/config/ minerva_cvm.gflags' nutanix@cvm\$ allssh 'genesis stop minerva_cvm' nutanix@cvm\$ cluster start</pre>
Manage sudo authentication settings.	Note: Running the commands from a single FSVM applies the changes to all FSVMs on the cluster. Enable sudo authentication settings. <pre data-bbox="850 741 1382 793">nutanix@fsvm\$ afs infra.set_sudo_auth state=on</pre> Disable sudo authentication settings before performing Files upgrades, expanding a files cluster, and before using the following commands: genesis restart, genesis stop, cluster stop, and cluster restart_genesis. Disable sudo authentication settings. <pre data-bbox="850 1037 1382 1089">nutanix@fsvm\$ afs infra.set_sudo_auth state=off</pre> After performing file server activate, clone, or migrate operations, re-enable sudo authentication. <pre data-bbox="850 1188 1382 1241">nutanix@fsvm\$ afs infra.set_sudo_auth state=on</pre>
Enable SMB signing.	Run the following command on any FSVM in the cluster. <pre data-bbox="850 1331 1395 1388">nutanix@fsvm\$ afs smb.set_conf "server signing" "mandatory" section=global</pre>

Secure Sockets Layer (SSL) Certificate Settings

Files supports installing a self-signed or custom SSL certificate for certificate-based authentication. For file servers, these operations are only available through the nCLI and Files CLI commands.

Note: For steps on generating a certificate signing request, see [Generating a Certificate Signing Request with Subject Alternative Name for submission to Certificate Authority](#) in the *Security Guide*.

Table 6: SSL Certificate Settings on a File Servers

Description	Command or Settings
Generate a new self-signed SSL certificate.	Run the following command. Note: Replace <i>fs_uuid</i> with the universally unique identifier of the file server. <pre data-bbox="850 415 1338 474">nutanix@CVM\$ ncli file-server ssl-certificate-generate uuid=fs_uuid</pre>
Replace the existing self-signed SSL certificate with a new one.	Run the following command. <pre data-bbox="850 552 1370 636">nutanix@CVM\$ ncli file-server ssl-certificate-change-pfx-file-password uuid=fs_uuid</pre>
Get the current SSL certificate.	Run the following command. <pre data-bbox="850 714 1403 772">nutanix@CVM\$ ncli file-server get-ssl-certificate uuid=fs_uuid</pre>
Apply a custom SSL certificate on an FSVM.	Run the following command. • Replace <i>ca-path</i> with the CA certificate or chain file path. • Replace <i>cert-path</i> with the SSL certificate file path. • Replace <i>key-path</i> with the private key path. • Replace <i>value</i> with the type of private key (ECDSA_256, ECDSA_384, ECDSA_521, RSA_2048). <pre data-bbox="850 1192 1403 1312">nutanix@FSVM\$ afs misc.import_ssl ca_chain_file=ca_path cert_file=cert-path key_file=key-path key_type=value</pre>

Quality of Service (QoS) Setup

The table following details QoS setup on a file server using differentiate services code point (DSCP) values.

Table 7: QoS Settings

Description	Command or settings
Enable QoS. Note: The default management traffic value is 16, and the default data traffic value is 10.	<pre>nutanix@FSVM\$ afs net.enable_qos [management_dscp_value=value] [data_dscp_value=value]</pre>
Get the currently configured DSCP values.	<pre>nutanix@FSVM\$ afs net.get_qos</pre>
Modify the DSCP value of one or more traffic types.	<pre>nutanix@FSVM\$ afs net.edit_qos [management_dscp_value=value] [data_dscp_value=value]</pre>
Disable QoS on all traffic types and delete the existing configuration.	<pre>nutanix@FSVM\$ afs net.disable_qos</pre>

Rsyslog

The table following provides details on rsyslog daemon configuration for log forwarding.

Table 8: Rsyslog

Description	Command or settings
Add a rsyslog server.	<pre>nutanix@FSVM\$ afs rsyslog.add_server server_name=server_name ip_address=server_ip port=server_port network_protocol=TCP/UDP relp_enabled=true/false</pre>
Modify the properties of the rsyslog server.	<pre>nutanix@FSVM\$ afs rsyslog.update_server server_name=configured_server_name ip_address=server_ip port=server_port network_protocol=TCP/UDP relp_enabled=true/false</pre>
Remove the configured rsyslog server.	<pre>nutanix@FSVM\$ afs rsyslog.remove_server server_name=configured_server_name</pre>
Add forwarded modules and their debug-level to the rsyslog server. Note: Files only supports the syslog module.	<pre>nutanix@FSVM\$ afs rsyslog.add_server_module server_name=configured_server_name module_name=SYSLOG_MODULE level=ALERT/ CRITICAL/ DEBUG/ EMERGENCY/ ERROR/ INFO/ NOTICE/ WARNING</pre>
Remove the module and forwarding level form the rsyslog server.	<pre>nutanix@FSVM\$ afs rsyslog.remove_server_module server_name=configured_server_name module_name=SYSLOG_MODULE</pre>
Set the status for rsyslog forwarding. (You can disable rsyslog forwarding completely without removing the configured server details.) Note: Setting the <i>false</i> status ends forwarding and removes the configuration.	<pre>nutanix@FSVM\$ afs rsyslog.set_status enable=true/false</pre>
Get the status of rsyslog forwarding.	<pre>nutanix@FSVM\$ afs rsyslog.get_status</pre>
Get the current rsyslog configuration.	<pre>nutanix@FSVM\$ afs rsyslog.list</pre>
Get the configured modules and their forwarding level for a configured rsyslog server.	<pre>nutanix@FSVM\$ afs rsyslog.list_modules server_name=server_name</pre>
Set up transport layer security (TLS) parameters of the rsyslog server. Note: Not setting the auth-mode to <i>anon</i>, requires specifying permitted peers in a comma-separated list.	<pre>nutanix@FSVM\$ afs rsyslog.set_tls auth_mode=anon, x509/certvalid, x509/ fingerprint, x509/name ca_chain_path=PEM_encoded_CA_certificate_file_absolut permitted_peers=permitted_peers</pre>
Get the configured TLS parameters.	<pre>nutanix@FSVM\$ afs rsyslog.get_tls</pre>
Disable all configured TLS parameters, and disable TLS on forwarded packets.	<pre>nutanix@FSVM\$ afs rsyslog.disable_tls</pre>

TROUBLESHOOTING

Disable distributed file system (DFS) support on a witness share

You are unable to use a Nutanix Files witness share with Windows 2019 Failover Cluster Services.

DFS is enabled by default on the witness share.

Procedure

1.

Note: This remedy is only possible for standard shares.

Disable DFS on the witness share.

```
nutanix@fsvm$ afs share.dfs share-name disable
```

2. Get the fully-qualified domain name (FQDN) of the FSVM hosting the share.

```
nutanix@fsvm$ afs share.owner_fsvm share-name
```

3. Use the FQDN of the specific FSVM hosting the share (not the file server) for the witness share.

Invalid Mounts After Authentication Change

Clients cannot mount share.

Changes to the authentication method.

Procedure

Remount the authentication type as a value for parameter sec.

```
host$ -o sec=authentication-type
```

For example, use `-o sec=krb5` for Kerberos 5.

Client Access Denial (NFS Protocol)

Linux client experiences a "permission denied" error while accessing NFS shares.

A user management (authentication) change on an existing file server.

Procedure

This problem might be fixed by restarting the RPC-GSSAPI service on the clients. For example, enter the following command on a Linux CentOS 6 client:

```
nutanix@fsvm$ sudo service rpcgssd restart
```

The command syntax to restart the RPC-GSSAPI service varies among different Linux versions.

Clients Cannot Mount Shares

Clients in the same subnet as the Controller VM or in the storage network of the file server cannot mount the shares of the file server.

The file server was configured with two separate networks for the client-side and storage-side networks.

Procedure

To allow clients in the same subnet as the Controller VM or storage network to mount shares, configure the file server with the same network for both the client-side and storage-side networks.

Client Side Network Mapping

The file server client network does not map to a site on the AD and does not specify the client side.

The file server's client network does not map to any site on the AD. Files cannot find the local domain controllers in a multi-site AD environment and uses a geographically remote DC, which, can result in delayed domain operations.

Procedure

1. In a multi-site environment, map the Files client network to a local site in the AD.
2. In a single-site environment or with only a single geographic location, ignore the warning.

Connecting to Authentication Services

The file server cannot connect with the AD server or it cannot contact the LDAP server for the given domain.

The file server cannot reach the given domain name with the specified DNS server list. Possible reasons include spelling mistakes in the domain name, incorrect DNS name servers, or connectivity issues with the domain controller servers.

Procedure

1. Check the DNS server addresses, domain name, and status of the domain controllers.
2. Verify the DNS entries for the given domain name.

Constraint Violation

Domain controller reused an operation due to a possible constraint violation.

- Incorrect SPN configuration.
- An SPN is not unique in the forest and the conflict results in failure.

Procedure

1. Ensure that Files related SPN entries are not present in the forest.
2. Ensure that the domain controllers do not have any replication issues.

Disable Directory Leasing on Snapshot

Open file and directory handles caused by directory leasing prevent snapshot deletion.

Procedure

1. Display the directory leasing configuration for a share.

```
nutanix@fsvm$ afs smb.get_conf "directory leasing disabled on snapshot"  
section=share_name
```

2. Disable directory leasing and reset the default configuration.

```
nutanix@fsvm$ afs smb.del_conf "directory leasing disabled on snapshot" section=share_name
```

Note: Use the following command to re-enable directory leasing when needed.

```
nutanix@fsvm$ afs smb.set_conf "directory leasing disabled on snapshot" "yes"  
section=share_name
```

DNS Missing SRV Records

SRV records not found on the specified DNS servers.

The specified DNS servers do not resolve the SRV records with the appropriate domain controller names.

Procedure

Add the domain controller SRV records for the required protocols and services.

Domain Controller Issues

Cannot find a writable and reachable domain controller.

Files cannot discover an active non-RODC (writable) LDAP domain controller (at the site or domain level).

Procedure

Ensure that one writeable domain controller is working in the given domain.

Finding IP Addresses

Fetch the IP addresses for all FSVMs.

Various causes.

Procedure

Enter the following command from any FSVM: `nutanix@fsvm$ afs misc.fsvmips`

Output lists the IP addresses for the FSVMs in the node.

Identifying the Share Owner

Identify the owner of a share or export.

The Files UI does not list the name of share and export owners.

Procedure

Perform one of the following commands to identify the share owner.

- » List the share owners of a standard share or export.

```
nutanix@fsvm$ afs share.owner_fsvm share-name
```

- » List the share owners of a distributed share or export by specifying the *share-name* and the *share-path* (including the name of the top-level directory TLD).

```
nutanix@fsvm$ afs share.owner_fsvm share-name path=share-path
```

Invalid Credential

Invalid user name or password.

Files cannot authenticate on the AD using the given user name and password combinations.

Procedure

1. Ensure that the user name and password are correct.
2. Verify that the user is not expired, locked, or disabled.

NLM Locks

Unable to get Network Lock Manager (NLM) locks from Mac client.

NLM recovery does not work over the User Datagram Protocol (UDP). Use the transmission Control Protocol (TCP) instead.

Procedure

1. Add the following lines to the /etc/nfs.conf file:

```
nfs.lockd.send_using_tcp = 1
nfs.statd.send_using_tcp = 1
```

2. Restart services.

```
user@host$ launchctl stop com.apple.lockd ; launchctl start com.apple.lockd
```

```
user@host$ launchctl stop com.apple.statd ; launchctl start com.apple.statd
```

Network Cannot Expand

You changed the network ID and the file server cannot expand.

Update the UUIDs.

Procedure

1. Log into the CVM.
2. To update the file server network, enter the following:

```
nutanix@cvm$ afs infra.update_file_server_network
```

3. To update a specific file server within the network, enter the following:

```
nutanix@cvm$ afs infra.update_file_server_network fs_name
```

NFS Cache Expiry During Share Migration

You get the following job error while migrating a primary or secondary NFS share: LeaderJobWorkFlow: NFS cache expiry time was reduced during migration. Unable to reset that to default value. Please refer migration doc for same.

Running the command to reset NFS cache expiry time.

Procedure

Update the NFS cache configuration.

```
nutanix@fsvm$ afs nfs.remove_config_param block=EXPORT export_name=target_export_name \
```

```
param=Attr_Expiration_Time restart=False
```

Replace *target_export_name* with the name of the target export.

Restore Target Share Access Upon Canceling Migration Plan

Target share is in *Read Only* state.

Migration Plans which were cancelled leave the target share in Read Only state.

Procedure

1. Trigger another iteration of migration on a cancelled plan to restore the write access on the target share once the iteration is completed.
2. Or you can restore *Write* access on the share without having to run another iteration of migration.

```
nutanix@fsvm$ afs share.client_blocking_filter share_name delete Read-only ip_list=*  
vendor_name=kMigration
```

For example, if the target share name is *share1*.

```
nutanix@fsvm$ afs share.client_blocking_filter share1 delete Read-only ip_list=*  
vendor_name=kMigration
```

NTLM Authentication Issues

Authentication might be unsuccessful for NTLM when contacting read-only domain controllers (RODC).

The list of allowed password replication must include the machine account name or file server name. To resolve, follow troubleshooting steps on a domain controller.

Procedure

1. Add the host name.

```
C:\>repadmin /prp add RODC_host_name allow machine_account_DN
```

2. View the list of added names.

```
C:\>repadmin /prp view RODC_host_name allow
```

For example, viewing the list of added names would appear similar to the following:

```
C:\>repadmin /prp view MINDC03 allow
```

Output looks similar to the following:

```
Allow list (msDS-RevealOnDemandGroup):  
RODC "CN=MINDC03,OU=Domain Controllers,DC=automation,DC=nutanix,DC=com":  
CN=MNRVATST124803,CN=Computers,DC=automation,DC=nutanix,DC=com  
CN=Allowed RODC Password Replication Group,CN=Users,DC=automation,DC=nutanix,DC=com
```

Share Copying

Copy operation interrupted while using Microsoft robocopy to copy large files to Files shares.

Various causes (for example network bandwidth issues).

Procedure

Use robocopy with the */z* option. This option resumes any interrupted copy operation.

Stale Statistics

Windows client experiences stale statistics.

Default cache entry time is approximately 5 minutes.

Procedure

To change the default cache entry time of 5 minutes, log into the FSVM and run the following :

```
nutanix@fsvm$ afs smb.set_conf "stats cache ttl" "value" section=global
```

Time Difference

A time difference exists between Files and the domain controller.

Files uses Kerberos protocol for authentication on the AD. Kerberos is a time sensitive protocol and cannot sync the correct time when the client and servers are out of sync for several minutes.

Procedure

Use the same NTP server for the domain controller and Files.

Unsuccessful Authentication

Authentication might be unsuccessful for the NT LAN manager (NTL) when contacting read-only domain controllers (RODC).

The list of allowed password replication must include the machine account name or file server name. To resolve, follow troubleshooting steps on a domain controller.

Procedure

1. Add the host name.

```
C:\>repadmin /prp add RODC_host_name allow machine_account_DN
```

2. View the list of added names.

```
C:\>repadmin /prp view RODC_host_name allow
```

For example, viewing the list of added names would appear similar to the following:

```
C:\>repadmin /prp view MINDC03 allow
```

Output looks similar to the following:

```
Allow list (msDS-RevealOnDemandGroup):
RODC "CN=MINDC03,OU=Domain Controllers,DC=automation,DC=nutanix,DC=com":
CN=MNRVATST124803,CN=Computers,DC=automation,DC=nutanix,DC=com
CN=Allowed RODC Password Replication Group,CN=Users,DC=automation,DC=nutanix,DC=com
```

COPYRIGHT

Copyright 2025 Nutanix, Inc.

Nutanix, Inc.

1740 Technology Drive, Suite 150


San Jose, CA 95110

All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. Nutanix and the Nutanix logo are registered trademarks of Nutanix, Inc. in the United States and/or other jurisdictions. All other brand and product names mentioned herein are for identification purposes only and may be trademarks of their respective holders.

Nutanix Objects User Guide

Objects 5.2

November 5, 2025



NUTANIX

Contents

Nutanix Objects Overview.....	6
Terminology Reference for Nutanix Objects.....	6
Salient Features of Nutanix Objects.....	7
Usage of Nutanix Objects.....	9
Advantages of Nutanix Objects.....	10
Fault Tolerance vs Storage Efficiency.....	11
Larger Erasure Coding Strips.....	13
Use Cases and Recommendations for Network File System on Nutanix Objects.....	13
Network File System and Simple Storage Service Interoperability.....	14
Nutanix Objects Workflow.....	15
 Enabling Nutanix Objects.....	 17
 Finding Nutanix Objects Version.....	 19
 Nutanix Objects Security Hardening Instructions.....	 20
Protecting Nutanix Objects Network with Flow Networking (Microsegmentation).....	22
 Nutanix Objects Prism Central Disaster Recovery.....	 24
 Nutanix Objects Life Cycle Manager Upgrades.....	 25
Installing Nutanix Objects (First-time Users).....	26
Microservices Platform.....	26
Finding the Microservices Platform Version.....	26
Finding the Primary and Secondary Microservices Platform Clusters.....	27
 Role-based Access Control for Nutanix Objects.....	 28
Prerequisites for Configuring Role-based Access Control for Nutanix Objects.....	28
Role-Based Access Control Workflows for Nutanix Objects.....	29
Nutanix Objects Permissions.....	29
Creating Full Administrator Role for Nutanix Objects.....	30
Creating View-only Administrator Role for Nutanix Objects.....	31
Creating Bucket Administrator Role for Nutanix Objects.....	32
Creating Infra Administrator Role for Nutanix Objects.....	33
Creating a Custom Role for Nutanix Objects.....	34
 Nutanix Objects Prerequisites and Limitations.....	 36
Nutanix Objects Deployment Prerequisites - AHV and ESXi.....	36
ESXi Prerequisites for Nutanix Objects Deployment.....	37
Port Requirements for Nutanix Objects.....	38
Nutanix Objects Network Configuration.....	39
AHV Configuration for Nutanix Objects Deployment.....	40

ESXi Configuration for Nutanix Objects Deployment.....	44
Shared Versus Single Network Configuration for Nutanix Objects.....	48
Nutanix Objects Limitations.....	48
Limitations of Network File System.....	49
Configuring Virtual Private Cloud Modes for Nutanix Objects Without Network Address Translation.....	51
Configuring Virtual Private Cloud Modes for Nutanix Objects with Network Address Translation.....	54
Federation Overview.....	57
Federated Namespace Prerequisites.....	58
Creating Federated Namespace.....	58
Adding External Object Store.....	60
Removing External Object Store.....	61
Managing Federated Namespace.....	62
Viewing Federated Namespaces.....	62
Deleting Federated Namespace.....	63
Errors in Federated Namespace.....	64
Object Store Service Deployment.....	65
Object Store and Domain Naming Conventions.....	65
Prerequisites for Enabling FT1n:2d on an Object Store.....	66
Prerequisites for Nutanix Objects Network Segmentation of an Object Store.....	66
Creating or Deploying an Object Store on Prism Central.....	67
Viewing Object Store Deployments.....	72
Access Nutanix Objects Endpoints.....	74
Viewing Object Store Summary.....	74
Deleting an Object Store.....	75
Object Store Expansion.....	76
Prerequisites for Expanding Storage for an Object Store in Nutanix Objects.....	77
Secondary Cluster Addition or Removal Guidelines.....	77
Expanding Storage for an Object Store.....	78
Scaling Out an Object Store.....	80
Managing FQDN and SSL Certificate.....	81
Adding FQDNs.....	82
Setting up SSL Certificate for an Object Store.....	83
Deploying Object Store on Nutanix Cloud Cluster.....	85
Deploying Object Store at a Dark Site (Offline Deployment).....	86
Types of Nutanix Objects Users.....	89
Directory Configuration and Access Key Generation.....	90
Nutanix Objects IAM-High Availability Overview.....	90
IAM-High Availability Prerequisites.....	90
IAM-High Availability Limitations.....	91
Configuring Directories.....	91
Generating Access Key for API Users.....	94

Viewing API Users.....	96
Managing API Keys.....	96
Deleting API Users.....	97
Bucket Creation, Operations, and Bucket Policy Configuration.....	99
Bucket Naming Conventions.....	99
Creating and Configuring an S3 Bucket in Nutanix Objects.....	99
Creating and Configuring an NFS Bucket in Nutanix Objects.....	101
Managing NFS Allowlist in Nutanix Objects.....	104
Bucket Policy Configuration in Nutanix Objects.....	105
Object Versioning in Nutanix Objects.....	105
Lifecycle Policies in Nutanix Objects.....	106
Cloud Tiering in Nutanix Objects.....	109
Legal Hold for Nutanix Objects.....	117
WORM Bucket in Nutanix Objects.....	118
Configuring a Bucket for Static Website Hosting in Nutanix Objects.....	120
Cross-Origin Resource Sharing (CORS) in Nutanix Objects Overview.....	122
Adding Tags to a Bucket in Nutanix Objects.....	124
Viewing Buckets in Nutanix Objects.....	125
Bucket Summary in Nutanix Objects.....	126
Updating a Bucket in Nutanix Objects.....	127
Sharing a Bucket in Nutanix Objects.....	128
Bucket Access Policies in Nutanix Objects.....	131
Listing the Shared Buckets.....	134
Viewing Bucket Users in Nutanix Objects.....	135
Editing User Access in Nutanix Objects.....	136
Deleting a Bucket in Nutanix Objects.....	137
Nutanix Objects Streaming Replication.....	139
Types of Bucket Replication.....	139
Nutanix Objects Streaming Replication Guarantees and Topologies.....	140
Replication Topologies Supported for Nutanix Objects as Destination.....	141
Replication Topologies Supported for S3 Destination.....	144
Prerequisites for Replication to Nutanix Objects.....	145
Adding Remote Prism Central as Availability Zone.....	145
Setting up IAM Synchronization with a Different Prism Central.....	146
Prerequisites for Replication to S3.....	147
Prerequisites for Creating a Replication Relation.....	148
Creating Replication Relation for Buckets.....	149
Delete Marker Replication Scenarios.....	153
Sample Replication Configurations.....	154
Combination of Versioned and Non-Versioned Nutanix Objects Replication.....	157
Pausing a Replication Rule for a Bucket.....	158
Resuming a Replication Rule for a Bucket.....	158
Deleting a Replication Rule of a Bucket.....	159
Viewing Bucket Replication Statistics.....	159
Achieving Fault Tolerance for IAM on Nutanix Objects.....	161
Monitoring and Alerts in Nutanix Objects.....	164
Viewing Performance of Object Stores.....	164
Viewing Performance of Buckets.....	165
Viewing Object Store Usage.....	166
Assigning Quota Policy to a User.....	167

Viewing Nutanix Objects Alerts.....	170
Nutanix Objects Specific Alerts.....	172
Nutanix Objects Notifications.....	182
Notification Types for Nutanix Objects.....	182
Configuring Event Notifications in Nutanix Objects.....	187
Creating Notification Rules for Data Events in Nutanix Objects.....	192
Nutanix Objects Browser.....	196
Browser Version Compatibility for Nutanix Objects Browser.....	196
Administrator Workflow: Nutanix Objects Browser.....	197
Launching the Nutanix Objects Browser.....	198
Supported Operations in Nutanix Objects Browser.....	199
Bucket Operations in Nutanix Objects Browser.....	199
Object Operations in Nutanix Objects Browser.....	213
User and Key Management in Nutanix Objects Browser.....	229
Nutanix Objects Prometheus Exporter.....	237
Nutanix Objects CRUD Operations by Using S3 APIs.....	239
Nutanix Objects Authentication.....	239
Nutanix Objects Supported and Unsupported APIs.....	239
Nutanix Objects Supported APIs.....	240
Nutanix Objects Unsupported APIs.....	244
Buckets and Nutanix Objects Tagging APIs Overview.....	244
API Operations Supported for Buckets and Nutanix Objects Tagging.....	245
Nutanix Objects S3 Select API Overview.....	246
Supported SQL Functions by Nutanix Objects S3 Select.....	247
Server-Side Encryption with Customer-Provided Keys Supported APIs.....	252
Nutanix Objects REST Error Responses.....	254
List of Error Codes for Nutanix Objects.....	254
Nutanix Objects Integration with Backup Applications.....	257
Troubleshooting Nutanix Objects.....	258
Shutting Down Nutanix Objects VMs.....	258
Starting the Nutanix Objects VMs.....	260
Detection of Slow Connections by Nutanix Objects.....	261
Prism Element Clusters or Subnets Not Listed in the User Interface.....	261
Object Stores Not Listed While Creating a Replication Rule.....	262
Manually Updating the DNS Server for Nutanix Objects Cluster.....	262
Copyright.....	265

NUTANIX OBJECTS OVERVIEW

Nutanix Objects is a software-defined Object Store Service. This service is designed with an Amazon Web Services Simple Storage Service (AWS S3) compatible REST API interface capable of handling petabytes of unstructured and machine-generated data. Nutanix Objects addresses storage-related use cases for backup, long-term retention, and data storage for your cloud-native applications by using standard S3 APIs. You no longer have to introduce an external, separately managed storage solution.

Nutanix Objects is deployed and managed as part of the Nutanix Unified Storage. It enables users of the Nutanix platform to store and manage unstructured data on top of a highly scalable hyperconverged architecture. In comparison to cloud-hosted solutions, this on-premises model offers more consistent control over the costs associated with storing objects, along with more transparency about the location of those objects.

You can manage objects by using Prism Central or the S3-compatible REST APIs after an administrator has authorized the applications and users to access buckets accordingly.

For more information on Nutanix Objects architecture, see [Solutions Documentation](#) and [Nutanix Bible](#)

Terminology Reference for Nutanix Objects

The following are the terms that you frequently encounter when you are using Nutanix Objects:

Table 1: Terminology Reference

Terminology	Description
Nutanix Objects™ (Nutanix Objects)	Nutanix Objects is deployed and managed as a part of Nutanix Unified Storage that delivers secure S3-compatible object storage at massive scale to hybrid cloud environments.
Bucket	An organizational unit exposed to the users and contains the objects. A deployment might have one or more buckets.
Object	Object represents the data uploaded by the user or application. The actual unit (blob) of storage and the item interfaced by using the Application Programming Interface (API) (For example, GET or PUT).
S3	The term used to describe the Amazon Web Services (AWS) interface. This term is now used synonymously for an object service. S3 is also used to describe the object API that you use to interact with an object store.
Storage Network	A Virtual Local Area Network (VLAN) is required for communication between Nutanix Objects Services.
Public Network	A VLAN used to access the object store endpoints externally.
Microservices Platform (MSP)	A platform based on Kubernetes where all the Nutanix Objects microservices run.

Terminology	Description
AHV IPAM	IP Address Management (IPAM) is a feature of AHV that allows it to assign IP addresses automatically to VMs by using DHCP. You can configure each virtual network with a specific IP address subnet, associated domain settings, and groups of IP address pools available for assignment to VMs.
Worker VMs	Virtual machines created during object store deployment that host various containerized Nutanix Objects services. Also, worker VMs are referred to as Nutanix Objects nodes.
Nutanix Objects Browser	A User Interface (UI) that allows the users to directly launch the object store instance in a web browser and perform bucket and object level operations.

Salient Features of Nutanix Objects

The salient features of Nutanix Objects are as follows:

Write-Once-Read-Many

Nutanix Objects enables you to create Write-Once-Read-Many (WORM) buckets to prevent anyone (including an administrator) from modifying or deleting data while the policy is active. WORM policies help you comply with strict regulations often mandated by the healthcare, financial, and government sectors. You can apply WORM policies to any bucket to prevent any updates to the data until the WORM policy expires. Nutanix provides a 24-hour grace period for letting you test the WORM policy. The bucket creator can undo the applied policy within the allowed grace period of 24 hours. After the grace period, no one can delete the policy or reduce the retention period. The option to extend the retention period is available. You must enable versioning on a WORM bucket so that you can update an object. The updated object gets stored as a new version. Versioning ensures that older versions of the object get retained and the data never gets overwritten.

For more information, see [WORM Bucket in Nutanix Objects](#) on page 118.

Immutability

The retrieved data must be the same as originally written. This means that you can retrieve the stored object (and its associated metadata) in the future without any modification. If you enabled WORM, the object remains unchanged and is also permanent.

Object Versioning

When you enable object versioning, uploading multiple copies of the same object creates multiple versions. You can use the versioning feature to protect data from accidental overwrite or deletion or as an option for reverting to a previous state. You can also turn off object versioning on a bucket at any time and set the objects in the bucket to expire after a specified time. By default, there is no limit on the number of versions for an object, considering the storage space is available.

For more information, see [Object Versioning in Nutanix Objects](#) on page 105.

Data Life Cycle Management

You can use age-based data retention policies to enforce compliance with strict data retention regulations that dictate the time to store specific data. For example, the Health Insurance Portability and Accountability Act (HIPAA) regulation mandates you to retain medical data for six years from the creation date. You can set a data retention policy to delete all data created in a specific bucket six years after its creation date.

Note: The WORM policy supersedes any set retention policy.

You can set retention policies at the bucket level where objects in a bucket expire after a certain amount of time you specify. If you do not set retention policies, the only limit to how many versions can be maintained depends on the storage space.

For more information, see [Lifecycle Policies in Nutanix Objects](#) on page 106.

Multipart Upload

Multipart upload allows you to reduce slow upload times by breaking large pieces of data into chunks. Then, the system can handle the data separately to increase upload speeds if you upload them simultaneously. You can also use multipart uploads to prevent losing progress during an upload. For example, if there is a connectivity loss, most applications retry only the unsuccessful chunk. Hence, you do not have to upload the entire object again.

For more information, see [Nutanix Objects Supported APIs](#) on page 240.

Data-at-Rest Encryption with Native Key Management

Nutanix Objects provides a FIPS 140-2 compliant data-at-rest encryption solution. To deliver this capability, Nutanix Objects uses the underlying AOS encryption capability. You can set encryption at an entire cluster level, always encrypting all data. With the native key management, the Nutanix cluster manages the keys, so the solution requires no additional device management or third-party costs.

Identity and Access Management

Native Identity and Access Management (IAM) functionality ensures that you have access only to the buckets and objects you created and granted access permissions. Each user gets a pair of access and secret keys that can be used by their applications to access Nutanix Objects. You can also generate access and secret keys for an Active Directory group. Administrators can revoke and regenerate keys at any time.

For more information, see [Directory Configuration and Access Key Generation](#) on page 90.

Multi-protocol Access

Nutanix Objects allows you to create buckets using both S3 and NFS protocols. NFS protocol support is natively implemented over Nutanix Objects and built on the same foundation that powers the S3 protocol.

For more information, see the following references:

- [Use Cases and Recommendations for Network File System on Nutanix Objects](#) on page 13
- [Network File System and Simple Storage Service Interoperability](#) on page 14
- [Limitations of Network File System](#) on page 49
- [Creating and Configuring an NFS Bucket in Nutanix Objects](#) on page 101
- [Creating and Configuring an S3 Bucket in Nutanix Objects](#) on page 99

Cloud Tiering

Cloud tiering enables you to move objects to another S3-compatible object store bucket for saving storage space in the Nutanix Objects cluster. Tiering can help you to save costs by sending the infrequently accessed objects to platforms such as AWS S3. The supported endpoints are AWS S3, Azure Blob Storage, Google Cloud Platform (GCP) and a different Nutanix Objects instance.

For more information, see [Cloud Tiering in Nutanix Objects](#) on page 109.

Nutanix Objects Streaming Replication

Nutanix Objects streaming replication enables automatic and asynchronous copying of source buckets to the target buckets in different Nutanix Objects instances. This provides you with the ability to replicate a single source bucket to a maximum of three destination buckets. These destination buckets can be on Nutanix Objects on the same Prism Central cluster or on a remote Prism Central cluster, or a Cloud Storage Endpoint.

For more information, see [Nutanix Objects Streaming Replication](#) on page 139.

Server-Side Encryption with Customer-Provided Keys (SSE-C)

Server-side encryption with customer-provided keys (SSE-C) allows you to encrypt specific objects at rest using encryption keys that you provide with each request. Unlike the data-at-rest encryption with native key management, which encrypts all data at the cluster level using system-managed keys, SSE-C gives you full control over both the encryption keys and the objects you choose to encrypt. Nutanix Objects encrypts the data before storing it but does not retain the encryption key, ensuring that only you can decrypt the data upon retrieval.

Note: SSE-C does not support Network File System, federation, replication, or S3 select.

For more information, see [Server-Side Encryption with Customer-Provided Keys Supported APIs](#) on page 252.

Usage of Nutanix Objects

Following are examples of solutions you can implement by using Nutanix Objects:

- **Backup:**

You can integrate Nutanix Objects with the backup applications such as Commvault, HYCU, Veeam, and Veritas. You can create backups to protect your data with a simple, scalable, and cost-effective active archive solution. You can start with small storage and scale to petabytes of storage to deliver great performance. Nutanix Objects supports the multipart upload API, with which you can reduce slow upload times by breaking data into chunks and uploading documents, images, and videos to the global namespace.

- **Long-term Retention:**

You can use Nutanix Objects for long-term data retention. Built-in object versioning provides storage protection and searches your data without the problem of tape systems. Versioning maintains previous copies of the object to avoid data loss from overwrites or deletes.

Some organizations, especially in the health, legal, and government sectors, are required to comply with the strict regulations mandating requirements, such as the minimum period for which data must be available or who can alter data. Nutanix Objects allows you to store data over a long period with features designed to comply with such strict regulations. For example, Write once read many (WORM) buckets prevent anyone (including an administrator) from modifying or deleting data while the policy is active.

- **DevOps:**

Easy application access to your data in a global namespace using simple PUT and GET commands makes Nutanix Objects the perfect fit for your dev-ops data. Nutanix Objects offers broad support for most programming environments, and you can access it using standard internet networking infrastructures, such as Hypertext Transfer Protocol (HTTP) and Hypertext Transfer Protocol Secure (HTTPS) protocols. With Nutanix Objects, you can easily access the objects with Simple Storage Service (S3)-compatible HTTP REST API requests from different locations. The bucket-level metrics

allow you to track resource utilization with ease. Also, Identity and Access Management (IAM) support enables secure access to Nutanix Objects resources and services.

DevOps and IT operations can use the S3-compatible interface for cross-geography, cross-team collaboration, and agile development.

Consider the following scenario. Joe and Adam work together as developers on the same team for an image-processing application called XYZ. Joe, residing in New York, develops code for the data creation and storage of XYZ. Adam, residing in Seattle, performs the testing of XYZ. Several buckets dedicated to the application XYZ are stored in the Nutanix Objects store instance deployed in the New York office. After completing the programming for the latest feature, Joe stores the latest version of XYZ in a bucket. Adam can use GET request automation scripts to pull the latest version of the code and store the latest test results in another bucket using PUT requests.

Thus, the team works more efficiently and quickly. Also, it demonstrates the immutable capabilities of object storage, where the data of Joe remains unchanged while Adam performs testing.

Advantages of Nutanix Objects

Following are the advantages of object storage with Nutanix Objects:

No Silos

Nutanix provides file services (Files) and block services (Volumes) as part of the Acropolis Distributed Storage Fabric (DSF). You can also add Nutanix Objects to the solution, thus allowing block, file, and object storage solutions to coexist with no silos. You can deploy and manage these features in a single environment.

Security-First Approach

Nutanix integrates security into every step of its solution stack from the early stages of development. For example, the stack conforms to Security Technical Implementation Guides (STIGs), which maintain a security baseline configuration based on common standards established by the National Institute of Standards and Technology (NIST). STIGs use machine-readable code to automate self-healing and compliance with the security standards for AOS and AHV. Nutanix also complies with *SEA 17a-4*, which specifies requirements for data retention and accessibility. Nutanix Objects also supports Data-at-Rest Encryption.

Capacity Optimization

Nutanix Objects leverages data-capacity optimization technologies such as compression and Nutanix Erasure Coding (EC-X) in the background. Inline EC is the default setting on Nutanix Objects deployments. A three-day post-process policy is set up in parallel for data that spills past the inline EC process.

In addition to compression savings, EC-X increases the usable storage capacity in a cluster with no overhead to the active path write.

In a cluster with the redundancy factor 2, two copies of the data get replicated among all nodes for resilience. Checksums of the data are stored with the metadata to ensure validity if corruption occurs. Therefore, the cluster (redundancy factor 2) uses half of its raw storage capacity to store copies of its data. EC-X performs the OR operation on these copies of data to compute a parity block. The original data blocks and the parity form an erasure code strip. This process reduces the number of actual data copies needed to protect the environment from a single node failure. If objects are being deleted regularly, it might be best to disable EC due to the garbage generated.

Suppose you have three data blocks (A, B, and C) stored in a four-node cluster, with each data block distributed across three nodes. Without EC-X, you might have the original three blocks of data plus three more copies of those blocks, which sum up to a total of six blocks of data. With EC-X enabled, you can reduce the storage footprint to four blocks of data that include the original three blocks of data plus their parity. These blocks are present in different nodes. If a node fails, the system can reconstruct the data from the data blocks and parity data present on the available

nodes. For example, If the node hosting block B fails, the system can reconstruct its data from blocks A, C, and P (parity).

Additionally, when the disks are not getting used heavily, the system regularly performs background checksum validations and data restoration to maximize protection against data corruption. These techniques (combined with a high tolerance for disk and node failures in a self-healing and distributed system) form a resilient and reliable enterprise object-storage platform.

Note: Erasure Coding can only be enabled on clusters with four or more nodes (redundancy factor - 2) or six or more nodes (redundancy factor - 3).

For more information on Erasure Coding best practices and requirements, see [Prism Element Web Console Guide](#).

For more information on Fault Tolerance and Storage Efficiency, see [Fault Tolerance vs Storage Efficiency](#) on page 11.

Ease of Use and Simple Management

Nutanix Objects is compatible with HTTP S3 REST APIs for object and bucket storage. Nutanix Objects supports all the necessary methods (including GET, PUT, POST, DELETE, and LIST) to perform basic object operations.

Nutanix Objects inherits the simplicity of use and management of Prism Central and provides the managing, monitoring, and reporting capabilities that are simple to administer.

For example, Prism Central automatically provides checks for interoperability and compatibility during a software update task. Alerts inform users about specific conditions, such as object storage components hitting capacity thresholds and reduced high availability. The Prism user interface reports on the size of the object store, the number of buckets and objects, average throughput, latency, and the number of GET and PUT operations.

Cost

With Nutanix Objects, the cost to store a large amount of unstructured data is less compared to traditional storage solutions. Compared to public cloud providers, Nutanix Objects does not charge for data ingress or egress.

Fault Tolerance vs Storage Efficiency

In Nutanix Objects, there is a tradeoff between how many faults a cluster can tolerate (for example, the number of drives or the number of node failures) and how efficiently a cluster can store the data.

For example, with Fault Tolerance 1 (FT1), an AOS cluster can tolerate the failure of one node or one disk. Without any capacity optimizations, such as compression or Erasure Coding, two copies of the data must be maintained (redundancy factor of 2), resulting in an overhead of 2x. In the case of Fault Tolerance 2 (FT2), the AOS cluster can tolerate the failure of two nodes or two disks. Similarly, without capacity optimizations, three copies of the data need to be maintained (redundancy factor of 3), leading to an overhead of 3x.

Both fault tolerance modes can be represented as FT1n:1d (Fault Tolerance of 1 node and 1 disk failure) and FT2n:2d (Fault Tolerance of 2 nodes and 2 disk failures), indicating how many node or disk failures the cluster can withstand.

Currently, the fault tolerance (FT) guarantees apply equally to both node and disk failures, based on the assumption that these failures occur with the same likelihood. However, in dense node configurations—typically those with 20 or more drives—there are significantly more disks in a cluster than there are nodes. As a result, the likelihood of experiencing multiple disk failures is greater than that of multiple node failures.

FT1n:2d offers some capacity savings compared to FT2, even without any capacity optimizations. However, these savings significantly increase when Erasure Coding is enabled. In traditional FT2 mode, each node can only store one member of the Erasure Code Strip. In contrast, FT1n:2d allows for up to two members of the Erasure Code Strip to be stored on the same node, provided they are on different disks.

This capability enables larger strips to be formed in a smaller cluster, thereby enhancing overall storage efficiency.

Fault Tolerance Mode Comparison Without Erasure Coding

Without Erasure Coding (EC), the only savings achieved between FT2 and FT1n:2d comes from the reduction in reserved rebuild capacity that needs to be allocated. With this in mind, the following table illustrates the overhead associated with different fault tolerance schemes, assuming no capacity optimizations (i.e., no compression, deduplication, or erasure coding). The values under each fault tolerance scheme represent the ratio of physical capacity to available logical capacity for various cluster sizes, with lower values indicating better efficiency.

In the case of FT1, as the cluster size increases, the ratio of physical to logical capacity approaches 2, which is the minimum if we are storing two copies of the data. Similarly, for FT2, as the cluster size grows, the ratio approaches 3, which is the minimum if we are storing three copies of the data.

FT1n:2d helps us create a new mode between FT1 and FT2. This mode offers better Fault Tolerance than FT1 and has less overhead than FT2.

Number of Nodes	(Physical Capacity / Available Logical Capacity) without EC		
	FT1n:1d (FT1)	FT1n:2d	FT2n:2d (FT2)
5	2.5x	3.75x	5.0x
6	2.4x	3.6x	4.5x
8	2.29x	3.43x	4.0x
12	2.18x	3.33x	3.6x
16	2.13x	3.2x	3.42x
32	2.06x	3.10x	3.2x

Fault Tolerance Mode Comparison with Erasure Coding

As noted above, with EC, one can form large strips on a smaller cluster with FT1n:2d. The following table summarizes the effective logical capacity that different FT schemes offer if all the data is erasure-coded. Note that these calculations assume that the largest strip sizes possible today are 4,1 and 4,2.

One can see that for FT1, as the cluster size increases, the ratio of physical to logical capacity approaches 1.25x, which is the lower bound if we are forming strips of size 4,1. Similarly, one can see that for FT2, as the cluster size increases, the ratio of logical to physical capacity approaches 1.5x, which is the lower bound if we are forming strips of size 4,2.

With FT1n:2d, it is clear that we are able to carve out a mode between FT1 and FT2, which offers better fault tolerance than FT2 with some additional storage overheads. These overheads are significantly lower than that offered by the FT2 scheme, especially for mid-sized clusters.

Number of Nodes	Physical Capacity / Available Logical Capacity) with EC		
	FT1n:1d(FT1)	FT1n:2d	FT2n:2d (FT2)
5	1.67x	1.88x	5.0x (no EC)
6	1.5x	1.8x	3.0x
8	1.43x	1.71x	2.0x
12	1.36x	1.63x	1.8x
16	1.33x	1.6x	1.71x
32	1.29x	1.55x	1.60x

Larger Erasure Coding Strips

Increased the Erasure Coding (EC) strip size for large clusters to improve storage efficiency.

EC strips are represented as N/K, where N is the number of information or data egroups, and K is the number of parity blocks.

The strip size of the EC can be configured as a container parameter for each container.

Note: Larger EC strips are enabled in a cluster that has fault tolerance of one node or two disk failures (FT1n:2d) enabled.

FT1n:2d is enabled on a newly deployed object store when the following criteria are met:

- The cluster has a minimum of five nodes.
- Fault tolerance (FT2) or a redundancy factor of 3 (RF3) is enabled on the cluster.
- The AOS version must be 6.8 or higher.
- Each node in the dense node platform hosts a minimum of 20 HDDs or SSDs.
- There are no non-internal containers (default, management, objects-lite, and self-service) with FT2 present.

Table 2: Wider EC Strip Configuration Requirements

Effective EC Strip Configuration	Node Requirement
6/2	5 nodes
8/2	6 nodes
10/2	7 nodes
12/2	8 nodes or more

Use Cases and Recommendations for Network File System on Nutanix Objects

Use Cases

Network File System (NFS) protocol support provides multiprotocol access (NFS 3.0, Simple Storage Service (S3)) to data stored within Nutanix Objects buckets. NFS protocol support is natively implemented over Nutanix Objects and builds on the same foundation that powers the S3 protocol.

NFS access over Nutanix Objects is ideally suited for large-scale, read-heavy workloads with sequential accesses where data is ingested once and minimally or never modified later.

Some ideal use cases are:

- **Backup applications:** Use Nutanix Objects as a large-scale NFS repository for backups while you migrate the underlying storage to Nutanix Objects.
- **Analytics:** Use Nutanix Objects multiprotocol access for in place analytics as you ingest data through object or S3 interface and access through file or NFS interfaces that analytic systems require.

Recommendations

Nutanix recommends the following:

- Not to use NFS access over Nutanix Objects for use cases that require modification to data including file edits and renames. For example, end user computing, home shares, app data that is frequently modified like presentations or computer-aided design (CAD) files.
- To have up to 100 subdirectories within a directory and a total of up to 100 million objects within a NFS bucket.
- To use `nconnect=1` mount option for NFS buckets.

Network File System and Simple Storage Service Interoperability

This section explains how objects in the Simple Storage Service (S3) namespace get mapped to files and directories in the Network File System (NFS) namespace and vice-versa.

Network File System to Simple Storage Service

Nutanix Objects in the NFS namespace get mapped to files and directories in the S3 namespace in the following ways:

- All the files and symbolic links created in the NFS namespace appears in the S3 namespace as objects.
- Any directory created from the NFS namespace will not appear in the S3 namespace, as the S3 protocol does not have any notion of directories.
- Any S3 operation like ObjectHead, ObjectGet or ObjectDelete on the directories fails with the error message `NfsDirectoryOperationNotAllowed` (Operation prohibited on an NFS directory). Therefore, a file present inside directories and subdirectories appears as a single object in the S3 namespace, as shown in the following table:

Table 3: Network File System to Simple Storage Service

NFS	S3
• dir1/ • dir2/ • file	dir1/dir2/file

Simple Storage Service to Network File System

Nutanix Objects in the S3 namespace get mapped to files and directories in the NFS namespace in the following ways:

- Nutanix Objects created from the S3 protocol appear as files and directories in the NFS namespace based on the object name. If the object name contains a directory-like hierarchy, then the object is stored in a hierarchical namespace as identified by its name.

The following table shows an example where an object `a/b/c` created from S3 protocol appears in the NFS namespace as `a/directory`, which contains the subdirectory `b/`, which in turn contains the file `c`.

Table 4: Simple Storage Service to Network File System

S3	NFS
a/b/c	• a/ • b/ • c

Note:

- These implicit directories are created when an object containing directory hierarchy is accessed from the NFS namespace.
- The implicit directories continue to exist in the NFS namespace after the object containing the directory hierarchy has been deleted.
- Folders created from the Nutanix Objects Browser also appear as a directory in the NFS namespace.

For more information on the Limitations of NFS on Nutanix Objects, see [Limitations of Network File System](#) on page 49.

Nutanix Objects Workflow

This section describes the basic workflow of Nutanix Objects. The subsequent sections provide detailed information about each step.

About this task

The Nutanix Objects workflow is as follows:

Procedure

1. Enable Nutanix Objects from Prism Central. For more information, see [Enabling Nutanix Objects](#) on page 17.
2. Deploy object store on a desired cluster. For more information, see [Object Store Service Deployment](#) on page 65.
3. Generate the access keys. For more information, see [Generating Access Key for API Users](#) on page 94.
4. Set up the Secure Sockets Layer (SSL) certificates for the object store. For more information, see [Setting up SSL Certificate for an Object Store](#) on page 83.
5. Access the object store endpoints from the third-party clients or Nutanix Objects Browser. For more information, see [Access Nutanix Objects Endpoints](#) on page 74 and [Nutanix Objects Browser](#) on page 196.
6. Create buckets using either Simple Storage Service (S3) or Network File System (NFS) protocol. For more information, see [Bucket Creation, Operations, and Bucket Policy Configuration](#) on page 99 and [Bucket Operations in Nutanix Objects Browser](#) on page 199.

7. Upload objects and perform object operations using Nutanix Objects Browser or S3 APIs. For more information, see [Supported Operations in Nutanix Objects Browser](#) on page 199 and [Nutanix Objects Supported APIs](#) on page 240.
8. Expand the object store if the storage is getting full. For more information, see [Expanding Storage for an Object Store](#) on page 78.

ENABLING NUTANIX OBJECTS

To create an object store, you first need to enable Nutanix Objects in Prism Central, then add a license in Prism Element or Prism Central, depending on the versions in use. With the new Prism Central experience, users with the Nutanix Objects Admin role cannot enable the Nutanix Objects service. Nutanix Objects service can only be enabled by Super Admin or Prism Admin users.

Before you begin

A non-admin user can enable objects only after the administrator creates an access control policy on a role in Prism Central for the non-admin user.

Note: Administrators and non-admin users are Prism Central users with the following roles:

- Administrator—A Super Admin or a Prism Admin in Prism Central.
- Non-admin user—A Prism Central user without any administrator privileges.

For more information on the built-in roles in Prism Central, see [Built-in Role Management](#) in the *Security Guide*.

About this task

To enable Nutanix Objects in Prism Central, follow these steps:

Procedure

1. Log on to Prism Central as an administrator.
2. In the Application Switcher, click **Admin Center**.
3. Click **Market place** in the navigation bar.
4. In the Nutanix Apps section, click **Get** for the Objects application.
The app details page appears.
5. Click **Deploy**.
It takes 4-5 minutes to complete the deployment.
6. In the Application Switcher, click **Objects**.
A welcome page appears with prerequisite details to create an object store.
7. Click **Download Creation Checklist**.
Enabling of Object Store Services is performed only once. After you enable Nutanix Objects, ensure that you perform LCM inventory and upgrade the MSP and Nutanix Objects Manager to the latest versions before you start with deployment. For more information, see [Nutanix Objects Life Cycle Manager Upgrades](#) on page 25.
A list of prerequisites for deploying an object store is downloaded.

8. (Only for ESXi clusters) Click **vCenter Registration** before deploying the object store on an ESXi cluster.

To deploy object stores on the ESXi clusters, you need to provide the vCenter credentials and configure the IPAM for the ESXi networks. For more information, see [Managing vCenter for Nutanix Objects Service](#) on page 46.

A non-admin user can perform vCenter Registration only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions:

- Create Object Store
- View Object Store

Administrators and non-admin users are Prism Central users with the following roles:

- Administrator—A Super Admin or a Prism Admin in Prism Central.
- Non-admin user—A Prism Central user without any administrator privileges.

For more information on the built-in roles in Prism Central, see [Built-in Role Management](#) in the *Security Guide*.

9. Click **Next**.

10. Click **Create Object Store**.

For more information on creating the object store, see [Creating or Deploying an Object Store on Prism Central](#) on page 67.

You can now create an object store.

What to do next

After deployment, the Nutanix Objects app appears on the **My Apps** page. You can also select **Objects** in the Application Switcher to access and use the app.

FINDING NUTANIX OBJECTS VERSION

You can find the installed versions of Nutanix Objects Manager and Nutanix Objects Service from Prism Central.

About this task

To find the version of Nutanix Objects, follow these steps:

Procedure

1. Log on to Prism Central.
2. On the Application Switcher, go to **Admin Center > LCM > Inventory**.
3. Click **Perform Inventory**.
The **Installed Versions** list displays the version of Nutanix Objects Manager and Nutanix Objects Service.

NUTANIX OBJECTS SECURITY HARDENING INSTRUCTIONS

This chapter describes the security hardening instructions for Nutanix Objects.

Before performing Nutanix Objects hardening configurations, ensure that Prism Central and Prism Element are hardened. For more information, see [Security Management Using Prism Element](#) and [Security Management Using Prism Central](#) in the *Nutanix Security Guide*.

Networking

The following are the instructions for hardening the network for Nutanix Objects:

- Ensure only the Nutanix Objects Simple Storage Service (S3) endpoint is exposed to users.
- A firewall protects Prism Central and Prism Element access and is not open to end users.
- The Nutanix Objects storage network is only accessible to Prism Element and Prism Central, and a firewall must protect this network.

Note: Use the storage network only for Nutanix Objects deployment. Do not use this storage network for deploying any workload or any other application.

- For the Nutanix Objects replication, open the required ports between the source and destination Nutanix Objects cluster IP addresses only. All other access must be protected.

Firewall Rules

Firewall rules must be configured to protect the infrastructure of Nutanix Objects. For more information, see [URL and Port Requirements](#).

The following is the list of networks:

- Nutanix Objects Storage Network
- Nutanix Objects Public Network
- Internet connection
- Other services
- Prism Central
- Prism Element

Flow Networking Rules

If Flow Networking is enabled on a setup, rules can be configured as security policies (application policy) manually or automatically instead of the external firewall.

For more information on configuring the security policy when Flow Networking is enabled, see [Protecting Nutanix Objects Network with Flow Networking \(Microsegmentation\)](#) on page 22.

If Flow Networking is not enabled, the same rules can be applied through an external firewall.

The following is the list of inbound traffic allowlist:

- Controller Virtual Machine (CVM) IP addresses (including secondary/segmented)
- Prism Central IP address

- Pod network range, usually 10.100.0.0/16
- All IP addresses, such as 0.0.0.0/0 for ports TCP 80 and 443.
Nutanix recommends not to allow port 80.

The following is the list of outbound traffic allowlist:

- CVM IP addresses (including secondary/segmented)
- Prism Central IP address
- Pod network range, usually 10.100.0.0/16
- All IP addresses such as 0.0.0.0/0 for port TCP 53, 7100, 5553 and UDP 53, 123

Nutanix Objects Bucket Settings

The following are the instructions for hardening Nutanix Objects bucket settings:

- Disable Hypertext Transfer Protocol (HTTP) port on the network firewall to restrict Nutanix Objects S3 endpoint access over HTTP protocol.
- Enable Write-Once-Read-Many (WORM) for buckets with sensitive data.

You can create WORM buckets to prevent anyone (including an administrator) from modifying or deleting data while the policy is active. For more information, see [WORM Bucket](#).

- Enable Nutanix Objects versioning.

When object versioning is enabled, uploading new data to the same object creates new versions. For more information, see [Creating and Configuring an S3 Bucket in Nutanix Objects](#) on page 99.

In order to prevent storage consumption from growing unbounded, set up a lifecycle policy to expire old versions and markers with an appropriate interval. For more information, see [Rules for Expiration based Lifecycle Policy](#) on page 107.

- Enable Data-at-Rest Encryption with Native Key Management.

Nutanix Objects provides a FIPS 140-2 compliant data-at-rest encryption solution. To deliver this capability, Nutanix Objects use the underlying AOS encryption capability. You can set encryption at an entire cluster level, always encrypting all data. For more information, see [Data-at-Rest Encryption](#) in the *Nutanix Security Guide*.

- Enable Identity and Access Management.

Review and delete unused or extra access keys for buckets.

- Configure replication.

Configure replication on buckets to keep at least one copy as backup. For more information, see [Creating Replication Rules between Buckets](#).

- Set up Secure Sockets Layer (SSL) Certificate.

Upload trusted Certificate Authority certificate to Nutanix Objects cluster to ensure secure communication between Nutanix Objects S3 endpoint and end users. For more information, see [Setting up SSL Certificate for an Object Store](#).

- On Prism Central, the Nutanix Objects user and Prism Central admin user is a root or administrator as they can perform all the operations on the buckets as well as generate new API keys.
- Prism Central supports role-based access control (RBAC) to configure and provide customized access to the users based on their assigned roles for the objects entities.
- Configure RBAC policies to provide Nutanix Objects access to Prism Central users and use those user credentials to perform operations on Nutanix Objects from Prism Central. For more information, see [Role-based Access Control for Nutanix Objects](#).
- Configure notifications.
You can send the completed events logs to the configured endpoints in your Nutanix Objects instance. For more information, see [Configuring Events Notification](#).
- Ensure that the tiering endpoint credentials are protected and not shared anywhere.

Remember: Never save your credentials in plain text form on your computer. Securely store your credentials in encrypted form by using credential storage applications.

Protecting Nutanix Objects Network with Flow Networking (Microsegmentation)

This section describes steps to protect Nutanix Objects Network with Flow Networking.

About this task

To protect the Nutanix Objects Network with Flow Networking, follow these steps:

Procedure

1. Enable Flow Networking microsegmentation.
For more information on enabling Flow Networking microsegmentation, see [Enabling Microsegmentation](#) in the *Flow Networking Microsegmentation Guide*.
2. Update the **AppType** category to **Objects_security**.
For more information, see [Modifying a Category](#) in the *Prism Self Service Administration Guide*.
3. Assign the newly created category to Nutanix Objects cluster VMs.
For more information, see [Assigning a Category](#) in the *Prism Self Service Administration Guide*.

4. Create an Application Security Policy on *AppType:Objects_security* with a relatable name like *Objects_security_policy*, and add the following data to **Inbounds** and **Outbounds**.
For more information, see [Application Security Policy Configuration](#) in the *Flow Networking Microsegmentation Guide*.

The following is the list of inbound traffic allowlist::

- CVM IP addresses (including secondary/segmented)
- Prism Central IP address
- Pod network range, usually 10.100.0.0/16
- All IP addresses such as 0.0.0.0/0 for port TCP 443.

The following is the list of outbound traffic allowlist::

- CVM IP addresses (including secondary/segmented)
- Prism Central IP address
- Pod network range, usually 10.100.0.0/16
- All IP addresses such as 0.0.0.0/0 for port TCP 53, 7100, 5553 and UDP 53, 123

NUTANIX OBJECTS PRISM CENTRAL DISASTER RECOVERY

Nutanix Objects includes built-in disaster recovery capabilities to recover automatically without any manual intervention.

This capability improves resilience and reduce manual intervention during Prism Central disaster recovery (PCDR).

Note: This topic is applicable for Prism Central version pc.7.3 or later.

Nutanix Objects Manager performs the following self-recovery operations after a PCDR event is triggered:

1. The Nutanix Objects Manager automatically runs a series of system prechecks.

Note: The Nutanix Objects user interface access is temporarily disabled during the precheck and recovery stages, with restoration expected in under an hour.

2. If the prechecks pass, the system automatically recovers without manual intervention. If the prechecks fail, the system does not trigger automatic recovery, and manual support is required.

If prechecks fail, contact Nutanix Support to restore Nutanix Objects functionality.

3. The Nutanix Objects user interface reflects the current recovery status, including any failure notifications.

Note: Data remains secure; however, users may intermittently lose access during recovery progress or failure.

NUTANIX OBJECTS LIFE CYCLE MANAGER UPGRADES

Nutanix Objects is upgraded using the Life Cycle Manager (LCM) feature through Prism Central.

As part of the Prism Central upgrades module, LCM handles updates for key components including Nutanix Objects Manager and Nutanix Objects Services.

Upgrades occur in the following order:

1. Prism Central
2. Prism Element
3. MSP Controller. For more information, see [Microservices Platform](#) on page 26.
4. Nutanix Objects Manager
5. Nutanix Objects Service

If Nutanix Objects is not enabled, upgrading the Prism Central instance will automatically upgrade the Nutanix Objects Manager. However, if Nutanix Objects is enabled, Nutanix Objects Manager must be upgraded manually.

For more information, see [Performing Firmware and Software Updates in a Connected Site Setup from Prism Central](#) in the *Life Cycle Manager Guide*.

Nutanix Objects Manager

Nutanix Objects Manager is a containerized service that runs on a Prism Central VM. It handles user input for object store deployment, validates inputs, manages certificates, deploys Nutanix Objects Services, and acts as an interface between Prism Central and object store backups.

A single Nutanix Objects Manager can manage one or more object stores. In a scale-out Prism Central VM, the service runs on each Prism Central node to provide high availability.

You can upgrade the Nutanix Objects Manager after MSP is upgraded. When multiple versions of Nutanix Objects Manager are available, always select the latest version. You can upgrade MSP, Nutanix Objects Manager, and Nutanix Objects Service together. When multiple upgrade options are selected, the upgrade happens serially, not in parallel.

Note:

- LCM disables the Nutanix Objects Manager upgrade if there is any unresolved Nutanix Objects Service upgrade failure.
- During Nutanix Objects Manager upgrade, input and output operations continue without disruption. However, the user interface becomes temporarily unavailable for statistics and management.

For more information, see [Performing Firmware and Software Updates in a Connected Site Setup from Prism Central](#) in the *Life Cycle Manager Guide*.

Nutanix Objects Service

Nutanix Objects Service provides the object store interface and handles storing and retrieving objects. It stores objects and metadata on selected Prism Element clusters. The service components are containerized and run on the Kubernetes platform, with each instance offering a single global namespace.

During deployment, the system creates the required VMs to run Kubernetes pods and the load balancer.

You can upgrade the Nutanix Objects Service after upgrading the Nutanix Objects Manager. Multiple Nutanix Objects Service instances appear separately in the upgrade list. You can upgrade them individually

or together. You can upgrade MSP, Nutanix Objects Manager, and Nutanix Objects Service together. When multiple upgrade options are selected, the upgrade happens serially, not in parallel.

Note:

- The Nutanix Objects Service becomes available for upgrade only after deployment.
- During Nutanix Objects Service upgrade, input and output operations experience disruption as internal services upgrade. The upgrade process takes around 15 to 30 minutes.

For more information, see [Performing Firmware and Software Updates in a Connected Site Setup from Prism Central](#) in the *Life Cycle Manager Guide*.

Installing Nutanix Objects (First-time Users)

Install Nutanix Objects by upgrading Prism Central through Life Cycle Manager (LCM), enabling Nutanix Objects, and updating to the latest version through the LCM interface.

About this task

To install and update Nutanix Objects in your environment, follow these steps:

Procedure

1. Perform an inventory and upgrade to a compatible version of Prism Central by using the LCM feature in Prism Central.

For more information, see [Performing Inventory with the Life Cycle Manager](#) section in the *Life Cycle Manager Guide*.

2. Enable Nutanix Objects.

For more information, see [Enabling Nutanix Objects](#) on page 17.

3. In LCM, perform inventory and update Nutanix Objects to the latest available version.

For more information, see [Nutanix Objects Life Cycle Manager Upgrades](#) on page 25.

The latest Nutanix Objects version is now installed.

Microservices Platform

Microservices Platform (MSP) is a Kubernetes-based platform that runs all Nutanix Objects microservices.

The MSP Controller, a service on the Prism Central VM, manages the MSP lifecycle. You must upgrade the MSP Controller before upgrading Nutanix Objects Manager and Nutanix Objects Service. The first object-store cluster deployed on Prism Central acts as the primary MSP or primary object store. All other clusters are secondary MSP clusters or secondary object stores.

For Nutanix Objects deployments, always upgrade to the latest MSP Controller version. Upgrade from MSP 1.0.8 to MSP 2.2.1 version to make Nutanix Objects visible for upgrade in dark site mode. For more information, see [Performing Inventory with the Life Cycle Manager](#) section in the *Life Cycle Manager Guide*.

Note: The primary object store cluster hosts all the common services, such as Identity and Access Management (IAM). You cannot delete the primary cluster without first deleting all secondary clusters. This note is applicable only when the IAM-high availability feature is not enabled. For more information, see the [Nutanix Objects IAM-High Availability Overview](#) on page 90 section.

Finding the Microservices Platform Version

Find the Microservices Platform (MSP) version by using the command line interface (CLI).

About this task

To find the MSP version, follow these steps:

Procedure

1. Log in to Prism Central VM.
2. Find the MSP Controller version:

```
admin@pcvm$ mspctl controller version
```

Note: If the `mspctl controller version` command does not work, use `mspctl controller_version` command since you might be using an older version of MSP (2.2.0 or earlier).

The system displays the MSP Controller version.

Finding the Primary and Secondary Microservices Platform Clusters

Find the primary and secondary Microservices Platform (MSP) clusters by using the command line interface (CLI).

About this task

To find the primary and secondary MSP clusters, follow these steps:

Procedure

1. Log on to Prism Central VM.
2. List the clusters:

```
admin@pcvm$ mspctl cluster list
```

If the `cluster_type` value is `primary_msp`, the cluster is a primary MSP cluster. All other clusters are secondary.

Note: If the `mspctl cluster list` command does not work, use `mspctl cluster_list` command since you might be using an older version of MSP (2.2.0 or earlier).

A list of cluster appears.

ROLE-BASED ACCESS CONTROL FOR NUTANIX OBJECTS

Prism Central supports role-based access control (RBAC) that allows you to configure and provide customized access to the users based on their assigned roles for the objects entities.

The access policies and permissions are user-based and apply only to the actions performed on the Prism Central web console. These policies and permissions do not apply to the actions performed on Nutanix Objects Browser or any S3-compatible client application, even if the same users have Prism Central access and Nutanix Objects Browser access.

From the Prism Central roles dashboard, you can define and assign the following additional roles to users or user groups:

- A set of predefined built-in roles (system roles)
- Additional custom roles

Nutanix Objects support the following built-in roles (system roles) that are defined by default:

- Nutanix Objects Admin—permissions required to create and manage object stores.
- Nutanix Objects Viewer—permissions required to perform read-only operations on object stores.

Note: For more information on Role Based Access Control, see the following references:

- *Controlling User Access (RBAC)* in the *Security Guide*.
- *Built-in Role Management* in the *Security Guide*
- *Operations Management* in the *Prism Central Guide*.

Prerequisites for Configuring Role-based Access Control for Nutanix Objects

Make sure that your system meets the following requirements before you start configuring role-based access control for Nutanix Objects.

- Prism Central version must be pc.2021.9 or later.
- Nutanix Objects Manager version must be 3.4 or later.
- Nutanix Objects Service version must be 3.4 or later (recommended version).
- Microservices Infrastructure must be enabled in Prism Central.

For information on enabling Microservices Infrastructure, see [Enabling Microservices Infrastructure Manually](#) in the *Prism Central Infrastructure Guide*. For more information on upgrading Nutanix Objects Manager and Nutanix Objects Service, see [Nutanix Objects Life Cycle Manager Upgrades](#) on page 25.

Only administrators (Super Admin or a Prism Admin in Prism Central) can create roles and access control policies for objects.

For more information on the built-in roles in Prism Central, see [Built-in Role Management](#) in the *Nutanix Security Guide*.

Note: To view the Prism Central version, log on to the Prism Central web console. In the Application Switcher, click **Infrastructure**. In the left navigation bar, select **Prism Central Settings > Prism Central Management**. The Prism Central version is listed in the **Prism Central Summary** section.

Role-Based Access Control Workflows for Nutanix Objects

Role-based access control (RBAC) allows users to configure and provide different levels of access to the users based on the role of the user within an organization.

You can create the following roles and assign them to the users based on your requirements:

- Full administrator
- View-only administrator
- Buckets administrator
- Infra administrator
- Custom

Note: For more information on role-based access control, see the following references:

- [Built-in Role Management](#) in the *Nutanix Security Guide*.
- [Controlling User Access \(RBAC\)](#) in the *Nutanix Security Guide*.

Nutanix Objects Permissions

Administrators can use the Prism Central RBAC feature to create access control policies and attach them to Nutanix Objects users.

A non-admin user can perform object-related tasks only after an administrator creates an access control policy on a role in Prism Central for the non-admin user. The administrators and non-admin users are Prism Central users with the following roles:

- Administrator—A Super Admin or a Prism Admin in Prism Central.
- Non-admin user—A Prism Central user without any administrator privileges.

Note: For more information on the built-in roles in Prism Central, see [Built-in Role Management](#) in the *Nutanix Security Guide*

The following table lists the permissions that an administrator can grant to a non-admin user to perform various objects tasks. These permissions allow the users to perform the following tasks:

Table 5: Nutanix Objects Permissions

Permission	Tasks Users Can Perform
Add Multicluster	Add Prism Element clusters to expand the storage.
Create Endpoint	Create remote endpoints for tiering.
Create Object Store	Create object stores.
Create Quota Policy	Create quota policies for objects.
Delete Object Store	Remove object stores.

Permission	Tasks Users Can Perform
Delete Quota Policy	Remove the quota policies created by the respective user for objects.
Download Certificate	Download the CA-signed certificates for an object store.
Edit Buckets	Edit the buckets in an object store.
Regenerate Self Signed Certificate	Regenerate self-signed certificates for an object store.
Replace Certificates	Replace certificates for an object store.
Scale Out	Add more resources to an existing object store cluster.
Set Capacity	Add additional storage capacity to an existing object store cluster.
Set Notification Endpoint	Add the notification endpoints.
Update Endpoint	Update the existing tiering endpoints.
Update Multicenter	Remove or modify the existing expanded storage (Prism Element cluster).
Update Quota Policy	Update the quota policies.
View Buckets	View the buckets in an object store.
View Endpoint	View the existing tiering endpoints.
View Multicenter	View the Prism Element clusters added for expanded storage.
View Notification Endpoint	View the notification endpoints.
View Object Store	View the object stores.
View Quota Policy	View the quota policies.

Creating Full Administrator Role for Nutanix Objects

You can create a full access administrator policy for users or user groups and grant full access rights to all operations on Nutanix Objects instances using role-based access control (RBAC) from the Prism Central web console.

About this task

Note:

- Only users with administrator privileges create a full access policy for Nutanix Objects instances.
- It might take some time (around five minutes or so) for the changes made to the policy to be effective.

To create a full access policy for Nutanix Objects instances, follow these steps:

Procedure

1. Log on to the Prism Central web console as an Admin user.

2. In the Application Switcher, select **Admin Center > IAM > Roles**.
3. Select the Nutanix Objects Admin role, and then click **Manage Assignment**.
The **Role Assignment** page appears.
4. Click **Add New** to create a new role.
5. From the **Select Users or Groups** section, select the required Active Directory group from the pull-down list, and then select one or more users or user groups using the **Search User** field.

Note: You can only add users who are part of an Active Directory group.

6. From the **Entities** section, select object store as the **Entity Type**, and then select the following:
 - Individual Entity
 - Select **All Nutanix Objects** or enter the name of the Nutanix Objects instance on which you need to grant custom permissions for the user.
7. Click **Save**.
An administrator user with full access is created.

Creating View-only Administrator Role for Nutanix Objects

You can create a view-only administrator role for users or user groups and grant complete access to read-only operations on Nutanix Objects instances using role-based access control (RBAC) from the Prism Central web console.

About this task

Note: Only users with administrator privileges can do this task.

To create a view-only admin access policy for Nutanix Objects instances, follow these steps:

Procedure

1. Log on to the Prism Central web console as an Admin user.
2. In the Application Switcher, select **Admin Center > IAM > Roles**.
3. Select the Nutanix Objects Viewer role, and then click **Manage Assignment**.
The **Role Assignment** page appears.
4. Click **Add New** to create a new role.
5. From the **Select Users or Groups** section, select the required Active Directory group from the pull-down list, and then select a user or user groups using the **Search User** field.

Note: You can add only those users who are part of a Windows Active Directory group.

6. From the **Entities** section, select object store as **Entity Type**, and then select the following:
 - **Individual Entity**
 - Select **All Nutanix Objects** or enter the name of the objects instance on which you need to grant read-only access to the user.
7. Click **Save**.
A view-only administrator access control policy for one or more users or user groups is created.

Creating Bucket Administrator Role for Nutanix Objects

You can create a bucket administrator policy to manage all bucket activities on Nutanix Objects instances using role-based access control (RBAC) from the Prism Central web console.

About this task

Note: Users with only administrator privileges can do this task.

To create a buckets admin policy, follow these steps:

Procedure

1. Log on to the Prism Central web console as an Admin user.
2. In the Application Switcher, select **Admin Center > IAM > Roles**.
3. Click **Create Role**.
The **Roles** page appears.
4. Enter the following details on the **Roles** page:
 - a. **Role Name:** Enter a name for the role.
 - b. **Description:** Enter a brief description of the role.
5. From the **Object Store** section, click **Change** next to the **Set Custom Permissions** option.
The **Custom Object Store Permissions** window appears.
(Optional) You can also enter *Object Store* in **Filter Entities** to go to the object store section.
6. Select the following permissions from the list and click **Save**:
 - Edit Buckets
 - View Buckets
 - View Object Store
 - When you select certain permissions, other additional permissions might automatically get added to the permission list. Read the on-screen instructions to understand the additional permissions that are automatically granted to the new custom role.
 - Users with Edit Bucket permission can perform bucket operations, such as create, update, and delete on the buckets available in the object stores that are assigned to the respective users.
7. Click the newly created role, and click **Manage Assignment**.
The **Role Assignment** page appears.
8. Click **Add New**.
You can assign the role to more users or user groups.
9. From the **Select Users or Groups** section, select the Windows AD group, and then select the user or user groups using the **Search User** field.
 - You can add only those users who are part of an Active Directory group.
 - When permission is granted to a role, the specific permission is not always inherited by the access control policy that is created for that role. This role needs to be assigned to a specific Nutanix Objects instance.

10. From the **Entities** section, select object store as the **Entity Type**, and then select the following:
 - Individual Entity
 - Select All Nutanix Objects or enter the name of the Nutanix Objects instance on which you need to grant custom permissions for the user.
11. Click **Save**.

Some bucket actions, such as setting a lifecycle policy for tiering, require additional permissions. For more information on the required permissions, see [Tiering](#).

The Buckets Admin access control policy is saved and listed on the Roles page. The user now has access to perform bucket management tasks (create, update, and delete buckets) on all buckets in that Nutanix Objects instance.

Creating Infra Administrator Role for Nutanix Objects

An administrator can provide access to users or to user groups to manage all the existing Nutanix Objects instances without providing any write access to buckets on the objects instances from the Prism Central web console.

About this task

Note: Users with only administrator privileges can perform this task.

To create an infra administrator access control policy, follow these steps:

Procedure

1. Log on to the Prism Central web console as an Admin user.
2. In the Application Switcher, select **Admin Center > IAM > Roles**.
3. Click **Create Role**.

The **Roles** page appears.
4. Enter the following details on the **Roles** page:
 - a. **Role Name:** Enter a name for the role.
 - b. **Description:** Enter a brief description of the role.
5. From the object Store section, select **Set Custom Permissions** and click **Change**.

The **Custom Object Store Permissions** window appears.

(Optional) You can also enter *Object Store* in **Filter Entities** to go to the **Object Store** section.
6. Select the following permissions from the list and click **Save**:

Add Multicloud, Create Quota Policy, Delete Quota Policy, Download Certificate, Regenerate Self Signed Certificate, Replace Certificates, Scale Out, Set Capacity, Set Notification Endpoint, Update Multicloud, Update Quota Policy, View Buckets, Edit Buckets, View Multicloud, View Notification Endpoint, View Object Store, and View Quota Policy.

When you select certain permissions, other additional permissions might automatically get added to the permission list. Read the on-screen instructions to understand the additional permissions that are automatically granted to the new custom role.
7. Click the newly created role, and click **Manage Assignment**.

The **Role Assignment** page appears.

8. Click **Add New** to assign the role to users or user groups.
9. From the **Select Users or Groups** section, select the Windows AD group, and then select the users or user groups using the **Search User** field.
 - You can add only those users who are part of an Active Directory group.
 - When permission is granted to a role, the specific permission is not always inherited by the access control policy that is created with that role. This role needs to be assigned to a specific Nutanix Objects instance.
10. From the **Entities** section, select object store as the **Entity Type**, and then select the following:
 - Individual Entity
 - Select **All Nutanix Objects** or enter the name of the objects instance on which you need to grant custom permissions for the user.

When permission is granted to a role, the specific permission is not always inherited from the access control policy that is created with that role. This role needs to be assigned to a specific Nutanix Objects instance.
11. Click **Save**.

The Infra Administration access control policy is saved and listed on the **Roles** page. The user now has access to manage existing Nutanix Objects instances. However, the user cannot perform any bucket update operations, create any objects instances, or delete any objects instances.

Creating a Custom Role for Nutanix Objects

An administrator can create custom roles and provide access to users or user groups from the Prism Central web console.

About this task

To create a custom access control policy (custom role), follow these steps:

Procedure

1. Log on to the Prism Central web console as an Admin user.

Only users with administrator privileges can do this task.

For information on creating a custom role, see [Creating a Custom Role](#) in the *Security Guide*.
2. In the Application Switcher, select **Admin Center > IAM > Roles**.
3. Click **Create Role**.

The **Roles** page appears.
4. Enter the following details on the **Roles** page:
 - a. **Role Name**: Enter a name for the role.
 - b. **Description**: Enter a brief description of the role.
5. From the Object Store section, select **Set Custom Permissions** and click **Change**.

The **Custom Object Store Permissions** window appears.

You can also enter *Object Store* in **Filter Entities** to go to the **Object Store** section.

6. Select the required permissions and click **Save**.

When you select certain permissions, other additional permissions might automatically get added to the permission list. Read the on-screen instructions to understand the additional permissions that are automatically granted to the new custom role.

7. Select the newly created role, and then click **Manage Assignment**.

The **Role Assignment** page appears.

8. Click **Add New** to assign the role to users or user groups.

9. From the **Select Users or Groups** section, select the required Active Directory group from the pull-down list, and then select one or more users or user groups using the **Search User** field.

- You can add only those users who are part of an Active Directory group.
- When permission is granted to a role, the specific permission is not always inherited by the access control policy that is created with that role. This role needs to be assigned to a specific Nutanix Objects instance.

10. From the **Entities** section, select object store as the **Entity Type**, and then select the following:

- **Individual Entity**
- Select **All Nutanix Objects** or enter the name of the Nutanix Objects instance on which you need to grant custom permissions for the user.

When permission is granted to a role, the specific permission is not always inherited from the access control policy that is created with that role. This role needs to be assigned to a specific Nutanix Objects instance.

11. Click **Save**.

The custom access control policy is saved and listed in the **Roles** page. The user now has access to manage existing Nutanix Objects instances granted to him in the policy.

NUTANIX OBJECTS PREREQUISITES AND LIMITATIONS

Ensure that you have met the deployment prerequisites, port requirements, and configured the network for AHV or VMware ESXi before proceeding with an object store deployment. This section also lists the Nutanix Objects and Network File System (NFS) limitations.

Nutanix Objects Deployment Prerequisites - AHV and ESXi

Before deploying Nutanix Objects on AHV or VMware ESXi, carefully review this section to ensure all prerequisites are met. The requirements listed apply to both online (internet-accessible) and offline (dark site) deployments unless specified otherwise.

General Requirements

Ensure that your environment conforms to the following requirements before running Nutanix Objects:

- The hypervisor must be either AHV or ESXi. ESXi is supported only with Nutanix Objects 3.0 or later versions.
- The Prism Element version must be 5.11.2 and later, and the Prism Central version must be 5.17.1 and later.
- The recommended browser is Google Chrome.
- A cluster running AHV or ESXi must have at least one node present.

Nutanix Objects utilizes up to 12 vCPUs for each AHV or ESXi node, which includes ten worker nodes and two load balancers. For dense all-flash nodes, Nutanix recommends using 20 vCPUs consisting of 16 worker nodes and four load balancers. If you need assistance with configuring all-flash nodes, please contact Nutanix Support.

- Make sure that no AHV, ESXi host, Prism Element, or Prism Central upgrades are in progress while deploying Nutanix Objects.
- Ensure the object store domain is specifically for the object store deployment.

For example, if the main domain is `mycompany.com`, the object store domain can be a subdomain such as `objectstore.mycompany.com`.

- Ensure the proxy can access the guest VM. Nutanix recommends enabling Pulse for sites that have Internet access.
- Nutanix recommends having a reliable, high-speed internet connection for online installations. Image downloads timeout after 90 minutes, resulting in a failed deployment.
- Nutanix recommends upgrading to the latest version of the Microservices Platform (MSP) Controller for deployment in a dark site.

For more information on upgrading MSP Controller, see [Microservices Platform](#) on page 26.

- Ensure that the Prism Element can access the Life Cycle Manager (LCM) web server where Nutanix Objects is deployed at a dark site.
- Ensure the LCM web server is accessible through the proxy if it is configured on Prism Central for dark site deployment.
- Allow Prism Central and Prism Element to access the web server through port 80 for dark site deployment.

- To expand the object store cluster, upgrade the MSP version to 2.0 or later and the Nutanix Objects version to 3.0 or later.
- To scale the object store cluster, ensure that you upgrade to version 2.1 or later.

Network Requirements

Ensure to configure the following network requirements before running Nutanix Objects:

- Ensure you meet the URL and port requirements. For more information, see [Port Requirements for Nutanix Objects](#) on page 38.
- Configure Domain Name Servers (DNS) on both Prism Element and Prism Central.
- Configure Network Time Protocol (NTP) servers on both Prism Element and Prism Central.
- Set up the Virtual IP address and data services IP address on the Prism Element for deploying Nutanix Objects. Also, configure the DSIP on the Prism Element cluster where Prism Central is deployed.
- Ensure that the communication between Nutanix Objects (Storage or Public network) and Prism Central must be direct and not through a proxy server.
- **AHV:** Ensure that the Virtual Local Area Networks (VLANs) required for internal Nutanix Objects services and external access to the object store endpoints are configured correctly on Prism Element. For more information on the guidelines, see [Nutanix Objects Network Configuration](#) on page 39.
- **ESXi:** For more information, see [ESXi Prerequisites](#).
- Ensure that both Prism Element and Prism Central have Internet connectivity for online deployment. If Internet access is not available, deploy the object store on a dark site. For more information, see [Deploying Object Store at a Dark Site \(Offline Deployment\)](#) on page 86.
- Basic VLAN, Open Virtual Network (OVN) VLAN, and virtual private cloud (VPC) subnets are supported.

Note: Network segmentation only supports basic VLAN.

- The best practice is to keep Nutanix Objects Storage and Nutanix Objects Public network on separate networks.

Maximum Configurations

For more information on Nutanix Objects configuration limits, see [Nutanix Configurations Maximums](#).

ESXi Prerequisites for Nutanix Objects Deployment

Before deploying Nutanix Objects on VMware ESXi, review this section to ensure that you meet all prerequisites.

Unless specified otherwise, the requirements listed apply to both online (internet-accessible) and offline (dark site) deployments.

Network Prerequisites

Ensure to configure the following network requirements before running Nutanix Objects:

- All ESXi hosts must have access to the network designated for Nutanix Objects deployment.
 - When you set up the network using a standard vSwitch, all hosts must have the same network name and VLAN configuration.
 - When you configure the network using the VMware Distributed Virtual Switch (DVS), all hosts must be part of it.
- Add nodes only from a single Prism Element cluster to an ESXi cluster in VMware vCenter server.
- Configure the VMware NSX network on all hosts in both the primary and secondary VMware vSphere Metro Storage Clusters (vMSC).
- Meet the Nutanix Objects ESXi IP Address Management (IPAM) requirements. For more information, see [ESXi Configuration for Nutanix Objects Deployment](#) on page 44.

vCenter and ESXi Prerequisites

Make sure that the ESXi cluster meets the following requirements:

- Register ESXi clusters with a vCenter server.
- Enable VMware Distributed Resource Scheduler (DRS) and high availability solutions on the cluster. For more information, see *Create a Cluster* in the *VMware vSphere Guide*.
- Map the AOS cluster that is used for deploying Nutanix Objects to a single ESXi cluster.
- The ESXi hypervisor must be version 6.5 or later.
- The Nutanix Objects version must be 3.0 or later.
- The MSP version must be 2.0 or later.
- The VMware NSX-T version must be 2.4 or later. For more information, see KB [8545](#).

Prism Central VM Prerequisites

Ensure that the Prism Central VM meets the following requirements:

- The Prism Central VM is hosted and registered on the Prism Element cluster. Volume must be provisioned from the hosted Prism Element cluster to the MSP controller for image conversion.

Note: Prism Central deployments in a non-Nutanix ESXi environment do not support Nutanix Objects.

- Allow the proxy configuration to include the vCenter server IP address, and connect the vCenter server to Prism Central VM.

Worker VMs, Prism Central VM, and Prism Element VM must connect directly to vCenter without using a proxy server.

Port Requirements for Nutanix Objects

Port requirements for Nutanix Objects instance deployment.

For more information on the required ports and the port diagram for your Nutanix Objects deployment, see [Ports and Protocols Guide](#) for Nutanix Objects.

Note:

- URLs are not required for the dark site deployment.

- Internal network IP ranges, such as 10.100.x.x and 10.200.x.x, might appear in firewall logs when traffic from the inter-Microservices Platform (MSP) routes through the firewall.

Nutanix Objects Network Configuration

The Nutanix Objects architecture uses the following virtual networks:

- **Nutanix Objects Storage Network:** Helps internal communication among various VMs and nodes of the object store, handling both data and management traffic when Nutanix Objects network segmentation is disabled, and only data traffic when it is enabled.
- **Nutanix Objects Public Network:** Allows external clients to access the object store.
- **Nutanix Objects Management Network:** Specifically isolates management traffic. It handles all communication from the Prism Central instance to the Microservices Platform (MSP) virtual machine, ensuring that this traffic remains securely segmented on its own dedicated network. The traffic destined to the management network of the Prism Element also originates from this network.

Note: The Nutanix Objects management network is used only when network segmentation is enabled for AHV clusters. This network option is not available for ESXi clusters.




Figure 1: Nutanix Objects Network Architecture - Using Two Separate Networks (Storage and Public)

Tip:

Key recommendations for Nutanix Objects networks:

- You can have two virtual networks, each for Nutanix Objects storage network and Nutanix Objects public network, but it is not mandatory. You can have the Nutanix Objects storage network and the Nutanix Objects public network on the same virtual network. However, Nutanix recommends configuring the Nutanix Objects storage network and the Nutanix Objects public network on different virtual networks for production deployments.
- Nutanix recommends that you configure the Nutanix Objects storage network to use the same network as the Controller VM (CVM) or the hypervisor. A single network enables the traffic between Nutanix Objects and CVM to flow within the same network, avoiding cross-network hops that can limit performance in some deployments. The traffic between Nutanix Objects

and the CVM is influenced by the capabilities of the underlying AOS and is significant in a dedicated objects deployment.

- When enabling network segmentation, Nutanix recommends using the same network for management plane communication as Prism Element.
- To use different networks for the Nutanix Objects storage network and CVM network, ensure that the network bandwidth between the top-of-rack switch and the L3 device is fast enough to avoid network congestion. Alternatively, you can enable L3 functionality on the top-of-rack switch.

Link Aggregation Control Protocol and Link Aggregation

Nutanix recommends that you configure link aggregation with link aggregation control protocol (LACP) and balance-transmission control protocol (TCP). Nutanix Objects is a network-heavy workload and requires a high network bandwidth. You can achieve a faster network by aggregating multiple physical links into one logical network interface that can be used for all traffic.

You require LACP and link aggregation to fully use the bandwidth from multiple links. In Open vSwitch (OVS), dynamic link aggregation is achieved through LACP, and load balancing is implemented using balance-TCP.

For more information on how to enable, turn off, and verify LACP in AHV, see [Enabling LACP and LAG \(AHV Only\)](#) in the *AHV Networking Guide*.

In ESXi-based Nutanix Objects deployments, you can also use the vSphere distributed switch with Load Based Teaming (LBT).

AHV Configuration for Nutanix Objects Deployment

Configure and manage virtual networks through the Prism Element web console and use these virtual networks to deploy an object store through Prism Central.

Nutanix Objects Storage Network Requirements

An object store uses the Nutanix Objects storage network, a private virtual network, to communicate between services. The number of IP addresses required for deploying your Nutanix Objects store instance varies according to the deployment size.

For more information, see [IP Address Consumption for an Object Store](#) on page 41. For more information on network configurations and enabling AHV IPAM, see [Network Configuration for VM Interfaces](#) in *Prism Web Console Guide*.

Requirements of a virtual network used for an object store:

- The virtual network must have AHV IPAM enabled.
- Sufficient IP addresses must be available in the IPAM dynamic host configuration protocol (DHCP) pool.
- Two static IP addresses outside the DHCP pool are required for each object-store. Later, while deploying the object store from the Prism Central UI, use the static IP addresses for the Nutanix Objects storage network configuration based on the object store storage and resource requirements.

For more information, see [Table 6: IP Address Consumption Based on Deployment Size for an Object Store Without Segmented Network](#) on page 42 and [Table 7: IP Address Consumption Based on Deployment Size for an Object Store With Segmented Network](#) on page 43.

Note: Nutanix Objects internal services use the 10.100.0.0/16 and 10.200.0.0/16 subnets. Ensure that you do not assign internal interface IP addresses from these subnets.

Nutanix Objects Public Network Requirements

An object store uses the Nutanix Objects public network, a virtual network, to allow access from external clients. For more information, see the [IP Address Consumption for an Object Store](#) on page 41 section.

Requirements of a virtual network to access an object store externally:

- AHV IPAM must be enabled, but no IP addresses are required in the IP pool.

While configuring the virtual network, ensure that you specify details such as **Virtual Local Area Network Identifier (VLAN ID)**, **Network IP Address/Prefix Length**, **Gateway IP address**, and **Domain Name System (DNS) Servers**.

- While deploying an object store from the Prism Central UI, use up to four static IP addresses for the Nutanix Objects public network configuration based on the object store storage and resource requirements.

Nutanix Objects Management Network Requirements

An object store uses the Nutanix Objects management network, a virtual network (with AHV IPAM enabled), to segment the network traffic and handle all management plane traffic. For more information, see the [IP Address Consumption for an Object Store](#) on page 41 section.

Requirements of a virtual network to segment the network of an object store:

- The VLAN for Nutanix Volumes service on a Prism Element cluster that has Nutanix Volumes network segmentation must be routable to the VLAN used for the Nutanix Objects storage network.
- When using the same VLAN for Nutanix Objects storage network and the Nutanix Volumes service on a Prism Element cluster that has Nutanix Volumes network segmentation, ensure that the range of IP addresses defined for the Nutanix Objects storage network does not overlap with the VLAN used for the Nutanix Volumes service.
- When deploying network segmentation enabled object store, the static IP addresses must be from the Nutanix Objects management network and not from Nutanix Objects storage network .
- When network segmentation is enabled, additional IP addresses are required for each worker VM and load balancer VM.

For more information, see [Table 6: IP Address Consumption Based on Deployment Size for an Object Store Without Segmented Network](#) on page 42 and [Table 7: IP Address Consumption Based on Deployment Size for an Object Store With Segmented Network](#) on page 43.

For more information on the prerequisites, see [Prerequisites for Nutanix Objects Network Segmentation of an Object Store](#) on page 66.

IP Address Consumption for an Object Store

Nutanix Objects requires the following IP addresses for its virtual networks:

Nutanix Objects Public Network

Nutanix Objects uses up to four static IP addresses to expose the Nutanix Objects Simple Storage Service (S3) endpoint, depending on the selected worker nodes. Clients connect to these IP addresses for all S3 requests. Each object store supports a maximum of four public IP addresses.

Nutanix Objects Storage Network

Ensure that you meet the following IP address requirements for Nutanix Objects storage network:

- **From DHCP pool:**
 - One IP address for each worker VM
 - One IP address for each load balancer to communicate with Nutanix Objects worker VMs
 - Two IP addresses for high availability of internal Microservices Platform (MSP) services
- **Outside DHCP pool (Static IP address):** Two static IP addresses for MSP DNS and API servers

Note: When network segmentation is enabled in an object store, static IP addresses (outside DHCP pool) from the Nutanix Objects storage network are not required.

Nutanix Objects Management Network

Ensure that you meet the following IP address requirements for Nutanix Objects management network:

- **From DHCP pool:**
 - One IP address for each worker VM
 - One IP address for each load balancer to communicate with Nutanix Objects worker VMs
- **Outside DHCP pool (Static IP address):** Two static IP addresses for internal MSP services

Table 6: IP Address Consumption Based on Deployment Size for an Object Store Without Segmented Network

Number of Worker Nodes	AHV Managed Networks					
	Nutanix Objects Storage Network					Nutanix Objects Public Network
	From DHCP Pool				Outside DHCP Pool	
	Number of IP Addresses Required for Worker VMs	Number of IP Addresses Required for the Load Balance VMs to Communicate with the Worker VMs	Number of IP Addresses for High Availability of Internal MSP Services	Number of IP Addresses Required for IPAM DHCP Server	Number of Static IP Addresses Required for MSP, DNS, and API Server	
1	1	1	2	1	2	1
2	2	1	2	1	2	1
3	3	2	2	1	2	2
4	4	3	2	1	2	3
5	5	4	2	1	2	4

Number of Worker Nodes	AHV Managed Networks					
	N	4 (without performance configuration) or N-1 (with performance configuration)	2	1	2	4 (without performance configuration) or N-1 (with performance configuration)

Table 7: IP Address Consumption Based on Deployment Size for an Object Store With Segmented Network

Number of Worker Nodes	AHV Managed Networks							
	Nutanix Objects Storage Network				Nutanix Objects Public Network	Nutanix Objects Management Network		
	From DHCP Pool					From DHCP Pool	Outside DHCP Pool	
	Number of IP Addresses Required for Worker VMs	Number of IP Addresses Required for the Load Balance VMs to Communicate with the Worker VMs	Number of IP Addresses for High Availability of Internal MSP Services	Number of IP Addresses Required for IPAM DHCP Server		Number of IP Addresses Required for Worker VMs	Number of IP Addresses Required for the Load Balance VMs to Communicate with the Worker VMs	Number of Static IP Addresses Required for MSP
1	1	1	2	1	1	1	1	2
2	2	1	2	1	1	2	1	2
3	3	2	2	1	2	3	2	2
4	4	3	2	1	3	4	3	2
5	5	4	2	1	4	5	4	2
N	N	4 (without performance configuration) or N-1 (with performance configuration)	2	1	4 (without performance configuration) or N-1 (with performance configuration)	N	4 (without performance configuration) or N-1 (with performance configuration)	2

ESXi Configuration for Nutanix Objects Deployment

For ESXi clusters, you can perform the Nutanix Objects ESXi IP Address Management (IPAM) configuration.

You can do the Nutanix Objects ESXi IPAM configuration from the **Nutanix Objects service > vCenter Management** available in the Prism Central web console. For more information, see [Managing vCenter for Nutanix Objects Service](#) on page 46.

Nutanix Objects requires you to add the IPAM range for the ESXi networks that you plan to use for object store deployment.

Nutanix Objects uses these ESXi networks for two purposes:

- To deploy Nutanix Objects VMs that host the various Nutanix Objects services (also referred to as **Nutanix Objects Storage Network**).
- To deploy load balancer VMs that provide object store endpoint to the Simple Storage Service (S3) clients (also referred to as **Nutanix Objects Public Network**).

Note: Nutanix Objects management network is used only when network segmentation is enabled for AHV clusters. This network option is not available for ESXi clusters.

The following section describes the Nutanix Objects Storage (internal) and Nutanix Objects Public (external) networks in more detail:

Nutanix Objects Storage Network

An object store uses this private virtual network to communicate between services. The number of static IP addresses required for deploying your Nutanix Objects store instance varies according to the deployment size. For more information, see the *IP Address Consumption* section.

Requirements of a virtual network used for an Object Store:

- Add the IPAM range for the ESXi networks in the **vCenter Management** page.
- Sufficient IP addresses must be available. For more information, see *IP Address Consumption Based on Deployment Size*.
- Subnet, Gateway, and DNS IP address to be used for the ESXi network. The provided values must be valid.

Note: Nutanix Objects uses the 10.100.0.0/16 and 10.200.0.0/16 subnets for internal services. Avoid subnet conflicts when assigning internal interface IP addresses.

Nutanix Objects Public Network

An object store uses this virtual network to allow access from external clients.

Requirements of a virtual network to access an object store externally:

Up to four static IP addresses that can either be part of or outside the IPAM range. Later, while deploying the object store from Prism Central, use the static IP addresses for the Nutanix Objects Public Network configuration based on the object store storage and resource requirements.

You can have two virtual networks, each for Nutanix Objects Storage Network and Nutanix Objects Public Network, but it is not mandatory. You can have the Nutanix Objects Storage Network and the Nutanix Objects Public Network on the same virtual network. However, Nutanix recommends to have the Nutanix Objects Storage Network and the Nutanix Objects Public Network on different virtual networks for production deployments.

Note:

- If the **Nutanix Objects Public Network** is different from the **Nutanix Objects Storage Network**, then only subnet and Gateway values are needed. IPAM range and DNS IP address values are optional.
- All the IP addresses might not be used during the deployment. The number of IP addresses used depends on the size of your deployment. The unused IP addresses are reserved for future usage.

You can view a list of deployed object stores and the general and networking details of the object stores. For example, you can view the Nutanix Objects Public IP addresses for your deployment. For more information, see [Viewing Object Store Deployments](#) on page 72.

IP Address Consumption

This section describes the IP address consumption based on the deployment size.

Nutanix Objects Public Network

Up to four static IP addresses are required to expose the Nutanix Objects S3 Endpoint based on the selected worker nodes. The client connects to these IP addresses for all S3 requests. Nutanix Objects supports a maximum of four public IP addresses for each object store. The IP addresses can be within or outside the range of the IPAM network.

Nutanix Objects Storage Network

The IP addresses are consumed from the IPAM range configured for the ESXi networks as follows:

- Each Worker VM requires one IP address.
- One IP address for each load balancer for communications with Nutanix Objects workers.
- Two IP addresses for High Availability of internal Microservices Platform (MSP) services.
- Two static IP addresses required for MSP Domain Name System (DNS) and Application Programming Interface (API) servers.

Table 8: IP Address Consumption Based on Deployment Size

Number of Worker Nodes	Nutanix Objects Storage Network					Nutanix Objects Public Network
	Nutanix Objects ESXi IPAM Range Configured for the ESXi Network					
	Number of IP addresses required for Worker VMs	Number of IP addresses required for the Load Balance VMs to communicate with the Worker VMs	Number of IP addresses for high availability of internal MSP services	Number of IP addresses required for IPAM DHCP Server	Number of static IP addresses required for MSP, DNS, API Server	Number of static IP addresses required for Load Balancers to communicate with external S3 clients
1	1	1	2	1	2	1
2	2	1	2	1	2	1

Number of Worker Nodes	Nutanix Objects Storage Network					Nutanix Objects Public Network
3	3	2	2	1	2	2
4	4	3	2	1	2	3
5	5	4	2	1	2	4
N	N	4 (without performance configuration) or N-1 (with performance configuration)	2	1	2	4 (without performance configuration) or N-1 (with performance configuration)

Managing vCenter for Nutanix Objects Service

To deploy Object Stores on ESXi clusters, you need to provide the vCenter credentials and configure the IP Address Management (IPAM) for ESXi networks.

About this task

For AHV clusters, IPAM is configured in the Prism Element. For ESXi clusters, you can perform the IPAM configuration in the Prism Central as described in this section.

Managing vCenter for object service consists of the following two steps:

- Add the vCenter IP address and login credentials in the Object service within Prism Central to create a trust relationship. Nutanix does not store the login credentials after the connection is established between the vCenter and Prism Central.
- Add IPAM for one or more ESXi networks.

When deploying the object store instance, you need to define the Nutanix Objects Storage Network and Object Public Network. You can use these pre-configured networks for Nutanix Objects internal and public access.

To register the vCenter in the Prism Central, follow these steps:

Procedure

1. Log on to the source Prism Central web console, and click the **Entity menu > Services > Objects**
2. Click **vCenter Management > Add vCenter**.




Figure 2: vCenter Management Page in Nutanix Objects

3. In the **vCenter** page, enter the IP address and login credentials of the vCenter.

The details you enter are used to generate a certificate and build a trust relationship between the vCenter and Prism Central.

Nutanix does not store the login credentials after the connection is established between the vCenter and Prism Central.

Note: For a user with non-administrator privileges to perform vCenter registration, the administrator must create an access control policy for the non-administrator user on a role that must have the following minimum permissions:

- Create Object Store
- View Object Store

4. Select one the following options:

» **Next:** Click this option to continue adding IPAM details.

» **Save & Close:** Click this option to add the IPAM details later.

In the **vCenter Management** page, you can click **Configure Network** in the **Actions** column to proceed with adding the IPAM for ESXi networks.

IPAM is added for the ESXi network.

5. Click **Add Network**.

The **Add Network** page appears.

6. In the **Add Network** page, follow these steps:

a. In the **Data Center**, select the datacenter where your ESXi cluster is located.

The **ESXi Network** drop-down will be populated with a list of supported ESXi networks belonging to the datacenter.

b. In the **ESXi Network**, select the ESXi network you want to use for deployment.

c. Enter the IP address range, subnet mask, Gateway IP address, and DNS IP address.

The IP address range you provide will be used for the Nutanix Objects Storage Network. The IP address range and DNS IP address are optional if the IPAM will be used only for the Nutanix Objects Public Network.

d. Click **Add** to complete your IPAM configuration.

If you want to add more networks, click **Add Network** on the **Configure** page and enter the details. It is recommended to use separate networks for the Nutanix Objects Storage Network and Nutanix Objects Public Network.

e. Click **Update**.

In the **vCenter Management** page, in the **Network** column, you can view the IPAM count.

In the **Actions** column, you can click the **Delete** option to delete the vCenter.

Even if you delete a vCenter, the IPAM details will still be available. If you add the deleted vCenter again, the IPAM details previously added get recovered.

What to do next

You can start with the deployment of the object store on the ESXi clusters. For more information, see [Object Store Service Deployment](#).

Shared Versus Single Network Configuration for Nutanix Objects

You can configure Nutanix Objects to use a single network for all services or separate the subnets used for internal communications and client communications.

You can have two virtual networks, each for Nutanix Objects Storage Network and Nutanix Objects Public Network, but it is not mandatory. You can have the Nutanix Objects Storage Network and the Nutanix Objects Public Network on the same virtual network. However, it is recommended to have the Nutanix Objects Storage Network and the Nutanix Objects Public Network on different virtual networks for production deployments.

Note: If you choose to deploy a network segmentation-enabled object store, you need to have the Nutanix Objects storage, public, and management networks on separate virtual networks; shared networks are not supported.




Figure 3: Nutanix Objects Network Architecture - Using Two Separate Networks (Storage and Public)




Figure 4: Nutanix Objects Network Architecture - Using One Network

Nutanix Objects Limitations

The following section lists the limitations for Nutanix Objects.

If the Nutanix Objects Public network is identical to the Prism Central VM network, and the Nutanix Objects Storage network is different, a firewall might block TCP handshakes between Prism Central and the Nutanix Objects Storage network. This might lead to connectivity issues.

To address this limitation, you must either ensure that the Nutanix Objects Public network is different from the Prism Central network or configure the Nutanix Objects Storage network to align with both the Prism Central and the Nutanix Objects Public network.

System Limitations

- Once an object store gets deployed, you cannot change the Data Services, Controller VM, Microservices Platform (MSP), and Prism Central IP addresses.
- Prism Central and Prism Element de-registration and reregistration is not supported.

Limitations of Network File System

This section lists the limitations of multiprotocol access.

- Network File System (NFS) adapter in Nutanix Objects does not support Windows NFS client.
- You can enable NFS access only at the time of bucket creation.
- NFS-enabled buckets are exposed as NFS shares, which the NFS clients can mount.
- You can perform updates (only multiprotocol access configurations), delete, and share actions on NFS buckets.

Note: You can delete NFS-enabled buckets only if they do not have any object and any directory explicitly created from the NFS protocol.

- You cannot enable other S3 bucket features, such as Lifecycle Policies, Versioning, Write once read many (WORM), Replication, Static Website, CORS, and Notifications on NFS-enabled buckets.
- You can only create symbolic links through NFS. Hard links are not supported.
- You can rename files and links; however, you cannot rename a directory. After renaming, the file handle of the renamed file changes.
- In addition to the Simple Storage Service (S3) object naming convention, the following are some cases that are not allowed for an object name in the NFS-enabled bucket:
 - Object name cannot be . or ..
 - Object name cannot start with ./ or ../
 - Object name cannot contain //, ./ or ../ pattern.
 - Object names cannot end with ./ or ../
 - If an object name contains a directory hierarchy, then each component in the object name separated by / cannot be more than 255 bytes in length.
- You cannot delete objects with / as a suffix or objects-browser-created folders from the NFS protocol, which appear as directories in the NFS namespace.
- If the objects are created from the S3 protocol in such a way that they end up creating a conflict in the object and directory name, then only the file is visible in the NFS namespace, and the directory is hidden. Only upon renaming or removing the file to some other name, the directory become visible again.

For example, suppose we create two objects dir1/dir2/file and dir1/dir2 from the S3 protocol. As we can see, dir2 is a directory as well as a file, so it creates a conflict. Now, while traversing this namespace

from the NFS protocol, once we are inside the `dir1/` directory, only the object `dir2` would be visible and not the subdirectory `dir2/`. So, Nutanix recommends not to use conflicting names for objects from the S3 protocol while creating them.

CONFIGURING VIRTUAL PRIVATE CLOUD MODES FOR NUTANIX OBJECTS WITHOUT NETWORK ADDRESS TRANSLATION

Nutanix Objects supports Virtual Private Clouds (VPCs) in underlay, mixed (underlay and overlay), and overlay modes for No-Network Address Translation (NAT).

Before you begin

You must meet the following requirements for configuring VPC modes with No-NAT for Nutanix Objects:

- Prism Central version must be pc.2024.3 and later.
- Prism Element version must be 6.10 and later.
- Nutanix Objects version must be 5.1 and later.
- MSP version must be 2.6.0.0 and later.
- AHV version must be 10.0 and later.

About this task

To configure VPC for Nutanix Objects with No-NAT, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. In the Application Switcher, select **Infrastructure**.
3. In the navigation bar, click **Network & Security > Subnets**.
The **Subnets** page opens the list of subnets.
4. Create an external subnet network with No-NAT.
For more information, see [Creating a Subnet](#) in the *Flow Virtual Networking Guide*.
5. Create an underlay network.
For more information, see [Creating a Subnet](#) in the *Flow Virtual Networking Guide*.
6. In the navigation bar, click **Network & Security > Virtual Private Clouds**.
The **Virtual Private Clouds** page opens the list of VPCs.
7. Click **Create VPC**.
For more information, see [Creating Virtual Private Cloud](#) in the *Flow Virtual Networking Guide*.
The **Create VPC** page appears.

8. Create VPC and associate the external subnet.
 - a. In **Name**, enter the VPC name.
 - b. In **External Connectivity**, click **Associate External Subnet**.
For more information, see the *External Connectivity* and *Associate External Subnet window details* rows in the table in [Creating Virtual Private Cloud](#) topic of the *Flow Virtual Networking Guide*.
 - c. In **Subnet Type**, select **VLAN**.
 - d. In **Static Routes**, check **Set this subnet as the default next hop for outbound traffic**.
 - e. In **Destination Prefix(es)**, enter 0.0.0.0/0.
 - f. In **SNAT IP/ Router IP**, select **IP Assignment Mode** as **Custom Defined**.
 - g. In **Custom SNAT IP / Router IP**, enter four IP addresses which are configured as next hop.
 - h. In **External Gateway Configuration**, enter 4 as **Number of Active Hosts**.
 - i. Click **Save**.
The external subnet is setup and the **Create VPC** page appears.
 - j. Enter the externally routable IP addresses.
Externally routable IP addresses is the flow VPC network.
 - k. In **Domain Name Servers (DNS)**, add the DNS server IP addresses.
For more information, see the *Other details on the Create VPC page* row in the table in [Creating Virtual Private Cloud](#) topic of the *Flow Virtual Networking Guide*.
 - l. Click **Create**.
The VPC is now created and associated with the external subnet.
9. On the **Virtual Private Clouds List** page, click the Flow Virtual Networking transit VPC name where you want to create an overlay network and click **Create Subnet**.
The **Create Subnet** page appears.
10. Configure the overlay network.
For more information, see [Creating a Subnet](#) in the *Flow Virtual Networking Guide*.
 - a. In **Name**, enter the name of the subnet.
 - b. In **Type**, confirm that the type is Overlay.
 - c. In **Network IP Address / Prefix**, enter the network IP address and prefix, and the gateway IP address.
This network IP address and prefix would be the same as the externally routable IP addresses provided in Step 8 j.
 - d. Click **Add IP Pool** and enter the first and last IP address of the range, then click the checkbox under **Actions**.
 - e. Expand and configure **Domain Settings**.
 - f. Click **Create**.
The external subnet, underlay, and overlay networks are created and listed in the Virtual Private Clouds list.
All three external subnet, underlay, and overlay networks are created and listed in the Virtual Private Clouds list.

What to do next

After creating all the networks, you can start the Nutanix Objects deployment. For more information, see [Object Store Service Deployment](#) on page 65. For more information on configuring VPC modes with NAT, see [Configuring Virtual Private Cloud Modes for Nutanix Objects with Network Address Translation](#) on page 54.

CONFIGURING VIRTUAL PRIVATE CLOUD MODES FOR NUTANIX OBJECTS WITH NETWORK ADDRESS TRANSLATION

Nutanix Objects supports Virtual Private Clouds (VPCs) in mixed mode for Network Address Translation (NAT).

Before you begin

You must meet the following requirements for configuring VPC modes with NAT for Nutanix Objects:

- Prism Central version must be pc.7.3 or later.
- Prism Element version must be 6.10 or later and compatible with Prism Central version pc.7.3.
- Nutanix Objects version must be 5.2 or later.
- MSP version must be 2.7.0.0 or later.
- AHV version must be 10.0 or later and compatible with Prism Central version pc.7.3.

About this task

Mixed mode includes Nutanix Objects storage network in underlay network and Nutanix Objects public network in overlay network.

To configure VPC for Nutanix Objects with NAT, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. From the [Application Switcher Function](#), select **Infrastructure**.
3. In the navigation bar, click **Network & Security > Subnets**.
The **Subnets** page opens the list of subnets.
4. Create an external subnet network with NAT.
For more information, see [Creating a Subnet](#) in the *Flow Virtual Networking Guide*.
5. Create an underlay network.
For more information, see [Creating a Subnet](#) in the *Flow Virtual Networking Guide*.
6. In the navigation bar, click **Network & Security > Virtual Private Clouds**.
The **Virtual Private Clouds** page opens the list of VPCs.
7. Click **Create VPC**.
For more information, see [Creating Virtual Private Cloud](#) in the *Flow Virtual Networking Guide*.
The **Create VPC** page appears.

8. Create VPC and associate the external subnet.
 - a. In **Name**, enter the VPC name.
 - b. In **External Connectivity**, click **Associate External Subnet**.
For more information, see the *External Connectivity* and *Associate External Subnet window details* rows in the table in [Creating Virtual Private Cloud](#) topic of the *Flow Virtual Networking Guide*.
 - c. In **Subnet Type**, select **VLAN**.
 - d. In **Static Routes**, check **Set this subnet as the default next hop for outbound traffic**.
 - e. In **Destination Prefix(es)**, enter 0.0.0.0/0.
 - f. In **SNAT IP/ Router IP**, select **IP Assignment Mode** as **Custom Defined**.
 - g. In **Custom SNAT IP / Router IP**, enter four IP addresses which are configured as next hop.
 - h. In **External Gateway Configuration**, enter 4 as **Number of Active Hosts**.
 - i. Click **Save**.
The external subnet is setup and the **Create VPC** page appears.
 - j. In **Domain Name Servers (DNS)**, add the DNS server IP addresses.
For more information, see the *Other details on the Create VPC page* row in the table in [Creating Virtual Private Cloud](#) topic of the *Flow Virtual Networking Guide*.
 - k. Click **Create**.
The VPC is now created and associated with the external subnet.
9. On the **Virtual Private Clouds List** page, click the Flow Virtual Networking transit VPC name where you plan to create an overlay network and click **Create Subnet**.
The **Create Subnet** page appears.
10. Configure the overlay network.
For more information, see [Creating a Subnet](#) in the *Flow Virtual Networking Guide*.
 - a. In **Name**, enter the name of the subnet.
 - b. In **Type**, confirm that the type is Overlay.
 - c. In **Network IP Address / Prefix**, enter the network IP address and prefix, and the gateway IP address.
 - d. Click **Add IP Pool** and enter the first and last IP address of the range, then click the checkbox under **Actions**.
 - e. Expand and configure **Domain Settings**.
 - f. Click **Create**.
The external subnet, underlay, and overlay networks are created and listed in the Virtual Private Clouds list.

11. Request a floating IP address.

- a. From the [Application Switcher Function](#), select **Infrastructure**.
- b. Click **Network & Security > Floating IPs**.
The **floating IPs details** page appears.
- c. Click **Request Floating IP**.
- d. In the **Number of Floating IPs** field, enter the required number of floating IP addresses.
For example, if you plan to assign floating IP addresses to two VMs, type 2 in this field. This floating IP address must match the number of Nutanix Objects public IP addresses required for the Nutanix Objects deployment later.
- e. Select the assignment type as **Assign Floating IPs** and select the floating IP address.
- f. Select the VPC and enter Nutanix Objects public IP addresses required for the Nutanix Objects deployment.
Repeat this step for all public IP addresses associated with the deployment of Nutanix Objects.
- g. Click **Submit**.
The floating IP addresses are assigned.

What to do next

After creating all the networks and assignment of floating IP addresses, you can start the Nutanix Objects deployment. For more information, see [Object Store Service Deployment](#) on page 65. You can also configure VPC modes without NAT. For more information, see [Configuring Virtual Private Cloud Modes for Nutanix Objects Without Network Address Translation](#) on page 51.

FEDERATION OVERVIEW

In the context of a single Nutanix Objects instance, *namespace* is a set of unique buckets hosted by that instance. Federation provides a way to create a single namespace across multiple individual Nutanix Objects instances. One federation instance corresponds to one such namespace.

A federated namespace can span across multiple Nutanix Objects instances that potentially reside in different geographical locations.

Each Nutanix Objects instance has a default local namespace that is independent of any other Nutanix Objects instance. In addition, Nutanix Objects instance can join multiple federated namespaces by becoming a member of each federation.

Any Nutanix Objects instance that joins a federation as a member can contribute to the federated namespace. A federated namespace can be accessed from any of the members using the Federation FQDN in the **Host** field of the HTTP requests.

Federated namespace supports the following properties:

- Federated namespace enforces unique bucket names within the namespace.
For example, if a client creates a bucket *backup* in a federated namespace, then another client that later tries creating the bucket *backup* would fail with the error message `BucketAlreadyExists` or `BucketAlreadyOwnedByYou`.
- All clients, irrespective of which federation member they choose to access the federated namespace from, see the same consistent list of buckets.
- Any buckets and objects that are part of the federated namespace can be accessed from any of the members.

A federation instance is supported in the backend using a subset of the federation members called **Core Members**. Core members, together, run a distributed consensus service in the backend to provide fault tolerance for the federated namespace.




Figure 5: Federation Overview

Federated Namespace Prerequisites

Before creating a federated namespace, review this section carefully to ensure you have met the prerequisites.

The following are the prerequisites:

- You must have at least one object store instance managed by a single or multiple Prism Central instances.
- For Nutanix Objects 4.0, it is necessary to upgrade Prism Central and Prism Element to compatible versions, with a minimum required Prism Central version of 2022.9.
- Nutanix Objects Manager on all Prism Central instances must be upgraded to compatible versions.
- Established availability zones pairing between Prism Central instances.
- Set up Identity and Access Management (IAM) replication across Prism Central instances.

Note: For IAM pairing, if the Nutanix Objects Manager version is 5.0 and above on either the source or the target, then the other side should also have Nutanix Objects Manager version 5.0 and above.

For more information on creating a federated namespace, see [Creating Federated Namespace](#) on page 58.

Creating Federated Namespace

Nutanix federation namespace is a mechanism that enables multiple independent systems to share a common namespace or naming system. This allows different systems to refer to the same resources using the same names, even if those resources are located in different systems and locations.

Before you begin

To create a federated namespace, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. In the Application Switcher, select **Objects > Federated Namespaces**.
3. Click **Federated Namespaces**.

Note:

- Ensure to update the SSL certificates with the federation Fully Qualified Domain Names (FQDN) for each object store that will be included in the federated namespace.
- Ensure your DNS is configured to map the federation FQDN to the object store public IP addresses so that applications can access the federation namespace.
- The federation FQDN should not be a sub-domain of the Prism Central domain.

4. Click **Create Federated NameSpace**.

After clicking **Get Started** on the prerequisites pop-up window that provides the necessary details, you will be directed to a new window for creating a federated namespace.

5. Update the fields as indicated:

- a. **Name:** Enter the name of federated namespace.

The federation name must meet the following guidelines.

- Must be unique across all existing federations.
- Must start with a letter.
- Can have alphanumeric or hyphen characters. Hyphen can only be used in the middle of the name.
- Must not contain special characters.
- Must be a minimum of 1 and a maximum of 16 characters long.

- b. **Domain:** Enter the domain.

The federation domain must meet the following guidelines:

- It must contain at least one dot.
- It may contain either alphanumeric or hyphen characters.
- It must not start or stop with a hyphen or a dot.
- When selecting a federation fully qualified domain name (FQDN), it's important to avoid using a subdomain that belongs to any of the MSP subdomains that are connected to the Prism Central web console.

- c. **Object Stores:** Select and add at least one object store instance. It is recommended to have at least three instances from different clusters to set up federation. You can select all the required object store instances that run the federation services.

Note: Once you select the Object stores, the visual representation on the right changes accordingly.

- d. **Federation core member:** A federation instance is supported in the backend using a subset of the federation members, called **Core Members**. Core members, together, run a distributed consensus service in the backend to provide fault tolerance for the federated namespace.

You can select and designate core members from the list of Object stores.

Note: To ensure higher fault tolerance in case of cluster failure, you can add 1, 3, or 5 core members from the same or different clusters.

- e. **Federation members:** All object stores participating in the federation are called federation members.

6. Click **Setup**.
The federation setup is ready.

☐	Name	Version	Domain	Worker Nodes	Usage (Logical)	Buckets	Objects	Notifications	Objects Public IPs
☐	OSS1 Federated	5.1		1	1.36 TiB / 50 TiB	232	6715K	Disabled	10.123.9.210
☐	OSS2 Federated	5.0		1	119.85 GiB / 940 TiB	183	4.47K	Disabled	10.123.9.230
☐	OSS3 Federated	5.0.11		1	59 GiB / 380 TiB	243	1M	Disabled	10.123.9.202

Figure 6: Federated Object Stores

What to do next

After creating a federated namespace, you can view and perform operations on the namespace. For more information, see [Viewing Federated Namespaces](#) on page 62 section.

Adding External Object Store

External object stores are AWS regions linked to an AWS account, allowing accessible buckets to be added to the federated global namespace.

Before you begin

To add an external object store to a federated namespace, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. In the Application Switcher, select **Objects > Federated Namespaces**.
3. Select and click the required federated namespace.

4. Click **Add member** > **Add External Object Store**.

After clicking **add external endpoint** on the prerequisites pop-up window that provides the necessary details, you will be directed to a new window for adding general details of external endpoint.

The screenshot shows a web form titled "Add External Object Store". The form is divided into two main sections: "General Details" and "Credentials".

General Details:

- Name:** A text input field with the placeholder "Please enter a name for the External Object Store".
- Provider:** A dropdown menu with "AWS" selected.
- Service Host:** A text input field with the placeholder "https://s3.us-east-1.amazonaws.com".

Credentials:

- A note: "Please provide AWS credentials with:
 - The 's3:ListAllMyBuckets' permission.
 - Owner-level access to all buckets you intend to add through this Endpoint.
- Access Key:** A text input field with the placeholder "Please enter Access key".
- Secret Key:** A text input field with the placeholder "Please enter Secret key".

To the right of the form is a network diagram showing a laptop icon connected to a box labeled "https://s3.us-east-1.amazonaws.com". Below this box is another box labeled "PC_10123/14", and below that is a box labeled "OS1 (Core Member)".

Figure 7: Adding external endpoint

5. Update the fields as indicated:

- Name:** Enter the name of External Object Store.
- Provider:** Select AWS.
- Service Host:** Service host is the complete URL of your endpoint provider where the buckets are located.
- Credentials:** Provide AWS credentials with s3:ListAllMyBuckets permission and Owner-level access to all buckets you intend to add through this Endpoint .
- Access key:** provide an access key.
- Secret Key:** provide a secret key.

6. Click **Save**.

The external endpoint is added.

What to do next

After adding an external object store, you can add buckets from the endpoint. For more information, see [Adding External Bucket](#) on page 202 section.

Removing External Object Store

Before you begin

All buckets associated with the external object store must be removed first. For more information, see [Deleting a Bucket in Nutanix Objects](#) on page 137.

About this task

To remove the external object store, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. In the Application Switcher, select **Objects > Federated Namespaces**.
3. Select the external object store you want to remove.
New window with External Summary opens.
4. Click **Remove External Object Store > Remove**.
External object store is removed.

Managing Federated Namespace

After creating a federated namespace, you can view details and manage the object stores.

About this task

To manage and add object store, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. In Application Switcher, select **Objects > Federated Namespaces**.
A list of existing federated namespaces appears.
3. Select a **Namespace**, and click **Manage Namespace**.
The **Federated Namespace** page appears.
4. (Optional) To remove an object store, first select an object store, and then click **Remove Object Store**.
The object store is removed.
5. (Optional) To add an object store to the existing federation, click **Add Object Store**.
The Add Object Store window appears. Select and add the required object stores.

Note:

- You can add between 1 and 20 object stores. Overall a federation can have maximum 128 members.
- Core members cannot be added to a federation.
- A core member cannot be removed.

What to do next

You can view the object store and check if they are part of the local namespace or federated namespace, see [Viewing Object Store Summary](#) on page 74.

Viewing Federated Namespaces

You can click any federated namespace to view all the details of the Summary and Object stores.

The following table lists and defines the fields used in summary and object stores:

Table 9: Summary

Term	Description
Name	Displays the name of the object store.
Domain	Displays the domain of the object store.
FQDN(s)	Displays a complete address to access the Federation.
Object Stores	Displays the number of object stores.
Buckets	Displays the number of buckets in an object store.
Object Browser	Displays a link to access the federated namespace.

Table 10: Object Stores

Term	Description
Name	Name of the object store.
Buckets	Number of buckets for the particular object store
Objects	Number of objects for that particular object store.
Capacity Usage (logical)	Displays the object storage capacity usage.
Version	Displays the version of Nutanix Objects in which the object store was created.
Public IPs	Displays the public IP addresses of object store. These are presented as links that can be used to launch the Nutanix Objects Browser.
PC Location	Displays the location of the object store.

Deleting Federated Namespace

You can delete an existing federated namespace.

About this task

Note:

- A healthy federation can be deleted
- Bulk deletion is not allowed.
- If a federation service fails to start, it can be deleted.

To delete a federated namespace, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. In Application Switcher, select **Objects > Federated Namespaces**.

3. Select the required **Federated Namespace** to delete.

Note: Clear out all the objects and buckets stored in the selected object store.

The **Federated Namespace** page appears.

4. Click **Delete Federation**.

The federated namespace is deleted.

What to do next

You can view the different error messages in federated namespace. For more information, see [Errors in Federated Namespace](#) on page 64.

Errors in Federated Namespace

When an error occurs while creating a federated namespace, the description lists the type of error.

Following are the different types of error seen in a federated namespace.

Table 11: Errors

Error type	Description	Steps to recover
Addition Failure	Inability to add a new object store to an existing federation.	1. Check the object store. 2. Get back the object store to a healthy state.
Backend Unreachable	When services of some of the object stores that are a part of the federation are down.	1. Check the object store. 2. Get back the object store to a healthy state.
Creation Failed	When the service cannot be created.	1. Check the object store. 2. Get back the object store to a healthy state.
Deletion Failure	Inability to delete an object store as it is hosting buckets or objects.	1. Check the object store. 2. Clear out all the objects and buckets stored in the selected object store.
Transport error	When core members are unreachable.	1. Check the object store. 2. Get back the services to a healthy state.

OBJECT STORE SERVICE DEPLOYMENT

Nutanix Objects is a highly available and distributed Object Store Service which can store petabytes of data.

To start using Nutanix Objects, you need to deploy the Object Store Service. For more information, see [Creating or Deploying an Object Store on Prism Central](#) on page 67 or [Deploying Object Store on Nutanix Cloud Cluster](#) on page 85. You can also perform offline deployment (site without Internet access) of an object store. For more information, see [Deploying Object Store at a Dark Site \(Offline Deployment\)](#) on page 86.

You can deploy multiple Nutanix Objects instance on a single Prism Element registered to a Prism Central, or you can register multiple Prism Elements to one Prism Central and deploy as many object stores on each of these Prism Elements provided you have sufficient storage.

You can do the following operations after deploying Nutanix Objects instance:

1. Configure the directory and generate the access key. For more information, see [Directory Configuration and Access Key Generation](#) on page 90.
2. Create buckets within the object store. For more information, see [Creating and Configuring an S3 Bucket in Nutanix Objects](#) on page 99
3. Upload objects and meta-data to the buckets by using the Simple Storage Service (S3) APIs. For more information, see [Nutanix Objects Supported APIs](#) on page 240.

Note:

- After registering Prism Element to Prism Central, wait for 10 minutes before starting the deployment.
- Parallel Object Store Service deployments are not supported.
- If your deployment fails due to precheck failures, you can resume the deployment after fixing the configuration.
- For objects containers, the Erasure Coding delay will be reduced from 7 days to 3 days for old and new deployments.

Object Store and Domain Naming Conventions

Established guidelines for names of object stores and domains.

Object Store Naming Guidelines

The name of an object store must conform to the following rules:

- Be unique across all existing object store names in Nutanix Objects.
- Begin with a letter, and end with a letter or number.
- Can contain alphanumeric or hyphen characters.
- Not contain any special character other than a hyphen.
- Minimum of 1 and a maximum of 16 characters long.

Note: You cannot change the name after creating the object store.

Domain Naming Guidelines

The name of the domain must conform to the following rules:

- Must include at least one period.
- Can consist of alphanumeric characters, hyphens, or underscores.
- Cannot start or end with a hyphen, underscore, or period.

Prerequisites for Enabling FT1n:2d on an Object Store

Fault tolerance for 1 node or 2 disk failures (FT1n:2d) is enabled on a newly deployed object store when you meet the following criteria:

- The cluster must consist of a minimum of five nodes.
- The fault tolerance/redundancy factor 3 must be enabled on the cluster.
- The AOS version must be 6.8 or later.
- Each node in the dense node platform must host a minimum of 20 HDDs.
- No non-internal (default, management, and self-service) FT2 containers must be present.
- When adding another cluster using the multicluster option in an FT1n:2d cluster with redundancy factor 3 enabled, ensure that the new cluster also has redundancy factor 3 enabled. Attempts to add a redundancy factor 2 cluster will be automatically rejected. For more information on fault tolerance, see [Fault Tolerance vs Storage Efficiency](#) on page 11.

Prerequisites for Nutanix Objects Network Segmentation of an Object Store

Nutanix Objects Network segmentation is enabled on an object store when you meet the following criteria:

- The AOS version must be 7.3 or later.
- The Prism Central version must be pc.7.3 or later.
- The Prism Element cluster must consist of a minimum of three nodes.
- Ensure that Nutanix Volumes network segmentation is enabled on the Prism Element cluster.

For more information, see [Isolating Service-Specific Traffic](#) and [Nutanix Volumes](#) in the Nutanix Security Guide.

- Ensure that the Nutanix Volumes network segmentation is configured on a different VLAN than Prism Central to maintain a stable connection between Prism Central and Prism Element.
- When enabling network segmentation for the Nutanix Volumes service on a Prism Element cluster that already has Nutanix Object Stores deployed (without network segmentation), make sure that the subnet IP address used for the Nutanix Volumes service network segmentation is different from the subnet IP address of any of the AHV subnets that are used as the storage network for any of the existing object stores.

Warning: To avoid I/O failures and potential cluster recovery issues, do not use the same subnet IP address for Nutanix Volumes network segmentation and the AHV subnet associated with object store storage networks.

- Nutanix recommends using the same network for management plane communication as Prism Element for a network segmentation enabled Nutanix object store.

- Ensure that for a network segmentation enabled Nutanix object store, the Nutanix Objects storage, public, and management networks each have a different AHV subnet.
- Ensure that for a network segmentation enabled Nutanix object store, the Nutanix Objects management network and the Nutanix Objects storage network are not routable to each other.
- When deploying a Nutanix object store without network segmentation on a Prism Element cluster that has Nutanix Volumes network segmentation enabled, ensure that the subnet IP address used for the AHV subnet of the storage network of the object store is a different subnet IP address than the one used for Nutanix Volumes service network segmentation on the Prism Element cluster.
- Ensure that Nutanix Volumes network segmentation is not disabled after deploying a network-segmented object store on a Prism Element cluster, even if the Prism Element user interface permits this action.

Warning: Disabling network segmentation for Nutanix Volumes in a network-segmented object store environment is not supported and might result in recovery failures.

Creating or Deploying an Object Store on Prism Central

Before you create a bucket and upload objects, you must deploy an object store.

Before you begin

Make sure that you meet the following prerequisites:

- Prerequisites for deploying an object store. For more information, see [Nutanix Objects Deployment Prerequisites - AHV and ESXi](#) on page 36.
- Prerequisites for enabling fault tolerance for one-node or two-disk failures (FT1n:2d) on a newly deployed object store. For more information, see [Prerequisites for Enabling FT1n:2d on an Object Store](#) on page 66.
- Prerequisites for enabling network segmentation. For more information, see [Prerequisites for Nutanix Objects Network Segmentation of an Object Store](#) on page 66.
- Ensure that you have the right user permissions. The access control policy must grant the following minimum permissions to a non-admin user for creating or deploying an object store:
 - **Create Object Store** access.
The user who creates an object store instance automatically receives full access to that instance. An administrator can later modify or revoke this access.
 - **View Cluster** access to the Prism Element instance on which you deploy the object store.
 - **View Subnet** access to the managed networks used for deployment if the Prism Element operates on AHV. This permission is applicable only for AHV.

For more information, see [Nutanix Objects Permissions](#) on page 29 and [Role-based Access Control for Nutanix Objects](#) on page 28.

About this task

You can deploy multiple object stores on a single Prism Element cluster or on each Prism Element cluster that is registered to a Prism Central instance, provided that you have enough storage. However, you cannot deploy a single object store across multiple Prism Element clusters.

Note:

- Multicloud is not supported when network segmentation is enabled for Nutanix Objects.

- In a network segmentation-enabled object store deployment, replication traffic defaults to the Nutanix Objects management network. To route replication traffic through the Nutanix Objects storage network, contact Nutanix Support for configuration assistance.

To create an object store, follow these steps:

Procedure

1. Log on to Prism Central.
2. In the Application Switcher, click **Objects**.
3. Click **Create Object Store**.
The **Create Object Store: Prerequisites** window appears.
4. Click **Confirm**.
The **Create Object Store** window appears.
5. In the **General Details** section of the **Object Store Details**, follow these steps and click **Next**:
 - a. In **Object Store Name**, enter the name of the object store.
For guidelines on choosing a compliant name, see [Object Store and Domain Naming Conventions](#) on page 65.
 - b. In **Domain**, enter the domain name.
This domain name is the default domain name for all the object store in that cluster.
For guidelines on choosing a compliant name, see [Object Store and Domain Naming Conventions](#) on page 65.
 - DNS configuration must be manually added on the client side to fully resolve the fully qualified domain name (FQDN) of Nutanix Objects to the Nutanix Objects public IP address.
 - When the Controller Microservices Platform (CMSP) is enabled, you cannot enter a domain name for the object store deployment. The CMSP domain is used instead. To use a different

domain, add Fully Qualified Domain Names (FQDNs) for the object store. For more information, see [Managing FQDN and SSL Certificate](#) on page 81.

- c. In **Cluster**, select the cluster for the object store deployment.
 - Use an encrypted cluster for encrypting the bucket.
 - Object stores deployed on one or two-node clusters cannot be scaled out after deployment.
- d. (Optional) To enable the Nutanix Objects fault tolerance for one-node or two-disk failures (FT1n:2d), click **Enable FT1n:2d objects qualification**.

- e. In **Worker Nodes**, add the number of worker nodes.

VM configuration guidelines:

- Each VM is assigned 10 vCPUs and a DHCP IP address.
- A minimum of 10vCPUs and 32 GiB of memory is required.
- Each click on plus for worker nodes adds 10vCPUs and 32 GiB of memory.
- vCPU and memory are linked. You must use multiples of 10 for vCPUs and multiples of 32 for memory.

Important:

- The number of configured worker nodes must not exceed the total nodes in the AOS cluster.
- Resources added cannot be reduced later.
- Actual performance relies on various factors.

- f. (Optional) To optimize high-performance workloads on all-flash clusters, select the following options:
 - » **Dense All-Flash Clusters:** If a cluster is a dense all-flash setup, enable the high-performance configuration through the Prism Central user interface. If you require better performance on these clusters, select **Scale up Worker Nodes**. This option is disabled by default.
 - » **Hybrid Clusters:** If a cluster is a hybrid setup, enable the maximum number of load balancers (N-1) when the number of MSP worker nodes (N) is six or greater. If your workloads on Nutanix Objects require more concurrent connections or higher throughput, select **Scale Out Load Balancer**. This option is also disabled by default.
- If you close the window by clicking the **X** while you create an object store, the system does not save any entered information or the precheck status.
- Hover over the help icon for more information about each field.
- The diagram updates automatically to show the necessary worker nodes, load balancers, and resources based on the number of selected worker nodes.
- The system highlights active components in the diagram.
- A new container with name `msh-<uuid>` is created for Nutanix Objects deployment on ESXi. This container downloads VM images for MSP worker nodes.

6. In the **Storage Network** section, follow these steps:

- a. In **Storage Network**, select the Nutanix Objects storage network that is used for internal communication (data and management) between various VMs and nodes of the object store.

For more information on the IP address requirements according to the deployment size, see [Nutanix Objects Network Configuration](#) on page 39.

- b. (Optional) To segment the network for AHV clusters, check the **Segment Management Traffic** checkbox and select the management network from the **Management Network** dropdown list.

Segmented networks separate data and management traffic. The management network handles all management plane traffic. You can enable Nutanix Objects network segmentation only if Nutanix Volumes network segmentation is enabled. If Nutanix Volumes network segmentation is disabled on the Prism Element cluster, the **Segment Management Traffic** checkbox appears grayed out.

Note: Network segmentation is not supported for ESXi clusters.

- c. (Only for AHV) In **Object Store Internal IPs (2)**, do one of the following:

- » If network segmentation is disabled, enter two Nutanix Objects storage network IP addresses, separated by a comma.
- » If network segmentation is enabled, enter two Nutanix Objects management network IP addresses, separated by a comma.

If network segmentation is disabled, **Object Store Internal IPs (2)** are within the Nutanix Objects storage network. However, if network segmentation is enabled, **Object Store Internal IPs (2)** are within the Nutanix Objects management network and Nutanix Objects storage and management network must be different.

For ESXi, these two internal IP addresses are not required and are selected automatically from the IPAM range configured for the ESXi networks.

7. In the **Public Network** section, follow these steps and click **Save & Continue**:

- a. In **Public Network**, select a public network to allow external clients access to the object store.

This VLAN must have up to four IP addresses in the usable IP address range. This network can be the same as the storage network. For more information, see [Shared Versus Single Network Configuration for Nutanix Objects](#) on page 48.

- b. In the **Public Network Static IPs**, enter the public access IP addresses (one for each load balancer) separated by a comma or as an IP address range.

For example, if one Load Balancer is used, then only one IP address is required. Also, you can enter the IP addresses in a range, 10.2.3.1-10.2.3.4, or separated by a comma 10.2.3.1, 10.2.3.2, 10.2.3.3, 10.2.3.4.

AHV: These IP addresses are within the Nutanix Objects public network and used to access the object store.

ESXi: The public access IP addresses can be either within or outside the range of the IPAM network.

For more information on network configurations, see [Nutanix Objects Network Configuration](#) on page 39.

8. (Optional) To continue with the deployment later, click **Save for Later**.

The object store is saved in the list of object stores.

9. (Optional) To complete the deployment later for the saved object store configuration, select the object store and click **Actions > Complete Deployment**.

Note: Deployment takes at least 30 minutes.

Prechecks start before the deployment begins. A list of checks performed is displayed in the UI. Also, a VM image named *predeployment_port_vm*, and two VMs named *predeployment_objects_public* and *predeployment_objects_storage* are created.

10. (Optional) To see a summary of the object store, click **Show Summary** in the right pane.

11. Depending on the results of the precheck, follow these steps:

- » If the prechecks pass, click **Download Report** to download the report, then click **Create Object Store** to start with the object store deployment.

The report contains the name of the check, status, and message. All prechecks must pass before you can start the deployment.

- » If the prechecks fail, an error message is displayed in the UI. Click **Download Report** to download the report. A Fail status is displayed next to the check name with a message. The object store is saved and you can complete the deployment after fixing the failed checks.

After the prechecks pass, you can view the status of the object store deployment. You can view the deployment progress in percentage and each step in the grid by hovering over the loading icon.

Warning: Do not delete MSP VMs (created with a specific prefix in Nutanix Objects) from vCenter or Prism Central, because their deletion goes undetected. You can identify the MSP VMs by their names. The naming conventions are as follows:

- MSP Worker VM: *deployment_name-XXXXXX-default-N*
- Load Balancer MSP VM: *deployment_name-XXXXXX-XXXXXXXXXX-envoy-N*

What to do next

After deploying the object store, you can perform the following steps:

- Configure directory and generate access key. For more information, see [Directory Configuration and Access Key Generation](#) on page 90.
- Access the endpoint provided as part of Client Access Network (S3 Endpoints) by using HTTP or HTTPS. For more information, see [Access Nutanix Objects Endpoints](#) on page 74.
- Create and configure buckets. For more information, see [Creating and Configuring an S3 Bucket in Nutanix Objects](#) on page 99.
- Create objects using S3 APIs. For more information, see [Supported S3 APIs](#)
- Expand the object store storage. For more information, see [Expanding Storage for an Object Store](#) on page 78
- Scale out an object store. For more information, see [Scaling Out an Object Store](#) on page 80
- Share buckets. For more information, see [Sharing a Bucket in Nutanix Objects](#) on page 128.

You can also deploy Nutanix Objects on Nutanix Cloud Clusters. For more information, see [Deploying Object Store on Nutanix Cloud Cluster](#) on page 85.

Viewing Object Store Deployments

View a list of deployed object stores, and the general and networking details of the object stores.

Before you begin

A non-admin user can view an object store only after the administrator creates an access control policy on a role in the Prism Central user interface for the non-admin user. The specific role must have minimum permission to view an object store and the non-admin user must be assigned to specific Nutanix Objects instances.

Note: Administrators and non-admin users are Prism Central users with the following roles:

- Administrator—A Super Admin or a Prism Admin in Prism Central.
- Non-admin user—A Prism Central user without any administrator privileges.


For more information on the built-in roles in Prism Central, see [Built-in Role Management](#) in the *Security Guide*.

About this task

To view the object store, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. From the [Application Switcher Function](#), click **Objects**.



Name	Version	Domain	Worker Nodes	Usage (Signed)	Buckets	Objects	Notifications	Object Public IPs
v4055-174401162	5.2	misp.pc-bhj.nutanix.com	3	1776 TiB / 120 TiB	9	27.89M	Enabled	10.40.242.3 10.40.242.4
v4055-1744017663	5.2	misp.pc-bhj.nutanix.com	3	32.72 TiB / 790 TiB	53	40.05M	Disabled	10.40.242.17 10.40.242.16

Figure 8: Object Store list

A list of existing object stores appears.

The following steps describe the fields that appear in the object store table. You can click on the name of an object store to open the **Object Store** page. You can click **View By** to view the **General**,

Usage, or **Networking** details of an object store. A dash (-) is displayed in the field when a value is not available or applicable.

General view:

- **Name:** Displays the name of the object store. A *Federated* tag is displayed next to the federated object stores.
- **Version:** Displays the version of Nutanix Objects in which the object store was created.
- **Domain:** Displays the domain of the object store.
- **Worker Nodes:** Displays the number of worker nodes of the object store.
- **Usage (Logical):** Displays the object store usage in GiB or TiB.
- **Buckets:** Displays the number of buckets in an object store.
- **Objects:** Displays the number of objects in an object store.
- **Notifications:** Displays notifications if any.
- **Objects Public IPs:** Displays the endpoints or the IP addresses used by the client.

You can access these endpoints by using HyperText Transfer Protocol (HTTP) and HyperText Transfer Protocol Secure (HTTPS) protocols. For more information, see [Access Nutanix Objects Endpoints](#) on page 74.

Networking view:

- **Name:** Displays the name of the object store.
- **Cluster:** Displays the cluster in which the object store is deployed.
- **Objects Public Network:** Displays the Virtual Local Area Network (VLAN) required for accessing object store endpoints externally.
- **Objects Public IPs:** Displays the public IP address of Nutanix Objects. This is used to access the Nutanix Objects Browser.
- **Objects Storage Network:** Displays the VLAN required internally for deploying Object Store Services on Prism Element.
- **Objects Storage Network Static IPs:** Displays the internal configured IP addresses.

Usage view:

- **Name:** Displays the name of the object store.
- **Usage (Logical):** Displays the object store usage in GiB or TiB.
- **Local Usage:** Displays the amount of data stored locally in the Nutanix Objects cluster.
- **Tiered Usage:** Displays the amount of data tiered out and stored remotely based on lifecycle policies.
- **Licensed Usage:** Displays the usage from a single object store accounted against the licensed capacity. It is calculated by adding **Local Usage** and **Tiered Usage**.

Note: Data tiered to another Nutanix Objects endpoint is not included.

For information on network configurations, see [Nutanix Objects Network Configuration](#) on page 39.

What to do next

You can view the summary of an object store and perform various object store operations. For information, see [Viewing Object Store Summary](#) on page 74.

Access Nutanix Objects Endpoints

Nutanix Objects endpoints are the entry point to Nutanix Objects. You can access these endpoints by using the Hypertext Transfer Protocol (HTTP) or Hypertext Transfer Protocol Secure (HTTPS) protocols through any third-party clients which support Simple Storage Service (S3) APIs.

Accessing Buckets and Nutanix Objects within an Object Store Instance

Nutanix Objects supports path-style and virtual hosted-style bucket access.

- **Path-Style Access:** The path-style syntax requires that you use the endpoint when attempting to access a bucket, and the request specifies a bucket by using the first slash-delimited component of the Request-URI path.

For example, if you have a bucket with the bucket name as **bucket-name** and the object name as **example.jpg**, and you want to use the path-style syntax. Following is the correct request:

```
PUT /bucket-name/example.jpg HTTP/1.1
Host: object-store-name.domain-name
```

- **Virtual Hosted-Style Access:** The virtual hosted-style syntax is used to address a bucket in a REST API call by using the HTTP Host header. This method requires the bucket name to be Domain Name System (DNS)-compliant.

For example, if you have a bucket with the bucket name as **bucket-name** and the object name as **example.jpg**, and you want to use the virtual hosted style. Following is the correct request:

```
PUT /example.jpg HTTP/1.1
Host: bucket-name.object-store-name.domain-name
```

For virtual hosted-style access, allow the Nutanix Objects fully qualified domain name (FQDN) in the DNS server with the wild card allowlist. For example, the following are the expected DNS entries for the Nutanix Objects endpoint.

objects.subdomain.example.com.	IN	A	192.168.5.101
objects.subdomain.example.com.	IN	A	192.168.5.102
objects.subdomain.example.com.	IN	A	192.168.5.103
objects.subdomain.example.com.	IN	A	192.168.5.104
*.objects.subdomain.example.com.	IN	A	192.168.5.101
*.objects.subdomain.example.com.	IN	A	192.168.5.102
*.objects.subdomain.example.com.	IN	A	192.168.5.103
*.objects.subdomain.example.com.	IN	A	192.168.5.104

Viewing Object Store Summary

You can view the summary of an object store. The summary comprises of configuration, namespaces, and usage of the object store.

About this task

To view the object store summary, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. In the Application Switcher, click **Objects**.

3. Click the **Object Stores** tab, and then click the name of the object store from the table to view its summary.
4. Click the **Summary** tab.
You can view the object store configurations, federated namespaces, and usage.

What to do next

You can do the following from the object store summary page:

- Managing FQDNs and SSL certificates. For more information, see [Managing FQDN and SSL Certificate](#) on page 81.
- Scaling out an object store. For more information, see [Scaling Out an Object Store](#) on page 80.
- Expanding an object store. For more information, see [Expanding Storage for an Object Store](#) on page 78.
- Deleting an object store. For more information, see [Deleting an Object Store](#) on page 75.
- Launching Nutanix Objects Browser. For more information, see [Launching the Nutanix Objects Browser](#) on page 198.

Deleting an Object Store

You can delete both successful and failed object store deployments. However, for deletion of successful deployment, ensure to first delete the objects and buckets within that object store.

Before you begin

A non-admin user can delete an object store only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions and the non-admin user must be assigned to specific objects instances: Delete Object Store and View Object Store.

Note: Administrators and non-admin users are Prism Central users with the following roles:

- Administrator—A Super Admin or a Prism Admin in Prism Central.
- Non-admin user—A Prism Central user without any administrator privileges.

For more information on the built-in roles in Prism Central, see [Built-in Role Management](#) in the *Security Guide*.

About this task

Note:

- Deployments in progress cannot be deleted.
- You cannot delete the primary cluster without deleting the secondary cluster as the primary object store cluster hosts all the common services. To delete the primary cluster, ensure all the secondary clusters are deleted.
- Multi-cluster containers are not deleted on secondary Prism Elements if the Nutanix Objects deployment is not successful.
- If a failure occurs while replacing an SSL certificate for an object store, you cannot delete that object store deployment; however, you can try replacing the SSL certificate again.

- Deletion of an object on the local object store may not immediately lead to space reclamation on the AWS S3. Space reclamation happens when all the source objects mapped to the corresponding objects in the AWS S3 get deleted.

To delete an object store deployment, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. In the Application Switcher, click **Objects**.
3. In the Object Stores table, select the object store which you want to delete, and then click **Actions > Delete**.
4. On the confirmation dialog box, click **Delete**.
A message appears to confirm the object store deletion.

Object Store Expansion

This section describes the expansion options available with object store.

Storage Expansion

If your existing object store storage is less than 85% full or you plan to use a different cluster for object store storage, you can expand the storage of your object store cluster by adding new nodes to the existing cluster or by adding additional clusters.

For example, if you have an object store deployed on a 10 TB cluster named *Cluster1* and your storage is almost or less than 8.5 TB full, then you can either add a node to the existing *Cluster1* or add a new cluster or existing clusters (such as *Cluster2* and *Cluster3*) if they have sufficient storage capacity. In this example, *Cluster1* is the existing primary cluster on which the object store was initially deployed and hosts the worker VMs for the object store. *Cluster2* and *Cluster3* are secondary clusters that are later added for capacity expansion.

In a multicluster setup, clusters are managed using a two-tier system: a preferred tier and a secondary tier.

- Clusters that use less than 85% of their physical capacity quota (where the quota is defined as a percentage of total physical capacity) start in the preferred tier. Once the usage of the cluster exceeds 85%, it is moved to the secondary tier.

For example, if the quota is set at 60% of the physical capacity (for example, 100TB), the object store cannot exceed a physical storage consumption of 60% (which equals 60TB). The 85% threshold for tier migration is applied to the defined capacity, meaning 85% of 60TB equals 51TB. Therefore, if the storage consumption goes beyond 51TB, the cluster will be moved to the secondary tier.

- When selecting a cluster for storage allocation, the preferred tier is prioritized. If the preferred tier is empty, the secondary tier is used.
- Within each tier, storage is allocated in a round-robin manner among the available clusters.
- Storage allocation from any single cluster cannot exceed 95% of its total physical capacity.

Quotas can only be configured for secondary clusters and are specified as a percentage of the total physical capacity. The primary cluster is assigned a default quota of 100%.

Scale-Out (Compute and Memory Expansion)

Scaling out an object store enables you to add more resources to an existing object store cluster. You can scale out the CPU count and memory, and optionally, add more storage in addition to the current usage. Scale out uses an extra 10 vCPUs and 32 GiB memory for the newly added node.

Prerequisites for Expanding Storage for an Object Store in Nutanix Objects

Make sure you meet the following requirements before expanding storage for an object store:

- Ensure to upgrade to Microservices Platform (MSP) 1.0.5 or later and Nutanix Objects 2.0 or later before expanding the object store cluster.
- Ensure that the new clusters are registered to the same Prism Central cluster where the object store is deployed.
- Both the primary and secondary clusters must be running AOS versions later than 5.11.2.
- Data services IP address must be configured in each secondary cluster.
- Firewall must be running on the Controller Virtual Machines (CVMs) in the cluster.
- Maximum latency required for adding secondary clusters is less than 5 milliseconds.
- Up to 4 secondary clusters are supported.
- VMware ESXi and Nutanix AHV clusters are supported. Microsoft Hyper-V cluster is not supported.

Secondary Cluster Addition or Removal Guidelines

Key considerations and operational rules for adding, updating, or removing secondary clusters in a Nutanix Objects object store.

Secondary Cluster Addition Guidelines

Following are the guidelines for adding a secondary cluster in an object store:

- Add one secondary cluster at a time to expand the object store.
- Do not add secondary clusters if the primary cluster is full.
- Add secondary clusters before primary storage reaches 80% capacity.
- If a secondary cluster is added to an empty object store and becomes unreachable:
 - No alert is shown in Prism Element.
 - The cluster is hidden in the UI.
 - After objects are added, an alert is triggered and the unreachable cluster is marked in the UI within a minute.

Cluster Removal and Update Guidelines

Following are the guidelines for removing or updating a secondary cluster in an object store:

- Use **Update Multi-Cluster** only to update capacity or delete an added cluster.
- You cannot remove or reduce the limit of a cluster while adding another cluster.
- You can add a cluster while removing or reducing the limit of another.

Timing Considerations

The following are the timing implications of various actions, such as removal or modification of storage limits:

- If the secondary cluster is empty, then the removal of that cluster takes up to three hours.
- If the secondary cluster has some data, and the storage limit is reduced or the cluster is removed, then starting the data migration process might take up to seven hours.
- When you reduce the limit of the secondary cluster and if the used capacity of the cluster is less than the updated limit capacity, then no data migration takes place and limit is changed without any delay.

Operational Restrictions

The following are the operational restriction for updating or removing the secondary clusters:

- When adding a cluster, you cannot remove another cluster or reduce the limit of an existing cluster. However, while removing a cluster or reducing its limit, you can add another cluster or increase the limit of an existing secondary cluster.
- If you cancel removing the cluster or decreasing the limit, then the last updated limit remains the same and any data migrated to other clusters is not migrated back to this cluster.
- While the limit reduction for one cluster is in progress, you can increase the limit of another cluster, but you cannot decrease the limit of that cluster.

Expanding Storage for an Object Store

Expand object store storage by adding nodes or clusters if usage is less than 85 percentage or if using a different cluster for storage.

Before you begin

See [Prerequisites for Expanding Storage for an Object Store in Nutanix Objects](#) on page 77.

Non-admin users can expand object store storage only if assigned a role with specific permissions in Prism Central. Minimum required permissions:

- View, Add, and Update Multiclustert
- View Object Store

In addition to the preceding permissions, the user must have *View Cluster* permission on the Prism Element cluster that needs to be added to the object store.

For more information on the built-in roles in Prism Central, see [Built-in Role Management](#) in the *Security Guide*.

About this task

To expand the object store storage, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. In the Application Switcher, click **Objects**.
3. Click the **Object Stores** tab, and then click the name of the object store for which you plan to expand the storage.
The object store page appears.

4. Click **Clusters**.

Caution: Once a primary cluster has been successfully added to an object store, it cannot be removed.

The table lists the Usage, Max Usable and Free Capacity of the cluster. All the capacities listed are physical capacities, not logical.

- **Usage (Physical):** Physical capacity used by this object store on the cluster.
- **Max Usable (Physical):** Maximum physical capacity on the cluster can be used by this object store. This is calculated as the total physical capacity of the cluster * any limit set for this cluster capacity used by other workloads. Your available capacity might be less than you planned for as other workloads are taking up the space.
- **Free Capacity (Physical):** Additional capacity the object store can consume on this cluster within the specified limit. It is possible that there might not be free capacity available for the object store to consume even if current consumed capacity is less than max usable since there might be other workloads consuming capacity from this cluster.

For example, Cluster A is a primary cluster with 10 TB as the total capacity. No hard limit can be set on this cluster as it is a primary cluster. So, the *Max Usable Physical Capacity* of this object store cluster is 10 TB. However, the current object store usage is 5 TB and other workloads usage is 2 TB. So, the additional *Physical Free Capacity* available for the object store is 3 TB. However, in case of Cluster B which is the secondary cluster with total capacity as 15 TB, the hard limit is 50%. So, the *Max Usable Physical Capacity* of this object store cluster is 50 % * 15 TB = 7.5 TB. However, the current object store usage is 2 TB and other workloads usage is 10 TB. So, the additional *Physical Free Capacity* available for the object store is 5.5 TB, but as the other workloads consume 2.5 TB of the *Max Usable Physical Capacity* of the object store, so the remaining *Physical Free Capacity* available for the object store is only 3 TB. So, other workloads can consume the *Max Usable Physical Capacity* of the object store; however, an object store cannot go beyond the set limit.

Your primary cluster where the object store is deployed will be displayed in the table and this cluster cannot be removed. You can also view the free and used storage space.

5. Click **Add Clusters**.

For more information on guidelines, see [Secondary Cluster Addition Guidelines](#) on page 77.

After adding 4 secondary clusters, the **Add Clusters** button is automatically disabled.

A list of clusters registered to the Prism Central will be displayed. Hypervisor type, total physical capacity and free physical capacity of the clusters will also be displayed. However, the clusters that are already added as secondary clusters to an object store will not be displayed.

6. Once you select the cluster, under **Set up hard limit** section, select the usage limit in percentage, or select **Custom** and enter a custom limit for an object store and click **Done**.

This limits the object store to use a maximum capacity on the selected cluster. If you exceed the limit, an alert is generated.

You can change the limit (increase or decrease) once a secondary cluster is added. Select the cluster, and then click **Update Limit**. You cannot update the limit for the primary cluster. You can also remove

the secondary clusters once added. Select the cluster, and then click **Remove**. The removed cluster can be added back to the multicluster.

For more information on secondary clusters considerations, see [Secondary Cluster Addition or Removal Guidelines](#) on page 77.

A new cluster is added to the object store. If any secondary cluster addition fails, you can remove that cluster. You can also see the usage of these clusters in the **Usage** tab. For more information, see [Viewing Object Store Usage](#) on page 166.

Note: It will take a minute to add a new cluster.

Scaling Out an Object Store

Scaling out an object store enables you to add more resources to an existing object store cluster. You can add a worker node, and optionally, add more storage in addition to the current usage. An additional 10 vCPUs and 32 GiB memory get added for each worker node.

Before you begin

- Make sure that you have at least a three-worker node cluster for performing scale-out, but it cannot exceed the maximum number of back-end AOS nodes.
- Make sure that physical resources are available.

A non-admin user can scale out an object store only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions and the non-admin user must be assigned to specific objects instances:

- Scale-Out
- View Object Store

Note: Administrators and non-admin users are Prism Central users with the following roles:

- Administrator—A Super Admin or a Prism Admin in Prism Central.
- Non-admin user—A Prism Central user without any administrator privileges.

For more information on the built-in roles in Prism Central, see [Built-in Role Management](#) in the *Security Guide*.

About this task

Note:

- Nutanix Objects version lower than Nutanix Objects 2.1 does not support scale-out. Upgrade Nutanix Objects to the latest version to use the scale-out feature.
- You can scale out one node or VM at a time.
- During scale out of an object store, no disruption happens. You can launch the object store during scale-out.
- If physical resources (VMs) are deployed, rollback is not supported. However, if the deployment of physical resources fails or if the deployment fails prior to deploying physical resources and your cluster is not scaling, you can roll back. For rolling back scale out of the object store, contact Nutanix Support at <http://portal.nutanix.com>.

You can perform compute scale out (adding worker nodes) and storage scale out (adding additional storage capacity) for an object store.

To scale out an object store, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. In the Application Switcher, click **Objects**.
3. Click the **Object Stores** tab, and then from the object store table, select the object store that you want to scale out.
4. If you want to add a worker node (compute scale out), follow these steps:
 - a. Click **Scale out**.
The **Compute Scale Out** window appears.
 - b. Click **Add Nodes**.

The object store is now scaling out. This process takes about 5 to 10 minutes. You can track step by step deployment progress with scale out workflow progress for the object store.

Once the object store scale out is completed, new node is added to the object store. The object store takes an additional 10 vCPUs and 32 GiB of memory for the worker node.
5. If you want to set usage alert to the current usage, follow these steps:
 - a. In the **Actions** list, click **Set Usage Alert**.
The **Set Usage Alert** page appears.
 - b. Enter the limit you want to add to the current usage and click **Add** to complete.

Note: Nutanix Objects generates an alert if the logical usage of your object store reaches 90% of the specified value. For more information on alerts, see [Viewing Alerts](#).

Managing FQDN and SSL Certificate

You can add multiple fully qualified domain names (FQDNs), download the Certificate Authority (CA) certificate, and set up or replace the secure sockets layer (SSL) certificate for your object store.

A non-admin user can manage FQDN and SSL certificates only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions and the non-admin user must be assigned to specific objects instances:

- Regenerate Self Signed Certificate
- Replace Certificates
- Download Certificate
- View Object Store

Note: Administrators and non-admin users are Prism Central users with the following roles:

- Administrator—A Super Admin or a Prism Admin in Prism Central.
- Non-admin user—A Prism Central user without any administrator privileges.

For more information on the built-in roles in Prism Central, see [Built-in Role Management](#) in the *Security Guide*.

Adding FQDNs

You can create multiple fully qualified domain names (FQDNs) for an object store. The FQDN that is used when creating an object store is considered the default FQDN, while the rest of the FQDNs are categorized as alternate FQDNs.

About this task

To add a FQDN, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. In the Application Switcher, click **Objects**.
3. Click the **Object Stores** tab, and then from the object stores table, select the object store for which you want to create a FQDN.
4. Click **Manage FQDNs and SSL Certificates**.

5. In the **New FQDN** field, type the FQDN and then click **+FQDN**.

Following are the guidelines for naming a FQDN:

- Must contain at least two dots.
- Can contain either alphanumeric, hyphen, or underscore characters.
- Cannot start or end with a hyphen or underscore or dot.
- Duplicate domain names are not allowed.
- A FQDN cannot be at a sub-domain of another FQDN.




Figure 9: Adding an FQDN

You can select one or more FQDNs, and click **Delete** to delete the FQDNs. However, you cannot delete the default FQDN (the FQDN used while creating an object store).

The new FQDN is listed in the table.

Note: A warning message is shown if the new FQDN is missing in the DNS configuration of the certificate. To add the FQDN to the certificate do either of the following:

- Regenerate the secure sockets layer (SSL) certificate to add all the newly added domains to the certificate.
- Replace the SSL certificate by importing a Certificate Authority (CA)-signed certificate. Make sure to add all the domains to that certificate.

6. Click **Save**.

A confirmation dialog box appears to replace the SSL certificate.

7. Click **Yes, Replace**.

The new FQDN is added and the SSL certificate remains the same.

Note: The object store will be unreachable for 2-3 minutes, and you will not be able to perform any operations on that object store.

What to do next

You can also regenerate or replace the SSL certificate. For more information, see [Setting up SSL Certificate for an Object Store](#) on page 83.

Setting up SSL Certificate for an Object Store

By default, self-signed secure sockets layer (SSL) certificates are generated. If you have strong security requirements, you can replace the default certificate for the object store to securely connect to the object

store while you use the Hypertext Transfer Protocol Secure (HTTPS) protocol. You can replace the certificates either by regenerating a self-signed certificate or if you have a Certificate Authority (CA)-signed certificate, then you can import your private key and certificate files. When you replace the existing certificate, it removes the web browser certificate error warnings.

Before you begin

Ensure that you have met the following requirements:

- The private key should be Rivest–Shamir–Adleman (RSA) key type with key size 2048- or 4096-bit. Contents of the private key can be in Public-Key Cryptography Standards (PKCS)#1 standard, unencrypted PKCS#8 standard, and Privacy Enhanced Mail (PEM) format.
- The provided public certificate must be signed by the provided CA.

Note: If you want the server to return the server certificate and the chain of intermediate certificates, upload the server certificate and the chain of intermediate certificates as a public certificate in a single file.

- The public certificate must have the fully qualified domain name (FQDN) of the Object Store Service along with the wild card in either CN or SAN. For example, if the object store name is `objects-2021` and the domain is `companyname.com`, then the FQDN should be `objects-2021.companyname.com`. Then, the certificate must have `*.objects-2021.companyname.com`, `objects-2021.companyname.com`.

About this task

To set the SSL certificate, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. In the Application Switcher, click **Objects**.
3. Click the **Object Stores** tab, and then from the object stores table, select the object store for which you want to set the certificate.
4. Click **Manage FQDNs and SSL Certificates**.
5. Under SSL Certificates, click **Replace SSL Certificate**.

6. In the **Replace SSL Certificate** window, select one of the following:

Figure 10: Replace SSL Certificate

- a. **By regenerating self signed certificate:** Uses RSA 2048 bit as the private key type.
A self-signed certificate is a certificate signed by the same entity that verifies the certificate.
- b. **By importing key and certificate:** Upload your private key and certificate files.

- **Private Key:** Click **Upload** to upload the private key.

This key is used to decrypt the message.

- **Public Certificate:** Click **Upload** to upload the public certificate.

A public key certificate is an electronic document used to prove the ownership of a public key. The certificate includes information about the key, information about the identity of its owner (called the subject), and the digital signature of an entity that has verified the contents of the certificate (called the issuer).

- **CA Certificate/Chain:** Upload the CA certificate.

A certificate chain is the certificate of a particular CA, plus the certificates of any higher CAs up through the root CA.

You can also add or delete the FQDNs while regenerating or replacing the SSL certificate. For more information, see [Adding FQDNs](#) on page 82.

7. Click **Save**.

The default SSL certificate is replaced.

Deploying Object Store on Nutanix Cloud Cluster

Perform the tasks described in this section to get started with Nutanix Cloud Cluster (NC2) on Nutanix Objects.

Before you begin

Ensure that you have registered for NC2 on Amazon Web Services (AWS) from the My Nutanix portal. For more information, see [Getting Started with NC2](#).

About this task

To deploy an object store on Nutanix Cloud Cluster (NC2), follow these steps:

Note: Nutanix Cloud Clusters (NC2) currently support only single-node deployments. For assistance with multi-node deployments, please contact Nutanix Support.

Procedure

1. Enable Nutanix Objects.

For more information, see [Enabling Nutanix Objects](#) on page 17.

2. Create an object store using the IP addresses from the subnet created on the Prism Element web console.

Select only the required number of IP addresses instead of using all 255.

3. Ensure that the name servers are configured correctly. Copy the name servers from the Prism Element web console and add them to the Prism Central web console.

4. Using the load balancer, log on to Prism Central.

5. You can get the IP address from the subnets created on Prism Element and deploy the object store.

Note: You can provide any free IP address in the subnet.

Object store is successfully deployed.

6. Create a load balancer for the Nutanix Objects endpoint IP address to access the Nutanix Objects Browser.

Deploying Object Store at a Dark Site (Offline Deployment)

Dark site deployment is a process for deploying Nutanix Objects at a site without Internet access.

Before you begin

Make sure you satisfy the deployment prerequisites before starting the deployment. For more information, see [Nutanix Objects Deployment Prerequisites - AHV and ESXi](#) on page 36.

About this task

To deploy an object store in a dark site, follow these steps:

Procedure

1. Deploy Prism Central 2023.1.0.1 or later version.
2. Check the Compatibility matrix to find the latest Nutanix Objects Manager Version from the [Support Portal](#).
3. Download the latest **Objects LCM bundle** from the [Support Portal](#).
4. Check the Compatibility Matrix to find the latest MSP version from the [Support Portal](#).
5. Download the latest MSP dark site bundle from the [Support Portal](#).
6. Download the latest LCM Framework Bundle from the [Support Portal](#).

7. Set up the Dark Site web server, and copy the latest bundles for Nutanix Objects, MSP, and LCM to the web server in the folder named darksite.

For more information on viewing example of setting up a web server, see [Setting Up a Local Web Server](#) in the *LCM Dark Site Guide*.

8. Run the following commands to untar all the three bundles within the same location (darksite folder).

```
tar -xvf lcm_dark_site_bundle_2.6.0.1.36249.tar.gz
```

```
tar -xvf lcm_darksite_msp-2.4.3.2.tar.gz
```

```
tar -xvf objects-4.0.tar.gz
```

9. Login to the Prism Central and navigate to the **Admin Center** from the application switcher.

- a. Enable the Nutanix Marketplace.

- b. Get the Nutanix Objects application.

This step enables the Nutanix Objects Service on the Prism Central.

10. Once the Nutanix Objects Service is enabled, navigate to **Objects** from the Application Switcher. The Nutanix Objects landing page appears.

11. To configure the dark site server in LCM, follow these steps:

- a. Navigate to the **Admin Center** from the Application Switcher, and then click **LCM > Settings**.

- b. Under Source, select **Dark Site (Local Web Server)** from the drop down, and enter the web server URL from Step 8 where all dark site bundles are untarred .

For example, `http://IP address of the web server/darksite/`

darksite is the folder where all the bundles are untarred.




Figure 11: LCM Dark Site Settings

12. SSH to Prism Central and run the following command to check the AirGap enabled status:

```
nutanix@PCVM~$ mspctl controller airgap get
```

13. Login to the Prism Central UI and navigate to **LCM > Inventory**, and then click **Perform Inventory**.
14. Once LCM Inventory completes, upgrade to latest Nutanix Objects Manager version should be available (if current Nutanix Objects manager is not the latest).
15. Select **Nutanix Objects Manager**, and click **View Upgrade Plan**, and then click **Apply Updates**. Once the Nutanix Objects Manager is upgraded to the latest version, the setup is ready to start the Nutanix Objects deployment.
16. SSH to Prism Central and run the following command to ensure MSP Controller version on Prism Central is same as shown in Step 4.

```
nutanix@PCVM~$ mspctl controller version
```

17. Go to the **Objects** app from the Application switcher, and then click **Create Object Store** to start the deployment.

TYPES OF NUTANIX OBJECTS USERS

There are two ways to manage Nutanix Objects.

- **Prism Central:** Prism Central users can access Nutanix Objects by using the Prism Central web console. They have the ability to create user accounts and perform all operations except for object-specific actions such as PUT, DELETE, copy, and list objects in any bucket (own or others). Prism Central users can also access a bucket by using the S3 APIs. These users can view buckets of all the API users, and can also share buckets of any API user with any other users.

The following additional types of Prism Central users with varying access levels, as described below, can be defined using role-based access control (RBAC).

- **Nutanix Objects-only admin** - Admin user having access to only Nutanix Objects service within Prism Central. Within Nutanix Objects, you can create the following three different users:
 - Full access to Nutanix Objects service (Nutanix Objects only—full access)
 - Read-only access to Nutanix Objects service (Nutanix Objects only—view access)
 - Restricted access to Nutanix Objects service (user can customize the access)
- **API:** S3 API users (IAM users) cannot access Nutanix Objects by using the Prism Central web console. They access buckets and perform operations only by using the S3 APIs. This includes S3-compatible applications. The API users have unconditional access to their own buckets and limited or no access to buckets of other users based on the shared policy. S3 API users are added using the Nutanix Objects GUI. For more information, see [Generating Access Key for API Users](#) on page 94.

DIRECTORY CONFIGURATION AND ACCESS KEY GENERATION

You can configure the directory, add people, and generate access keys for the people. You can use these directories to search for people who can have access to the service. Only users with an access key can share buckets.

Nutanix Objects IAM-High Availability Overview

The Identity and Access Management-High Availability (IAM-HA) feature is introduced to ensure the high availability of Nutanix Objects IAM across data centers. It provides the ability to withstand Prism Central or primary datacenter failures. Earlier, Nutanix Objects IAM was deployed only on the first Nutanix Objects cluster, and the rest of the Nutanix Objects clusters were dependent on it. IAM-HA ensures that each Nutanix Objects cluster is operational on its own, without depending on the other Nutanix Objects clusters, by deploying IAM on each Nutanix Objects cluster. With IAM-HA, any new deployments will use IAMv2 (the same version of IAM as Prism Central).

IAM-HA is enabled in two stages:

Note: The following two stages are not applicable if there are no Nutanix Objects clusters.

Stage 1: Nutanix Objects Manager is Upgraded: When the Nutanix Objects Manager is upgraded to the latest version, it performs migration of IAM directory service configuration and IAM users from older IAM in the primary Nutanix Objects cluster to the IAM in Prism Central. During this stage, the following limitations apply:

- You cannot update users or keys (such as adding or deleting users or keys) from UI during migration. Only read operations can be performed.
- Operations such as deploy, upgrade, replace certificate, scale-out, set usage alert fails in prechecks.
- S3 APIs are not affected during this stage and they continue to allow all operations.

Stage 2: Nutanix Objects Service Upgrade:

During the Nutanix Objects Service upgrade, in the primary Nutanix Objects cluster, the existing older IAM service is deleted, and IAMv2 is deployed. In the secondary Nutanix Objects cluster, the IAM service was not available previously, so IAMv2 has been newly deployed.

Once the IAM-HA is enabled, the following points are applicable:

- During IAM-HA enablement on clusters, all usernames are converted to lowercase.
- Any newly created username is in lowercase.
- If any existing user has a mixed case in the username, it is converted to lowercase when a new key is added to that user.

Note: This point is applicable only to the IAM-HA cluster deployed before Nutanix Objects 5.0. Any user created or migrated after enabling IAM-HA with Nutanix Objects 5.0 always has lowercase user names.

- Bucket policy evaluation is case-insensitive.
- IAM users cannot be deleted. However, once all keys are deleted, the user disappears from the **Access Keys** page in the user interface.

IAM-High Availability Prerequisites

This section lists the prerequisites for the IAM-HA feature.

The following are the prerequisites for the IAM-HA feature to be enabled automatically:

- Prism Central version must be pc.2024.1 or later.
- The MSP version must be MSP 2.5.0 or later.
- Make sure to open the additional ports needed for the IAM-HA feature. For more information, see [Ports and Protocols Guide](#).

IAM-High Availability Limitations

This section lists the limitations of the IAM-High Availability (IAM-HA) feature.

IAM-HA is not enabled under the following conditions:

- If the same user exists in ObjectsIAM (IAM currently used by Nutanix Objects) and Prism Central IAMv2 but with a different user UUID.
- If the same directory service is configured in both ObjectsIAM (IAM currently used by Nutanix Objects) and Prism Central IAMv2 but with different names. To enable IAM-HA, the directory service name must be the same.

Note: For a successful IAM-HA migration, the directory service names must match in letter case. For example, NTNXdomain and ntnxdomain are considered different directory services.

- If multiple users in ObjectsIAM (IAM currently used by Nutanix Objects) have the same user name but in a different case (user1@ntnx.com vs USER1@ntnx.com), only one such user can exist for IAM-HA to be enabled. To enable IAM-HA, delete any one user from the ObjectsIAM.
- If any stale Active Directory users exist in ObjectsIAM. To enable IAM-HA, delete any stale Active Directory users.

Note: Stale AD users are users belonging to a directory service that is no longer configured in ObjectsIAM (IAM currently used by Nutanix Objects).

- If any of the Active Directory users present in IAM are disabled or do not exist in the Active Directory, IAM migration fails, and IAM-HA is not enabled.

Configuring Directories

If IAM-high availability is disabled, you can configure directories in Nutanix Objects. If enabled, use Prism Central to manage directory settings.

About this task

If IAM-high availability is disabled, you can add directories that Nutanix Objects can use to search for users who might have access to the service. You can also configure multiple Active Directory servers in the user interface and perform searches across one or more Active Directories. However, if IAM-high availability is enabled, you can only view the directories; and cannot configure directories through the Nutanix Objects user interface. You can go to the **Directory Configuration** page on Prism Central to configure the directory.

Non-admin users can configure directories only if assigned a role with **View Object Store** and **Create Object Store** permissions by the admin user in Prism Central. Admins are Super Admins or Prism Admins in Prism Central; non-admins are Prism Central users without admin privileges.

For more information on the built-in roles in Prism Central, see [Built-in Role Management](#) in the *Security Guide*.

About this task

To configure the directories, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. In the Application Switcher, click **Objects**.
3. Click **Access Keys**.

Note: Deletion of the directory is not allowed if a user has been created and access keys are present.




Figure 12: View Directories if IAM-High Availability is Enabled

4. Click **Configure Directories**.
5. Click **+ Add Directory**.




Figure 13: Directory Configuration If IAM-High Availability is Disabled

If you have already added the directory earlier, you will see a list of directories in this screen, and you can edit or remove the directory from the **Edit** or **Remove** link next to the directory name.

6. In the **Add Directory** window, select any one of the following:

a. **Active Directory:** To add a directory through Active Directory, enter the following directory details and service account credentials:

- **Name:** Enter the name of the Active Directory server.
- **Domain:** Enter the domain that represents the top of the Active Directory tree and uses Domain Name System (DNS) to define its namespace.
Usually, the domain is the DNS name of the company.
- **Directory URL:** Enter the Active Directory URL with the port to access the Active Directory server.
For example, ldap://10.2.3.111:389
- **Username:** Enter the username for accessing the Active Directory server to retrieve the user details.
- **Password:** Enter the password for accessing the Active Directory server to retrieve the user details.

Note: To access Nutanix Objects, no expiry must be set on the Active Directory account.

b. **LDAP:** To add a directory through Lightweight Directory Access Protocol (LDAP), enter the following directory details, user and group details hierarchy (or tree under which the user details for generating access keys have to be retrieved), and service account credentials:

- **Name:** Enter the name of the OpenLDAP server.
- **Domain:** Enter the domain that represents the top of the LDAP tree that uses DNS to define its namespace.
Usually, the domain is the DNS name of the company.
- **Directory URL:** Enter the domain that represents the LDAP URL with port to access OpenLDAP server.
For example, ldap://10.2.3.111:389
- **User Object Class:** Enter the LDAP object class value that defines users in the directory service. When the user is created, this list of user object classes is added to the attributes list of the user.
- **User Search Base:** Enter the location or the search starting point in the LDAP tree, which locates the users.
For example, OU=people.
- **Username Attribute:** Enter the attribute names that are searched to retrieve users from the LDAP tree.
- **Group Object Class:** Enter the LDAP object class value that defines groups in the directory service to which the users belong.
When the group is created, this list of group object classes is added to the attribute list of the user.
- **Group Search Base:** Represents the location or the search starting point in the LDAP tree under which groups are located.
For example, OU=people
- **Group Member Attribute:** Enter the member attribute that specifies the group memberships.

- **Group Member Attribute Value:** Enter the group entries for which the memberships are specified by using the member attribute.

These member attributes can have member attribute values specifying group membership in Distinguished Names (DNs). Member attribute values are used for group membership resolution.

- **Username:** Enter the username for accessing the openLDAP server to retrieve the user details.
- **Password:** Enter the password for accessing the openLDAP server to retrieve the user details.

Note: To access Nutanix Objects, no expiry must be set on the LDAP account.

What to do next

You can now generate an access key for the API users. For more information, see [Generating Access Key for API Users](#) on page 94.

Generating Access Key for API Users

You can add people (users) from the directory or through email address, and generate access key for the users. Only users with access keys can share buckets.

Before you begin

A non-admin user can generate access keys only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions: View Object Store and Create an Object Store.

Note: Administrators and non-admin users are Prism Central users with the following roles:

- Administrator—A Super Admin or a Prism Admin in Prism Central.
- Non-admin user—A Prism Central user without any administrator privileges.

For more information on the built-in roles in Prism Central, see [Built-in Role Management](#) in the *Security Guide*.

About this task

To generate access keys for the users, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. In the Application Switcher, click **Objects**.
3. Click **Access Keys**.
4. Click **+ Add People**.

5. In the **Add People** window, select any of the following options:

- a. **Search for people in a directory service:** Select to add people from the directory.

For more information on adding a directory, see [Configuring Directories](#) on page 91.

You can use the Active Directory (AD) group to generate key pairs. Nutanix Objects IAM generates key pairs for each user as a separate file (inside a single zip file). The administrator can distribute these individual key pair files to the end users.

- b. **Add people not in a directory service:** Add the email addresses of the people. Also, you can add a display name for the user.

Adding the display name is optional and can contain up to 255 characters.

Click **+Add** to add multiple users.

6. Click **Next**.

In this page, you generate and download the access keys. You can also add a tag to the access keys for key management.




Figure 14: Generate and Download Keys Page

The **Generate and Download Keys** page appears.

7. (Optional) Select **Apply tag to keys** check box and enter a tag name for the access keys.

If you added multiple users, the same tag applies to all. The tags can contain up to 255 characters.

Tags cannot be changed later. If tags are not set, auto-generated tags apply in version pc.2024.3 with Nutanix Objects Manager 5.0.

8. Click **Generate Keys**.

The keys are generated for the selected people.

9. To download the keys, click **Download Keys**.

The keys are downloaded.

Caution: For Google Chrome, Microsoft Edge, and Internet Explorer, you can directly download the keys. For Safari and Firefox, after you click **Download keys**, a new tab opens listing the keys. You must copy and paste the keys at the desired location manually from the tab. You no longer have access to the keys after you close the tab.

Viewing API Users

You can view the list of people who can have access to Nutanix Objects.

Before you begin

For a user with non-administrator privileges to view API users, the administrator must create an access control policy for the non-administrator user on a role that must have the minimum permission to View Object Store.

About this task

Note: If the Identity and Access Management-High Availability (IAM-HA) feature is enabled, the users with 0 access keys are not displayed on the **Access Keys** page.

To view the list of API users of Nutanix Objects, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. In the Application Switcher, click **Objects**.
3. Click **Access Keys**.

Note: You can add keys for multiple users, but you cannot delete multiple users at the same time.

The table displays the list of API users.

Managing API Keys

If you lose or forget your access keys, you cannot retrieve the same key. However, you can add the key and get access. You can also delete a key.

Before you begin

A non-admin user can manage API keys only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions:

- View Object Store
- Create an Object Store

Note: Administrators and non-admin users are Prism Central users with the following roles:

- Administrator—A Super Admin or a Prism Admin in Prism Central.
- Non-admin user—A Prism Central user without any administrator privileges.

For more information on the built-in roles in Prism Central, see [Built-in Role Management](#) in the *Security Guide*.

About this task

To add or delete the access keys, follow these steps:

Procedure

1. Log on to the Prism Central web console.

2. In the Application Switcher, click **Objects**.
3. Click **Access Keys**.
The table displays the list of API users.
4. Click **Manage** against a user.



Figure 15: Manage Keys Page

The **Manage Keys** page appears. This page provides you with the option to add or delete access keys.

5. To manage a key, you can do one of the following:
 - » **Add Key:** Click this button to generate the access key for the user.
 - » **Apply tag to keys:** Select this option if you want to associate a tag with the access key.

Note: You can add one key at a time and up to five keys for a user.

- » **Delete:** Select the access key and click this button to delete the access key.

The access key of the user gets deleted and has no access to the object.

Note:

- You can delete one key at a time.
- If all keys for a user are deleted, then that user is not displayed on the **Access Keys** page. If you want to add a key to that user, you can generate an access key. For more information, see [Generating Access Key for API Users](#) on page 94.

Deleting API Users

This section describes the procedure to delete an API user.

Before you begin

A non-admin user can delete API users only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions:

- View Object Store
- Create an Object Store

Note: Administrators and non-admin users are Prism Central users with the following roles:

- Administrator—A Super Admin or a Prism Admin in Prism Central.
- Non-admin user—A Prism Central user without any administrator privileges.

For more information on the built-in roles in Prism Central, see [Built-in Role Management](#) in the *Security Guide*.

About this task

Note:

- You cannot delete multiple API users at the same time.
- A deleted user can work up to 3 hours after deletion.

To delete an API user, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. In the Application Switcher, click **Nutanix Objects**.
3. Click **Access Keys**.
The table displays the list of API users.
4. Select the user you want to delete, and then click **Delete User**.
The **Delete User** button appears at the top after you select a user.
The user is deleted.

BUCKET CREATION, OPERATIONS, AND BUCKET POLICY CONFIGURATION

Before you can upload objects to a bucket, you must create a bucket in an object store. You can create a bucket using both S3 and NFS protocols. You can also configure a S3 bucket while creating it except applying the WORM policy.

A user with Edit Buckets permissions can perform all bucket operations, such as create, update, and delete on all the buckets on the assigned Object Store.

You can enable Lifecycle Policies, Versioning, WORM, Replication, Static Website, CORS, and Notifications to an S3 bucket. These S3 features cannot be enabled in a multi-protocol access bucket (NFS). However, you can also filter, update, share, and delete a bucket, and view all the existing buckets list created through both protocols.

Bucket Naming Conventions

Naming a bucket requires adherence to specific rules.

The name of a bucket must conform to the following rules:

- Be a unique Domain Name System (DNS)-compliant name within a deployed bucket instance.
- Can contain alphanumeric, dot, or hyphen characters.
- Begin with a lowercase letter or a number.
- Cannot contain uppercase and special characters.
- Minimum of 3 and a maximum of 64 characters long.

Note: You cannot change the bucket name after creating the bucket.

Creating and Configuring an S3 Bucket in Nutanix Objects

You can create and configure the bucket settings. There are multiple options available for configuring a bucket and all are accessible at the time of creating a bucket. However, you cannot configure a WORM bucket while creating a bucket. You can edit WORM policies only after creating a bucket. You cannot enable multi-protocol access on the S3 bucket.

Before you begin

Ensure that the bucket names are unique for all the users.

A non-admin user can create and configure S3 buckets only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions and the non-admin user must be assigned to specific objects instances:

- Edit Buckets
- View Buckets
- View Object Store

Note:

- Administrators and non-admin users are Prism Central users with the following roles:

- Administrator—A Super Admin or a Prism Admin in Prism Central.
- Non-admin user—A Prism Central user without any administrator privileges.

For more information on the built-in roles in Prism Central, see [Built-in Role Management](#) in the *Security Guide*.

- Users with Edit Bucket permission can perform a create bucket operation on the buckets available in the object stores that are assigned to the respective users.

About this task

Note: While creating a new versioned bucket, the system will automatically create a lifecycle policy to expire the delete markers if they are the only remaining object versions.

To create and configure a bucket, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. In the Application Switcher, click **Objects**.
3. In the object stores table, click the name of the object store in which you want to create a bucket. The **Object Store** page appears.
4. Click **Buckets > Create Bucket**.

Note: You can also create a bucket without enabling versioning and life cycle policies.

5. On the **General Settings** section, type a name for your bucket and select a namespace.
For more information on namespace, see [Federation Overview](#) on page 57 .
For more information on naming buckets, see [Buckets Naming Conventions](#).
6. (Optional) On the **Object Versions** section, configure the following:
For more information on versions, see [Object Versioning](#).
 - a. **Enable Versioning:** Select this check box to enable versioning on objects and to keep all the versions on the same bucket.
To apply the life cycle policy with versioning enabled, see [Rules for Expiration based Lifecycle Policy](#) on page 107.

Note: When you select versioning, you are able to recover objects from accidental deletion or overwriting.

7. (Optional) On the **Lifecycle Policies** section, select the following:

For more information, see [Lifecycle Policies](#).

a. **Expire current objects version after::** Select to type a time period after which the current version of the object expires.

You can specify the number in days, months, or years.

b. **Expire previous objects versions after:** Select this check box to enter a time period to delete all the previous versions of the objects. This option is available only if versioning is enabled.

You can specify the number in days, months, or years.

Note:

- If versioning is not enabled, the current object deletes permanently. When you enable versioning, the current object becomes a past object.
- Multi-protocol access cannot be enabled on the S3 bucket. If you want to create buckets with multi-protocol access, see [Creating and Configuring an NFS Bucket in Nutanix Objects](#) on page 101.

8. Click **Save**.

The bucket is created successfully.

What to do next

After creating a bucket, you can perform object operations from the Nutanix Objects Browser or the S3 APIs. For more information, see [Object Operations in Nutanix Objects Browser](#) on page 213 and [Supported S3 APIs](#). You can also configure object-tiering lifecycle rules. For more information, see [Creating a Lifecycle Rule in Nutanix Objects](#) on page 114.

Creating and Configuring an NFS Bucket in Nutanix Objects

You can create and configure the bucket with multi-protocol access. If you are enabling multi-protocol access on a bucket, S3 features such as Lifecycle Policies, Versioning, WORM, Replication, Static Website, CORS, and Notifications cannot be enabled on that bucket. You can perform update (only multi-protocol access configurations), delete, or share operations on these buckets.

Before you begin

Make sure you are aware of the uses cases, recommendations, NFS-S3 interoperability, and limitations of NFS on Nutanix Objects. See [Use Cases and Recommendations for Network File System on Nutanix Objects](#) on page 13, [Network File System and Simple Storage Service Interoperability](#) on page 14 and [Limitations of Network File System](#) on page 49.

A non-admin user can create and configure NFS buckets only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions and the non-admin user must be assigned to specific objects instances:

- Edit Buckets
- View Buckets
- View Object Store

Note:

- Administrators and non-admin users are Prism Central users with the following roles:
 - Administrator—A Super Admin or a Prism Admin in Prism Central.
 - Non-admin user—A Prism Central user without any administrator privileges.

For more information on the built-in roles in Prism Central, see [Built-in Role Management](#) in the *Security Guide*.

- Users with Edit Bucket permission can perform a create bucket operation on the buckets available in the object stores that are assigned to the respective users.

About this task

- If you deployed object store prior to Nutanix Objects 3.3 version, then you will not get the option to enable NFS for the bucket of that object store. Similarly, if you upgrade to later versions of Nutanix Objects, then the option to enable NFS will not be available.
- Ensure that the bucket names are unique for all users.
- As the Nutanix Objects-NFS does not support NLM (Network Lock Manager), no lock option is required while mounting the NFS bucket.
- The total and available bytes returned in the FSSTAT response denote the logical capacity and logical available space and not the physical capacity of the cluster which also takes RF2 into consideration.

To create and configure a bucket with multi-protocol access, follow these steps:

Procedure

- Log on to the Prism Central web console.
- In the Application Switcher, click **Objects**.
- In the object stores table, click the name of the object store in which you want to create a bucket. The **Object Store** page appears.
- Click **Buckets > Create Bucket**.

The screenshot shows the 'Create Bucket' dialog box. It includes a 'Lifecycle Policies' section with a dropdown menu. Below that is the 'Multiprotocol Access' section, which is currently disabled, with a note stating 'Access cannot be enabled or disabled once the bucket is created.' The 'Owner and Default permissions for S3 written objects' section contains fields for 'Owner' and 'Group'. The 'Advanced Settings' section at the bottom includes fields for 'Name', 'Access Key ID', and 'Secret Key'. At the bottom right, there are 'Cancel' and 'Create' buttons.

Figure 16: Create Bucket Window

5. On the **General Settings** section, type a name for your bucket and select a namespace.
For more information on namespace, see [Federation Overview](#) on page 57.
For more information on naming buckets, see [Buckets Naming Conventions](#).
Versioning and lifecycle policies cannot be enabled on buckets with multi-protocol access.
For more information on creating and configuring buckets for S3 features, see [Creating and Configuring an S3 Bucket in Nutanix Objects](#) on page 99
6. On the **Multiprotocol Access** section, select **Enable NFS v3 Access**.
Access cannot be enabled or disabled after the bucket is created.
7. For owner and default permissions for S3 written objects, do the following.
For files written using NFS protocol, these settings are inherited from the client.
 - a. In the **Owner** section, enter the **UID** and **GID**.
Any object (file or directory) created from S3 protocol have UserID (UID) and GroupID (GID) in the NFS namespace.
 - b. In the **Default Access Permissions** section, by default read, write, and execute permissions are set for the files and directories for the owner, group (other users in a group), and others (users that are not part of any group).
The following are the default permissions for files:
 - **Owner:** Read, Write
 - **Group:** Read
 - **Others:** ReadThe following are the default permissions for directories:
 - **Owner:** Read, Write, Execute
 - **Group:** Read, Execute
 - **Others:** Read, ExecuteYou can change these permissions as needed.
8. In the **Advanced Settings** section, select any one of the following squash options.
 - » **None:** Select this option if you do not want to convert the UID and GID of the users on the server.
 - » **Root Squash:** Select this option if the user has root privileges and you want to convert the UID and GID to an anonymous UID and GID on the server. The anonymous UID and GID are automatically generated, however, you can change it.
 - » **All Squash:** Select this option if you want to map all users to a single identity. This will convert the UID and GID of all users to an anonymous UID and GID on the server.
9. Click **Create**.
The bucket is created successfully.

10.

Note: Before mounting a bucket, add the client to the NFS allowlist. Only the client present in the NFS allowlist will be given access to the NFS buckets.

For more information on adding and managing clients to NFS allowlist, see [Managing NFS Allowlist in Nutanix Objects](#) on page 104.

Mount the bucket from the client VM.

```
$ sudo mount -tnfs -o nfsvers=3,proto=tcp,soft -v objectstore-endpoint-ip-address:/bucketname path/to/mount
```

Here *objectstore-endpoint-ip-address* is the Nutanix Objects Public IP address which is used to send any S3 request. You can find this Public IP address in the Object Store table. For more information, see [Viewing Object Store Deployments](#) on page 72. *bucketname* is the name of the bucket.

For example,

```
$ sudo mount -tnfs -o nfsvers=3,proto=tcp,soft -v 10.45.53.75:/test-nfs-bucket home/folder-1/mnt-point
```

The bucket is mounted successfully.

What to do next

After creating a bucket, you create directories, files, and symbolic links from the NFS namespace. You can also perform object operations from the Nutanix Objects Browser or the S3 APIs. For more information, see [Object Operations in Nutanix Objects Browser](#) on page 213 and [Supported S3 APIs](#).

Managing NFS Allowlist in Nutanix Objects

You need to add the IP addresses of the client VMs that are allowed to access an NFS bucket in an object store.

About this task

Note: Make sure that you add the required client IP addresses to the allowed list before you mount a bucket.

To add clients to the NFS, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. In the Application Switcher, click **Objects**.
3. In the object stores table, click the name of the object store in which the bucket exists. The **Object Store** page appears.
4. Click **Settings > NFS Clients Allowlist**, and then click **Add Client**. The **Add NFS Clients** window appears.
 - a. Enter the IP address of the client VM in Classless Inter-Domain Routing (CIDR) format and click **Add**.
 - b. Click **Save**.
The newly added IP address is listed in the **Added Clients** list.

Note: The **Add** button appears only when you add the first IP address to the allowlist. To add more clients or to manage the existing clients, click **Manage Clients**.

5. Click **Manage Clients** to add or remove the clients.

The **Manage NFS Clients** window appears.

- a. (Optional) Enter the IP address of the client VM in Classless Inter-Domain Routing (CIDR) format and click **Add** to add client.
- b. (Optional) Select one or more IP addresses from **Client IP(s)** and click **Remove** to remove the clients.
- c. Click **Save**.
The IP address is removed from the **Added Clients** list.

What to do next

After adding the required clients to the NFS allowlist, you can mount the buckets and then perform tasks, such as creating files, directories, or symbolic links. For more information on mounting the bucket, see [Creating and Configuring an NFS Bucket in Nutanix Objects](#) on page 101.

Bucket Policy Configuration in Nutanix Objects

You can configure multiple policies in a bucket such as object versioning, lifecycle policies, WORM, static website hosting and Cross-Origin Resource Sharing (CORS). Object versioning and lifecycle policies are accessible at the time of creating a bucket, however, you can apply WORM policy only after creating a bucket.

Note: Lifecycle Policies, Versioning, WORM, Replication, Static Website, CORS, and Notifications cannot be enabled for buckets created using NFS protocol.

Object Versioning in Nutanix Objects

Object versioning enables you to keep multiple versions of an object in one bucket. By default, versioning is disabled for a new bucket. You can enable versioning while creating a bucket or editing a bucket. See [Creating and Configuring an S3 Bucket in Nutanix Objects](#) on page 99.

With the object versioning option, you can enable versioning on objects of that particular bucket. Object versioning varies depending on the lifecycle policies applied to the object.

For more information on the lifecycle rules for object versioning, see [Rules for lifecycle policy](#).

Note:

- Versioning cannot be enabled for the buckets created using the NFS protocol.
- For a versioned bucket, the number of objects shown for each bucket is indicative of the number of versions of all the objects present in the bucket.
- You cannot disable but can suspend the object versioning at any time.

When you suspend versioning, the accumulation of the new object versions is stopped and previous object versions are retained.

A non-admin user can enable versions on objects only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions and the non-admin user must be assigned to specific objects instances:

- Edit Buckets
- View Buckets
- View Object Store

Note:

- Administrators and non-admin users are Prism Central users with the following roles:

- Administrator—A Super Admin or a Prism Admin in Prism Central.
- Non-admin user—A Prism Central user without any administrator privileges.

For more information on the built-in roles in Prism Central, see *Built-in Role Management* in the *Security Guide*.

- Users with Edit Bucket permission can perform object versioning on the buckets available in the object stores that are assigned to the respective users.

Lifecycle Policies in Nutanix Objects

Lifecycle policy enables you to create or update a set of rules that define actions that Nutanix Objects applies to a group of objects.

With these policies, you can expire objects when no longer required or move them to a low-cost storage tier to preserve for a longer-term.

A non-admin user can create lifecycle policies only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions and the non-admin user must be assigned to specific objects instances:

- Edit Buckets
- View Buckets
- View Object Store

Note:

- Administrators and non-admin users are Prism Central users with the following roles:

- Administrator—A Super Admin or a Prism Admin in Prism Central.
- Non-admin user—A Prism Central user without any administrator privileges.

For more information on the built-in roles in Prism Central, see [Built-in Role Management](#) in the *Security Guide*.

- Users with Edit Bucket permission can apply lifecycle policies on the buckets available in the object stores that are assigned to the respective users.
- Lifecycle policies cannot be applied for the buckets created using NFS protocol.
- Rules or any updates to the rules get applied to the new objects that you create and will not apply to the objects existing before the rule creation or update.
- Starting from Nutanix Objects 5.2 the non- retroactive lifecycle policies only apply to new or modified objects added after you create the policy. The retroactive lifecycle policies apply to all existing objects, plus the new objects that are added later. On the older buckets, you have an option in the UI to enable retroactive lifecycle policies.

To change to retroactive, see [Evaluating Lifecycle Policy](#) on page 108.

You can apply these policies while creating a bucket. You can create multiple rules within a lifecycle policy. This means that different objects within the bucket can have different rules based on prefixes and tags.

Example: You can create rule 1 to expire the current versions of the objects with the tag value as a dev. Similarly, you can create multiple rules with different tiering and expiring configurations and apply them to other objects using prefixes and tags.

With lifecycle policies, you can configure a lifecycle policy rule to:

- Automatically delete objects after a specified number of days or months or years from the date of object creation.
- Tier objects to an S3-compatible object storage bucket after a specified number of days or months, or years from the date of tiering rule creation.
- Expire the current version and previous versions of an object independently. This means that you can set different expiration durations for the current version and the previous versions.
- Tier the current version and previous versions of an object independently on buckets with retroactive lifecycle policies enabled.
- Expire the incomplete multi-part uploads of an object.
- Apply to all objects or a subset of objects based on prefixes, tags, or both.

For example, if you want to store log files or business transaction records for a fixed period and after that period, you want to delete them.

Note: You cannot recover the objects once it is deleted.

To create or update a lifecycle rule, see [Creating a Lifecycle Rule in Nutanix Objects](#) on page 114.

If retroactive lifecycle policies are not enabled on the bucket, lifecycle policies with the scope of a tag are not retroactive. Only tags that were applied along with Object PUT are considered. Tags applied to an object after its upload is not considered by lifecycle policies.

With retroactive lifecycle policies enabled, all rules (with both tag and prefix filters) are retroactively applied to the objects

Rules for Expiration based Lifecycle Policy

Following are the rules for expiration based lifecycle policy:

- If you apply lifecycle policy **Expire current objects after # days/months/years** on an object with versioning enabled, it deletes the current version of the objects after the specified time, and does not delete any past versions of the objects.
- If you apply lifecycle policy **Expire previous objects versions after # days/months/years** on an object with versioning enabled, it deletes all the past versions after the specified time. This specified time gets calculated from the day the object version becomes non-current or past. This operation does not delete the current version.

Note: This policy cannot be configured on a bucket with suspended versioning.

- If you apply both the lifecycle policies **Expire current objects after # days/months/years** and **Expire previous objects versions after # days/months/years** on an object with versioning enabled, it deletes all the past versions based on the time specified in **Expire previous objects versions after # days/months/years**, and the current version expires based on the time specified in **Expire current objects after # days/months/years**.
- If you apply expiration based lifecycle policy **Expire or Abort incomplete multipart uploads after # days/months/years** on an object, it deletes the parts associated with the multipart uploads after the specified time.

- If you apply any lifecycle policy to a WORM bucket with versioning enabled, the lifecycle policy is not applicable until the WORM retention period has elapsed.
- When creating a new versioned bucket, the system automatically creates a lifecycle policy to expire the delete markers if they are the only remaining object versions.

Rules for Tiering based Lifecycle Policy

Following are the rules for the lifecycle policy when you enable tiering to a given endpoint:

- If you apply lifecycle policy **Tier current objects after # days/months/years** on an object with versioning enabled, it tiers out the current version of the objects after the specified time, and does not tier out any past versions of the objects.
- If you apply lifecycle policy **Tier previous objects versions after # days/months/years** on an object with versioning enabled, it tiers out all the past versions after the specified time. This specified time gets calculated from the day the object version becomes non-current or past. This operation does not tier out the current version.

Note: This policy cannot be configured on a bucket with suspended versioning.

- If you apply both the lifecycle policies **Tier current objects after # days/months/years** and **Tier previous objects versions after # days/months/years** on an object with versioning enabled, it tiers all the past versions based on the time specified in **Tier previous objects versions after # days/months/years**, and the current version tiers based on the time specified in **Tier current objects after # days/months/years**.

Evaluating Lifecycle Policy

Enabling Retroactive policy enforcement will apply your latest configured lifecycle policy to all existing objects.

About this task

Note: All the lifecycle policies on buckets created after upgrading to Nutanix Objects 5.2 will be retroactive by default. On older buckets, there is an option to enable retroactive lifecycle policy management on the bucket.

To enable lifecycle policy evaluation, do the following:

Procedure

1. Log on to the Prism Central web console.
2. In the Application Switcher, click **Objects**.
3. In the object stores table, click the name of the object store that contains the bucket where you want to edit the lifecycle configuration.
The **Object Store** page appears.
4. Click the **Buckets** tab.
5. In the **Buckets** table, click the bucket that you want to edit the lifecycle configuration.
The **Bucket** page appears.
6. Click the **Lifecycle** tab.
You can view the rules that you defined while creating the bucket.
7. In the **Policy Enforcement** table, click **Change to Retroactive**
The **Change enforcement to Retroactive** page appears.

8. Enter the **Bucket Name** in the text field.

9. Click **Change**.

In the **Lifecycle** tab, the **Policy Enforcement** Type is **Retroactive**.

Retroactive Policy Limitation

The following section lists the limitations of enabling retroactive policy.

- If an object is already tiered to a remote end point and the policy is changed to use a different endpoint, the object will not be re-tiered.
- S3 imposes a limit of 10,000 parts per object. Objects that exceed this limit cannot be tiered out.
- Large objects that exceed the maximum supported size for S3-compatible or Azure endpoints are not eligible for tiering.

For example, Azure supports a maximum single object PUT size of 5000 MiB, while Nutanix supports up to 5120 MiB. As a result, objects between 5000 MiB and 5120 MiB will not be tiered when using Azure as the target endpoint.

- Retroactive tiering increases API costs for remote endpoints because operations are performed on a per-object basis rather than in bulk.

Cloud Tiering in Nutanix Objects

Cloud tiering enables you to move objects to another S3-compatible object store bucket for saving storage space in the Nutanix Objects cluster. Tiering can help you to save costs by sending the infrequently accessed objects to platforms such as AWS S3, Microsoft Azure Blob Storage, and Google Cloud Platform (GCP). The supported endpoints are AWS S3, Azure Blob Storage, Google Cloud Platform (GCP) and a different Nutanix Objects instance.

A non-admin user can configure a tiering endpoint only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions and the non-admin user must be assigned to specific objects instances:

- Create Endpoint
- Update Endpoint
- View Endpoint
- View Object Store

Note: Administrators and non-admin users are Prism Central users with the following roles:

- Administrator—A Super Admin or a Prism Admin in Prism Central.
- Non-admin user—A Prism Central user without any administrator privileges.

For more information on the built-in roles in Prism Central, see [Built-in Role Management](#) in the *Security Guide*.

Cloud tiering is managed through lifecycle policies. You can configure multiple lifecycle rules for different objects within a bucket.

Cloud tiering configuration consists of the following steps:

- **Step 1: Configure a remote endpoint in the object store**

You can configure remote endpoints using the **Cloud Bucket Endpoints** page in the object store instance. The objects within the bucket get moved to the endpoint according to the rules you configure. For more information, see [Configuring a Cloud Bucket Endpoint in Nutanix Objects](#) on page 111.

- **Step 2: Create lifecycle rules for a bucket**

You can create a tiering rule for a bucket where you define the scope (all objects or a subset of objects), select the remote endpoint, and specify the details for the rule. For more information, see [Creating a Lifecycle Rule in Nutanix Objects](#) on page 114.

Points to Note

Before you start with object tiering, note the following points:

- The following points are strongly recommended:
 - Only encrypted data is stored on buckets for which tiering to the cloud is enabled.
 - Object store admins enable audit trails for the S3 bucket or other storage endpoints to ensure data is not being accessed or tampered with external malicious entities and that all access is only coming from the object store instance.
 - Admins follow recommended security best practices of AWS or other storage endpoints while setting up buckets for tiering.
- Tiering lifecycle policies are non-retroactive. The policy gets applied to the new objects that you create. The policy will not apply to existing objects.
- Tiering lifecycle policies get applied to both versioned and non-versioned objects in the same way. A separate non-current version transition action in lifecycle policies is not supported.
- Do not perform write operations on the configured-endpoint bucket. Ensure that the endpoint bucket gets used only for the object store instance.
- Removing a configured endpoint in the object store instance is not supported.
- There is a N:N relationship between the source bucket and the configured endpoint bucket. You can create multiple tiering-lifecycle rules for different objects within a bucket and use a separate endpoint for each tiering rule. Also, an endpoint bucket can be a destination to many Object Store buckets.
- Nutanix Objects within a WORM-enabled bucket will continue to adhere to their WORM property even after getting tiered out to the endpoint bucket.
- Only the object data gets tiered to the endpoint bucket. The metadata of the object is not tiered.
- The behavior to access tiered objects using the `Object Get` method remains the same. If you send a request to retrieve an object from the Object store and the object is already tiered out, the Object store fetches the data from the endpoint bucket and fulfills your request.
- Users with **Edit Bucket** permission can create lifecycle rules bucket operation on the buckets available in the Object stores that are assigned to the respective users.
- Tiered objects are not readable on the tiering endpoint as they are moved in an unstructured form, without metadata.
- You can check if the objects are being tiered on the source Bucket **Summary > Tiering**. Also tiering endpoint bucket will grow in size and show data blocks.
- Once the specified time period has elapsed, an object becomes eligible for tiering out. However, it is important to note that the actual movement of the object may occur up to one day later.
- When reading a tiered object, the Nutanix Objects cluster retrieves the object data from the tiering endpoint cluster and presents it to the client, without storing it locally that is no caching is performed.

Configuring a Cloud Bucket Endpoint in Nutanix Objects

You can configure a cloud bucket endpoint for both tiering and replication. The infrequently used objects are moved to the configured endpoints according to the configured lifecycle rule for tiering. The supported endpoints for tiering are AWS S3, Microsoft Azure Blob Storage, Google Cloud Platform (GCP), and a different Nutanix Objects instance.

Before you begin

Make sure the bucket used for replication must be created with the recommended settings. For more information, see [Cloud Bucket Endpoint Recommended Settings](#) on page 113.

About this task

Note:

- Cloud tiering supports GCP endpoint as *Other S3 compatible* endpoint type.
- For tiering the objects to GCP, you need to create a bucket in GCP with the following default configurations. This bucket will be used as the tiering endpoint.
 - **Location type:** Multi-region
 - **Default storage class:** Standard
 - **Public access prevention:** Off
 - **Access control:** Uniform
 - **Protection tools:** None
- Ensure that the tiering endpoint has a Certificate Authority (CA) signed certificate. Note that self-signed certificates are not supported.
- To enable tiering and replication to the cloud, the endpoint must be connected to the storage network. For example, tiering to S3 requires the S3 ports (80 and 443) to be open and accessible, and the endpoint `s3.amazonaws.com` to be reachable from the Nutanix Objects storage network.

To configure a cloud bucket endpoint, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. In the Application Switcher, click **Objects**.
3. In the object stores table, click the name of the object store to configure the cloud bucket endpoint. The **Object Store** page appears.
4. Click **Endpoints > Cloud Bucket Endpoints**.
5. To open the **Add Cloud Bucket Endpoint** page, click **Create**.

If no endpoints are configured, the **Add Endpoint** button appears in the center of the page. Click **Add Endpoint** to start an endpoint configuration.

6. (Only for Nutanix Objects as the endpoint) Enter the following details:
 - a. **Description:** Enter a description that identifies your cloud bucket endpoint.
 - b. **Endpoint Provider:** Select **Nutanix** as the endpoint type from the pull-down list.
 - c. **Service Host:** Enter the complete URL of the endpoint that you want to use for tiering the objects.
For example, the Nutanix Objects public IP address or the domain name of an Object store instance.

Example of a Nutanix Objects endpoint: example.buckets.company.com.
 - d. **Bucket Name:** Enter the name of the bucket within the service host to which you want the objects to tier out.
 - e. **Access Key:** Enter the access key of the bucket owner.
 - f. **Secret Key:** Enter the secret key of the bucket owner.
 - g. **Skip SSL Certificate Validation:** Select this check box to skip SSL certificate validation.
7. (Only for AWS as the endpoint) Enter the following details:
 - a. **Description:** Enter a description that identifies your cloud bucket endpoint.
 - b. **Endpoint Provider:** Select **AWS** as the endpoint type from the pull-down list.
 - c. **Service Host:** Enter the complete URL of the endpoint that you want to use for replication or tiering the objects.

The following AWS S3 URL formats are supported:
 - s3.region.amazonaws.com
 - s3-region.amazonaws.com
 - s3.region.amazonaws.comExample of an AWS S3 endpoint: s3.us-east-1.amazonaws.com.
 - d. **Bucket Name:** Enter the name of the bucket within the service host to which you want the objects to tier out or replicated.
 - e. **Access Key:** Enter the access key of the bucket owner.
 - f. **Secret Key:** Enter the secret key of the bucket owner.
8. (Only for Azure Blob Storage as the endpoint) Enter the following details:
 - a. **Description:** Enter a description that identifies your cloud bucket endpoint.
 - b. **Endpoint Provider:** Select **Azure** as the endpoint type from the pull-down list.
 - c. **Container Name:** Enter the name of the container within the service host to which you want the objects to tier out.
 - d. **Account Name:** Enter the Azure account name where the containers are located.
 - e. **Secret Key:** Enter the secret key of the container owner.

9. (Only for Other S3 Compatible as the endpoint) Enter the following details:
 - a. **Description:** Enter a description that identifies your cloud bucket endpoint.
 - b. **Endpoint Provider:** Select **Other S3 Compatible** as the endpoint type from the pull-down list.
 - c. **Service Host:** Enter the complete URL of the endpoint that you want to use for tiering the objects.
The AWS S3 URL format is `s3.region.amazonaws.com`
Example of an AWS S3 endpoint: `s3.us-east-1.amazonaws.com`.
 - d. **Bucket Name:** Enter the name of the bucket within the service host to which you want the objects to tier out.
 - e. **Access Key:** Enter the access key of the bucket owner.
 - f. **Secret Key:** Enter the secret key of the bucket owner.

10. Click **Save** to complete configuring the cloud bucket endpoint.

You can also update the access and secret keys of the bucket for the configured endpoints. For example, if the bucket owner generates new access and secret keys, the owner can update the configured endpoint with the new keys using the **Update** option.

Select the endpoint and click **Update** to open the **Update Endpoint** page. Update the required fields and click **Save** to complete the workflow.

What to do next

After you configure an endpoint, you can view the cloud bucket configured endpoints, create lifecycle rules for tiering objects within your bucket, or create a replication relation. For more information, see [Viewing Cloud Bucket Endpoints](#) on page 117, [Creating a Lifecycle Rule in Nutanix Objects](#) on page 114 or [Creating Replication Relation for Buckets](#) on page 149.

Cloud Bucket Endpoint Recommended Settings

This section describes the permission settings required on the destination S3 bucket configuration and lock configuration.

Permissions Required for the Destination S3 Bucket Configuration

- Creating an endpoint involves testing read and write permissions by adding and deleting objects on the destination bucket. Therefore, the destination bucket should not have WORM (Write Once Read Many) enabled when creating an endpoint, as deleting it will fail. After creating the endpoint, you can enable WORM on the destination bucket. If the destination bucket is already WORM-enabled, you see the following error message: Method Not Allowed: An object from the object-lock enabled bucket can not be modified or deleted unless the retention period is elapsed. If you are unable to remove WORM to configure the endpoint, contact Nutanix support for assistance with configuring it.
- The destination bucket must have all the object-level (such as HEAD, GET, PUT, DELETE object, WORM, tags-related permission, and so on) and bucket-level permissions (such as HEAD, tags, and so on) same as the source bucket.
- The lock permission on the source and destination must be the same.

For example, the following is a safe configuration with only access to a particular bucket:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "S30013200001BbN8ZAAVqa",
```

```

        "Effect": "Allow",
        "Action": [
            "s3:*"
        ],
        "Resource": [
            "arn:aws:s3:::bucket-name",
            "arn:aws:s3:::bucket-name/*",
        ]
    }
]
}

```

Caution: Changing these permissions after replication rule creation might cause replication operation failure.

Creating a Lifecycle Rule in Nutanix Objects

You can create a lifecycle rule for a bucket where you define the scope (all objects or a subset of objects), select the remote endpoint, and specify the details for the rule.

Before you begin

Make sure to configure a tiering endpoint before creating object-tiering lifecycle rules.

A non-admin user can create lifecycle rules for a bucket in cloud tiering only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions and the non-admin user must be assigned to specific objects instances:

- Edit Buckets
- View Buckets
- View Endpoint
- View Object Store

Note:

- Administrators and non-admin users are Prism Central users with the following roles:
 - Administrator—A Super Admin or a Prism Admin in Prism Central.
 - Non-admin user—A Prism Central user without any administrator privileges.

For more information on the built-in roles in Prism Central, see [Built-in Role Management](#) in the *Security Guide*.

About this task

When creating a bucket, you can enable versioning, set a rule to delete past versions, and add a lifecycle policy to expire the current objects after the specified time. Once the bucket gets created, you can go to the bucket page and configure object-tiering lifecycle rules.

To create a lifecycle rule to tier out objects, do the following:

Procedure

1. Log on to the Prism Central web console.
2. In the Application Switcher, click **Objects**.

3. In the object stores table, click the name of the object store that contains the bucket where you want to create tiering rules.
The **Object Store** page appears.
4. Click the **Buckets** tab.
5. In the **Buckets** table, click the bucket that you want to tier out.
The **Bucket** page appears.
6. Click the **Lifecycle** tab.
You can view the rules that you defined while creating the bucket.
7. Click **Create Rule** to open the **Create Rule** page.
You can also import lifecycle rule by clicking **Create Rule > Import XML**. Type or copy the lifecycle rule in XML format.
8. In the **Scope** page, do the following, and then click **Next**:
 - a. In the **Name** box, enter a name that identifies the rule you are creating.
 - b. In the **Scope** list, select **All objects** or **Tags/Prefix**.
 - Select **All objects** to apply the tiering rule to all the objects within the bucket.
 - Select **Tags/Prefix** to apply the tiering rule to specific objects. You can filter objects by entering a prefix, tags, or both. Nutanix Objects with the prefix and tags you specified get filtered and the tiering rule applies to those objects.

Nutanix Objects is a flat namespace and directories are a file construct. Directories may be simulated in Nutanix Objects but they are not true directories.

Lifecycle policy prefix should include the entire directory hierarchy.

Some examples of prefix and explanation which files (objects) it will apply to:

- **log** - Only files/folders in the root of the bucket that start with "log".
- **serverA**- All files in folder "serverA", which is in the root of the bucket.
- **serverB/log** - Files in folder "serverB" that start with "log".
- **server*/log**- Variable prefixes (for matching multiple folders) are not supported.

9. In the **Configure Rule** page, do the following, and then click **Next**.

a. To create a tiering rule, select the **Tiering** check box.

1. In the **Tier** list, select **current version** according to your requirement.

For a retroactive lifecycle policies enabled bucket, **Previous version** is available for tiering.

2. In the **Endpoint** list, select the endpoint where you want to tier out the objects.

3. Enter a time period after which you want to move the objects to the selected endpoint. You can specify the number in days, months, or years.

Tiering time must be less than expiration time.

The screenshot shows the 'Configure Rule' page with three tabs: 'Scope' (checked), 'Configure Rule' (active), and 'Review'. The 'Tiering' section is expanded, showing a checkbox for 'Tiering' which is checked. Below it, there are three dropdown menus: 'Tier' (set to 'Current version'), 'Endpoint' (set to 'Test Endpoint 1'), and 'Tier After' (set to '1 days'). The 'Expiration' section is also expanded, showing a checkbox for 'Expiration' which is checked. Below it, there are two dropdown menus: 'Expire' (set to 'Current version') and 'After last creation date' (set to '4 days'). At the bottom of the 'Expiration' section, there is a button labeled '+ Add Action'. At the bottom of the page, there are three buttons: '< Back', 'Cancel', and 'Next'.

Figure 17: Creating Tiering Rule

b. To create an expiration rule, select the **Expiration** check box.

1. In **Expire**, select **Current version**, **Previous version**, **Multipart uploads**, or **Delete markers** according to your requirement.

The **Previous version** and **Delete markers** options appear only for a version-enabled bucket.

The expiration of delete marker and the current version cannot be configured in a single rule.

Tip: It is recommended to create a rule to automatically expire delete markers to enhance the performance of your object stores.

2. Enter a time period after which you want the objects to expire. You can specify the number in days, months, or years.

You can click **Add Action** to add multiple expiration rules. You can create expiration rules for the current version, previous version, and multipart uploads.

10. On the **Review** page, check your configurations. Then, click **Done** to create the rule.

The rule you just created gets enabled and appears in the Rules table.

- You can select a bucket, and update, delete, disable, and enable the rule using the **Actions** drop-down.
- You can also export multiple rules to an XML file by clicking **Export to XML**.

Viewing Cloud Bucket Endpoints

You can view all the cloud bucket endpoints configured for both replication and tiering.

About this task

To view all the configured cloud bucket endpoints, follow this steps:

Procedure

1. Log on to the Prism Central web console.
2. In the Application Switcher, click **Objects**.
3. In the object stores table, click the name of the object store to configure the cloud bucket endpoint. The **Object Store** page appears.
4. Click **Endpoints > Cloud Bucket Endpoints**.
The endpoints are listed in the Cloud Bucket Endpoint table.

Tiering Statistics

Cloud tiering supports endpoints such as, Nutanix Objects, Other S3 Compatible, and Azure Blob Storage. You can view the amount of object data moved to the endpoint bucket and the amount of pending data. You can also view the statistics for the source bucket (Object Store bucket) and the endpoint bucket.

Object Store Bucket - Statistics

In the object store instance page, click the **Bucket** tab. Click the bucket name, and then click the **Summary** tab to view its tiering statistics.

You can view the following information:

- **Space pending reclamation:** The amount of the deleted data and metadata left out due to incomplete operations.
- **Total size of objects marked for tiering (pending):** The amount of the object data (eligible for tiering) that is in the process of tiering.
- **Tier out object size:** The amount of object data that has been tiered from this source bucket to the endpoint buckets.

Note: If all tiering endpoints are in idle state for a long period of time, such as a day or two, then the usage statistics displays zero usage. However, the actual usage statistics reflects in the user interface some time (approximately 30 minutes or so) after the next tiering task starts.

The tiering statistics are updated in the user interface after a tiering task is completed. However, when large amount of data is tiered to an endpoint, you might experience a delay in seeing the updated statistics in the user interface.

Legal Hold for Nutanix Objects

Legal hold allows you to lock an object indefinitely. There is no expiration date. A legal hold remains active until an authorized user explicitly removes it. Implementing a legal hold prevents the object data from being modified or deleted.

There are various cases where you would want to implement a legal hold on your object data. For example, put an indefinite lock on objects that you want to preserve for legal cases, audits, or compliance purposes.

Note the following points about the legal hold functionality in Nutanix Objects:

- A user with the write permission on the bucket can apply a legal hold to the objects within the bucket.

- Only an administrator or bucket owner can remove the legal hold applied to the objects.
- A legal hold can be applied while writing an object or to an existing object also.
- If a bucket gets created using the Prism GUI and a user with write permissions applies a legal hold to the objects within the bucket, an administrator user needs to be set up to remove the legal hold. Contact Nutanix support to set up an administrator user with the necessary permissions. It is recommended that the user should create a bucket. In this case, the user will be the bucket owner and will have permission to remove the legal hold.
- If you are replicating a bucket created in Object v3.1 (with the legal hold applied) to a bucket created in Nutanix Objects v3.0, the legal hold attributes will not get replicated.

A legal hold can be applied to the objects using the AWS S3 APIs. For more information, see [Nutanix Objects Supported APIs](#) on page 240.

WORM Bucket in Nutanix Objects

Write-once-read-many (WORM) buckets protect your data and metadata. You can configure a WORM bucket to allow the creation of new objects and prevent overwrites or deletion of the existing content for a particular retention period. By default, versioning is not enabled on a bucket. When you apply the WORM policy on a bucket, you can choose to enable versioning. In some industries, regulations or compliance rules mandate long-term records retention, sometimes for more than 7 years. For example, in the financial and health services industry, you must maintain the records in their original state, which cannot be overwritten or erased.

A non-admin user can perform WORM bucket operations only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions and the non-admin user must be assigned to specific objects instances:

- Edit Buckets
- View Buckets
- View Object Store

Note:

- Administrators and non-admin users are Prism Central users with the following roles:
 - Administrator—A Super Admin or a Prism Admin in Prism Central.
 - Non-admin user—A Prism Central user without any administrator privileges.

For more information on the built-in roles in Prism Central, see [Built-in Role Management](#) in the *Security Guide*.

- Users with Edit Bucket permission can create WORM buckets on the buckets available in the Object stores that are assigned to the respective users.

When you increase the retention period of a bucket, the new retention period applies to the existing objects as well as the newly added objects.

When you apply any lifecycle policy to a WORM bucket with or without versions enabled, the policy is not applicable until the WORM retention period has passed.

Note:

- WORM can not be enabled for the buckets with NFS protocol support enabled.
- You cannot enable the WORM policy while creating a bucket.

- You can set 0 days as the minimum and 100 years as the maximum limit for the retention period.

Operations on a WORM Bucket in Nutanix Objects

This topic lists the valid and invalid operations on a WORM bucket. You can refer to this list before performing any operation on a WORM bucket.

Valid Operations

Following are the valid operations:

- Delete objects.

Note: This operation does not delete the data, but creates a delete marker on top of the existing versions of objects.

- Delete the delete marker.
- Delete the target version of the object after the retention period.
- Create a new version of the object and retain old versions.
- Extend the retention period.

Invalid Operations

Following are the invalid operations:

- Enable WORM policy while creating a bucket.
- Enable WORM policy without specifying the retention period.
- Reduce the retention period.
- Delete the targeted version before the retention period.
- Change the version state.
- Does not support legal hold and governance mode.

Applying WORM Policy to Buckets from Prism Central

You can apply WORM policy only after the bucket is created. A WORM bucket allows you to create new objects and prevents you to overwrite or delete the existing content for a particular retention period.

About this task

Warning: You cannot modify or delete objects inside a WORM bucket for the specified time period.

To apply WORM policies to buckets, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. In the Application Switcher, click **Objects**.
3. In the object stores table, click the name of the object store which contains the bucket. The **Object Store** page appears.
4. Click the **Buckets** tab.

5. In the Buckets table, select the bucket to apply the WORM policy.
6. Click **triple dot horizontal icon > Configure WORM**.
The **Configure WORM on *bucketname*** window appears.

Note:

- If versioning is disabled on a bucket, you can enable versioning while enabling WORM. However, once the WORM is enabled, you cannot change the versioning state.
- You cannot suspend versioning on WORM buckets.

Caution:

- For legal compliance reasons, the setting becomes permanent after 24 hours. You can disable the WORM policy within the first 24 hours of the grace period.
- You can only edit the retention period to increase the length of retention. You cannot decrease the retention period.

7. (Optional) Click the **Enable Version** check box to enable versioning.
8. Click the **Enable WORM** checkbox.
9. Type the retention period (in years or months or days) in the **Retention Period** field by entering a number, and then selecting the time period from the drop-down menu.
10. Click **Enable WORM**.
This procedure successfully applies the WORM policy to the bucket.

Configuring a Bucket for Static Website Hosting in Nutanix Objects

You can use Nutanix Objects to host a static website that has individual web pages with static content. To host a static website on Nutanix Objects, you can configure a bucket for website hosting, and then upload your website files (objects) to the bucket. When you configure a bucket as a static website, you enable static website hosting, and optionally, add an index document and an error page. You can upload files (such as index documents and error pages) to the bucket using an S3 browser. The S3 browser uses the S3 protocols by providing access and the secret key. You can also choose to redirect to a website. Once you have configured your bucket as a static website, you can access the bucket through the object store endpoints for your bucket.

Before you begin

A non-admin user can configure static website hosting only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions and the non-admin user must be assigned to specific objects instances:

- Edit Buckets
- View Buckets
- View Object Store

Note:

- Website hosting implicitly grants anonymous public read access of all objects.

- Administrators and non-admin users are Prism Central users with the following roles:
 - Administrator—A Super Admin or a Prism Admin in Prism Central.
 - Non-admin user—A Prism Central user without any administrator privileges.

For more information on the built-in roles in Prism Central, see [Built-in Role Management](#) in the *Security Guide*.

- Users with Edit Bucket permission can configure buckets for static website hosting on the buckets available in the Object stores that are assigned to the respective users.

About this task

Note:

- Static website hosting cannot be configured for the buckets created using the NFS protocol.
- Once you configure the static website for a bucket, you cannot turn off this feature from the Nutanix Objects user interface. To turn off the static website for buckets, contact Nutanix Support.

To configure a bucket for static website hosting, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. In the Application Switcher, click **Objects**.
3. In the object stores table, click the name of the object store which contains the bucket.
The **Object Store** page appears.
4. Click the **Buckets** tab.
5. In the **Buckets** table, select the bucket to configure for static website hosting.
6. Click **tripple dot horizontal icon > Static Website**.
The **Configure Static Website** window appears.
7. By default, the endpoint is auto-populated when you click **Save** at the last step.
For example, when an endpoint auto populates, the URL will be in the format `objectstorename.domain/bucketname`. For example, `testobjectstore.nutanix.com/teamobjects`. However, if they have set up the DNS correctly, then you can also access the website with `bucketname.objectstorename.domain` endpoint using HTTP or HTTPS. For example, `https://teamobjects.testobjectstore.nutanix.com`.
8. Click the **Host Website or Redirect** check box.
 - » **Use this bucket to host a website:** Select this option to use the bucket to host the website. Optionally, you can enter the name of the index document (for example, `myindex.html`) and an error page.

An index document is a web page that Nutanix Objects returns when you request the root of a website. It is the default page that loads when you are not requesting any specific page. After you enable static website hosting for your bucket, you can upload an HTML file with the index document name (for example, `myindex.html`) to your bucket. For example, if you specify no object in the URL, then the website loads the index page (`myindex.html`) that you have configured. If you have not

configured an index document, then website access to the root will return an access denied error message.

A custom error page is a web page that Object returns when an error occurs. For example, if you are trying to load an object that does not exist, the website loads the error page that you have configured.

- » **Redirect:** Select this option to enter a website URL to redirect to that website. For example, when you try to access the bucket endpoint, you will be redirected to this website. The protocol used is either HTTP or HTTPS.

9. Click **Save**.

An endpoint is auto-generated when you click **Save**. This endpoint will be the object store endpoint for your bucket and is used as the website address.

You can now use your bucket as a static website. You can use the endpoint to test your static website.

Cross-Origin Resource Sharing (CORS) in Nutanix Objects Overview

Cross-Origin Resource Sharing (CORS) allows a web application loaded in one domain to access the restricted resources that are requested from another domain. With CORS support in Nutanix Objects, you can create rich web applications and allow cross-origin sharing of resources from Nutanix Objects.

For example, if you upload an image (for example, `image1.png`) in a bucket (for example, *first-bucket*) in *domain1* that contains some security-related information, and you are not allowed to make access to `image1.png` from a website (for example, `www.example.com`) on *domain2*. Then you can configure a CORS policy for that *first-bucket* to allow `www.example.com` to access the resources of the *first-bucket*.

Note:

- CORS cannot be configured for the buckets created using NFS protocol.
- When you configure a bucket for static website hosting, public can have only read access to that bucket. POST, PUT, and DELETE requests on the bucket will be denied.

For configuring CORS for a bucket, create an XML document with the following. The document size is limited to 64 KB.

- Rules that identify the origins that you will allow to access your bucket.
- The operations (HTTP methods) that will support each origin.
- Other operation-specific information.

When Nutanix Objects receives a cross-origin request for a bucket, it checks the CORS configuration on the bucket and uses the first `CORSRule` rule that matches the incoming browser request to enable a cross-origin request. You can add up to 100 rules to the configuration.

Following are the conditions to match the rules.

- The `Origin` header of the request must match `AllowedOrigin` elements.
- The request method (for example, GET, PUT, HEAD, and so on) or the `Access-Control-Request-Method` header for a pre-flight `OPTIONS` request must be one of the `AllowedMethod` elements.
- Every header specified in the `Access-Control-Request-Headers` request header of a pre-flight request must match an `AllowedHeader` element.

Following is an example of CORS configuration with two `CORSRule`:

```
<CORSConfiguration>
  <CORSRule>
```

```

<AllowedOrigin>http://www.example.com</AllowedOrigin>
<AllowedMethod>PUT</AllowedMethod>
<AllowedMethod>DELETE</AllowedMethod>
<AllowedHeader>*</AllowedHeader>
</CORSRule>
<CORSRule>
  <AllowedOrigin>*</AllowedOrigin>
  <AllowedMethod>GET</AllowedMethod>
</CORSRule>
</CORSConfiguration>

```

- The first `CORSRule` allows cross-origin PUT and DELETE requests whose origin is `http://www.example.com` origin. The rule also allows all headers in a pre-flight `OPTIONS` request through the `Access-Control-Request-Headers` header. So, in response to any pre-flight `OPTIONS` request, Nutanix Objects will return any requested headers.

Note: Other than pre-flight `OPTIONS` request, no other requests are denied that fail the CORS policy checks.

- The second rule allows cross-origin GET requests from all the origins.
The `*` wild-card character refers to all the origins.

Configuring CORS on a Bucket in Nutanix Objects

Cross-Origin Resource Sharing (CORS) allows a web application loaded in one domain to access the restricted resources that are requested from another domain.

Before you begin

A non-admin user can configure CORS only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions and the non-admin user must be assigned to specific objects instances:

- Edit Buckets
- View Buckets
- View Object Store

Note:

- Administrators and non-admin users are Prism Central users with the following roles:
 - Administrator—A Super Admin or a Prism Admin in Prism Central.
 - Non-admin user—A Prism Central user without any administrator privileges.

For more information on the built-in roles in Prism Central, see [Built-in Role Management](#) in the *Security Guide*.

- Users with Edit Bucket permission can configure CORS on the buckets available in the object stores that are assigned to the respective users.

About this task

You set this configuration on a bucket so that the bucket can service cross-origin requests.

To configure CORS for a bucket, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. In the Application Switcher, click **Objects**.
3. In the object stores table, click the name of the object store which contains the bucket.
The **Object Store** page appears.
4. Click the **Buckets** tab.
5. In the **Buckets** table, select the bucket to configure CORS.
6. Click **tripple dot horizontal icon > CORS**.
The **Configure CORS** window appears.
7. Type or copy and paste a configuration file, or edit an existing configuration.
The configuration file must be an XML file.
8. Click **Save**.
The CORS configurations are saved for the bucket.

Adding Tags to a Bucket in Nutanix Objects

A tag is a label that you assign to a bucket, and it consists of a key and a value pair that you can define. You can add or remove tags from your buckets.

About this task

To add tags to a bucket, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. In the Application Switcher, click **Objects**.
3. In the object stores table, click the name of the object store which contains the bucket.
4. Click the **Buckets** tab, and from the **Buckets** table, select the bucket to add tags.
5. Click **tripple dot horizontal icon > Tags**.
The **Add Tags** window appears.
6. To add a new tag, click **+ add tag**.
The existing tags will be auto-populated. You can click the delete icon to delete a tag.

Note:

- The maximum number of tags allowed per bucket is 50.
- Tag keys can be up to 128 Unicode characters in length, and tag values can be up to 256 Unicode characters in length.
- Tag keys and values are case-sensitive.
- Tag key should be unique and is a mandatory field, but tag value is optional.

A new row appears to fill the key value of the tag.

7. Enter the key and value.
8. Click **Save**.
The changes are saved.

Viewing Buckets in Nutanix Objects

The buckets view allows you to view the list of buckets in the object store and access detailed information about each bucket.

Before you begin

A non-admin user can view buckets only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions and the non-admin user must be assigned to specific objects instances:

- View buckets
- View object store

Note:

- Administrators and non-admin users are Prism Central users with the following roles:
 - Administrator—A Super Admin or a Prism Admin in Prism Central.
 - Non-admin user—A Prism Central user without any administrator privileges.

For more information on the built-in roles in Prism Central, see [Built-in Role Management](#) in the *Security Guide*.

- Users with View Buckets permission can view all the buckets on an assigned object store.

About this task

To view the list of buckets, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. In the Application Switcher, click **Objects**.
3. In the object stores table, click the name of the object store in which you want to create a bucket.
The **Object Store** page appears.

4. Click the **Buckets** tab.

A list of buckets appears in a tabular view.

The following list describes the fields that appear in the buckets list. A dash (-) is displayed in a field when a value is not available or applicable.

Note:

In Prism Central, the object store only displays the federated buckets that it hosts and not all buckets that are part of the federation. To view all the buckets in a federation namespace, you need to use the Object Browser or any S3 client.

- **Name:** Displays the name of the bucket. Click the name to display the bucket summary.
- **Size:** Displays the size of the bucket.
- **Number of Objects:** Displays the number of objects in a bucket.
- **Versioning:** Displays if versioning is enabled or disabled in a bucket.
- **WORM:** Displays if WORM is enabled or disabled in a bucket.
- **Outbound Replication:** Displays the outbound replication status. The following statuses are displayed:
 - Enabled—If a replication rule is enabled.
 - Disabled—If all the replication rules are disabled.
 - None—If no replication rule is created on the bucket.
- **Multiprotocol Access:** Displays if multiprotocol access is enabled or disabled for the bucket.
- **Notifications** Displays if notifications are enabled or disabled.
- **Static Website & CORS:** Displays if static website or CORS is configured for a bucket.

You can also launch the Nutanix Objects Browser for this object store, click **Launch Nutanix Objects Browser**.

Bucket Summary in Nutanix Objects

This page displays the bucket properties and tiering details.

To view the **Summary** page of a bucket, click the name of the bucket in the buckets table, and then click **Summary**. See [Viewing Buckets in Nutanix Objects](#) on page 125.




Figure 18: Bucket Summary

Updating a Bucket in Nutanix Objects

You can update the bucket settings after creating the buckets and adding the objects to the buckets.

Before you begin

A non-admin user can update buckets only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions and the non-admin user must be assigned to specific objects instances:

- Edit Buckets
- View Buckets
- View Object Store

Note:

- Administrators and non-admin users are Prism Central users with the following roles:
 - Administrator—A Super Admin or a Prism Admin in Prism Central.
 - Non-admin user—A Prism Central user without any administrator privileges.

For more information on the built-in roles in Prism Central, see [Built-in Role Management](#) in the *Security Guide*.

- Users with Edit Bucket permission can perform an update bucket operation on the buckets available in the object stores that are assigned to the respective users.

About this task

Note:

- You cannot disable versioning (if enabled) but you can suspend it.
- You cannot edit multiple buckets at a time.

To edit a bucket, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. In the Application Switcher, click **Objects**.
3. In the object stores table, click the name of the object store which contains the bucket.
A new object store window appears.
4. Click the **Buckets** tab.
5. In the **Buckets** table, select the bucket for which you want to change the settings.

6. Click **Update**.
The **Update Bucket** window appears.

Update Bucket

Lifecycle rules for objects can be managed via Bucket details -> Lifecycle

General Settings

Name: test

NameSpace: Local Namespace

Object Versions

Enable versioning if you wish to safeguard data against accidental deletions or overwrites, or for applications that are versioning aware. If you enable versioning, consider applying lifecycle policies to this bucket to remove older versions of objects after a specified length of time. [Learn More](#)

☒ Enable versioning

Lifecycle Policies

System will also automatically create a lifecycle policy to expire delete markers if they are the only remaining object version.

Multiprotocol Access

Recommended for read-heavy workloads with sequential accesses, e.g. Backup targets, log archives, large media files, etc. Access cannot be enabled or disabled once the bucket is created. [Learn More](#)

☒ Enable NFS v3 access

Cancel Save

Figure 19: Update Bucket Window

7. Edit the settings that you want to change.
For more information on the bucket settings, see [Bucket Configuration](#).
8. Click **Save**.
The changes are saved.

Sharing a Bucket in Nutanix Objects

You can share a bucket with multiple users that have access keys.

Before you begin

A non-admin user can share a bucket only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions, and the non-admin user must be assigned to specific objects instances:

- Edit Buckets
- View Buckets
- View Object Store

Note:

- Administrators and non-admin users are Prism Central users with the following roles:
 - Administrator—A Super Admin or a Prism Admin in Prism Central.
 - Non-admin user—A Prism Central user without any administrator privileges.

For more information on the built-in roles in Prism Central, see [Built-in Role Management](#) in the *Security Guide*.

- Users with Edit Bucket permission can share a bucket available in the object stores that are assigned to the respective users.

- You can share only one bucket at a time.

About this task

To share a bucket, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. In the Application Switcher, click **Objects**.
3. In the object stores table, click the name of the Object Store that contains the bucket. The **Object Store** page appears.
4. Click the **Buckets** tab, and in the **Buckets** table, select the bucket that you want to share with other users.

Name	Size	Configuration	Number of Objects	Accessing	Version	Lifecycle Policies	Replication Policies	Metadata Access	Permissions	Owner	Created
...	...	...	...	...	...	...	...	...	...	...	...
...	...	...	...	...	...	...	...	...	...	...	...
...	...	...	...	...	...	...	...	...	...	...	...
...	...	...	...	...	...	...	...	...	...	...	...
...	...	...	...	...	...	...	...	...	...	...	...

Figure 20: Sharing the Bucket

5. Click **triple dot horizontal icon > Share**. The **Bucket Access Configuration** window appears.
6. The name of the bucket owner is visible under **Owner**.
You cannot edit the name of the bucket owner. Removing or changing the permission for the user who is a bucket owner does not affect the bucket. The owner is still allowed to do all the operations on a bucket.

7. Type the email address of the users and set the required permission for that user. You can also select **Anyone with Link** as a user to share a bucket for public access and set the required permission. When you select **Anyone with Link** as a user, a public bucket access link is provided. You can copy and share the link with any user to directly access the bucket from the web browser. You can only add users who have access keys.

Caution: Creating permissions with **Anyone with Link** as a user, allows anonymous users to use the bucket without requiring authentication.




Figure 21: Share Bucket Window

To generate access keys, see [Generating Access Key for API Users](#) on page 94.

Nutanix supports the following list of permissions on a bucket:

- **Read Only:** Provides read-only access to the user.
- **Full Access:** Provides all access to the user.
- **Custom:** Provides customizing of varied levels of access to the user. Click **Set** under Permissions to set custom permission for a user.

To know more about Bucket Permissions, see [Bucket Access Policies](#).

8. (Optional) To add more users, click **+ Add User**.
- If you add **Anyone with Link** as **Users** and grant them **Full Access**, the user will be able to perform all operations without any authentication.
 - While adding a new user, auto suggestions are provided.
 - You can enter multiple user names by providing space in between each user's name.
9. (Optional) Click the delete icon to remove an existing set of users and permission.
10. Click **Save**.
- It is not possible to save a policy with an empty set of permissions. It is required to have at least one user and permission.
 - Modifying or revoking permissions for a user who is the bucket owner does not affect the bucket. The owner retains unrestricted access and can continue to perform all operations on the bucket.
- The bucket is now shared with the listed users with the allotted permissions.

What to do next

You can list the buckets that are shared with you.

For more information, refer to [Listing the Shared Buckets](#) on page 134.

Bucket Access Policies in Nutanix Objects

This section describes the authorization and access policies implemented in Nutanix Objects.

Nutanix supports the following list of policies on a bucket:

- **Read Only:** Provides read-only access to the user.
- **Full Access:** Provides all access to the user.
- **Custom:** Provides customizing of varied levels of access to the user.

Note: To evaluate the policy for a user, the union of the user specific policy and the anonymous policy is computed. For example, if the user specific policy contains one set of S3 actions and the anonymous policy for the bucket contains another set of S3 actions, then the resulting policy for the user for that bucket will be union of both set.

Table 12: Access Roles

Access Role	Description
User context with respect to buckets	
Owner	The user who creates the bucket. The owner can grant and revoke access to the non-admin users on the bucket. Owner can implicitly perform all operations.
Admin	PC admin who can manage the object-store from the UI.
Shared user	The non-owner user who gains access to the bucket based on the bucket policy assigned. A shared user can apply a bucket policy and grant access to various operations to any other user, provided they have the necessary access to perform PutBucketPolicy API.
Other (non-shared) users	The bucket is not shared with them and they do not have access to the bucket.
Anonymous	Refers to someone who does not authenticate themselves, such as static website users.
User privilege levels (IAM)	
Admin user	Admin user will be implicitly allowed all the bucket resource related operations as mentioned in Permissions for Nutanix Objects Operations on page 132 .
Standard user	Any user who is not an admin. They have full control on the bucket which they create.

Permissions for Nutanix Objects Operations

This section describes the operations that can be performed by the owners and the admin users.

Owner Privileges

The owner can perform any S3 operations on a bucket. They have full control over the buckets they create and can grant access to the buckets to other users (non-admin users).

Any user can make changes as long as they have the necessary access permissions specified in the policy to perform the desired operation.

Bucket Owner has all actions permissions.

Table 13: Operations Performed by an Owner

AbortMultipartUpload
DeleteBucket
DeleteBucketWebsite
DeleteObject
DeleteBucketPolicy
DeleteObjectTagging
DeleteObjectVersion
DeleteObjectVersionTagging
GetAccelerateConfiguration
GetAnalyticsConfiguration
GetBucketAcl
GetBucketCORS
GetBucketLocation
GetBucketLogging
GetBucketNfsConfiguration
GetBucketNotification
GetBucketObjectLockConfiguration
GetBucketPolicy
GetBucketPolicyStatus
GetBucketPublicAccessBlock
GetBucketRequestPayment
GetBucketTagging
GetBucketVersioning
GetBucketWebsite
GetEncryptionConfiguration
GetInventoryConfiguration
GetLifecycleConfiguration

GetMetricsConfiguration
GetObject
GetObjectLegalHold
GetObjectRetention
GetObjectTagging
GetObjectVersionTagging
ListBucket
ListBucketMultipartUploads
ListBucketVersions
ListMultipartUploadParts
PutBucketCORS
PutBucketNfsConfiguration
PutBucketNotification
PutBucketObjectLockConfiguration
PutBucketPolicy
PutBucketTagging
PutBucketVersioning
PutBucketWebsite
PutEncryptionConfiguration
PutInventoryConfiguration

Admin Privileges

An admin user is implicitly granted permission to perform all operations related to bucket resources.

An admin can perform the following actions:

Table 14: Operations Performed by an Admin

DeleteBucket
DeleteBucketWebsite
DeleteBucketPolicy
GetAccelerateConfiguration
GetAnalyticsConfiguration
GetBucketAcl
GetBucketCORS
GetBucketLocation
GetBucketLogging
GetBucketNfsConfiguration
GetBucketNotification

```
GetBucketObjectLockConfiguration
GetBucketPolicy
GetBucketPolicyStatus
GetBucketRequestPayment
GetBucketTagging
GetBucketVersioning
GetBucketWebsite
GetEncryptionConfiguration
GetInventoryConfiguration
GetLifecycleConfiguration
GetMetricsConfiguration
GetReplicationConfiguration
ListBucket
ListBucketMultipartUploads
ListBucketVersions
PutBucketCORS
PutBucketNfsConfiguration
PutBucketNotification
PutBucketObjectLockConfiguration
PutBucketPolicy
PutBucketTagging
PutBucketVersioning
PutBucketWebsite
PutInventoryConfiguration
PutLifecycleConfiguration
PutReplicationConfiguration
```

Listing the Shared Buckets

Listing the Shared Buckets feature extends the S3 ListBuckets API and displays the buckets owned by or shared with the current user.

The current behavior of the S3 API for ListBuckets is to list the buckets that the current owner (caller of the API) owns. If you have created a bucket, then you are the owner of that bucket. For example, when an admin creates a bucket using the Nutanix Objects user interface, the owner of all such buckets is the admin user. You (owner of the buckets) can choose to share buckets with one or more IAM users with *Read Only* or *Full Access* or *Custom* permissions. IAM users with whom a bucket is shared can access the bucket and its objects subject to the permissions granted. However, the Vanilla S3 API for ListBuckets does not allow for the discovery (listing) of buckets that were shared with them.

In earlier versions of Nutanix Objects, if you share a bucket, you have to inform the user which bucket is shared with them. However, with Nutanix Objects 2.1 and later versions, *Listing the Shared Bucket* feature extends the S3 ListBuckets API and displays the buckets owned by or shared with the current user. If you delete a bucket that is shared, the bucket will automatically be removed from the list.

Note:

- This feature is a departure from the S3 specifications and is enabled by default. To disable the feature, contact Nutanix Support.
- This feature is only enabled for sharing relations created using Nutanix Objects 2.1 or later versions. You cannot list the buckets created using earlier versions of Nutanix Objects. A workaround to enable this feature retroactively is to have the owner of the bucket remove all the sharing relations and recreate them using the latest Nutanix Objects version.

For example, an Admin created three buckets B0, B1, and B2. B1 and User2 were shared with User1 with any set of S3 actions. Following are the outputs of the ListBuckets API call before and after the introduction of *Listing the Shared Buckets* feature.

Before the introduction of *Listing the Shared Bucket* feature:

```
Admin B0 B1 B2
User1
User2
```

After the introduction of *Listing the Shared Bucket* feature:

```
Admin B0 B1 B2
User1 B1
User2 B2
```

Viewing Bucket Users in Nutanix Objects

You can view all the users and their permissions for a particular bucket.

Before you begin

A non-admin user can view bucket users only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions and the non-admin user must be assigned to specific objects instances:

- View Buckets
- View Object Store

Note:

- Administrators and non-admin users are Prism Central users with the following roles:
 - Administrator—A Super Admin or a Prism Admin in Prism Central.
 - Non-admin user—A Prism Central user without any administrator privileges.

For more information on the built-in roles in Prism Central, see [Built-in Role Management](#) in the *Security Guide*.

- Users with View Buckets permission can view all the buckets on an assigned object store.

About this task

To view the usage of the bucket, follow these steps:

Procedure

1. Log on to the Prism Central web console.

2. In the Application Switcher, click **Objects**.
3. In the object stores table, click the name of the object store which stores the bucket.
The **Object Store** page appears.
4. Click the **Buckets** tab, and then from the Buckets list table, click the name of the bucket.
5. Click **User Access**.




Figure 22: List of Users

The table displays the list of users and their permissions.

You can edit the user access by clicking **Edit User Access**.

For more information on editing user access, see [Editing User Access in Nutanix Objects](#) on page 136 topic.

Editing User Access in Nutanix Objects

You can modify user access to grant different permissions, including read-only access, full access, or custom access to provide varied levels of access.

Before you begin

The non-admin user must be assigned to specific objects instances:

- Edit Buckets
- View Buckets
- View Object Store

Note:

- Administrators and non-admin users are Prism Central users with the following roles:
 - Administrator—A Super Admin or a Prism Admin in Prism Central.
 - Non-admin user—A Prism Central user without any administrator privileges.

For more information on the built-in roles in Prism Central, see [Built-in Role Management](#) in the *Security Guide*.

- If you access any of the buckets that are still utilizing the old-style policy, they will automatically be upgraded to the new policy. During the upgrade process, the set of access permissions will remain the same as it was before.

About this task

To edit user access, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. In the Application Switcher, click **Objects**.
3. In the object stores table, click the name of the Object Store which contains the bucket.
The **Object Store** page appears.
4. Click the **Buckets** tab, and then from the buckets list table, click the name of the bucket.
5. Click **User Access > Edit User Access** to edit.
The **Bucket Access Configuration** page appears.
6. Type the email address of the user and set the permission for that user.




Figure 23: Share Bucket Window

Nutanix supports the following list of permissions on a bucket:

- **Read Only:** Provides read-only access to the user.
- **Full Access:** Provides all access to the user.
- **Custom:** Provides customizing of varied levels of access to the user.

7. To add more users and permissions, click **+ Add User and Permissions**.

Note:

- Ensure that multiple users that share the same set of permissions, are added as a single entry. This approach reduces the overall policy size and enables the sharing of buckets with a larger number of users without reaching the policy document size limit.
- If you add **All users** as **Users** and grant them **Full Access**, the user will be able to perform all operations without any authentication.

8. Click the delete icon to remove an existing set of user and permission.
9. Click **Save**.
The bucket is now shared with the listed users with the allotted permissions.

Note: It is not possible to save a policy with an empty set of permissions. It is required to have at least one user and permission.

Deleting a Bucket in Nutanix Objects

The owner and admin can delete an empty bucket directly. However, non-empty buckets can be deleted only by admins, provided those buckets do not have any WORM or replication rules enabled, and are

not part of any Federation. An admin can also delete a bucket with versioning enabled; however, it permanently deletes all the versions of the bucket.

Before you begin

A non-admin user can delete a bucket only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions and the non-admin user must be assigned to specific objects instances:

- Edit Buckets
- View Buckets
- View Object Store

Administrators and non-admin users are Prism Central users with the following roles:

- Administrator—A Super Admin or a Prism Admin in Prism Central.
- Non-admin user—A Prism Central user without any administrator privileges.

For more information on the built-in roles in Prism Central, see [Built-in Role Management](#) in the *Security Guide*.

Note:

- Users with **Edit Bucket** permission can delete only the buckets they created in the object stores. They cannot delete buckets created by others.
- If the bucket deletion fails, the bucket will be marked as a read-only bucket until it is deleted.
- The space reclamation from AOS is not instant when you delete the non-empty buckets.

About this task

To delete a bucket, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. In the Application Switcher, click **Objects**.
3. In the object stores table, click the name of the object store which contains the bucket.
The **Object Store** page appears.
4. Click the **Buckets** tab, and then from the Buckets list table, click the check box next to the name of the bucket that you want to delete.
5. Click **Actions > Delete**.
While performing this action for a bucket on AWS S3, the cloud bucket is deleted from the federated global namespace. It continues to exist in AWS S3.
6. In the popup, click **Confirm**.
A message appears to confirm the deletion of the bucket.

What to do next

(optional) After removing all the S3 bucket, you can remove the external endpoint from the Prism Central. For more information refer [Removing External Object Store](#) on page 61

NUTANIX OBJECTS STREAMING REPLICATION

Nutanix Objects streaming replication enables automatic and asynchronous copying of source buckets to the target buckets in different Nutanix Objects instances.

You can now create up to three replication rules. This provides you with the ability to replicate a single source bucket to a maximum of three destination buckets. These destination buckets can be on Nutanix Objects on the same PC cluster or on a remote PC cluster, or a Cloud Storage Endpoint.

For more information, see [Types of Bucket Replication](#) on page 139.

In Nutanix Objects, for version-enabled buckets, you can also replicate the delete markers to a destination bucket. You can only replicate a non-versioned bucket to an S3 destination bucket.

Note: Nutanix Objects streaming replication is only supported for buckets created using the S3 protocol.

Nutanix Objects streaming replication helps you with the following:

- Maintain another copy of data for disaster recovery.
- Maintain another copy of data at different sites for local access. This helps to minimize latency in accessing data.

Types of Bucket Replication

The Nutanix Objects Streaming Replication feature supports two types of bucket replication: Replication to Nutanix Objects and Replication to S3.

Replication to Nutanix Objects Destination

You can replicate the data to a destination bucket in the same Prism Central instance or a remote Prism Central instance in Nutanix Objects. Replication is supported for both federated and non-federated buckets.

You can replicate the bucket to Nutanix Objects destination in the following ways:

- **Nutanix Objects instances within the same Prism Central instance**

Replicating a bucket to Nutanix Objects instance within the same Prism Central instance involves the single step of creating a replication rule.

For more information, see [creating a replication rule](#).

- **Nutanix Objects instances within different Prism Central instances** Replicates buckets to Nutanix Objects instance within a different Prism Central instance. This involves three high-level steps.
 - Add the remote Prism Central clusters as an availability zone. For more information, see [Adding the remote PC clusters as an availability zone](#).
 - Perform IAM synchronization between the source Prism Central cluster and remote Prism Central clusters. For more information, see [Setting up IAM Synchronization with a Different Prism Central](#) on page 146.
 - Create a replication rule. For more information, see [Creating a replication rule](#).

Replication to an S3 Destination

You can replicate the data from a non-versioned bucket to an S3 destination bucket. The destination bucket can be versioned or non-versioned.

For more information, see [Creating a replication rule](#).

Object user metadata is replicated, but the destination object has its own system metadata. For example, the creation time of the object at the destination might differ from the creation time of the source object.

After you set up the replication rule between the source and destination bucket, you can track the replication statistics for each bucket.

For more information, see [Viewing Replication Statistics for a Bucket](#).

Points To Remember for Replication to Nutanix Objects

Note the following points about Nutanix Objects streaming replication when buckets are replicated to Nutanix Objects:

- If the Nutanix Objects instances belong to different Prism Central clusters, a trust must be set up between the Prism Central clusters by adding each other as Availability Zones.
Availability Zone - A physically isolated site where you can replicate the data that you want to protect. A Prism Central instance represents an availability zone.
- For replication-enabled buckets, versioning, and WORM modifications are prevented.
- Nutanix Objects streaming replication between buckets relies on the connection (for example, a secure VPN connection) for encryption.
- Network Address Translation (NAT) performed by any device in between two Nutanix Objects instances is not currently supported.
- Proxy between Nutanix Objects instances is not supported.

Nutanix Objects Streaming Replication Guarantees and Topologies

This section covers the replication guarantees, how replication works, and supported replication topologies.

Replication Guarantees

The following object attributes are replicated:

- Object operations—Object PUT, Object Copy, Updates (PutTags, PutObjectLock), and Delete
- Object metadata—ETag, create and modification time, and lock property
- Version numbers, User metadata, and Tags

Note:

- Creation and modification timestamps are not replicated to non-Nutanix targets such as AWS S3 buckets or third-party S3-compatible endpoints.
- Source and destination buckets are independently managed and not replicated. Also, any changes to the bucket policies (for example, access or lifecycle policies) are not replicated.

How Replication Works

Note the following points about how the Nutanix Objects streaming replication works:

- Replication starts as soon as the object gets written on the source bucket.

- The replication completion time might vary depending on the object size and other factors such as available bandwidth, number of replications, and so on.
- Nutanix Objects replications might not strictly follow the same time order in which they are written on the source bucket.
- Replication of the versions of an object might not happen in sequential order, but they are replicated eventually.

Replication Topologies

See the [Replication Topologies Supported for Nutanix Objects as Destination](#) on page 141 and [Replication Topologies Supported for S3 Destination](#) on page 144 section.

Replication Topologies Supported for Nutanix Objects as Destination

This section lists the topologies that you can use for your Nutanix Objects replication scenarios.

Note:

- Ensure that applications do not perform conflicting write operations (objects with the same name) on the remote bucket while replication is enabled.
- The I/O operations performed on the remote bucket are unrestricted. You can perform read-and-write operations on a remote bucket.
- For an N:1 replication relation, a bucket can have a maximum of five inbound relationships. For example, Bucket A can be the destination bucket for a maximum of five source buckets.
- For a 1:N replication relation, a bucket can have a maximum of three outbound relationships. For example, Bucket A can be the source bucket for a maximum of three destination buckets.

Replication Topologies

The following are the topologies that you can use for your replication scenarios:

Single-Replication Relation

In this topology, you replicate objects one way from the source to the destination.




Figure 24: Single Replication - Unidirectional

Different buckets on the source Nutanix Objects instance can replicate to buckets belonging to different Nutanix Objects instances.




Figure 25: Single Replication - Different Nutanix Objects Instances

Bidirectional Replication Relation

In this topology, you can set up a bidirectional-replication relation between a pair of buckets. Nutanix Objects written on one side get replicated on the other side.

Independent replication relations between the buckets have to be created, that is, create a replication rule with Bucket A as the source and Bucket B as the destination and the other way around.

Note: Ensure that the application does not write conflicting objects (objects with the same name) on both buckets while replication is enabled.




Figure 26: Bidirectional Replication

Chain Replication Relation

In this topology, objects from Bucket A can be replicated to Bucket B, and objects that originated in Bucket B can be replicated to Bucket C




Figure 27: Chain Replication

N:1 Replication Relation

In this topology, a bucket can be a destination to many source buckets.




Figure 28: N:1 Replication

1:N Replication Relation

In this topology, a source bucket can replicate to a maximum of three destination buckets. You can have a maximum of three rules in this topology. By default, the priority of replication on all three destination buckets is equal.




Figure 29: 1:N Replication

Replication Topologies Supported for S3 Destination

This section lists the topologies that you can use for your S3 replication scenarios.

Replication Topologies

The following are the topologies that you can use for your S3 replication scenarios:

Note: These replication relations are unidirectional.

1:1 Replication Relation

In this topology, you replicate objects one way from the source to the destination.




Figure 30: 1:1 Replication - Unidirectional

N:1 Replication Relation

In this topology, a bucket can be a destination for many source buckets.




Figure 31: N:1 Replication - Unidirectional

Prerequisites for Replication to Nutanix Objects

The following are the prerequisites for bucket replication to Nutanix Objects:

- Nutanix Objects instances containing the source and destination buckets must be deployed with Nutanix Objects version 3.0 or later.
- For Nutanix Objects instances within different Prism Central instances replication, ensure to deploy Nutanix Objects instances with IAM users if you want the same users to have access to the replicated objects.
- For Nutanix Objects instances within different Prism Central instances replication, you must establish a secure VPN connection between the Prism Central clusters.

For Nutanix Objects instances within the same Prism Central cluster, you must establish a secure VPN connection between the Prism Element clusters.

- For IAM pairing, both the source and target must have Nutanix Objects Manager version 5.0 or later.
- Ensure that the firewall allows traffic to port 7100 on the load balancers for Nutanix Objects instances for replication.
- The pod network of the source Nutanix Objects cluster must not overlap with the public network of the target Nutanix Objects cluster. Also, the public network of the target Nutanix Objects cluster must not use the 10.100.x.x network range.

Adding Remote Prism Central as Availability Zone

For *Nutanix Objects instances within different PC instances* replication, you need to add the remote PC clusters as an Availability Zone (AZ) on the source PC cluster to establish a secure connection.

About this task

To add the remote PC as an availability zone, follow these steps:

Procedure

1. Log on to the Prism Central web console.

2. In the Application Switcher, select **Infrastructure** > **Administration** > **Availability Zones**.




Figure 32: Administration - Availability Zones Page

3. Click **Connect to Availability Zone**.
 4. In the **Availability Zone Type** list, click **Physical Location**.
 5. Enter the IP address and login credentials of the remote PC in the corresponding boxes.
 6. Click **Connect**.
- The remote PC gets added as an AZ in the source PC and the connectivity status is shown as **Reachable**.

Setting up IAM Synchronization with a Different Prism Central

To replicate Nutanix Objects, perform IAM synchronization between the source and remote Prism Central instances.

About this task

The IAM users of the source Nutanix Objects instance is replicated to the destination Nutanix Objects instance belonging to a different Prism Central with the same access key and secret key pair. The admin or bucket owner needs to provide permission to the users for the replicated buckets.

Note: For IAM pairing, if either the source or target uses Nutanix Objects Manager version 5.0 or later, the other must also be on version 5.0 or later.

A non-admin user can set up IAM synchronization with a different Prism Central only after an administrator assigns a role with the following permissions:

- Create Object Store
- View Object Store

For more information on the built-in roles in Prism Central, see [Built-in Role Management](#) in the *Security Guide*.

To perform the IAM pairing, follow these steps:

Procedure

1. Log on to the source Prism Central web console.

2. In the Application Switcher, click **Objects**.
3. Click the **Access Keys** tab.
4. Click **IAM Replication Settings**, and then click **Add IAM Pairing**.
5. In the **Target Prism Central** list, select the remote Prism Central that you added as an Availability Zone (AZ) in the source Prism Central.
 - Only the Prism Central instances added in **Administration > Availability Zones** appear in this list.
 - The target Prism Central IAM must have all the Active Directory configured in this source Prism Central IAM.

Figure 33: Add IAM Pairing

6. Click **Connect** to complete the pairing.

After you click **Connect**, all existing users get replicated to the target Prism Central IAM. You can monitor the progress of the replication on the **IAM Replication Settings** page. Once the IAM pairing is complete, all the future users and keys that are created and deleted get replicated to the target Prism Central IAM.

If any of the replication fails, go to the **IAM Replication Settings** page and click **Sync** to replicate any unreplicated users to the target Prism Central. The administrators can view the users for whom the replication failed and use the export option to download the list of errors.

Note:

- You cannot add more than five IAM replication targets.
- The **Sync** option does not replicate user or key deletions. For failed delete replications, log in to the target Prism Central and manually delete the users and keys.

Prerequisites for Replication to S3

This section lists the prerequisites for bucket replication to S3.

The prerequisites for replicating buckets to S3 are as follows:

- The source bucket must be non-versioned.

- While configuring the replication rule, the source and destination bucket must have the same WORM policies (such as retention period). Changing the WORM policy parameter on the destination bucket after setting the replication rule might result in unexpected behavior. For example, if the object is deleted on the source, the destination bucket might reject the delete request.
- You must set up a default lifecycle policy on the destination S3 bucket.
On AWS, the lifecycle policy can be set by the S3 bucket. Click **Management > Add Lifecycle Rule**. Select **Clean up incomplete multipart uploads** checkbox. Enter the number of days as expiry duration. Here, the number of days depends on the object size, connection bandwidth, workload, and system resources. For example, 30 days is a good starting point.
- Make sure that the destination bucket has the recommended settings. For more information, see [Cloud Bucket Endpoint Recommended Settings](#) on page 113.

Prerequisites for Creating a Replication Relation

Before you create a replication relation, ensure that you meet the following requirements:

- For remote replication of Prism Central clusters, ensure that the fully qualified domain names (FQDN) of the object store instances in the two different Prism Central clusters are different.
- If you intend to replicate data to a bucket on another Prism Central cluster, make sure to synchronize IAM between the source and the remote Prism Central cluster.
- Nutanix Objects does not automatically replicate the bucket access permissions to the destination bucket. You must manually assign them to the destination bucket.
- Lifecycle policies are not replicated and can be assigned independently to both source and destination buckets.
- When you replicate a bucket created in Nutanix Objects version 3.1 with a legal hold to a bucket created in Nutanix Objects version 3.0, the system does not replicate the legal hold attributes.
- The source and destination buckets must have the same versioning, WORM state, and WORM retention period.

Note: For S3 replication, the source bucket must be non-versioned, while the destination bucket can be either versioned or non-versioned.

- When you create a replication rule, only the objects that you add to the buckets or modify afterward is transferred to the destination bucket. For more information on replicating existing objects in the bucket, see [Creating Replication Relation for Buckets](#).
- Versioning and WORM modifications are not allowed for replication-enabled buckets. To make changes, you must first delete the replication rule, make the necessary edits on the buckets, then recreate the replication rule.
- When creating a replication rule for buckets located in different object storage instances, first establish a relationship with the object store that contains the destination bucket using the Prism Central user interface. For example, to replicate Bucket 2 in object store A to Bucket 6 in object store B, use the user interface for the first step. After that, you can create additional replication rules for any buckets in object store A to any other bucket in object store B using either the user interface or S3 API.
- A non-admin user can create replication rules between buckets only after the administrator creates an access control policy for the non-admin user on a role in Prism Central. The specific role must have the following minimum permissions, and the non-admin user must be assigned to specific Nutanix Objects instances:
 - Edit Buckets

- View Buckets
- View Object Store

Note: Administrators and non-admin users are Prism Central users with the following roles:

- Administrator: A Super Admin or a Prism Admin in Prism Central.
- Non-admin user: A Prism Central user without any administrator privileges.

For more information on the built-in roles in Prism Central, see [Built-in Role Management](#) in the *Security Guide*.

- Users with **Edit Buckets** permission can configure bucket replication on all the buckets on an assigned object store.
- If the target object store is on the same Prism Central cluster, the user must have **View Object Store** access to that target object store as well. However, if the target object store is on a different Prism Central cluster, the user is allowed to select any object store available on the target Prism Central cluster.

Creating Replication Relation for Buckets

You can create a replication relation for buckets that allows you to replicate data in a bucket to other buckets. You can replicate the data to a destination bucket in the same Prism Central or a remote Prism Central, or you can replicate the data in any S3 destination bucket.

Before you begin

- Ensure that you meet the prerequisites for creating a replication relation and replication to Nutanix Objects and S3. For more information, see [Prerequisites for Creating a Replication Relation](#) on page 148, [Prerequisites for Replication to Nutanix Objects](#) on page 145, and [Prerequisites for Replication to S3](#) on page 147.
- If you want to replicate the data to an S3 destination bucket, make sure you configure a Cloud Bucket Endpoint for the destination bucket. See [Configuring a Cloud Bucket Endpoint in Nutanix Objects](#) on page 111.

About this task

To create a replication relation for buckets, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. In the Application Switcher, click **Objects**.
3. In the object stores table, click the name of the object store which stores the bucket. The **Object Store** page appears.
4. Click the **Buckets** tab, and then from the buckets list table, select the source bucket.
5. Click **Actions > Replication Rules** or click the bucket name and then **Replication**.

6. Click **Outgoing Replication Relation**.
The **Create Replication Relation** page appears.

Create Replication Relation for

Target Rules

Important Note

IAM Replication Settings
If you plan to replicate data to a bucket on a different Prism Central, please ensure you go to IAM Replication Settings and type IAM users to the target Prism Central.

Bucket access permissions
This is not automatically replicated to destination buckets. You will need to manually assign them to the destination bucket.

Lifecycle policies
Lifecycle policies are not replicated and can be assigned independently on the source and the destination buckets.

Versioning and WORM
Versioning and WORM state must be the same between the source and destination buckets.

Other
Only objects added to the buckets after the replication rule has been created will be transferred to the destination bucket.

Destination

Endpoint Type
Nutanix Objects

Namespace
Search

Bucket
Please enter the bucket name

Cancel Next

Figure 34: Create Replication Relation Page

7. In the **Target tab**, select the **Endpoint Type** as Nutanix Objects or Cloud Storage Endpoint configured for the bucket.

For more information, see [Viewing Cloud Bucket Endpoints](#) on page 117 .Nutanix Objects supports only the S3 Cloud Bucket Endpoint for replication.

8. (Only for Nutanix Objects endpoint) Select **Nutanix Objects** from the **Endpoint Type** dropdown, enter the following details and click **Next**:

- a. **Namespace**: Select a namespace to replicate the object. You can select any local or federated namespace.

You can replicate the bucket on the same Prism Central cluster or on a remote Prism Central cluster. You can select **Local AZ** to replicate the source bucket to a destination bucket on the same Prism Central cluster. You can select a remote Prism Central from the **Federated** list to replicate your bucket to a destination bucket on the remote Prism Central cluster.

- b. **Bucket**: Enter the name of the destination bucket.

9. (Only for S3-compatible endpoint) Select **AWS** from the **Endpoint Type** dropdown, select the **Bucket Name** from the dropdown, and click **Next**.

Nutanix Objects supports only buckets configured for the AWS S3 endpoint for replication.

10. In the **Rules** tab, under the **Sync** section, select one of the following options:

- » **Yes:** Select this option if you want to replicate the existing objects upon rule creation.
- » **No:** Select this option if you do not want to replicate the existing objects upon rule creation.

Replication is skipped for objects that do not match the current versioning state of the bucket and generates an alert.

- **When versioning is enabled:** Nutanix Objects are replicated along with their available version history. However, replication is skipped for objects that do not have a version history, such as NULL versioned objects.
- **When versioning is disabled:** Only NULL versioned objects are replicated, and replication excludes any associated version history. In this case, objects that have versions and their respective version histories are not replicated.

For more information on replication alerts, see [Nutanix Objects Specific Alerts](#) on page 172.

If the user specifies filters, only the objects that comply with those filters will be replicated.

If the bucket already contains a mix of versioned and non-versioned objects, see [Combination of Versioned and Non-Versioned Nutanix Objects Replication](#) on page 157.

11. Under the **Rules** section, click **+ Add** to add a new rule for bucket replication.

A maximum of 1000 rules can be created across all destinations. For more information on the sample replication configurations, see [Sample Replication Configurations](#) on page 154.

a. **Name/ID:** Enter a valid name for the filter.

Note:

- Each rule must have a unique alphanumeric ID. For example, rule_10. The ID must be unique across all the rules set up among all the destinations.
- Name/ID can be a maximum of 255 characters.

b. **Priority:** Each rule must have a unique numeric priority. The priority must be unique across all the rules set up among all the destinations. The higher the number, the higher the priority of that rule.

c. (Optional) Select the **Replicate delete operations** check box to replicate delete markers to the destination bucket.

- If **Replicate delete operations** is enabled, **Tag** is disabled. Tag based filtering cannot be used with Replicate delete operations.
- Selecting the **Replicate delete operations** check box enables the `DeleteMarkerReplication` parameter and replicates delete markers to the destination bucket.
- The `DeleteMarkerReplication` status now appears in the Nutanix Objects user interface. You can also find the `DeleteMarkerReplication` status using S3 API. The relations

created using Nutanix Objects versions earlier than Nutanix Objects 3.5 return the correct `DeleteMarkerReplication` enabled status.

- Set `DeleteMarkerReplication` as `Disabled` when you create a replication relationship. This setting ensures that the replicated objects are not deleted from the destination bucket when the corresponding object is deleted from the source bucket.

However, if objects that are pending replication are deleted from the source, these objects are not replicated to the destination bucket. Hence, the objects are not be available on the destination bucket.

- You can create a maximum of three rules for the 1:N replication.
- For more information on the delete marker replication scenarios, see [Delete Marker Replication Scenarios](#) on page 153.
- When you have an overlapping set of rules where both rules have different values for `DeleteMarkerReplicationStatus`, only the delete operation is replicated if the highest priority matching rule has the `DeleteMarkerReplicationStatus` enabled. However, tag-based rules are not considered even if they have the highest priority.

For example, consider a scenario where the highest priority rule has tags and `DeleteMarkerReplicationStatus` disabled, but a lower priority matching rule (with no tags) has the delete marker replication enabled. In this case, the delete operation is replicated. If there is no rule matching, the delete operation is not replicated.

- d. (Optional) **Prefix**: You can add any prefix. You can replicate only the object whose name matches the prefix.

Prefix length must be less than 255 characters.

- e. (Optional) **Tag**: You can replicate only if the object tags match with the rule tags.

- A maximum of 255 tags per replication rule is allowed.
- The object must match with both prefix and tag if specified. If you do not provide a prefix or tag, all the objects in the bucket are replicated.
- If **Tag** is enabled, **Replicate delete operations** is disabled. Replicate delete operations cannot be used with Tag based filtering.
- Prefix or tag replication does not work retroactively. Replication takes place when an object is uploaded. If you edit the replication rule, all the objects uploaded after edit will be matched against the new rule.
- Replication Status (Paused/Enabled) is at an endpoint level. All the replication rules with the same destination must have the same replication status.
- When you upload an object, the object is matched against the filters set for the replication relations. If at least one of the filters matches, a `replication_info_map` entry is created for that object and replicates the object immediately. For the background path, it is possible that the replication filter no longer matches the object, therefore, before replicating the object in the background path, the object against the filters is matched one more time to ensure the decision of replication.
- A tag consists of a key-value pair where the value field is optional. Also, the keys provided must be unique within the replication rule. For example, user cannot specify `<key1, val1>` and `<key1, val2>`.

- f. Click **Add** to add the rule.
The list of **Rules** appears in new page.

12. (Optional) Select an existing replication relation, and click **Action > Add** to add multiple rules.

13. (Optional) Select an existing replication relation, and click **Action** > **Remove** to delete an existing replication relationship.

14. (Optional) Select an existing replication rule, and click **Action** > **Sync** to sync existing objects of already created replication rule.

The operation might have a significant impact on bandwidth and performance for several hours.

The status is changed to **Active** on completion of sync and an alert is raised.

For more information on replication alerts, see [Nutanix Objects Specific Alerts](#) on page 172.

15. Click **Save**.

The new replication rule appears in the **Replication Relation** table. The replication relation between the buckets is established. You can see the status in the **Status** column. The replication time might vary and depends on various factors, such as object size, workloads on the objects cluster, and so on. This replication rule is applied to the objects based on the **Sync** option selected in Step 10. The status is changed to **Active** on completion of sync and an alert is raised.

For more information on replication alerts, see [Nutanix Objects Specific Alerts](#) on page 172.

- You cannot pause, delete, or edit the replication rule while synchronization is in progress. You can only abort the synchronization.
- For disconnected Availability Zones: If you remove a Prism Central from the Availability Zones after setting up the IAM synchronization and bucket replication:
 - Successive IAM user additions and deletions are not replicated.
 - Existing bucket replication is not affected. However, you cannot create new replication rules to the object store deployed in the disconnected availability zone.

You can select the destination bucket and click **View** on the **Replication** page to view the replication rule details.

You can also view the overview or statistics of the replication rule by clicking the **Overview** tab. For more information, see [Viewing Bucket Replication Statistics](#) on page 159.

What to do next

You can also delete a replication rule. For more information, see [Deleting a Replication Rule of a Bucket](#) on page 159

Delete Marker Replication Scenarios

This section describes the delete marker behavior in a few replication scenarios.

Note: Only non-versioned source buckets are replicated to S3 destination buckets.

In the case of unidirectional replication, the following is the behavior of the delete marker:

When the `DeleteMarkerReplication` is enabled:

- For a non-versioned bucket, the delete operation is replicated to the destination bucket.
- For a versioned bucket, the delete marker is replicated to the destination bucket.

When the `DeleteMarkerReplication` is disabled:

- For a non-versioned bucket, the delete operation is not replicated to the destination bucket.
- For a versioned bucket, the delete marker is not replicated to the destination bucket.

Sample Replication Configurations

This section provides some examples of the replication configurations.

Replication Configuration with One Rule

The following basic replication configuration specifies one rule. This rule replicates all the objects in the source bucket to the destination bucket. For example, bucket-example in this case.

```
<ReplicationConfiguration>
<Role/>
<Rule>
  <ID>rule_id_1</ID>
  <Status>Enabled</Status>
  <Priority>1</Priority>
  <DeleteMarkerReplication>
    <Status>Enabled</Status>
  </DeleteMarkerReplication>
  <Destination>
    <Bucket>bucket-example:example_fqdn</Bucket>
  </Destination>
</Rule>
</ReplicationConfiguration>
```

To choose a subset of objects to replicate, you can add a filter. In the following configuration, the filter specifies an object key prefix. This rule applies to objects that have the prefix `Tax/` in their object names. Priority is irrelevant because there is only one rule.

```
<ReplicationConfiguration>
<Role/>
<Rule>
  <ID>rule_id_1</ID>
  <Status>Enabled</Status>
  <Priority>1</Priority>
  <Filter>
    <Prefix>Tax/</Prefix>
  </Filter>
  <DeleteMarkerReplication>
    <Status>Enabled</Status>
  </DeleteMarkerReplication>
  <Destination>
    <Bucket>bucket-example:example_fqdn</Bucket>
  </Destination>
</Rule>
</ReplicationConfiguration>
```

In the following configuration, the filter specifies an object key tag. This rule applies to objects that have the tag `vm1`, `us-west-1` in its list of tags.

Note: The `DeleteMarkerReplicationStatus` must always be **disabled** for tag-based rules.

```
<ReplicationConfiguration>
<Role/>
<Rule>
  <ID>rule_id_1</ID>
  <Status>Enabled</Status>
  <Priority>1</Priority>
  <Filter>
    <Tag>
      <Key>vm1</Key>
      <Value>us-west-1</Value>
    </Tag>
  </Filter>
  <DeleteMarkerReplication>
```

```

        <Status>Disabled</Status>
      </DeleteMarkerReplication>
      <Destination>
        <Bucket>bucket-example:example_fqdn</Bucket>
      </Destination>
    </Rule>
  </ReplicationConfiguration>

```

In the following configuration, the filter specifies one prefix and two tags. The rule applies to the subset of objects that have the specified key prefix and tags. Specifically, it applies to object that have the `Tax/` prefix in their key names and the two specified object tags in its list of object tags. Priority doesn't apply because there is only one rule.

Note: The `DeleteMarkerReplicationStatus` must always be **disabled** for tag-based rules

```

<ReplicationConfiguration>
<Role/>
<Rule>
  <ID>rule_id_1</ID>
  <Status>Enabled</Status>
  <Priority>1</Priority>
  <Filter>
    <And>
      <Prefix>Tax</Prefix>
      <Tag>
        <Key>IRS</Key>
        <Value>2024</Value>
      </Tag>
      <Tag>
        <Key>IRS2</Key>
        <Value>2023</Value>
      </Tag>
    </And>
  </Filter>
  <DeleteMarkerReplication>
    <Status>Disabled</Status>
  </DeleteMarkerReplication>
  <Destination>
    <Bucket>bucket-example:example_fqdn</Bucket>
  </Destination>
</Rule>
</ReplicationConfiguration>

```

Replication Configuration with Two Rules

In the following replication configuration:

- Each rule filters on a different key prefix so that each rule applies to a distinct subset of objects. In this example, we will replicate objects with the key names `Tax/doc1.pdf` and `Finance/project1.pdf`, but it does not replicate objects with the key name `Logs/snapshot1.txt`.
- Rule priority is irrelevant because the rules apply to two distinct sets of objects. The next example shows what happens when rule priority is applied.

```

<ReplicationConfiguration>
<Role/>
<Rule>
  <ID>rule_id_1</ID>
  <Status>Enabled</Status>
  <Priority>1</Priority>
  <Filter>
    <Prefix>Tax</Prefix>

```

```

</Filter>
<DeleteMarkerReplication>
  <Status>Enabled</Status>
</DeleteMarkerReplication>
<Destination>
  <Bucket>bucket-example:example_fqdn</Bucket>
</Destination>
</Rule>
<Rule>
  <ID>rule_id_2</ID>
  <Status>Enabled</Status>
  <Priority>2</Priority>
  <Filter>
    <Prefix>Finance</Prefix>
  </Filter>
  <DeleteMarkerReplication>
    <Status>Enabled</Status>
  </DeleteMarkerReplication>
  <Destination>
    <Bucket>bucket-example:example_fqdn</Bucket>
  </Destination>
</Rule>
</ReplicationConfiguration>

```

Replication Configuration with Two Overlapping Rules

In this configuration, the two rules specify filters with overlapping key prefixes with same destination, `Doc` and `Document`. Both rules apply to objects with the key name `Document1`. In this case, Nutanix Objects uses the rule priority to determine which rule to apply. The higher the number, the higher the priority.

```

<ReplicationConfiguration>
<Role/>
<Rule>
  <ID>rule_id_1</ID>
  <Status>Enabled</Status>
  <Priority>1</Priority>
  <Filter>
    <Prefix>Doc</Prefix>
  </Filter>
  <DeleteMarkerReplication>
    <Status>Enabled</Status>
  </DeleteMarkerReplication>
  <Destination>
    <Bucket>bucket-example:example_fqdn</Bucket>
  </Destination>
</Rule>
<Rule>
  <ID>rule_id_2</ID>
  <Status>Enabled</Status>
  <Priority>2</Priority>
  <Filter>
    <Prefix>Document</Prefix>
  </Filter>
  <DeleteMarkerReplication>
    <Status>Disabled</Status>
  </DeleteMarkerReplication>
  <Destination>
    <Bucket>bucket-example:example_fqdn</Bucket>
  </Destination>
</Rule>
</ReplicationConfiguration>

```

If you delete the object `Document1`, the delete will not be replicated to the destination as the rule with priority 2 has the **DeleteMarkerReplicationStatus** disabled.

In this configuration, the two rules specify filters with overlapping key prefixes with different destinations, `Doc` and `Document`. Since both the rules have different destinations, priority is not applicable even if the rules are overlapping. Both rules will be applied to objects with the key name `Document1` and object will be replicated to both the destinations.

```
<ReplicationConfiguration>
<Role/>
<Rule>
  <ID>rule_id_1</ID>
  <Status>Enabled</Status>
  <Priority>1</Priority>
  <Filter>
    <Prefix>Doc</Prefix>
  </Filter>
  <DeleteMarkerReplication>
    <Status>Enabled</Status>
  </DeleteMarkerReplication>
  <Destination>
    <Bucket>bucket-example:example_fqdn</Bucket>
  </Destination>
</Rule>
<Rule>
  <ID>rule_id_2</ID>
  <Status>Enabled</Status>
  <Priority>2</Priority>
  <Filter>
    <Prefix>Document</Prefix>
  </Filter>
  <DeleteMarkerReplication>
    <Status>Disabled</Status>
  </DeleteMarkerReplication>
  <Destination>
    <Bucket>bucket-example2:example_fqdn</Bucket>
  </Destination>
</Rule>
</ReplicationConfiguration>
```

Combination of Versioned and Non-Versioned Nutanix Objects Replication

If the bucket already contains a mix of versioned and non-versioned objects, follow this procedure.

About this task

To replicate a combination of versioned and non-versioned object, follow these steps:

Procedure

1. Replicate all objects that correspond to the current versioning state.
For more information, see [Creating Replication Relation for Buckets](#) on page 149.
2. Delete the replication relationship.
For more information, see [Deleting a Replication Rule of a Bucket](#) on page 159.
3. Change the versioning state of both the source and destination buckets.
4. Re-enable replication and perform another sync.
All objects are successfully replicated, regardless of their versioning state.

Pausing a Replication Rule for a Bucket

You can pause an active replication rule at any point in time. When you pause a replication, the objects are accumulated by the source. When the replication resumes, all the objects uploaded during the pause state are replicated to the destination bucket.

About this task

Note: Nutanix recommends not pausing the replication for more than 24 hours to avoid bloating the number of objects pending for replication. If the replication rule is paused for more than 24 hours and if at least one object was uploaded during the pause state, the alert `PauseReplicationWarningDurationTimeExceeded` is raised to resume the replication. For more information on the alerts, see [Nutanix Objects Specific Alerts](#) on page 172.

To pause a replication rule for buckets, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. In the Application Switcher, click **Objects**.
3. In the object stores table, click the name of the object store that stores the bucket.
The **Object Store** page appears.
4. Click the **Buckets** tab, then from the buckets list table, click the name of the source bucket.
5. Click the **Replication** tab.
6. To pause an active replication, select the replication rule and click **Actions > Pause**.
The replication rule is paused and the status is changed from **Active** to **Paused**. The objects are accumulated by the source.

What to do next

You can also resume or delete a replication rule. For more information, see [Resuming a Replication Rule for a Bucket](#) on page 158 and [Deleting a Replication Rule of a Bucket](#) on page 159

Resuming a Replication Rule for a Bucket

You can resume a paused replication rule at any point in time. When you pause a replication, the objects are accumulated by the source. When the replication is resumed, all the objects uploaded during the pause state are replicated to the destination bucket.

About this task

To resume a replication rule for buckets, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. In the Application Switcher, click **Objects**.
3. In the object stores table, click the name of the object store that stores the bucket.
The **Object Store** page appears.
4. Click the **Buckets** tab, then from the buckets list table, click the name of the source bucket.
5. Click the **Replication** tab.

6. To resume a paused replication, select the replication rule and click **Actions > Resume**.
The replication rule is resumed and the status changes from **Paused** to **Active**. The objects uploaded during the paused state are replicated to the destination bucket.

What to do next

You can also pause or delete a replication rule. For more information, see [Pausing a Replication Rule for a Bucket](#) on page 158 and [Deleting a Replication Rule of a Bucket](#) on page 159

Deleting a Replication Rule of a Bucket

This section provides the steps to delete a replication relation.

About this task

A typical scenario where you may require to delete a replication rule is if you want to modify the versioning and WORM state of a replication-enabled bucket. Since versioning and WORM changes are prevented, you need to delete the replication rule, perform the required edits on the bucket, and then create the replication rule again.

Note: Deleting a replication relation would cause all pending replications for that relationship to be dropped immediately. It is recommended to wait for the pending replications to complete, then start making changes by deleting the relation.

To delete a replication rule for a bucket, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. In the Application Switcher, click **Objects**.
3. Click the name of the object store where the bucket is deployed.
The **Object Store** page appears.
4. Click **Buckets**.
5. In the **Buckets** table, select the replication-enabled bucket.
6. Click **Actions > Replication Rules**.
The **Replication** page appears.
7. In the **Replication Rule** table, select the replication destination bucket.
8. Click **Delete**.
After the replication rule is deleted, new objects are not replicated to the destination bucket.
The replication rule is deleted.

What to do next

You can also pause a replication rule. For more information, see [Pausing a Replication Rule for a Bucket](#) on page 158.

Viewing Bucket Replication Statistics

This section describes the steps to view the replication statistics, such as the replication properties and the average bandwidth for a bucket.

Before you begin

A non-admin user can view replication statistics only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions and the non-admin user must be assigned to specific objects instances:

- View Buckets
- View Object Store

Note:

- Administrators and non-admin users are Prism Central users with the following roles:
 - Administrator—A Super Admin or a Prism Admin in Prism Central.
 - Non-admin user—A Prism Central user without any administrator privileges.

For more information on the built-in roles in Prism Central, see [Built-in Role Management](#) in the *Security Guide*.

- Users with View Buckets permission can view the replication statistics on all the buckets on an assigned object store.

About this task

To view the replication statistics, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. In the Application Switcher, click **Objects**.
3. In the object stores table, click the name of the object store which stores the bucket.
The **Object Store** page appears.
4. Click the **Buckets** tab, and then from the Buckets list table, select the replication-enabled bucket.
5. Click **tripple dot horizontal icon > Replication Rules**.
The **Replication** page appears.

6. Click the **Overview** tab.

The replication properties and the average bandwidth appears on the **Overview** page.




Figure 35: View Replication Overview

If the bucket has both inbound and outbound traffic, you can select either of the traffic types from the **Traffic Type** list and then select the corresponding source or destination bucket.

- **Outbound from this bucket**—Shows the statistics for replication of objects from the selected bucket (source) to the destination bucket.
- **Inbound to this bucket**—The selected bucket can be a destination to many source buckets. Inbound statistics show the data for the replication of objects to the selected bucket. You can select the source bucket and view the inbound statistics from that bucket.

Note: The updates to the destination bucket happen periodically. This might result in a lag in the destination bucket statistics display.

- **Destination Bucket**—The destination bucket for replication.
- **Destination Bucket Object Store**—The name of the object store that contains the destination bucket.
- **Last replication point**—Point-in-time up to which all the objects created on the source bucket have been replicated.
- **Number of objects pending replication**—Object count pending replication.
- **Objects size pending replication**—Total amount of data pending replication.
- **Average bandwidth**—The rate of the amount of data transferred from the source to the destination. For inbound relationships, the average bandwidth is the cumulative value of all the incoming data from the source buckets. The bandwidth graph helps you to visualize the progress of the replication.

What to do next

You can create more replication rules for buckets. For more information, see [Creating Replication Relation for Buckets](#) on page 149.

Achieving Fault Tolerance for IAM on Nutanix Objects

This section describes how you can use IAM replication and bucket replication to withstand data center failures.

Note: This section is applicable only if IAM-HA is not enabled. For more information, see [Nutanix Objects IAM-High Availability Overview](#) on page 90.

IAM Service on Nutanix Objects

At present, all the Object Stores (Nutanix Objects instance) in the same PC share a single IAM. This IAM instance resides in the first object store (primary) that is deployed in the PC. If that primary object store becomes unavailable, the IAM service will not be available for all secondary object stores of that PC.




Figure 36: IAM Service on Nutanix Objects

For example, you have one PC managing all your clusters that is *PC1*. This *PC1* has four object stores (OSS 1, 2, 3 and 4). However, *IAM 1* resides only in the first object store *OSS 1* (primary object store). All other secondary object stores on that PC (OSS 2, 3 and 4) are for back up of *OSS 1* and relies on *IAM 1* for authentication. If *OSS 1* goes down, then all other back ups also goes down. In this case, Nutanix recommends you the following IAM replication configuration.

Note: The Nutanix Objects cluster will work even when the PC fails.

Recommended IAM Replication Configurations

Nutanix recommends to transition to multiple PC, that is multiple IAM solution for achieving fault tolerance for IAM replication. You can have one or more PCs to manage IAM replication. The remote PC can have different object stores, and you can set up IAM replication between primary PC and remote PC.




Figure 37: Nutanix Objects IAM Service with Replication across Prism Centrals

For example, in the preceding image, you have three PCs managing all your clusters that is *PC1*, *PC2*, and *PC3*. *PC1* has four object stores (OSS 1, 2, 3 and 4) with *IAM 1* residing only in the first object store OSS 1 (primary object store). All other secondary object stores on that PC (OSS 2, 3 and 4) relies on *IAM 1* for authentication.

Similarly, *PC2* and *PC3* have different object stores respectively (OSS 5, 6, 7, 8 and OSS 9, 10, 11, 12) with *IAM 2* and *IAM 3* residing only in the first object stores OSS 5 and OSS 9 (primary object stores) of *PC2* and *PC3*.

Now as recommended, you use *PC2* and *PC3* for replication of *PC1*. In this case, even if any object store in *PC1* goes down, then replicated buckets will be available in *PC2* and *PC3* without any disruption.

MONITORING AND ALERTS IN NUTANIX OBJECTS

You can monitor the usage and the performance of object stores and buckets, and you can also view system-generated alerts for Nutanix Objects.

A non-admin user can view the usage and performance of object stores and buckets, and also view system-generated alerts for objects only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have view object store permission and the non-admin user must be assigned to specific objects instances.

Note: Administrators and non-admin users are Prism Central users with the following roles:

- Administrator—A Super Admin or a Prism Admin in Prism Central.
- Non-admin user—A Prism Central user without any administrator privileges.

For more information on the built-in roles in Prism Central, see *Built-in Role Management* in the *Security Guide*.

Note: For the object counts that are shown at the object store instance level, Nutanix Objects counts each upload of a multipart upload as a separate object until the object is finalized. However, for the object count at the bucket level, Nutanix Objects does not include the upload counts in the objects count since an object is considered uploaded in a bucket only after it is finalized. For example, if you have done a multipart upload of 10 objects (suppose 5 parts of each object), then the multipart count of the objects is shown as 50 in the object store instance level. However, at the bucket level, it is shown as 0 because the objects are not finalized or completely uploaded.

Viewing Performance of Object Stores

You can view and analyze the performance of an object store by placing the cursor anywhere on the horizontal axis to display the value at that time. You can also select the time interval (last 1 hour, 3 hours, 6 hours, 12 hours, and 24 hours) from the pull-down list above the graphs.

About this task

To view the performance of an object store, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. In the Application Switcher, click **Objects**.
3. In the object stores table, click the name of the object store to view its performance.
The **Object Store** page appears.

4. Click **Performance**.




Figure 38: Object Store Performance Graph

The graph shows the following information:

a. **Requests Per Second:** Displays the following graphs:

- The **Total** graph displays the total input and output requests in each second in the object store or bucket.
- The **Puts** graph displays the total input in each second
- The **Gets** graph displays the total output in each second.

b. **Throughput (MB per sec):** The Throughput graph displays granular read-and-write throughput in MB in each second.

You can see the total in, total out, gets, puts, NFS reads, and NFS writes in a bucket.

c. **Time to First Byte (GET Operations):** This graph displays the time taken to read the first byte from the object in milliseconds.

Viewing Performance of Buckets

You can view and analyze the performance of a bucket by placing the cursor anywhere on the horizontal axis to display the value at that time. You can also select the time interval (last 1 hour, 3 hours, 6 hours, 12 hours, and 24 hours) from the pull-down list above the graphs.

About this task

Note: The **Throughput** chart shows data for each connection. This data is not the cumulative throughput across all the clients and connections.

To view the performance of a bucket, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. In the Application Switcher, click **Objects**.

3. In the object stores table, click the name of the object store to view its performance.
The **Object Store** page appears.
4. Click the **Bucket** tab, and then click the name of the bucket from the Buckets list table, and then click **Performance**.




Figure 39: Bucket Performance Graph

The graph displays the performance of the bucket.

a. **Requests Per Second:** Displays the following graphs:

- The **Total** graph displays the total input and output requests in each second in the object store or bucket.
- The **Puts** graph displays the total input in each second
- The **Gets** graph displays the total output in each second.
- The **NFS Reads** graph displays the total NFS reads in each second.
- The **NFS Writes** graph displays the total NFS writes in each second.

b. **Throughput (MB per sec):** This graph displays granular read-and-write throughput in MB per second.

You can see the total in, total out, gets, puts, NFS reads, and NFS writes in a bucket.

c. **Time to First Byte (GET Operations):** This graph displays the time taken to read the first byte from the object in milliseconds.

Viewing Object Store Usage

The Usage tab displays the physical and logical storage usage. You can view the space used across all buckets in an object store.

About this task

To view the object store usage, follow these steps:

Procedure

1. Log on to the Prism Central web console.

2. In the Application Switcher, click **Objects**.
3. In the object stores table, click the name of the object store to view its usage.
The **Object Store** page appears.
4. Click **Usage > Usage Summary**.




Figure 40: Object Store Usage

The storage usage section displays the following information:

Physical Usage: Displays the physical capacity used by the object store on the cluster.

- You can switch between **Across Clusters** and **Cluster wise** for the required view in graph.
- Total used and total usable storage for the cluster.
- Space used across all buckets on the existing cluster displayed in a bar.
- Total and free physical capacity of the cluster.
- Logical capacity assigned for the object store.

Logical Usage: Displays the logical capacity used by the object store on the cluster. As the Redundancy Factor is 2 (RF-2), the data is stored in two backups.

- Data stored locally
- Data tiered to another S3 endpoint

Assigning Quota Policy to a User

Assigning a quota policy to a user Nutanix Objects to set soft and hard thresholds on storage limits of all buckets owned by the user and number of buckets created by the user within an object store. You can assign quota policy to multiple users at the same time. If a soft threshold is set and if any user exceeds the limit, they can still create buckets and objects, however an alert is raised in the Alerts page. However, if a hard threshold is set, then the user is restricted from exceeding the limit. For example, if a user has 20 TB storage and has six buckets, then storage for all the six buckets will be accounted. If you are an owner of a bucket and have shared your bucket with multiple users, then your storage quota will be used up even if other users use the shared bucket storage. Quota policies will be automatically removed within 24 hours after the user is deleted. If some multi-part operations are pending, those pending parts will count towards the quota.

Before you begin

A non-admin user can assign quota policy to a user only after the administrator creates an access control policy on a role in Prism Central for the non-admin user. The specific role must have the following minimum permissions:

- Create Quota Policy
- Delete Quota Policy
- View Quota Policy
- View Object Store

Note: Administrators and non-admin users are Prism Central users with the following roles:

- Administrator—A Super Admin or a Prism Admin in Prism Central.
- Non-admin user—A Prism Central user without any administrator privileges.

For more information on the built-in roles in Prism Central, see [Built-in Role Management](#) in the *Security Guide*.

About this task

- You cannot list the buckets only owned by a particular user. So, if you have exceeded your storage quota and try to delete a bucket shared with you to release the storage space, the quota level remains as is. You have to delete your own bucket to release the storage space.

For example, if user A has exceeded the storage quota and tries to list all the buckets, the bucket list will also contain the buckets that were shared with user A by user B. To release the storage space within user A quotas, if user A deletes objects in the buckets shared with user A by user B, the quota level remains as is.

You can run `get-bucket-acl` command on the bucket to list the user owned buckets, and then delete the data from that bucket.

- Quota policies are enforced for buckets or space usage across all namespaces, federated or non-federated alike.
- If federation is enabled, quotas for a particular owner/user are enforced only on the OSS cluster where the policy was created and not on other OSS clusters which are part of the same federation.

To assign a quota policy to a user, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. In the Application Switcher, click **Objects**.
3. In the object stores table, click the name of the object store to see the object store usage and to assign quota policy to the users.
The **Object Store** page appears.
4. Click **Usage > Quota Policy**.
5. Select the user for which you want to create the quota policy.
If you are creating quota policy for the first time, skip this step.

6. Click **+ Create Policy**.

Configure Quota

Select People

demo.user@nutanix.com

Quota Limit

☐ Storage Capacity Limit

Enter capacity limit GiB

Capacity limit can be larger than the current available storage

☐ Bucket Number Limit

Enter number Buckets

Quota Limit Enforcement

Once a limit is reached alert notifications will be generated on Prism Central.

☒ Hard Limit

No writes allowed to user owned buckets and/or new bucket creation.

☐ Soft Limit

No restriction on writes or bucket creation. Only alerts raised.

Cancel Save

Figure 41: Configure Quota

A **Configure Quota** window appears.

7. Select a user or multiple users from the **People** drop-down.

You can create quote policy for multiple users simultaneously, but you cannot assign multiple quotas to the same user.

8. In the **Quota Limit** section, select either one check box or both check boxes, and enter the limit.

- (Optional) Click the **Storage Limit** check box and enter the total storage capacity limit for the user. Capacity Limit cannot be larger than the current available storage.
- (Optional) Click the **Bucket Number Limit** check box and enter the number of buckets the user can create.

9. In the **Quota Limit Enforcement** section, select any of the following, and then click **Save**.

- » **Hard limit:** Once the user reaches the specified quota limits, the request is blocked.
- » **Soft limit:** Once the user has exceeded the specified quota limits, only alerts are raised and the request is not blocked.

You can always update the existing policy.

A list of users with the configured storage- and bucket-quota, and the current usage is displayed.

When the storage usage or bucket usage reaches 90% of the defined limit, you are notified with an exclamation mark.

A background task checks for any violation of quota policies and generates this alert. So, it can take about an hour to generate the alert. On a heavily loaded object store, it can take up to a day to generate an alert.

What to do next

You can view the alerts. For more information, see [Viewing Nutanix Objects Alerts](#) on page 170.

Viewing Nutanix Objects Alerts

You can view any informational, warning, and critical alerts in a tabular format. The table presents each alert with the color-coded severity level, a description of the alert, and the time stamp.

About this task

Note: Nutanix Objects-related alerts are now displayed in the Prism Central Alerts section as well. You see a banner at the top of the page to open the Prism Central alerts page.

To view object store alerts, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. In the Application Switcher, click **Objects**.
3. In the object stores table, click the name of the object store.
The **Object Store** page appears.

4. Click the **Alerts** tab.




Figure 42: Alerts List in the Nutanix Objects Alerts Page

The alerts table displays the following information:

Table 15: Alerts Table

Field	Description
Name	Name of the alert
Description	Description of the alert issue
Severity	Severity level of the alert: • Critical: Requires immediate attention • Warning: Needs attention soon • Info: Condition to be aware of.
Created	Number of days since the alert was created. Hover over the number of days to view the date and time when the alert was created.

Field	Description
Silenced	Suppress an alert.

You can also suppress multiple alerts by using **Silence** at the top-left of the screen.

Note:

- Silencing alerts cannot be done from Prism Central alerts page.
- Silencing of alerts is temporary.

You can use **Filters** at the top-right corner to filter the alerts you want to view. You can filter the alerts by severity levels (**Critical**, **Warning**, and **Info**) and state (**Active** and **Silenced**).

You can also view these alerts from the Prism Central **Alerts** page. The following image shows the list of Nutanix Objects alerts on Prism Central **Alerts** page. The Source Entity column for the Nutanix Objects alerts will be the object store name.

Title	Source Entity	Impact Type	Severity	Status	Created Time	Last Occurred	Cluster
Last Replication Sync Time Fo	final-oss-2	Availability	Critical	-	12/06/23...	12/06/23...	Prism Ce...
Unsynchronized clock	final-oss-2	Availability	Critical	-	12/06/23...	12/07/23...	Prism Ce...
Federation WAL usage is high	final-oss-2	Availability	Critical	-	12/06/23...	12/07/23...	Prism Ce...
Unsynchronized clock	final-oss-1	Availability	Critical	-	12/06/23...	12/06/23...	Prism Ce...
Federation WAL usage is high	final-oss-1	Availability	Critical	-	12/06/23...	12/06/23...	Prism Ce...

Figure 43: Alerts List in the Prism Central Alerts page

Nutanix Objects Specific Alerts

The following tables describe the Nutanix Objects specific alerts. Apart from the Nutanix Objects specific alerts, alerts that monitor the underlying infrastructure of the Nutanix Objects clusters are also triggered. For more information on these alerts, contact Nutanix Support at <https://portal.nutanix.com>.

Table 16: Nutanix Objects Alert - High Time to First Byte

Title	High Time to First Byte
Name	HighTimeToFirstByte
Description	The HighTimeToFirstByte alert appears when the Time To First Byte (TTFB) for all the HTTP GET operations in the past 10 minutes exceeds 1 second.
Alert Message	Get operations issued to the object store in the past 10 minutes have been showing TTFB of <value> msec

Cause	High network latency, improper sizing, and component crashes may generate HighTimeToFirstByte alert.
Impact	The response for Object GET requests becomes slow.
Resolution	Do the following: • Check for the networking issues impacting the nodes in the Prism Element cluster. • Check if the GETs in each second and PUTs in each second shown on the Performance page of the object store exceeds the estimated workload at the time of deployment. If there is an indication of an overload, consider reducing the workload on the Object Store Service (OSS). • Check for crashing services on the object store. • Check the Prism Element cluster on which you deployed the object store for any alerts that may be relevant to HighTimeToFirstByte alert.

Table 17: Nutanix Objects Alert - High Object Store Usage

Title	High Object Store Usage
Name	HighObjectStoreUsage
Description	The HighObjectStoreUsage alert appears when the total object store space usage exceeds the estimated capacity specified at the time of deployment.
Alert Message	Current object store usage <val> TB exceeds the provisioned capacity <val>TB
Cause	The following may cause the HighObjectStoreUsage alert to generate. • Missing lifecycle management policies • Performing too many object operations.
Impact	Object store gets overloaded, thus causing slow performance.

Resolution	Do the following: • Delete nonessential objects from the object store to reduce the number of objects in the object store. • If the object store usage is unexpectedly high based on the workload and available resources, check for misconfigured Lifecycle Policies on the buckets in the object that may be retaining objects longer than intended. • If the storage space on the Prism Element cluster outside of the object store is high, consider deleting nonessential VMs, snapshots, data, or adding extra nodes to free the storage space in the Prism Element cluster.

Table 18: Nutanix Objects Alert - High Error Rate

Title	High Error Rate
Name	HighErrorRate
Description	The HighErrorRate alert appears when the object store returns one or more HTTP 4XX or HTTP 5XX errors in each second for the last 10 minutes.
Alert Message	Operations issued to Nutanix Buckets have been failing with 5XX/4XX errors with observed error rate <val>/s over the past 10 minutes
Cause	Improper credentials and component crashes may generate the HighErrorRate alert.
Impact	The object store operations fail.
Resolution	Do the following: • Check the client applications for the correct access and secret key combination. • Check for any crashing services on the object store. • Check the Prism Element cluster on which you deployed the object store for any alerts that may be relevant to HighErrorRate alert.

Table 19: Nutanix Objects Alert - Endpoint <name> Unreachable

Title	Endpoint <cluster_name> Unreachable
Name	EndpointUnreachable
Description	The EndpointUnreachable alert appears when connectivity to the remote endpoint is down.

Alert Message	Connectivity to the remote endpoint <code><cluster_name></code> is down. Identify and fix the connectivity issue.
Cause	The network connectivity to the remote endpoint cluster is lost.
Impact	This issue can result in object GET or PUT failures.
Resolution	Check the network connectivity to the remote endpoint cluster and fix the connectivity issue.

Table 20: Nutanix Objects Replication Alert - Replication Pause Time For a Bucket Exceeded by 24 Hours

Title	Replication Pause Time For a Bucket Exceeded by 24 Hours
Name	PauseReplicationWarningDurationTimeExceeded
Description	The PauseReplicationWarningDurationTimeExceeded alert appears when the replication rule is paused for more than 24 hours and if at least one object was uploaded during the pause state.
Alert Message	Pause replication duration time from source bucket to destination bucket has exceeded warning pause time by duration.
Cause	The replication rule was paused for more than 24 hours and at least one object was uploaded during the pause state.
Impact	If the replication of a bucket is paused, it can cause potential data loss and bloat the number of objects pending for replication.
Resolution	Resume the replication to avoid bloating the number of objects pending for replication.

Table 21: Nutanix Objects Replication Alert - Last replication sync time for a bucket exceeded by 12 hours

Title	Last replication sync time for a bucket exceeded by 12 hours
Name	ReplicationWarningRPOTimeExceeded
Description	The ReplicationWarningRPOTimeExceeded alert appears when the replication of pending objects does not finish even after 12 hours of RPO.
Alert Message	Last sync time for bucket <code><bucket name></code> exceeded RPO time by <code><time_period_secs></code>
Cause	The cause can be a combination of many reasons, such as low network bandwidth, sizing, replication failures, and so on.

Impact	Replication for one bucket has not synced within expected time and can cause potential data loss.
Resolution	Do the following: • Check if any other replication alerts are generated. • Check the connectivity or if the available network bandwidth is sufficient. • Check for any alerts generated due to service crashes.

Table 22: Nutanix Objects Replication Alert - Last replication sync time for a bucket exceeded by 3 days

Title	Last replication sync time for a bucket exceeded by 3 days.
Name	ReplicationCriticalRPOTimeExceeded
Description	The ReplicationCriticalRPOTimeExceeded alert appears when the replication of pending objects does not finish even after 3 days of RPO.
Alert Message	Last sync time from bucket <source bucket name> to destination bucket <destination bucket> exceeded Critical RPO time by <time_period_secs>.
Cause	The cause can be a combination of many reasons, such as low network bandwidth, sizing, replication failures, and so on.
Impact	Replication of one bucket has not synced within expected time and can cause potential data loss.
Resolution	Do the following: • Check if any other replication alerts are generated. • Check the connectivity or if the available network bandwidth is sufficient. • Check for any alerts generated due to service crashes.

Table 23: Nutanix Objects Replication Alert - Replication Endpoint Storage Full

Title	Replication Endpoint Storage Full
Name	ReplicationEndpointStorageFull
Description	The ReplicationEndpointStorageFull alert appears when the object store instance where the remote bucket exists runs out of storage space.

Alert Message	Storage full on replication endpoint <endpoint name>.
Cause	The remote objects instance storage becomes full.
Impact	Replication to the remote site fails since the destination object store is out of space.
Resolution	Add capacity to the remote objects PE cluster or delete existing objects to create space.

Table 24: Nutanix Objects Replication Alert - Replication Endpoint Unreachable

Title	Replication Endpoint Unreachable
Name	ReplicationEndpointUnreachable
Description	The ReplicationEndpointUnreachable alert appears when the object store instance where the remote bucket exists is unreachable for 15 mins.
Alert Message	Replication to endpoint <endpoint name> has lost connectivity for the last 15 minutes
Cause	The following can be the reasons for this alert: • The network connectivity to the remote object store instance is lost. • The remote object store instance is down.
Impact	Replication of bucket to remote site fails since remote site is unreachable.
Resolution	Do the following: • Check the network connectivity to the remote objects instance. • Check if the remote objects instance is up and running.

Table 25: Nutanix Objects Replication Alert - Bucket Sync Completed

Title	Bucket Sync Completed
Name	kBucketSyncCompleted
Description	The system raises this alert to notify the user upon successful completion of the bucket sync operation.
Alert Message	Sync operation has completed for bucket: <i>source_bucket_name</i>
Cause	The synchronization of bucket replication has been completed, and the status has changed from "Running" to "Completed."
Impact	Alerts are displayed on both Prism Central and Nutanix Objects Alerts page.

Resolution	No resolution is required. The alert will automatically resolve itself after 7 days.

Table 26: Nutanix Objects Replication Alert - Few Nonconforming Nutanix Objects Skipped during Sync

Title	Few nonconforming objects skipped during sync
Name	kObjectsSkippedInSync_<bucket_name>
Description	The system raises this alert this alert when the bucket synchronization operation skips certain objects during replication. This usually occurs when the versioning state of the skipped objects does not match the versioning state of the destination bucket
Alert Message	Some objects were skipped due to differing version states.: Match the bucket version state and reissue sync to replicate all objects.
Cause	Some objects were skipped because they have different version states.
Impact	Alerts are displayed on both Prism Central and Nutanix Objects Alerts page.
Resolution	Match the bucket version state and reissue sync to replicate all objects.

Table 27: Nutanix Objects Replication Alert - Replication Sync Operation on Bucket Stalled

Title	Replication Sync operation on bucket stalled
Name	kBucketSyncProgressStalled
Description	The system raises this alert when the bucket synchronization operation is stalled for longer than expected, indicating it is unable to make progress.
Alert Message	Unexpected error occurred while syncing bucket. If Replication Sync does not complete, please contact Nutanix support.
Cause	An unexpected error occurred while syncing the bucket.
Impact	Alerts are displayed on both Prism Central and Nutanix Objects Alerts page.
Resolution	If replication sync does not complete, contact Nutanix Support.

Table 28: Nutanix Objects Replication Alert - Bucket Replication Sync has Pending Replication Entries

Title	Bucket Replication Sync has pending replication entries
Name	kBucketSyncPendingEntriesTimeout
Description	The system raises this alert when the bucket synchronization operation has pending replication entries for a long time.
Alert Message	Replication Sync operation is stuck for a while. Please contact Nutanix support.
Cause	The replication synchronization operation has been stuck for a while.
Impact	Alerts are displayed on both Prism Central and Nutanix Objects Alerts page.
Resolution	If replication sync does not complete, contact Nutanix Support.

Table 29: Nutanix Objects NFS Alert - High NFS Ops Drop Rate

Title	High NFS Ops Drop Rate
Name	HighNfsOpsDropRate
Description	The HighNfsOpsDropRate alert is triggered when some of the NFS ops are not being executed for past 10 minutes.
Alert Message	Due to high payload, few operations submitted to NFS are not executed for the past 10 minutes. Operations are dropped when the outstanding NFS operations have exceeded the threshold value of 1000, or when the QoS queue has reached its maximum capacity of 128 ops, or when the operation wasn't admitted to the queue within 10 seconds.
Cause	The following can be the reasons for this alert: • NFS payload is causing the number of ops to exceed the threshold value of 1000. • Overall high payload is causing QoS queue to reach its maximum capacity.
Impact	NFS ops are being dropped. Latencies may increase due to client retries.

Resolution	Do the following: • Reduce the overall payload on the impacted bucket. • Modify maximum number of concurrent NFS requests. For example, in CentOS, this can be done by changing <code>/proc/sys/sunrpc/tcp_slot_table_entries</code> value to 128.

Table 30: Nutanix Objects Too Many Delete Markers

Title	Too Many Delete Markers
Name	kTooManyDeleteMarkers
Description	The kTooManyDeleteMarkers alert is an early warning triggered when the number of objects and delete markers in a bucket exceed safe thresholds.
Alert Message	The average number of delete markers for each object is greater than 1 million.
Cause	This typically occurs when a bucket is misconfigured with versioning enabled (instead of disabled), causing applications to leave behind delete markers that are never cleaned up.
Impact	Accumulated delete markers can lead to kTimeouts and other performance issues when serving LIST-class S3 requests.
Resolution	Use life-cycle policies on versioned buckets to define retention and expiry policies on delete-markers.

Table 31: Nutanix Objects Too Many Versions

Title	Too Many Versions
Name	kTooManyVersions
Description	The kTooManyVersions alert is an early warning triggered when average number of versions for each unique object name is greater than a set threshold 1000.
Alert Message	The average number of versions for each object is greater than 1000.
Cause	This typically occurs when the average number of versions for each object is high.
Impact	Certain Nutanix Objects version APIs may experience performance degradation.

Resolution

Use life-cycle policies on versioned buckets to define retention and expiry policies on non-current versions of objects.

NUTANIX OBJECTS NOTIFICATIONS

Notifications for Nutanix Objects enables you to send the completed events logs to the configured endpoints in your Nutanix Objects instance. This helps you with centralized events log management, thus enabling you to monitor and analyze the logs and identify performance or configuration issues. TCP is the supported protocol for Nutanix Objects notifications.

The following endpoints are supported for logging the notification events:

- **Syslog**—System Logging Protocol is a standard protocol for sending events logs to the Syslog server. Enter the hostname or host IP address and port number of your Syslog server when configuring the endpoints. Only one endpoint is supported for Syslog. The Syslog server should be up and running when performing the endpoint configuration in your Nutanix Objects instance.
- **Nats-streaming**—A lightweight, reliable streaming platform built on top of the core NATS platform that provides persistent logs. Enter the hostname or host IP address and port number of your NATS streaming server when configuring the endpoints. Only one endpoint is supported for NATS streaming. NATS streaming server should be up and running when performing the endpoint configuration in your Nutanix Objects instance.

Note: The default topic used to create the NATS queue is *OSSEvents*. You can use this as the subject while using the NATS client to connect to the NATS streaming server.

- **Apache Kafka**—Kafka is a distributed system consisting of *servers* and *clients* that communicate through a high-performance TCP network protocol. It can be deployed on bare-metal hardware, virtual machines, and containers in on-premise as well as cloud environments. Enter the hostname or host IP address and port number of your Kafka server when configuring the endpoints. You can add up to 10 Kafka Brokers for an endpoint for Kafka. The Kafka server should be up and running when performing the endpoint configuration in your Nutanix Objects instance.

Notification Types for Nutanix Objects

There are two types of notifications for Nutanix Objects that can be enabled for either successful, or failed, or both successful and failed instance-level and data events.

- **Instance-level events**—All operations performed on a bucket. For example, create a bucket, update a bucket, delete a bucket, enable versioning, and so on. The following instance-level events are supported for successful, or failed, or both successful and failed events:

Note: All the instance-level events notifications are enabled by default once you configure the endpoints. You cannot enable or disable the notifications for a single instance-level event.

Table 32: S3 APIs for Only Successful Instance-level Events

Supported Instance-level Events	S3 APIs for Only Successful Instance-level Events
Create bucket	S3:BucketCreated:Put
Delete bucket	S3:BucketRemoved>Delete
Add bucket policy	S3:BucketPolicy:Put
Enable static website	S3:BucketWebsite:Put
Disable static website	S3:BucketWebsite>Delete

Supported Instance-level Events	S3 APIs for Only Successful Instance-level Events
Create bucket CORS	S3:BucketCORS:Put
Delete bucket CORS	S3:BucketCORS:Delete
Update bucket lifecycle policy	S3:BucketLifecyclePolicy:Update
Delete bucket lifecycle policy	S3:BucketLifecyclePolicy:Delete
Enable bucket lock (WORM)	S3:BucketLock:Enabled
Disable bucket lock (WORM)	S3:BucketLock:Disabled
Add bucket replication rule	S3:BucketReplication:Put
Delete bucket replication rule	S3:BucketReplication:Delete
Enable bucket versioning	S3:BucketVersioning:Enabled
Suspend bucket versioning	S3:BucketVersioning:Disabled
Enable bucket notification	S3:BucketNotify

Table 33: S3 APIs for Only Failed Instance-level Events

Supported Instance-level Events	S3 APIs for Only Failed Instance-level Events
Create bucket	S3:BucketCreatedError:Put
Delete bucket	S3:BucketRemovedError:Delete
Add bucket policy	S3:BucketPolicyError:Put
Enable static website	S3:BucketWebsiteError:Put
Disable static website	S3:BucketWebsiteError:Delete
Create bucket CORS	S3:BucketCORSError:Put
Delete bucket CORS	S3:BucketCORSError:Delete
Update bucket lifecycle policy	S3:BucketLifecyclePolicyError:Update
Delete bucket lifecycle policy	S3:BucketLifecyclePolicyError:Delete
Enable bucket lock (WORM)	S3:BucketLockError:Enabled
Disable bucket lock (WORM)	S3:BucketLockError:Disabled
Add bucket replication rule	S3:BucketReplicationError:Put
Delete bucket replication rule	S3:BucketReplicationError:Delete
Enable bucket versioning	S3:BucketVersioningError:Enabled
Suspend bucket versioning	S3:BucketVersioningError:Disabled

Supported Instance-level Events	S3 APIs for Only Failed Instance-level Events
Enable bucket notification	S3:BucketNotifyError

Table 34: S3 APIs for Both Successful and Failed Instance-level Events

Supported Instance-level Events	S3 APIs for Both Successful and Failed Instance-level Events
Create bucket	S3:BucketCreated:Put and S3:BucketCreatedError:Put
Delete bucket	S3:BucketRemoved:Delete and S3:BucketRemovedError:Delete
Add bucket policy	S3:BucketPolicy:Put and S3:BucketPolicyError:Put
Enable static website	S3:BucketWebsite:Put and S3:BucketWebsiteError:Put
Disable static website	S3:BucketWebsite:Delete and S3:BucketWebsiteError:Delete
Create bucket CORS	S3:BucketCORS:Put and S3:BucketCORSError:Put
Delete bucket CORS	S3:BucketCORS:Delete and S3:BucketCORSError:Delete
Update bucket lifecycle policy	S3:BucketLifecyclePolicy:Update and S3:BucketLifecyclePolicyError:Update
Delete bucket lifecycle policy	S3:BucketLifecyclePolicy:Delete and S3:BucketLifecyclePolicyError:Delete
Enable bucket lock (WORM)	S3:BucketLock:Enabled and S3:BucketLockError:Enabled
Disable bucket lock (WORM)	S3:BucketLock:Disabled and S3:BucketLockError:Disabled
Add bucket replication rule	S3:BucketReplication:Put and S3:BucketReplicationError:Put
Delete bucket replication rule	S3:BucketReplication:Delete and S3:BucketReplicationError:Delete
Enable bucket versioning	S3:BucketVersioning:Enabled and S3:BucketVersioningError:Enabled
Suspend bucket versioning	S3:BucketVersioning:Disabled and S3:BucketVersioningError:Disabled
Enable bucket notification	S3:BucketNotify and S3:BucketNotifyError

For more information, see [Configuring Event Notifications in Nutanix Objects](#) on page 187.

The following is the structure of the notification output for instance-level events that get published to the endpoints.

```
{ "EventType": "s3:BucketCreated:Put", "Key": "first", "Records":
[ { "eventError": "", "eventVersion": "2.0", "eventSource": "aws:s3", "awsRegion": "us-
```

```
east-1", "eventTime": "2022-02-09T07:40:13Z", "eventName": "s3:BucketCreated:Put", "userIdentity":
{"principalId": "upgrade@ntnx.com"}, "requestParameters":
{"sourceIPAddress": "10.44.91.64"}, "responseElements":
{"x-amz-id-2": "110000+1229+2129+240", "x-amz-request-
id": "110000+1229+2129+240", "x-ntnx-origin-endpoint": "0.0.0.0"}, "s3":
{"s3SchemaVersion": "1.0", "configurationId": "SuccessfulWriteManagementEvent", "OSSDomain": "rc4.eng
{"name": "frst", "ownerIdentity":
{"principalId": "upgrade@ntnx.com"}, "arn": "arn:aws:s3:::frst"}}}], "level": "info", "msg": "", "time":
```

- **Data events**—Data events are specific to data operations. These events can be enabled for each bucket. To enable notifications for successful, or failed, or both data events for a bucket, you need to create notification rules. You can define the scope of a notification rule by selecting **All Objects** or **Subset of objects**. For more information, see [Creating Notification Rules for Data Events in Nutanix Objects](#) on page 192.

The following is the structure of the notification output for data events that get published to the endpoints.

```
{"EventType": "s3:ObjectCreated:Put", "Key": "put", "Records":
[{"eventError": "", "eventVersion": "2.0", "eventSource": "aws:s3", "awsRegion": "us-
east-1", "eventTime": "2022-02-09T08:04:48Z", "eventName": "s3:ObjectCreated:Put", "userIdentity":
{"principalId": "upgrade@ntnx.com"}, "requestParameters":
{"sourceIPAddress": "10.44.91.64"}, "responseElements":
{"x-amz-id-2": "110000+1229+2791+297", "x-amz-request-
id": "110000+1229+2791+297", "x-ntnx-origin-endpoint": "0.0.0.0"}, "s3":
{"s3SchemaVersion": "1.0", "configurationId": "5836c513-51ef-45a9-9c84-6a0c01996e40", "OSSDomain": "r
{"name": "put", "ownerIdentity":
{"principalId": "upgrade@ntnx.com"}, "arn": "arn:aws:s3:::put"}, "object":
{"key": "csv.docx", "size": "31", "eTag": "c7a74379a77360a6b85926adb8b2aafb", "versionId": "1", "sequenc
```

The following data events are supported for successful, or failed, or both successful and failed events:

Note: You can enable notifications for all Nutanix Objects Create Events, Nutanix Objects Access Events, or Nutanix Objects Delete Events, or you can choose to enable notifications for one or more sub-categories of Nutanix Objects Create Events (such as Put, Post, Copy, Multipart Upload Complete), Nutanix Objects Access Events (such as Get, Head), or Nutanix Objects Delete Events (such as Delete, Delete Markers).

Table 35: S3 APIs for Only Successful Data Events

Supported Data Events	S3 APIs for Only Successful Data Events
All Create Events	s3:ObjectCreated:*
Put	s3:ObjectCreated:Put
Post	s3:ObjectCreated:Post
Copy	s3:ObjectCreated:Copy
Multipart Upload	s3:ObjectCreated:CompleteMultipartUpload
All Access Events	s3:ObjectAccessed:*
Get	s3:ObjectAccessed:Get
Head	s3:ObjectAccessed:Head
All Delete Events	s3:ObjectRemoved:*
Delete	s3:ObjectRemoved:Delete

Supported Data Events	S3 APIs for Only Successful Data Events
Delete Markers	s3:ObjectRemoved:DeleteMarkerCreated

Table 36: S3 APIs for Only Failed Data Events

Supported Data Events	S3 APIs for Only Failed Data Events
All Create Events	s3:ObjectCreatedError:*
Put	s3:ObjectCreatedError:Put
Post	s3:ObjectCreatedError:Post
Copy	s3:ObjectCreatedError:Copy
Multipart Upload	s3:ObjectCreatedError:CompleteMultipartUpload
All Access Events	s3:ObjectAccessedError:*
Get	s3:ObjectAccessedError:Get
Head	s3:ObjectAccessedError:Head
All Delete Events	s3:ObjectRemovedError:*
Delete	s3:ObjectRemovedError>Delete
Delete Markers	s3:ObjectRemovedError>DeleteMarkerCreated

Table 37: S3 APIs for Both Successful and Failed Data Events

Supported Data Events	S3 APIs for Both Successful and Failed Data Events
All Create Events	s3:ObjectCreated:* and s3:ObjectCreatedError:*
Put	s3:ObjectCreated:Put and s3:ObjectCreatedError:Put
Post	s3:ObjectCreated:Post and s3:ObjectCreatedError:Post
Copy	s3:ObjectCreated:Copy and s3:ObjectCreatedError:Copy
Multipart Upload	s3:ObjectCreated:CompleteMultipartUpload and s3:ObjectCreatedError:CompleteMultipartUpload
All Access Events	s3:ObjectAccessed:* and s3:ObjectAccessedError:*
Get	s3:ObjectAccessed:Get and s3:ObjectAccessedError:Get
Head	s3:ObjectAccessed:Head and s3:ObjectAccessedError:Head

Supported Data Events	S3 APIs for Both Successful and Failed Data Events
All Delete Events	s3:ObjectRemoved:* and s3:ObjectRemovedError:*
Delete	s3:ObjectRemoved>Delete and s3:ObjectRemovedError>Delete
Delete Markers	s3:ObjectRemoved>DeleteMarkerCreated and s3:ObjectRemovedError>DeleteMarkerCreated

Configuring Event Notifications in Nutanix Objects

Specify endpoints when configuring event notifications. By default, only successful events are logged. You can choose success, failure, both, or none.

Before you begin

You must set up and maintain your own Syslog, NATS, or Kafka servers.

About this task

Nutanix Objects supports event notifications through Syslog, NATS, or Kafka. To log operations, enable notifications to one of these services. Nutanix does not manage these services.

Non-admin users can configure event notifications only if an admin assigns them a role with specific permissions in Prism Central. Required permissions include **Set Notification Endpoint**, **View Notification Endpoint**, and **View Object store**. Admins are Super Admins or Prism Admins in Prism Central; non-admins are Prism Central users without admin privileges. For more information, see [Built-in Role Management](#) in the *Nutanix Security Guide*.

To configure events notifications, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. In the Application Switcher, click **Objects**.
3. In the object stores table, click the name of the object store to view its performance. The **Object Store** page appears.
4. Click **Endpoints > Notification Endpoints**.

The **Notification Endpoints** page appears to configure the endpoints by entering the endpoint server details. The available endpoints are **Syslog**, **Nats-streaming**, and **Kafka**. You can enable any one of these endpoints or all of them.

When you add a notification endpoint (Syslog, NATS, or Kafka) with correctly formatted fields, the operation succeeds even if the endpoint is unreachable or invalid. Errors during setup appear on the **Object Store list** page.

5. In the **Syslog** tab, turn on the **Enable** toggle button to configure Syslog as the endpoint.

Important: Syslog server must be up and running when performing the endpoint configuration in your Nutanix Objects instance.

The screenshot shows the 'Notification Endpoints' configuration page. At the top, there are three tabs: 'Syslog', 'Nats-Streaming', and 'Kafka'. The 'Syslog' tab is selected and highlighted with a red box. Below the tabs, there is a toggle switch labeled 'Enable', which is turned on and also highlighted with a red box. Underneath the toggle, there is a text input field labeled 'Host name and port' containing the text '10.48.'. Below this, there is a section titled 'Logged Events'. It contains a link 'View Events List' with a blue arrow pointing to it from a callout box that says 'Click to view the list of bucket events that get logged to the endpoints'. Below the link, there is a dropdown menu labeled 'Object Store Events' with the option 'Successful and Failed' selected. At the bottom of the page, there are two buttons: 'Discard Changes' and 'Save'.

Figure 44: Notification Endpoints: Syslog

- In the **Host name and port** box, enter the host name and port number of your Syslog server in the *Host name:Port* format.
- In the **Logged Events** section, select one of the following object store events for which you plan to enable notifications. By default, **Successful only** option is selected.
 - **Successful and Failed:** Logs both successful and failed object store events.
 - **Successful only:** Logs only successful object store events.
 - **Failed only:** Logs only failed object store events.
 - **None:** Logs none of the object store events.

You can click the **View Events List** link to view the list of bucket events that are logged to the endpoints.

6. In the **Nats-streaming** tab, turn on the **Enable** toggle button to configure Nats-streaming as the endpoint.

Important: Nats-streaming server must be up and running when performing the endpoint configuration in your Nutanix Objects instance.

The screenshot shows the 'Notification Endpoints' configuration page. At the top, there are three tabs: 'Syslog', 'Nats-Streaming' (which is selected and highlighted with a red box), and 'Kafka'. Below the tabs, there is an 'Enable' toggle switch, also highlighted with a red box. Underneath, there are two text input fields: 'Host name and port' containing '10.48' and 'Cluster ID' containing '_cluster'. Below these is the 'Logged Events' section, which includes a link 'View Events List' and a dropdown menu currently set to 'Successful and Failed'. A callout box points to the 'View Events List' link with the text 'Click to view the list of bucket events that get logged to the endpoints'. At the bottom of the form, there are two buttons: 'Discard Changes' and 'Save'.

Figure 45: Notification Endpoints: Nats-Streaming

- In the **Host name and port** box, enter the host name and port number of your NATS streaming server in the *Host name:Port* format.
- In the **Cluster ID** box, enter the cluster ID of the server that you used to start the NATS server.
- In the **Logged Events** section, select one of the following object store events for which you plan to enable notifications. By default, **Successful only** option is selected.
 - Successful and Failed:** Logs both successful and failed object store events.
 - Successful only:** Logs only successful object store events.
 - Failed only:** Logs only failed object store events.
 - None:** Logs none of the object store events.

You can click **View Events List** link to view the list of bucket events that are logged to the endpoints.

7. In the **Kafka** tab, turn on the **Enable** toggle button to configure Kafka as the endpoint.

When adding a Kafka endpoint, if the OSSEvents topic does not exist, the system creates it. If it already exists, ensure it matches your Kafka server configuration.

Important: Kafka server must be up and running when performing the endpoint configuration in your Nutanix Objects instance.

The screenshot shows the 'Notification Endpoints' configuration page. At the top, there are three tabs: 'Syslog', 'Nats-Streaming', and 'Kafka'. The 'Kafka' tab is selected and highlighted with a red box. Below the tabs, there is a section with a toggle switch labeled 'Enable', which is turned on and also highlighted with a red box. Underneath, there is a text input field labeled 'Host name and port' containing three entries: '10.48.1', '10.48.', and '10.48.'. Below this, there is a section titled 'Logged Events' with a link 'View Events List' and a dropdown menu currently set to 'Successful and Failed'. A callout box points to the 'View Events List' link with the text 'Click to view the list of bucket events that get logged to the endpoints'. At the bottom of the form, there are two buttons: 'Discard Changes' and 'Save'.

Figure 46: Notification Endpoints: Kafka

- a. In the **Host name and port** box, enter the host name and port number of your Kafka streaming server in the *Host name:Port* format.

Note: Up to 10 Kafka Brokers are supported for a Kafka endpoint.

For entering multiple brokers, hit space after entering a broker.

- b. In the **Logged Events** section, select one of the following object store events for which you plan to enable notifications. By default, **Successful only** option is selected.
- **Successful and Failed:** Logs both successful and failed object store events.
 - **Successful only:** Logs only successful object store events.
 - **Failed only:** Logs only failed object store events.
 - **None:** Logs none of the object store events.

You can click **View Events List** link to view the list of bucket events that are logged to the endpoints.

8. Click **Save** to complete the events notification configuration for Syslog, Nat-streaming, and Kafka. Notification is enabled for your Nutanix Objects instance. Also, the status of the notification endpoint is shown in the **Notifications** column. For example, if the NATS endpoint gets disconnected due to network issues, a red icon appears in the **Notifications** column. Hovering over the red icon displays Cannot connect to nats endpoint message.

Note: Changes to notification configuration might take up to 5 minutes to take effect.

9. (Optional) To undo an endpoint configuration, follow these steps:
 - a. Click **Endpoints > Notification Endpoints**. The **Notification Endpoints** page appears.
 - b. Clear the checkbox for the endpoint that you plan to remove, and click **Save**.
 - c. In the **Data Event Notification** page, delete all the notification rules corresponding to the endpoint that you removed before adding a new notification rule to the bucket.
For more information, see [Creating Notification Rules for Data Events in Nutanix Objects](#) on page 192

What to do next

You can now create notification rules for data events of the bucket, see [Creating Notification Rules for Data Events in Nutanix Objects](#) on page 192.

Creating Notification Rules for Data Events in Nutanix Objects

To track successfully completed data events, create notification rules for the buckets in Nutanix Objects.

About this task

Non-admin users can create notification rules for data events only if assigned a role with the following permissions by the admin user in Prism Central.

- View Notification Endpoint
- Edit Buckets

Users with this permission can create notification rules for data events across all buckets in the assigned object stores.

- View Buckets
- View Object Store

Admins are Super Admins or Prism Admins in Prism Central; non-admins are Prism Central users without admin privileges. For more information on the built-in roles in Prism Central, see [Built-in Role Management](#) in the *Security Guide*.

To create a notification rule for a bucket, follow these steps.

Procedure

1. Log on to the Prism Central web console.
2. In the Application Switcher, click **Objects**.
3. In the object stores table, click the name of the object store to view its performance.
The **Object Store** page appears.

4. Click the **Buckets** tab, and then from the list of buckets, select the bucket for which you plan to create a notification rule.
5. Click **Data Event Notification**.
The **Data Event Notification** page appears.
6. Click **Create Rule** to open the **Create Rule** page.
Configuring an endpoint is required to create notification rules. Configure endpoints that can be used to log both bucket events and data events. For more information, see [Configuring Event Notifications in Nutanix Objects](#) on page 187.
7. (Optional) If a rule is already created for the bucket, click **+ Add Rule**.

The screenshot shows the 'Create Rule' dialog box. It has a title bar with the text 'Create Rule' and a close button (X). The main content area is divided into several sections:

- Endpoint**: A dropdown menu with 'Nats-streaming' selected.
- Scope**: A dropdown menu with 'Subset of objects' selected.
- Filter with prefix**: A checkbox that is checked, followed by a text input field containing 'Please enter a prefix'.
- Filter with suffix**: A checkbox that is checked, followed by a text input field containing 'Please enter a suffix'.
- Data Events**: A section with two radio buttons: 'All Events' (selected) and 'Successful and Failed'.
- Objects Create Events**: A section with five checkboxes: 'All Create Events', 'Put', 'Post', 'Copy', and 'Multipart upload complete'.
- Objects Access Events**: A section with three checkboxes: 'All Access Events', 'Get', and 'Head'.
- Objects Delete Events**: A section with three checkboxes: 'All Delete Events', 'Delete', and 'Delete Markers'.

At the bottom right of the dialog, there are two buttons: 'Cancel' and 'Save'.

Figure 47: Create Rule

8. In the **Create Rule** page, follow these steps:

a. In the **Endpoint** list, select the endpoint where you plan the data events for this bucket to be logged. The endpoints you enable are listed here.

To update endpoints for your Nutanix Objects instance, see [Configuring Events Notification](#).

b. In the **Scope** list, select **All Objects** or **Subset of objects**.

- If you select **All Objects**, the notification rule applies to all the objects.
- Select **Subset of objects** to apply the notification rule to specific objects. You can filter Nutanix Objects by entering a prefix, suffix, or both. Nutanix Objects with the prefix and suffix you specified are filtered and the rule applies to those objects.

Note:

- You cannot proceed without providing a suffix or prefix for the scope **Subset of objects**.
- You can enter only one prefix and suffix. You can create another rule with a different suffix and prefix.
- The suffix-prefix pair in two rules corresponding to the same event must not overlap. For example, any object name *xyz.jpg* must be selected by a maximum of one rule corresponding to the event type.

c. In the **Data Events** section, click the **All Events** check box to select all data events, or select a check box next to the name of the events listed under **Objects Create Events**, **Objects Access Events**, and **Objects Delete Events**, and choose from the following that you plan to log to the endpoint.

- **Successful and Failed**: Logs both successful and failed object store events.
- **Successful only**: Logs only successful object store events.
- **Failed only**: Logs only failed object store events.

When you select the **All Events** check box, all the **Objects Create Events**, **Objects Access Events**, and **Objects Delete Events** are selected.

9. Click **Save** to complete, and then click **Done** to close the **Create Rule** page.

The notification rule you created appears in the list of data events notification rules on the **Data Event Notification** page. You can click **All Events** or **Number Events** in the Events column to view

the configured data events for that bucket. For example, click **8 Events** to view the 8 data events configured.




Figure 48: Configured Notification Events

Also, you can delete a notification rule. In the **Data Event Notification** page, select a notification rule from the list, and then click **Delete**.

After you create a notification rule for a bucket, the status in the **Notifications** column changes to **Enabled**.

NUTANIX OBJECTS BROWSER

The Nutanix Objects Browser is a User Interface (UI) that helps users to directly launch the object store instance in a web browser and perform bucket and object-level operations. This eliminates the steps of logging into Prism Central and launching the Nutanix Objects service.

The administrators use Prism Central to manage the object stores. When an administrator creates an object store, a default bucket (*objectsbrowser*) gets created. The static website hosting capability is already enabled on this default bucket. The Nutanix Objects Browser UI is hosted on this default bucket.

To launch the Nutanix Objects Browser, select the object store and click **Actions > Launch Objects Browser**. Nutanix Objects Browser will be opened in a new window. You can also launch the Nutanix Objects Browser from the buckets list page (See [Viewing Buckets in Nutanix Objects](#) on page 125).

Alternatively, if you have multiple load balancer IP addresses for the object store, you can click any of the IP addresses to launch the Nutanix Objects Browser for that object store.




Figure 49: Prism Central - Launch Nutanix Objects Browser

The administrator can also share the object store URL with the user to access the Nutanix Objects Browser UI. The URL can be formed using the Nutanix Objects Public IP address.

`http://objects-public-ipaddress/objectsbrowser`

Note: You can use http or https to access the Nutanix Objects Browser.

Refer to the following image to understand the URL formation.

Note:

- Refreshing the page logs out the user and cancels the pending and in-progress uploads.
- Any uploads taking more than 60 minutes gets terminated and retried three times before getting canceled.

Browser Version Compatibility for Nutanix Objects Browser

Supported browser versions for Nutanix Objects Browser.

Table 38: Browser Compatibility

Browser	Minimum Supported Version
Google Chrome	92
Microsoft Edge	92
Safari	15.4

Browser	Minimum Supported Version
Mozilla Firefox	90

Administrator Workflow: Nutanix Objects Browser

The administrator needs to perform the following steps to grant access to the Nutanix Objects Browser to a particular user.

About this task

Note:

- Nutanix Objects Browser does not store the access and secret keys.
- If the user refreshes the page with the refresh button of the web browser, they will need to provide the credentials again, and any unsaved changes will be lost (including any pending or in-progress uploads).

To grant access to Nutanix Objects Browser, follow these steps:

Procedure

1. Add the IAM user.

For more information, see [Generating Access Key for API Users](#) on page 94.

After the administrator adds the user, the access and secret keys are generated.

2. (Optional) Grant access permissions (read only, full access, or custom) to the buckets.


For more information, see [Sharing a Bucket in Nutanix Objects](#) on page 128.

Note: This step is optional and can be performed if the administrator wants to share bucket of another user with a new user. If the administrator does not perform this step, the buckets list will be empty for the IAM users. The users can still create buckets from the Nutanix Objects Browser UI.

3. Share the Nutanix Objects Browser URL with a user to launch the object store in a web browser.

4. Share the access and secret keys that you generated for the user.

The user would require to enter the keys in the login page of the Nutanix Objects Browser UI.



The image shows the login page of the Nutanix Objects Browser UI. At the top, there is a header with the Nutanix logo and the word "Objects". Below the header, there are two input fields: "Access Key" and "Secret Key". The "Access Key" field has a placeholder text "Provide the access key" and the "Secret Key" field has a placeholder text "Provide the secret key". Below these fields is a blue "Login" button. At the bottom, there is a link that says "Don't have a key ? Generate keys".

Figure 50: Nutanix Objects Browser Login Page

What to do next

The user can log on to the Nutanix Objects Browser UI and perform CRUD operations within the object store. For example, create buckets, upload objects, assign a tag to objects, and so on. See [Supported Operations in Nutanix Objects Browser](#) on page 199.

Launching the Nutanix Objects Browser

Access the Nutanix Objects Browser UI using the object store URL and the keys provided by the administrator.


About this task

To launch the Nutanix Objects Browser, follow these steps:

Procedure

1. Open the object store URL shared by the administrator in a web browser.

2. In the **Nutanix Objects Browser** login page, enter the **Access** and **Secret** keys to access the object store instance.

The image shows the login page for the Nutanix Objects Browser. At the top, there is a header with the Nutanix logo and the word "Objects". Below this, there are two input fields: "Access Key" and "Secret Key". The "Access Key" field has a placeholder text "Provide the access key". The "Secret Key" field has a placeholder text "Provide the secret key". Below these fields is a blue "Login" button. At the bottom, there is a link that says "Don't have a key ? Generate keys".

Objects

Access Key

Provide the access key

Secret Key

Provide the secret key

Login

Don't have a key ? [Generate keys](#)

Figure 51: Nutanix Objects Browser Login Page

Note: The Administrator shares the access and secret key.

The object store is opened in the Nutanix Objects Browser.

What to do next

You can now perform CRUD operations on buckets and objects. See [Supported Operations in Nutanix Objects Browser](#) on page 199. You can also change the mode from light to dark, logout, or see the version by clicking the user name at top-right corner.

Supported Operations in Nutanix Objects Browser

Users can perform various operations on buckets and objects in the Nutanix Objects Browser UI.

In Nutanix Objects Browser UI, users can create buckets, configure policies, enable versioning, and set quotas. Buckets serve as logical containers for objects, and users can manage access permissions, lifecycle rules, and replication settings. These operations help structure and secure object storage for scalable data management.

Bucket Operations in Nutanix Objects Browser

After you log into the Nutanix Objects Browser, all the buckets that the administrator shared with you are listed with creation date and owner information.

You can perform the following operations at the bucket level:

- **Create Bucket:** Allows you to create buckets. For more information, see [Creating an S3 Bucket Using Nutanix Objects Browser](#) on page 200.

- **Lifecycle:** Allows you to create lifecycle rules.

Click the name of the bucket, and then click the **Lifecycle** option at the left pane. For information, see [Creating Lifecycle Rules Using Nutanix Objects Browser](#) on page 208.

- You can use the **Actions** list to update bucket properties, host a static website, configure CORS, and delete a bucket.

- **Update:** Allows you to update the bucket properties.

To update bucket properties, select the bucket, and then click **Actions > Update**.

- **Delete:** Allows you to delete a bucket.

To delete a bucket, select the bucket, and then click **Actions > Delete**.

Note: You can only delete an empty bucket. In the case of a version-enabled bucket, the delete operation performed on an object is not permanent. The object gets removed from the list and moved to Recycle Bin and a Delete Marker gets created. For more information, see [Understanding Object Versions](#) on page 226.

- **Share:** Allows you to share the bucket.

To update bucket properties, select the bucket, and then click **Actions > Share**.

You can share a bucket with multiple users and provide permissions such as **Read only**, **Full Access**, and **Custom**. For more information see, [Sharing a Bucket Using Nutanix Objects Browser](#) on page 206.

- **Static Website:** Allows you to configure a bucket to host a static website.

You can configure a bucket for website hosting, and then upload your website files (objects) to the bucket. For more information see, [Configuring a Bucket for Static Website Hosting using Nutanix Objects Browser](#) on page 211.

- **CORS:** Allows you to configure CORS on a bucket.

This allows the bucket to service cross-origin requests. For more information see, [Configuring CORS on a Bucket using Nutanix Objects Browser](#) on page 212.

Creating an S3 Bucket Using Nutanix Objects Browser

You can create, modify, and delete a bucket using the Nutanix Objects Browser.

About this task

Note:

- Make sure that the bucket names are unique for all users.
- You cannot configure a WORM bucket while creating a bucket. You can edit WORM policies only after creating a bucket.
- You cannot enable multi-protocol access on an S3 bucket.
- NFS buckets cannot be created using Nutanix Objects Browser.
- While creating a new versioned bucket, the system will automatically create a lifecycle policy to expire the delete markers if they are the only remaining object versions.

To create and configure a bucket, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. In the Application Switcher, click **Objects**.
3. Click the name of the object store in which you want to create a bucket and launch the Nutanix Objects Browser.

For more information on launching the Nutanix Objects Browser, see [Launching the Nutanix Objects Browser](#) on page 198.

The **Object Store** page appears.

4. Click **Add Bucket > Create Bucket**.
The **Create Bucket** window appears.

Create Bucket

Additional lifecycle rules can be added via Bucket details -> Lifecycle

General Settings

Name

Object Versions

Enable versioning if you wish to safeguard data against accidental deletions or overwrites or for applications that are versioning aware. If you enable versioning, consider applying lifecycle policies to this bucket to remove older versions of objects after a specified length of time. [Learn More](#)

☒ Enable versioning

Lifecycle Policies

☒ Expire current objects version after years

☐ Expire previous objects version after days

System will also automatically create a lifecycle policy to expire delete markers if they are the only remaining object version.

Cancel Create

Figure 52: Create Bucket Window

5. On the **General Settings** section, type a name for your bucket.
For more information on naming buckets, see [Buckets Naming Conventions](#).
6. (Optional) On the **Object Versions** section, configure the following:
For more information on versions, see [Object Versioning](#).
 - a. **Enable Versioning:** Select this check box to enable versioning of objects.
This allows you to keep all versions of an object in the same bucket. For more information on versions, see [Object Versioning](#).
To apply life cycle policy with versioning enabled, see [Rules for Expiration based Lifecycle Policy](#) on page 107.

Note: Enabling versioning allows you to recover objects from accidental deletion and overwrite.

7. (Optional) On the **Lifecycle Policies** section, select the following:

For more information, see [Lifecycle Policies](#).

a. **Expire current objects version after::** Select to type a time period after which the current version of the object expires.

You can specify the number in days, months, or years.

b. **Expire previous objects versions after:** Select this check box to enter a time period to delete all the previous versions of the objects. This option is available only if versioning is enabled.

You can specify the number in days, months, or years.

Note:

- If versioning is not enabled, the current object is deleted permanently. When you enable versioning, multiple versions of the same object are maintained in a bucket.
- NFS buckets cannot be created using Nutanix Objects Browser.

8. Click **Save**.

The bucket is created successfully.

What to do next

After creating a bucket, you can create objects through S3 APIs and manage them. For more information, see [Supported S3 APIs](#).

You can also perform various actions on a bucket, such as configuring static websites and configuring CORS on a bucket.

- For more information on creating lifecycle rules, see [Creating Lifecycle Rules Using Nutanix Objects Browser](#) on page 208.
- For more information on configuring static websites, see [Configuring a Bucket for Static Website Hosting using Nutanix Objects Browser](#) on page 211.
- For more information on configuring CORS on a bucket, see [Configuring CORS on a Bucket using Nutanix Objects Browser](#) on page 212.

Adding External Bucket

You can add an existing external bucket to the federated namespace.

Before you begin

To add an external bucket, following permissions are needed:

- `"s3:ListBucket"`
- `"s3:*Object"`

.

About this task

To add an external bucket, follow these steps:

Procedure

1. Log on to the Prism Central web console.

2. In the Application Switcher, click **Objects**.
3. Click the name of the object store in which you want to add a bucket and launch the Nutanix Objects Browser.
For more information on launching the Nutanix Objects Browser, see [Launching the Nutanix Objects Browser](#) on page 198.
The **Object Store** page appears.
4. Click **Add Bucket > Add External Bucket** .
The **Add External Bucket** window appears.

Figure 53: Add External Bucket Window

5. On the **General Settings** section, type a name of existing external bucket.

Note: Nutanix recommends that external buckets use a differentiating common prefix in order to distinguish external buckets from other buckets on Nutanix Objects.

6. On the **Endpoint**, select an appropriate external object store.
7. Provide **Access Key** and **Secret Key**.
These credentials are specific to the bucket that is being added.
8. Click **Add**.
External bucket is added successfully.

What to do next

After adding an external bucket, you can share the bucket with same permissions with another user. For more information, see [Sharing a Bucket Using Nutanix Objects Browser](#) on page 206.

- By default, full access permissions are granted locally.
- Final permissions are governed by the bucket policies defined in S3.

After sharing a bucket using Nutanix Objects Browser, you can upload Nutanix Objects to S3 bucket, for more information see,

You can also perform various actions on a bucket, such as configuring static websites and configuring CORS on a bucket.

- For more information on creating lifecycle rules, see [Creating Lifecycle Rules Using Nutanix Objects Browser](#) on page 208.

Uploading Nutanix Objects to External Buckets

You can upload objects to external buckets.

About this task

To upload objects to external buckets, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. In the Application Switcher, click **Objects**.
3. Select the object store and launch the **Nutanix Objects Browser**.
4. Click the **Buckets** tab, and in the **Buckets** table, select the bucket to upload Object.
5. Click **Upload Nutanix Objects > Select Files** .
Upload Nutanix Objects window appears.
6. Click **Close**.
The objects are in the bucket.

What to do next

(optional) After uploading Nutanix Objects to the external bucket, you can remove the bucket and the endpoint. For more information, see [Deleting a Bucket in Nutanix Objects](#) on page 137.

Bucket Operations in Nutanix Objects Browser for External Bucket

After you log into the Nutanix Objects Browser, all the buckets that the administrator shared with you are listed with creation date and owner information.

You can perform the following operations at the bucket level:

- **Create Bucket:** Allows you to create buckets. For more information, see [Creating an S3 Bucket Using Nutanix Objects Browser](#) on page 200.
 - **Lifecycle:** Allows you to create lifecycle rules.
Click the name of the bucket, and then click the **Lifecycle** option at the left pane. For information, see [Creating Lifecycle Rules Using Nutanix Objects Browser](#) on page 208.
 - You can use the **Actions** list to update bucket properties, and delete a bucket.
 - **Update:** Allows you to update the bucket properties.
To update bucket properties, select the bucket, and then click **Actions > Update**.
 - **Delete:** Allows you to delete a bucket.
To delete a bucket, select the bucket, and then click **Actions > Delete**.
- Note:** External buckets are only removed from the namespace, not deleted from the S3 endpoint.
- **Share:** Allows you to share the bucket.
To update bucket properties, select the bucket, and then click **Actions > Share**.
You can share a bucket with multiple users and provide **Full Access** permission. For more information see, [Sharing a Bucket Using Nutanix Objects Browser](#) on page 206.

Updating a Bucket using Nutanix Objects Browser

After you create a bucket, you can update the bucket settings using Nutanix Objects Browser.

About this task

Note: Life cycle rules created for a bucket cannot be modified from here. To modify the lifecycle rules for a bucket, see [Creating Lifecycle Rules Using Nutanix Objects Browser](#) on page 208.

To update bucket settings, follow these steps:

Procedure

1. Log on to the Nutanix Objects Browser.
For more information on launching Nutanix Objects Browser, see [Administrator Workflow: Objects Browser](#) and [Launching the Objects Browser](#).
2. From the object store, select the bucket that you need to update.
3. Click **tripple dot horizontal** icon > **Update**.

Update Bucket

Lifecycle rules for objects can be managed via [Bucket details -> Lifecycle](#)

General Settings

Name

Object Versions

Enable versioning if you wish to safeguard data against accidental deletions or overwrites or for applications that are versioning aware. If you enable versioning, consider applying lifecycle policies to this bucket to remove older versions of objects after a specified length of time. [Learn More](#)

☒ Enable versioning

☐ Suspend versioning

Cancel Update

Figure 54: Update Bucket Window

4. On the **Object Versions** section, select one of the checkboxes to enable or disable versioning of objects in the bucket.
 - a. **Enable versioning:** Select this check box to enable versioning on objects and to keep all the versions of the object on the same bucket.

Note: Select this option to recover objects from accidental deletion or overwrite.

- b. **Suspend versioning:** Select this check box to disable versioning of objects in a bucket.
When you suspend versioning, accumulation of the new object versions is stopped. However, versions of objects already existing in the bucket are retained.
5. Click **Done**.
The updated settings are applied to the bucket successfully.

Sharing a Bucket Using Nutanix Objects Browser

You can modify user access to grant different permissions, including read-only, full, or custom.

Before you begin

The ability to perform the share bucket operation is limited to the owners of the bucket. If a non-owner attempts to update the policy, they must have the necessary permissions granted by the owner or administrator in the bucket policy. Without these permissions, the non-owner receives an Access Denied error message.

About this task

To share a bucket, follow these steps:

Procedure

1. Log on to the Nutanix Objects Browser.

For information on launching the Nutanix Objects Browser, see [Launching the Nutanix Objects Browser](#) on page 198.

All the buckets that the administrator shared with you are listed with the creation date and owner information in a new window.

2. Click the **Buckets** tab, and in the **Buckets** table, select the bucket to be shared.

3. Click **tripple dot horizontal icon > Share**.

The name of the bucket owner is visible under **Owner**. You cannot edit the name of the bucket owner. Removing or changing the permission for the user who is a bucket owner does not affect the bucket. The owner is still allowed to do all the operations on a bucket.

4. Type the email address of the users or select **Anyone with Link** as a user and set the required permission for that user.

You can only add users with access keys. When you select **Anyone with Link** as a user, a public bucket access link is provided. You can copy and share the link with any user to directly access the bucket from the web browser.

Caution: Creating permissions with **Anyone with Link** as a user grants unauthenticated access, exposing the bucket to anonymous users.




Figure 55: Share Bucket Window

Nutanix supports the following list of permissions on a bucket:

- **Read Only:** Provides read-only access to the user.
- **Full Access:** Provides all access to the user.
- **Custom:** Provides customizing of varied levels of access to the user. Click **Set** under Permissions to set custom permission for a user.

Modifying or revoking permissions for a user who is the bucket owner does not affect the bucket. The owner retains unrestricted access and can continue to perform all operations on the bucket.

For more information on Bucket Permissions, see [Bucket Access Policies](#).

For more information on generating access keys, see [Generating Access Key for API Users](#) on page 94.

5. (Optional) To add more users and permissions, click **+ Add User and Permissions**.
 - If you add **Anyone with Link** as **Users** and grant them **Full Access**, the user is able to perform all operations without any authentication.
 - To add a new user, you must enter the precise user name without any variations or errors.
 - You can enter multiple user names by providing space between each user's name.
6. (Optional) Click the delete icon to remove an existing set of users and permission.
7. Click **Save**.

The bucket is now shared with the listed users with the allotted permissions.

Note: You cannot save a policy with an empty set of permissions. It is required to have at least one user and permission.

Creating Lifecycle Rules Using Nutanix Objects Browser

Lifecycle policy enables you to create or update a set of rules that define actions that Nutanix Objects applies to a group of objects within a bucket. With these policies, you can expire objects when no longer required.

Before you begin

See [Lifecycle Policies in Nutanix Objects](#) on page 106.

About this task

Note: Rules or any updates to the rules get applied to the new objects that you create and do not apply to the objects existing before the rule creation or update.

To create a lifecycle rule, follow these steps:

Procedure

1. Log on to the Nutanix Objects Browser.
2. In the list of buckets, click the name of the bucket for which you want to create a rule.
3. Click **Lifecycle**.
4. Click **Create Rule**.

When you have no rules created, you can directly import the XML. The **Import XML** option is not available if you already have rules to avoid overwriting.

5. In the **Scope** section, do the following, and then click **Next**:

The screenshot shows a three-step wizard: 1. Scope, 2. Configure Rule, and 3. Review. The 'Scope' step is highlighted with a red box. Below the steps, there is a 'Name' field with the value 'rule1'. The 'Scope' dropdown menu is set to 'Tags/Prefix'. Under this, there are two checked options: 'Prefix' and 'Tag'. The 'Prefix' section has a note 'Only one prefix is supported' and a text input field with the placeholder 'Please enter'. The 'Tag' section has a note 'Enter one or upto ten tags' and a table with two columns: 'Key' and 'Value'. Both columns have text input fields with the placeholder 'Please enter'. Below the table is a button labeled '+ Add Tag'. At the bottom right of the form are 'Cancel' and 'Next' buttons.

Figure 56: Defining Scope for a Rule

- a. **Name:** Enter a name that identifies the rule you are creating.
- b. **Scope:** Select the scope of the rule to either all objects of that bucket, or to tags and prefixes.
 - » **Prefix:** You can enter only one prefix.
 - » **Tag:** You can enter up to 10 tags in key value pair.
Click **+ Add Tag** to add more tags.

6. In the **Configure Rule** section, do the following, and then click **Next**:




Figure 57: Configuring Rule

You cannot configure the tiering endpoints from the Nutanix Objects Browser. You can configure it from the Nutanix Objects UI and the same is visible as *Read Only* in the Nutanix Objects Browser UI.

- a. **Expire** Select which version to expire: **Current version**, **Previous version**, **Multipart uploads**, or **Delete markers** according to your requirement.
- The **Previous version** and **Delete markers** options appear only for a version-enabled bucket.
 - Duplicate expiration field is not allowed. For example, you cannot create two rules for the current version expiration.
 - Expiration of delete marker and the current version cannot be configured in a single rule.

Tip: It is recommended to create a rule to automatically expire delete markers to enhance the performance of your object stores.

Multipart uploads expiration should not be specified with tag-based filters.

- b. **after last creation date** Specify the number in days, months, or years after which that respective version of the object expires after last creation date of that object.

You can click **Add Action** to add up to three expiration rules. You can create expiration rules for the current version, the previous version, and multipart uploads.

Click **Delete Icon** to delete the rule.

7. In the **Review** section, review the scope and actions, and then click **Done**.

- You can select a bucket, and update, delete, disable, and enable the rule using the **Actions** drop-down.
- You can also export the multiple rules to an XML file by clicking **Export to XML**.

Actions listed are executed in sequence.

The rule you just created gets enabled and appears in the Rules table.

Configuring a Bucket for Static Website Hosting using Nutanix Objects Browser

You can use Nutanix Objects Browser to host a static website that has individual web pages with static content. To host a static website on Nutanix Objects, you can configure a bucket for website hosting, and then upload your website files (objects) to the bucket. When you configure a bucket as a static website, you enable static website hosting, and optionally, add an index document and an error page. You can upload files (such as index documents and error pages) to the bucket using the Nutanix Objects Browser or S3 browser. The S3 browser uses the S3 protocols by providing access and the secret key. You can also choose to redirect to a website. Once you have configured your bucket as a static website, you can access the bucket through the object store endpoints for your bucket.

About this task

To configure a bucket for static website hosting, follow these steps:

Procedure

1. Log on to the Nutanix Objects Browser.
2. In the **Buckets** table, select the bucket to configure it for static website hosting.
3. Click **triple dot horizontal icon > Static Website**.
The **Configure Static Website** window appears.

4. By default, the endpoint is auto-populated when you click **Save** at the last step.

For example, when an endpoint auto populates, the URL will be in the format `objectstorename.domain/bucketname`. For example, `objectstore.nutanix.com/teamobjects`. However, if they have set up the DNS correctly, then you can also access the website with `bucketname.objectstorename.domain` endpoint using HTTP or HTTPS. For example, `https://teamobjects.objectstore.nutanix.com`.

5. Click the **Host Website or Redirect** check box.

- » **Use this bucket to host a website:** Select this option to use the bucket to host the website. Optionally, you can enter the name of the index document (for example, `myindex.html`) and an error page.

An index document is a web page that Nutanix Objects returns when you request the root of a website. It is the default page that loads when you are not requesting any specific page. After you enable static website hosting for your bucket, you can upload an HTML file with the index document name (for example, `myindex.html`) to your bucket. For example, if you specify no object in the URL, then the website loads the index page (`myindex.html`) that you have configured. If you have not configured an index document, then a website access to the root will return an access denied error message.

A custom error page is a web page that Object returns when an error occurs. For example, if you are trying to load an object that does not exist, the website loads the error page that you have configured.

- » **Redirect:** Select this option to enter a website URL to redirect to that website. For example, when you try to access the bucket endpoint, you will be redirected to this website. The protocol used is either HTTP or HTTPS.

To remove the static website configuration from the Nutanix Objects Browser, uncheck **Host website or redirect**, and then click **Save**.

6. Click **Save.**

An endpoint is auto-generated when you click **Save**. This endpoint will be the object store endpoint for your bucket and is used as the website address.

You can now use your bucket as a static website. You can use the endpoint to test your static website.

Configuring CORS on a Bucket using Nutanix Objects Browser

Cross-Origin Resource Sharing (CORS) allows a web application loaded in one domain to access the restricted resources that are requested from another domain.

Before you begin

See [Cross-Origin Resource Sharing \(CORS\) in Nutanix Objects Overview](#) on page 122.

About this task

You set this configuration on a bucket so that the bucket can service cross-origin requests.

To configure CORS for a bucket, follow these steps:

Procedure

1. Log on to the Nutanix Objects Browser.
2. In the **Buckets** table, select the bucket to configure CORS.
3. Click **trippel dot horizontal icon** > **CORS**.
The **Configure CORS** window appears.
4. Type or copy and paste a configuration file, or edit an existing configuration.
The configuration file must be an XML file.
5. Click **Save**.
The CORS configurations are saved for the bucket.

Adding Tags to a Bucket Using Nutanix Objects Browser

A tag is a label that you assign to a bucket, and it consists of a key and a value pair that you can define. You can add or remove tags from your buckets using the Nutanix Objects Browser.

About this task

To add tags to a bucket from the Nutanix Objects Browser, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. In the Application Switcher, click **Objects**.

3. Click the name of the object store and launch the Nutanix Objects Browser.
For information on launching the Nutanix Objects Browser, see [Launching the Nutanix Objects Browser](#) on page 198.
The **Object Store** page appears.
4. In the **Buckets** table, select the bucket to add tags.
5. Click **tripple dot horizontal icon > Tags**.
The **Add Tags** window appears.
6. To add a new tag, click **+ add tag**.
The existing tags will be auto-populated. You can click the delete icon to delete a tag.

Note:

- The maximum number of tags allowed per bucket is 50.
- Tag keys can be up to 128 Unicode characters in length, and tag values can be up to 256 Unicode characters in length.
- Tag keys and values are case-sensitive.
- Tag key should be unique and is a mandatory field, but tag value is optional.

A new row appears to fill the key value of the tag.

7. Enter the key and value.
8. Click **Save**.
The changes are saved.

Bucket Summary in Nutanix Objects Browser

The Summary page allows you to view the list of various policies applied to the bucket.

To view the **Summary** page of a bucket in an object store, click the name of the bucket in the buckets table, and then click **Summary**.

Object Operations in Nutanix Objects Browser

You can perform the following operations on an object within a bucket in the Nutanix Objects Browser UI:

- Upload objects
For more information, see [Uploading an Object Using Nutanix Objects Browser](#) on page 217.
- Create new text object
For more information, see [Creating a Text Object with the Nutanix Objects Visualizer in Nutanix Objects Browser](#) on page 219.
- Create new folder
For more information, see [Creating a Folder in Nutanix Objects Browser](#) on page 220.
- Open an object: Open the objects and preview it.
Open objects in Nutanix Objects visualizer using a user-selected preview plug-in. The preview can be in the form of text, image, PDF, audio, or video.
- For more information, see [Nutanix Objects Visualizer](#) on page 214.

- Download an object
For more information, see [Nutanix Objects Visualizer](#) on page 214.
- Copy sharing link with other users: Generate a link to the object that you can share with other users. Other users can directly open the shared object using the link and perform actions depending on the permissions. Users are prompted to log on when they open the link.
- Add tags to an object
For more information, see [Adding Tags to an Object Using Nutanix Objects Browser](#) on page 221.
- View and manage the different versions of an object
For more information, see [Understanding Object Versions](#) on page 226.
- Query the CSV objects using the SQL query
For more information, see [Querying CSV Nutanix Objects Using S3 Select in Nutanix Objects Browser](#) on page 224.
- Delete an object
For more information, see [Deleting an Object Using the Nutanix Objects Browser](#) on page 225.
- Search an object
For more information, see [Searching an Objects Using Nutanix Objects Browser](#) on page 223.

Object Naming Conventions

The object key name is a sequence of UTF-8 characters not exceeding 1024 bytes.

You can use the following characters to create an object key name or prefixes:

- Alphanumeric characters—lowercase letters (a-z), uppercase letters (A-Z), and numbers (0-9)
- Special characters—forward slash (/), exclamation mark (!), hyphen (-), underscore (_), period (.), asterisk (*), single quote ('), open parenthesis ((), and close parenthesis ())

Consider the following limitations before you create an object key name:

- Only the UTF-8 character set is supported.
- The length of the object key name cannot exceed 1024 bytes.
- Special characters, such as pound (#) and percent (%) are not supported.

Note: Starting with the Nutanix Objects 5.1 release, Nutanix Objects no longer supports null bytes in object names, aligning this change with Amazon Web Services Simple Storage Service (AWS S3) and Microsoft Azure.

Nutanix Objects Visualizer

Within the Nutanix Objects Browser, you can directly interact with object content.

You can perform the following tasks directly within the Nutanix Objects Browser:

1. *Preview Content:* You can conveniently preview the content of the most common types of objects without the need for downloading.
2. *Edit object content:* For specific object types, you have the capability to edit content directly within the objects visualizer.
3. *Stream and view large objects:* You now have the option to stream and view audio/ video without the requirement for complete downloads.

4. *Integrated Help/Support:* When editing certain types of objects, you can access integrated help and support directly within the Nutanix Objects Visualizer in the form of syntax highlighting and autocomplete snippets.

Supported Object Types in Nutanix Objects Visualizer

Nutanix Objects visualizer supports following object types:

- Image
- Text Code
- Audio
- Video
- PDF

Image

This functionality enables the preview of common image formats, supporting the following: png, jpeg, jpg, webp, gif, apng, avif, svg, bmp, ico.

Supported operations:

- A user-friendly slider is available in the UI for zoom adjustments.
- *Rendering*
 - Images are rendered using the largest possible size based on the user's view port.
 - If image dimensions exceed the view port, they are scaled down, with the scaling reflected in the zoom slider.
 - The aspect ratio of the image is preserved during scaling to maintain image quality and prevent distortion.
 - Scaling is restricted to the image's actual resolution to prevent overstretching.
- Users can press and drag images on the screen to navigate and focus on specific areas of large images.
- *Download:* For animated formats (like gif, avif), the plugin renders only the first frame, displaying a static image in the preview. For more information, see [Downloading Objects Using Nutanix Objects Browser](#) on page 222.
- The plugin supports images up to 50 MB.

Text Code

This functionality facilitates the preview and editing of various text-based objects, supporting a broad range of formats: log, xml, json, txt, csv, html, htm, js, ts, jsx, tsx, sh, php, scss, less, sass, css, py, c, cpp, java, swift, h, dart, gitignore, go, proto, json5, kotlin, ps1, scala, sql, yaml, yml, md, tf, rust, r, ini, man, me, xsl, coffee, scpt, xbat, ws, vbs, jsf, ejs, mjs, crt, dhtml, chtml, rjs, shtml, public, asp, aspx, xss, rhtml, do, csr, phtml, xhtml, xht, jss, cgi, pl, jsp, cs.

Note: While editing large text objects, the performance might be degraded with an increase in object size.

Supported operations:

- xml, json, html, js, jsx, ts, tsx, php, css, py, c, cpp, java, go, proto, sql, yaml, yml, md objects feature color-coded syntax highlighting and collapsible sections for easy content navigation.

- You can initiate a search on content using `CMD + F` or `CTRL + F` for efficient find and replace operations.
- *Download*: Downloads the object. For more information, see [Downloading Objects Using Nutanix Objects Browser](#) on page 222.
- The plugin supports files up to 50 MB.

Audio

This feature enables the preview of audio with seamless streaming support. It supports a variety of audio formats, including mp3, ogg, and wav.

Supported operations:

- *Play*: Initiate audio playback.
- *Pause*: Halt audio playback.
- *Seek*: Navigate to specific points in the audio.
- *Streaming*: Audio content can be previewed without the need to download the entire object.
- *Download*: Downloads the object. For more information, see [Downloading Objects Using Nutanix Objects Browser](#) on page 222.
- There is no enforced size limit for preview since only partial content is loaded in the browser. Unused content is periodically garbage collected to optimize resource usage.

For more information on the limitations, see [Limitations for Audio and Video Preview in Nutanix Objects Visualizer](#) on page 217.

Video

This feature enables users to preview videos with seamless streaming support. It supports various video formats, including mp4, webm, mov, and 3gp.

Supported operations:

- *Play*: Initiate video playback.
- *Pause*: Halt video playback.
- *Seek*: Navigate to specific points in the video.
- Expand the video to full-screen mode.
- *Picture-in-Picture Mode (in few browsers)*: Enjoy video playback in a compact overlay, allowing multitasking.
- *Download*: Downloads the object. For more information, see [Downloading Objects Using Nutanix Objects Browser](#) on page 222.
- Video content can be previewed without the need to download the entire object. Partial data is streamed from object stores using content-range, rendering it directly in the browser. This feature enables users to preview even large videos without complete downloads.
- There is no enforced size limit for preview since only partial content is loaded in the browser. Unused content is periodically garbage collected to optimize resource usage.

For more information on the limitations, see [Limitations for Audio and Video Preview in Nutanix Objects Visualizer](#) on page 217.

PDF

This functionality enables the preview of PDF objects and is compatible with browsers that support native PDF rendering.

Supported operations:

- Efficiently renders a limited number of pages by default, loading additional pages dynamically as the user scrolls
- You can navigate directly to a desired page by specifying its page number.
- You can adjust the PDF scale using either a slider or options available in the user interface.
- You can click on and utilize links within PDFs.
- This plug-in can support PDF up to 100 MB.
- Compatibility for native and custom PDF viewers is ensured on the following browser versions:
 - Google Chrome version 92 and above.
 - Microsoft Edge version 92 and above.
 - Safari version 15.4 and above.
 - Mozilla Firefox version 90 and above.

Limitations for Audio and Video Preview in Nutanix Objects Visualizer

The following are the limitations for audio and video objects preview in Nutanix Objects Browser:

- Presigned URL Expiration: Nutanix Objects Visualizer relies on presigned URLs with a 15-minute expiration time for loading object previews. This constraint limits the time frame during which the plugin can fetch and process audio/video content.
- Audio/Video Operations Timeframe: The plugin supports all operations on audio and video content only within the 15-minute window of the presigned URL's validity. Beyond this timeframe, certain operations may be affected.
- Playback Disruption: Once the presigned URL expires, there is a potential for playback disruption. Specifically, audio operations may encounter difficulties when attempting to resume playback from parts that were not downloaded or were subject to garbage collection.
- Limited Support for URL Renewal: Currently, the plugin does not support the renewal of expired presigned URLs. This limitation impacts the ability to extend the timeframe for operations and playback beyond the initial 15-minute duration.
- Media Play Continuation: The plugin lacks built-in support for the seamless continuation of media playback after the expiration of the presigned URL. This may result in interruptions, particularly when dealing with large audio files.
- Handling Extended Video/Audio Durations: Similar challenges may arise when attempting to play video or audio files exceeding 15 minutes in duration for their full intended playback duration.

Uploading an Object Using Nutanix Objects Browser

After creating a bucket, you can use the Nutanix Objects Browser to select and upload files or folders containing multiple files to a bucket.

About this task

To upload files or folders to a bucket, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. In the Application Switcher, click **Objects**.
3. Click the name of the object store in which you want to create a bucket and launch the Nutanix Objects Browser.
For information on launching the Nutanix Objects Browser, see [Launching the Nutanix Objects Browser](#) on page 198.
The **Object Store** page appears.
4. Click the bucket name to which you want to upload files or folders.
5. (Optional) To select and upload a file, click **Upload Objects > Select Files**
The **Upload Objects** window appears.

Note:

- This page displays the upload progress, the upload status, the bucket name, the object name, the size of the object, and the actions that you can perform on the object. You can also search by object name, bucket name or prefix.
- Uploading objects is an asynchronous process. If the upload size is large, you can close the **Upload Objects** window and perform other operations in the Object Browser UI. You can check the upload status by clicking the status icon in the top-right corner of the page.
- Multi-part upload is used to upload large files (more than 1 GB). This feature enables to upload objects of any size to the object store. Depending upon the size of the file, upload may take a few minutes.
- If you refresh the page using the refresh button of the web browser, logout, or navigate to any non-application URL, any unsaved changes will be lost. This includes any pending or in-progress uploads. A warning message appears if you try to log out or refresh the page

while an upload is in progress. Keep the Nutanix Objects Browser open until the uploads are complete.




Figure 58: Upload Objects Window

- a. Click the **Cancel All Updates** to cancel the updates that are in progress.
 - b. Click the **Close** button or the **X** icon, to close the **Upload Objects** window.
6. (Optional) To select and upload a folder and the files in the folder, click **Upload Objects > Select Folder**
The **Upload Objects** window appears.
7. (Optional) Click **Summary** to view the bucket summary.

What to do next

After you upload objects to the bucket, you can add a tag to an object, download an object, view and manage the versions of an object, and also delete an object. For more information, see:

- [Object Operations in Nutanix Objects Browser](#) on page 213
- [Adding Tags to an Object Using Nutanix Objects Browser](#) on page 221
- [Understanding Object Versions](#) on page 226
- [Deleting an Object Using the Nutanix Objects Browser](#) on page 225

Creating a Text Object with the Nutanix Objects Visualizer in Nutanix Objects Browser

Create a text object in a bucket of an object store with the Nutanix Objects Visualizer in the Nutanix Objects Browser UI.

Before you begin

See [Nutanix Objects Visualizer](#) on page 214.

About this task

To create a text object in the Nutanix Objects Browser, follow these steps:

Procedure

1. Log on to the Prism Central web console.

2. From the [Application Switcher Function](#), click **Objects**.
3. Select the object store, click the **Triple Dot Horizontal** icon, and click **Launch Objects Browser**.
4. Enter the access and the secret key.

Note: The Administrator shares the access and secret key.

For information on launching the Nutanix Objects Browser, see [Launching the Nutanix Objects Browser](#) on page 198.

The **Object Store** page appears in the Nutanix Objects Browser.

5. Click the bucket name that contains the object.
6. Click **Create New > Text Object**.
A **New Text Object** dialog box appears.
7. Enter the object name and click **Create**.

An object can be any form of text, including code or other written content. You can assign any file extension to the object name, and the system treats it as a text file.

The text file opens up in Nutanix Objects Visualizer in edit mode.

8. Edit the text file and click **Save**.
The system saves the text file and switches to view mode.
9. (Optional) Do any of the following:
 - » **Edit:** Click to edit the file.
 - » **Download:** Click to download the file.
10. (Optional) To go back to the **Objects List** page, click the **Close** icon.
The newly created object appears in the **Objects List** page.

What to do next

You can perform other object-related tasks, such as creating a folder, uploading objects, managing versions, and deleting an object. For more information, see the following references:

- [Creating a Folder in Nutanix Objects Browser](#) on page 220
- [Object Operations in Nutanix Objects Browser](#) on page 213
- [Understanding Object Versions](#) on page 226
- [Deleting an Object Using the Nutanix Objects Browser](#) on page 225

Creating a Folder in Nutanix Objects Browser

Create folders to organize and manage your data within a bucket of an object store in the Nutanix Objects Browser UI.

About this task

You can create multiple folders to organize the objects in a bucket.

To create a folder in the Nutanix Objects Browser, follow these steps:

Procedure

1. Log on to the Prism Central web console.

2. From the [Application Switcher Function](#), click **Objects**.
3. Select the object store, click the **Triple Dot Horizontal** icon, and click **Launch Objects Browser**.
4. Enter the access and the secret key.

Note: The Administrator shares the access and secret key with the user.

For information on launching the Nutanix Objects Browser, see [Launching the Nutanix Objects Browser](#) on page 198.

The **Object Store** page appears in the Nutanix Objects Browser.

5. Click the bucket name that contains the object.

Note: You cannot create a new folder in NFS-enabled buckets.

6. Click **Create New > Folder**.

Note:

- Folder names can only contain alphanumeric characters and specific special characters, including exclamation marks, hyphens, spaces, underscores, dots, asterisks, single quotes, and parentheses.
- When you create a folder, Nutanix Objects Browser creates a 0-byte object named after the specified folder, followed by a slash (/).

A **New Folder** dialog box appears.

7. Enter the folder name and click **Create**.
The new folder is created and appears in the **Objects List** page.

What to do next

You can perform other object-related tasks, such as uploading objects, managing versions, and deleting an object. For more information, see the following references:

- [Creating a Text Object with the Nutanix Objects Visualizer in Nutanix Objects Browser](#) on page 219
- [Object Operations in Nutanix Objects Browser](#) on page 213
- [Understanding Object Versions](#) on page 226
- [Deleting an Object Using the Nutanix Objects Browser](#) on page 225

Adding Tags to an Object Using Nutanix Objects Browser

You can use the Nutanix Objects Browser to add tags to an object in a bucket.

About this task

To add tags to an object, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. In the Application Switcher, click **Objects**.

3. Click the name of the object store and launch the Nutanix Objects Browser.
For information on launching the Nutanix Objects Browser, see [Launching the Nutanix Objects Browser](#) on page 198.
The **Object Store** page appears.
4. Click the bucket name that contains the object.
5. To tag an object, select the object that you need to tag, and click **Actions > Tags**.
The **Add Tags** window appears.




Figure 59: Adding Tags

- a. Enter the key and value pair.
- b. Click **Add Tag** to add more tags.

Note: You can add up to 10 tags for an object.
Click the delete icon against a key-value pair to delete it.

- c. Click **Save**.

What to do next

You can perform other object-related tasks, such as uploading objects, managing versions, and deleting an object. For more information, see:

- [Object Operations in Nutanix Objects Browser](#) on page 213
- [Understanding Object Versions](#) on page 226
- [Deleting an Object Using the Nutanix Objects Browser](#) on page 225

Downloading Objects Using Nutanix Objects Browser

You can download an object from the Nutanix Objects Browser.

About this task

To download an object, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. From the [Application Switcher Function](#), click **Objects**.
3. Click the name of the object store and launch the Nutanix Objects Browser.
For information on launching the Nutanix Objects Browser, see [Launching the Nutanix Objects Browser](#) on page 198.
The **Object Store** page appears.
4. Click the bucket name that contains the object.
5. Click **Actions > Download**.
The object is downloaded.

What to do next

You can perform other object-related tasks, such as uploading objects, managing versions, and deleting an object. For more information, see the following references:

- [Object Operations in Nutanix Objects Browser](#) on page 213
- [Understanding Object Versions](#) on page 226
- [Deleting an Object Using the Nutanix Objects Browser](#) on page 225

Searching an Objects Using Nutanix Objects Browser

Perform a prefix-based search for an object in the search bar of the Nutanix Objects Browser.

About this task

You can perform the search on the **Objects List** page and the **Recycle Bin**. The search keyword is case-sensitive. You can enter any prefix as a search keyword and the objects starting with that keyword are listed. For example, if you search for the prefix *copy*, all objects whose names start with the keyword *copy* are listed.

You can also search for exact name matches. For example, searching for *copy.txt* returns only that specific file. Files with similar names, such as *copy.text*, do not appear in the results.

To search an object, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. From the [Application Switcher Function](#), click **Objects**.
3. Click the name of the object store and launch the Nutanix Objects Browser.
For information on launching the Nutanix Objects Browser, see [Launching the Nutanix Objects Browser](#) on page 198.
The **Object Store** page appears.
4. Click the bucket name that contains the object.
The **Objects List** page appears.

5. In the search bar, type the prefix or the exact name as the search keyword.

Note: The scope of search is limited to the current folder. For example, if you search from within the *BigData* folder, the search results appear only from the objects in the *BigData* folder, and not from the entire list of objects.

The system displays the search results.

What to do next

You can perform other object-related tasks, such as uploading objects, managing versions, and deleting an object. For more information, see the following references:

- [Object Operations in Nutanix Objects Browser](#) on page 213
- [Understanding Object Versions](#) on page 226

Querying CSV Nutanix Objects Using S3 Select in Nutanix Objects Browser

S3 select adds support for querying the CSV objects using SQL query and viewing the results. You can perform all S3 Select-compatible operations using the queries.

Before you begin

Read [Nutanix Objects S3 Select API Overview](#) on page 246 and [Supported SQL Functions by Nutanix Objects S3 Select](#) on page 247 sections.

About this task

To perform S3 Select compatible operations using SQL query, follow these steps:

Procedure

1. Launch the Nutanix Objects Browser.

For information on launching the Nutanix Objects Browser, see [Launching the Nutanix Objects Browser](#) on page 198.

The **Object Store** page appears.

2. Click the bucket name that contains the object.
3. Select the CSV object to run the query on, and click **Actions > Run SQL Query**.

Note:

- Querying of CSV objects with header information is only supported.
- Query with column numbers is not supported.

The **Run SQL Query on <CSV_object_name>** window appears.

4. Enter the SQL query in the **Query** field and click **Run Query**.

Note: You can only run the S3 Select supported queries in the **Query** field. See [Supported SQL Functions by Nutanix Objects S3 Select](#) on page 247.




Figure 60: Nutanix Objects Browser - Running SQL Query on the CSV Object

Note: The result is limited to the first 1000 rows.

The result is displayed in a tabular format in the **Result** section.

5. (Optional) Click the **Export** link to export the result.

What to do next

You can perform other object-related tasks, such as uploading objects, managing versions, and deleting an object. For more information, see the following references:

- [Object Operations in Nutanix Objects Browser](#) on page 213
- [Understanding Object Versions](#) on page 226

Deleting an Object Using the Nutanix Objects Browser

You can use the Nutanix Objects Browser to delete objects in a bucket that you do not need.

About this task

To delete objects from a bucket, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. From the [Application Switcher Function](#), click **Objects**.
3. Click the name of the object store and launch the Nutanix Objects Browser.
For information on launching the Nutanix Objects Browser, see [Launching the Nutanix Objects Browser](#) on page 198.
The **Object Store** page appears.
4. Click the bucket name that contains the object.
5. Select the objects that you need to delete, and click **Actions > Delete**.

6. Click **Confirm**, to confirm deletion of the objects.

Note:

- For versioned buckets, the delete operation performed on an object is not permanent and a delete marker is created. You can view the deleted objects, and if needed, retrieve them from the **Recycle Bin**.
- For non-versioned buckets, the delete operation performed on an object is permanent and the deleted object cannot be retrieved. Non-versioned buckets do not have a **Recycle Bin**.
- When you delete an object within a folder, the system lists the deleted object in the recycle bin under the folder name. For example, if the file "object1.txt" is in the "BigData" folder, it appears in the recycle bin as "BigData/object1.txt."

7. Click **Recycle Bin**, select the objects to be deleted permanently and click **Delete Permanently**.




Figure 61: Delete Nutanix Objects

Caution:

- The **Delete Permanently** action permanently deletes all versions of the selected objects. You cannot recover the deleted objects after you perform this operation.
- Any new version added to an object that is currently being deleted, gets automatically deleted.

What to do next

You can perform other object-related tasks, such as uploading objects and folders, managing versions, and adding tags to an object. For more information, see:

- [Object Operations in Nutanix Objects Browser](#) on page 213
- [Understanding Object Versions](#) on page 226
- [Adding Tags to an Object Using Nutanix Objects Browser](#) on page 221

Understanding Object Versions

Versions allow you to view and manage the different versions of an object. From the **Versions** page, you can revert, delete, and download the different versions of an object. These options are available only for version-enabled buckets.

For more information on versions, see [Object Versioning](#).

For more information on managing object versions, see [Managing Object Versions Using Nutanix Objects Browser](#) on page 227.

When you delete a version-enabled object from the **Bucket** page, it is removed from the objects list and moved to the **Recycle Bin**.

To view the latest version of an object, click the object name, and then click **View all versions**. The latest version is listed with a **Delete Marker** banner in the **Recycle Bin**.

You can permanently delete any version by selecting the object version, and then by clicking the **Delete** button. You can also permanently delete the selected objects including all its versions from the **Recycle Bin**. Select the objects, and then click **Delete Permanently**.

Note:

- Selection of objects is limited to a page-by-page basis. If you select all objects listed in the Recycle Bin, all the objects listed on the first page are deleted. For example, if the total number of objects in the Recycle Bin is 150, but the first page has 100 objects listed, then the 100 objects are deleted.
- For the selected objects, if any new version is added while the deletion is in progress, that version is also deleted.

To restore (revert) this object to any previous version, select the object, and then click **View all versions**. Select the version of the object, and then click **Revert**. Now, the reverted version is the latest version visible in the objects list. You can also select the latest delete marker and delete it to restore the object to the last version.



Figure 62: Managing Object Versions - Using Nutanix Objects Browser

Note: Versions with a *Delete Marker* banner can only be deleted permanently. After deletion, they cannot be reverted or downloaded.

Managing Object Versions Using Nutanix Objects Browser

You can use the Nutanix Objects Browser to view and manage the versions of an object. Version allows you to revert an earlier version of an object to its latest version. You can also delete or download any version of the object.

About this task

To manage objects versions, follow these steps:

Procedure

1. Log on to the Prism Central web console.
2. In the Application Switcher, click **Objects**.

- Click the name of the object store and launch the Nutanix Objects Browser.
For information on launching the Nutanix Objects Browser, see [Launching the Nutanix Objects Browser](#) on page 198.
The **Object Store** page appears.
- Click the bucket name that contains the object.
- Select the objects whose version you need to manage, and click **Actions > Versions**.




Figure 63: Nutanix Objects Versions

Note: The latest version of an object is indicated by appending the banner Latest appended to the object name.

- (Optional) Select a version that you need to revert as the latest version, and click **Revert**.




Figure 64: Revert Version

The above image depicts an example, where employeeDetails.xlsx (1) is reverted to the latest version indicated by employeeDetails.xlsx (4) [Latest].

Note:

- Using the revert feature, you can designate any earlier version of an object as its latest version.
- If you deleted an object from the **Buckets** page, where all the objects within the bucket are listed. On the **Versions** page, you can select a version that does not have the Delete Marker banner, and then click **Revert**. The object gets restored on the **Bucket** page.

- (Optional) Select a version of the object that you need to delete, and click **Delete**.
This action permanently deletes the selected version of the object.

8. (Optional) Select a version of the object that you need to download, and click **Download**.
You cannot download the object version that is marked with a **Delete Marker** banner. The object is downloaded to the default download folder of your browser.
9. Click **Close** or the **X** icon, to close the window.

User and Key Management in Nutanix Objects Browser

Nutanix Objects Browser also provides end users with access to the Object service directly.

With the Nutanix Objects Browser, you can also enable IT consumers within your organization, such as development, test, or devops team to manage their own access and secret keys directly through the Nutanix Objects Browser, without requiring intervention from a Prism Central Administrator.

Roles for User and Key Management in Nutanix Objects Browser

The following roles are available when configuring User and Key management in Nutanix Objects Browser.

1. *Prism Central Administrator*: Use the Prism Central administrator role to set up access control policies, delegating management of bucket access and secret keys to end users.
2. *Sub Admin*: Use the sub-admin role to manage access keys on behalf of a specific group of users. A user with a sub admin role optionally creates new users if additional permissions are granted.
3. *End user*: Use the end user role for direct access to Object storage through the browser interface or an external S3 client.

Prerequisites for User and Key Management in Nutanix Objects Browser

Deployment prerequisites, port configurations, and IAM roles requirements for user and key management in Nutanix Objects Browser.

Ensure that you meet the following requirements before configuring the user and key in the Nutanix Objects Browser:

- Required IAM roles with appropriate permissions are configured.
- The client system that opens the Nutanix Objects Browser must be able to reach the Prism Central IP over port 9440 for Active Directory/LDAP authentication.
- The Nutanix Objects Storage network must connect to Prism Central over port 8443 to perform IAM API operations.

Limitations for User and Key Management in Nutanix Objects Browser

The following limitations apply when you configure the user and key in Nutanix Objects Browser:

- Only Active Directory (AD) or LDAP users can log in to the Nutanix Objects for user and key management. Local users are not supported.
- Only the IAM users imported from AD or LDAP can be managed from Nutanix Objects Browser. Local IAM users cannot be managed from Objects Browser.
- Users can only view, manage, or delete the access keys that they create. Keys created using the Prism Central UI are owned by the admin and are not visible or editable by other users—even if those keys are associated with their username.
- A user who creates keys for another user becomes the owner of those keys. As a result, the recipient user cannot manage those keys in Nutanix Objects Browser.
- The permissions required for user and key management in Nutanix Objects Browser are different from those required for managing IAM users and keys from the Prism Central Objects page.

User and Key Management in Nutanix Objects Browser Work flow

Configuring user and key in Nutanix Objects Browser feature enables the end users to manage their own access key and secret keys.

The Nutanix Objects Browser user and key feature is an extension of the existing access key management capabilities using identity and access management (IAM) in the Prism Central interface. The primary use cases are the following:

1. Support for end users to manage their own access keys, this is applicable in both of the following cases
 - The end user already exists in IAM.
 - The end user does not exist in IAM.
2. Support for delegated sub-admins to manage access keys for other users.

Table 39: Nutanix Objects Browser Workflow for Configuring Permissions and Policies

Sequence	Tasks	Description
1	Configure roles for admins or the end user with specific set of permissions.	• For configuring role, see Creating a Custom Role for End User on page 230 • For a Sub-admin managing the keys, Creating a Role for Sub-Admin on page 233
2	Configuring access control policy	• For an existing user, see Creating an Authorization Policy for Configurable Access (for an existing user) on page 231 • For a non existing user, see Creating an Authorization Policy for Configurable Access (for a non existing user) on page 232 • For a Sub-admin managing the keys, Creating an Authorization Policy for a Sub-Admin on page 234
3	Managing IAM Keys	For more information, see Managing IAM Keys using LDAP on page 235

Creating a Custom Role for End User

You can create a custom role for an end user.

About this task

To create the custom role for an end user in the Prism Central UI, follow these steps:

Procedure

1. Log in to Prism Central as an administrator.
2. From the **Application Switcher Function**, select the **Admin Center** application and from the navigation bar, select **IAM**.
3. Select the **Roles** tab.
4. Click **Create Role > New Role**.
5. In the **Role Name** field, enter a unique name for the role.
6. (Optional) In the **Description** field, enter a description.
7. In the **Entity Type** filter the individual operations available for a specific resource type, select the Operation from the drop down list and in the search box enter the name of the resource.

Note: The following permissions are required:

- View User
- View User Key
- Create User Key
- Delete User Key

8. Click the plus icon to add an individual operations to the role.
9. Do one of the following:
 - » To create the role, click **Save**.
 - » To create the role and attach the authorization policy, click **Save & Create Authorization Policy**.

What to do next

- After creating a custom role for existing user, you can create authorization policy. For more information, see [Creating an Authorization Policy for Configurable Access \(for an existing user\)](#) on page 231 section.
- After creating a custom role for non-existing user, you can create authorization policy. For more information, see [Creating an Authorization Policy for Configurable Access \(for a non existing user\)](#) on page 232 section.

Creating an Authorization Policy for Configurable Access (for an existing user)

Create an authorization policy with full access to all entity types and instances for the added users in the associated role.

Before you begin

Create a custom role for an end user to generate their own user key. For more information, see [Creating a Custom Role for End User](#) on page 230.

About this task

To create an authorization policy, follow these steps:

Procedure

1. Log in to Prism Central as an administrator.
2. From the **Application Switcher Function**, select the **Admin Center** application and from the navigation bar, select **IAM**.
3. Select the **Authorization Policies** tab.
4. Click **+ Create Authorization Policy**.
5. (Optional) To edit the name, click the pencil icon next to the default name, edit the name, and click the checkmark icon.
6. On the **Select Role** search box, enter the built-in or custom role's name and select the role from the list of suggestions. The system displays role details for the selected role.
7. Click **Next**.
8. Select **Configure access: select entity type & instances** and follow these steps:
 - a. From the **Entity Type** dropdown list, select the entity type.
The list of available entities depends on the role that you select.
 - b. From the **Filters** dropdown list, select the filter type.
 - c. On the **Search** box, select the specific entities that belongs to the selected entity type.
9. To allow the assigned users to access any new entity instances created by them, select the **Allow user access to entities created by them** checkbox.
10. Click **Next**.
11. From the **Users** dropdown list, select the active directory the user belongs to.
12. On the **Search** box, enter the first few letters of the user or user group's name and select the correct user or user group from the list of suggestions.

Note: The **scope** and **identities** must have the same user.

13. Click **Save**.

What to do next

After creating an authorization policy, you can manage IAM keys. For more information, see [Managing IAM Keys using LDAP](#) on page 235

Creating an Authorization Policy for Configurable Access (for a non existing user)

Create an authorization policy with full access to all entity types and instances for the added users in the associated role.

Before you begin

Create a custom role for an end user to generate their own user key. For more information, see [Creating a Custom Role for End User](#) on page 230.

About this task

To create an authorization policy, follow these steps:

Procedure

1. Log in to Prism Central as an administrator.
2. From the **Application Switcher Function**, select the **Admin Center** application and from the navigation bar, select **IAM**.
3. Select the **Authorization Policies** tab.
4. Click **+ Create Authorization Policy**.
5. (Optional) To edit the name, click the pencil icon next to the default name, edit the name, and click the checkmark icon.
6. On the **Select Role** search box, enter the built-in or custom role's name and select the role from the list of suggestions. The system displays role details for the selected role.
7. Click **Next**.
8. Leave **Define Scope** empty and click **Next**.
9. From the **Users** dropdown list, select the active directory the user belongs to.
10. On the **Search** box, enter the first few letters of the user or user group's name and select the correct user or user group from the list of suggestions.
11. Click **Save**
End user is created in IAM.
12. Edit the newly created policy, follow these steps:
 - a. From the **Entity Type** dropdown list, select the entity type.
The list of available entities depends on the role that you select.
 - b. From the **Filters** dropdown list, select the filter type.
 - c. On the **Search** box, select the specific entities that belongs to the selected entity type.
 - d. To allow the assigned users to access any new entity instances created by them, select the **Allow user access to entities created by them** checkbox.
13. Click **Save**.

What to do next

After creating an authorization policy, you can manage IAM keys. For more information, see [Managing IAM Keys using LDAP](#) on page 235

Creating a Role for Sub-Admin

You can create a custom role for a sub-admin managing keys for a set of users.

About this task

To create a custom role for sub-admin in Prism Central, follow these steps:

Procedure

1. Log in to Prism Central as an administrator.

2. From the **Application Switcher Function**, select the **Admin Center** application and from the navigation bar, select **IAM**.
3. Select the **Roles** tab.
4. Click **Create Role > New Role**.
5. In the **Role Name** field, enter a unique name for the role.
6. (Optional) In the **Description** field, enter a description
7. In the **Entity Type** filter, the individual operations available for a specific resource type, select the Operation from the drop-down list, and in the search box, enter the name of the resource.

Note: The following permissions are required:

- View User
- View User Key
- Create User Key
- Delete User Key

The following permissions are optional and not required for managing existing users but these permissions are required for creating new users:

- Create User
- View Directory Service
- Search Directory Service

8. Click the plus icon to add an individual operations to the role.
9. Do one of the following:
 - » To create the role, click **Save**.
 - » To create the role and attach the authorization policy, click **Save & Create Authorization Policy**.

What to do next

After providing the required permissions in the IAM role, you can create an authorization policy. For more information, see [Creating an Authorization Policy for a Sub-Admin](#) on page 234.

Creating an Authorization Policy for a Sub-Admin

Create an authorization policy with full access to all entity types and instances for the added users in the associated role.

Before you begin

Create a custom role for sub-admin managing keys for a set of users. For more information, see [Creating a Role for Sub-Admin](#) on page 233.

About this task

To create an authorization policy, follow these steps:

Procedure

1. Log in to Prism Central as an administrator.
2. From the **Application Switcher Function**, select the **Admin Center** application and from the navigation bar, select **IAM**.
3. Select the **Authorization Policies** tab.
4. Click **+ Create Authorization Policy**.
5. (Optional) To edit the name, click the pencil icon next to the default name, edit the name, and click the checkmark icon.
6. On the **Select Role** search box, enter the built-in or custom role's name and select the role from the list of suggestions. The system displays role details for the selected role.
7. Click **Next**.
8. Select **Configure access: select entity type & instances** and follow these steps:
 - a. From the **Entity Type** drop-down list, select the entity type.
Select the list of users that a sub-admin can manage.
 - b. From the **Filters** drop-down list, select the filter type.
 - c. On the **Search** box, select the specific entities that belongs to the selected entity type.
9. To allow the assigned users to access any new entity instances created by them, select the **Allow user access to entities created by them** checkbox.
10. Click **Next**.
11. From the **Users** drop-down list, select the sub-admin user.
The user group selected will be able to manage the entities selected in the **Define Scope** (step 8).
12. Click **Save**.

What to do next

After creating an authorization policy, you can manage IAM keys. For more information, see [Managing IAM Keys using LDAP](#) on page 235

Managing IAM Keys using LDAP

You can use the object store URL to launch the Nutanix Objects Browser UI and use the Active Directory or LDAP credentials to manage IAM keys.

Before you begin

Create a custom role for an end user to generate their own user key. For more information, see [Creating a Custom Role for End User](#) on page 230.

Create a custom role for sub-admin to manage keys for a set of users. For more information, see [Creating a Role for Sub-Admin](#) on page 233.

About this task

To launch the Nutanix Objects Browser, follow these steps:

Procedure

1. Open the Object Browser URL shared by the administrator in a web browser.
2. In the Nutanix Objects Browser login page, click **Login with LDAP/AD**

Note: If you open the Objects Browser URL in the same browser where you're already logged in to Prism Central with an AD/LDAP account, you might be automatically signed in to Objects Browser with the same credentials. To avoid this, open Objects Browser in a different browser or use an incognito/private browsing window.

The Prism Central Login page appears.

3. Log on to Prism Central using the created policy credentials.
The **Access Keys** table appears with the current policy.
4. Click **Add Key**, generate the access key for the user.
 - a. (Optional) Click **Add Key** at the top to add new users.
 - b. (Optional) Click **Add Key** next to the desired user to add a key for that specific user.
5. You can do one of the following actions :
 - » **Add Key:** Click this button to generate the access key for the user.
 - » **Delete:** Select the access key and click this button to delete the access key.

Note: A logged-in user can manage only the access keys created by that user.

NUTANIX OBJECTS PROMETHEUS EXPORTER

The Nutanix Objects Prometheus Exporter provides a solution for Nutanix Objects observability. The Nutanix Objects Prometheus Exporter service runs on Prism Central. Using an external Prometheus server, stats can be periodically retrieved from the exporter service and visualized through Grafana dashboards. You can also use any other observability tools that can import stats in Prometheus format and store and visualize them.

The exporter can be used to monitor the following stats:

- Object stores and buckets stats:
 - Capacity consumed or usage
 - Number of objects
 - Number of buckets
- Performance stats (individual bucket level as well as cumulative object store level): Total requests/sec, GET/sec, PUT/sec, throughput.
- Object store resource consumption stats: CPU and memory consumption for object store services - Object Controller, Metadata service, Atlas, and Load Balancer.

List of stats at bucket level

- nutanix_objectstore_total_object_data_gets
- nutanix_objectstore_object_count
- nutanix_objectstore_total_object_data_puts
- nutanix_objectstore_storage_usage_bytes
- nutanix_objectstore_total_deletes
- nutanix_objectstore_total_batch_deletes
- nutanix_objectstore_total_batch_delete_objects
- nutanix_objectstore_total_batch_delete_with_versions
- nutanix_objectstore_total_delete_with_versions
- nutanix_objectstore_total_heads
- nutanix_objectstore_total_lists
- nutanix_objectstore_total_list_versions
- nutanix_objectstore_total_multipart_lists
- nutanix_objectstore_total_multipart_uploads
- nutanix_objectstore_total_posts
- nutanix_objectstore_total_requests_sec
- nutanix_objectstore_total_select_content_ops
- nutanix_objectstore_total_throughput_in_sec

- nutanix_objectstore_total_throughput_out_sec

List of stats at object store level

- nutanix_objectstore_total_object_data_gets
- nutanix_objectstore_total_batch_deletes
- nutanix_objectstore_total_batch_delete_objects
- nutanix_objectstore_total_batch_delete_with_versions
- nutanix_objectstore_total_delete_with_versions
- nutanix_objectstore_num_buckets
- nutanix_objectstore_num_objects
- nutanix_objectstore_total_object_data_puts
- nutanix_objectstore_total_deletes
- nutanix_objectstore_total_heads
- nutanix_objectstore_total_lists
- nutanix_objectstore_total_list_versions
- nutanix_objectstore_total_multipart_lists
- nutanix_objectstore_total_multipart_uploads
- nutanix_objectstore_total_posts
- nutanix_objectstore_total_requests_sec
- nutanix_objectstore_total_select_content_ops
- nutanix_objectstore_total_throughput_in_sec
- nutanix_objectstore_total_throughput_out_sec
- nutanix_objectstore_usage_bytes

List of resource stats for services of an object store

CPU and memory usage for Atlas, Object Controller and Metadata Server.

For more information on the integration steps with external Prometheus and Grafana, contact your Nutanix account team for further guidance.

NUTANIX OBJECTS CRUD OPERATIONS BY USING S3 APIS

Create and manage buckets and objects by using S3 APIs in Nutanix Objects.

Nutanix Objects provides S3-compatible object storage, enabling users to create, manage, and interact with buckets and objects using standard Amazon S3 APIs. You can interact with the service via REST API or AWS SDKs. All API interactions with Nutanix Objects are authenticated to ensure secure access. The object store supports Amazon S3 features such as tagging (key-value labels), S3 Select for partial content retrieval, and server-side encryption with customer-provided keys (SSE-C).

Nutanix Objects Authentication

You can send requests to Nutanix Objects by using the REST API or the Amazon Web Services Software Development Kit (AWS SDK) wrapper libraries that wrap the underlying S3 REST API. Every interaction with Nutanix Objects is authenticated. In this authentication process, the identity of the requester who is trying to access Nutanix Objects is verified with a signature value. The signature value is generated from the AWS access keys (access key ID and secret access key) of the requester. This AWS access keys and endpoint URL is provided by the administrator to the user.

If you are using the AWS SDK, the libraries compute the signature from the keys you provide. However, if you make direct REST API calls, the signature is computed from the request.

For creating the buckets and objects, you need the following information from the administrator:

- Endpoint URL (Static IP address and the Port number)
- Access key ID
- Secret access key

Once you get this information, you can import the SDK libraries, create a session and a client, and then you can start making the requests (for example, creating buckets and objects).

Note: Nutanix Objects supports the following:

- Signature Version 2 and Signature Version 4 (regular and pre-signed)
- Streaming signed payloads for PUT requests with Signature Version 4

Nutanix Objects Supported and Unsupported APIs

This section describes Nutanix Objects support for Amazon S3 API. The object store service is available on the following ports:

- **HTTP:** 80
- **HTTPS:** 443

Note:

- Transport Layer Security (TLS) 1.2 is supported on Nutanix Objects.
- Support for all CBC based ciphers are deprecated.

If you are using RSA certificate, the following ciphers are allowed.

TLS/1.2

IANA : TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 (OpenSSL : ECDHE-RSA-AES256-GCM-SHA384)
IANA : TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256 (OpenSSL : ECDHE-RSA-AES128-GCM-SHA256)

If you are using EC certificate, the following ciphers are allowed.

TLS/1.2

IANA : TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384 (OpenSSL : ECDHE-ECDSA-AES256-GCM-SHA384)

IANA : TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256 (OpenSSL : ECDHE-ECDSA-AES128-GCM-SHA256)

- The default region for Nutanix Objects is **us-east-1**.

Nutanix Objects Supported APIs

The following table lists the supported S3 API methods.

Table 40: Supported S3 APIs and Parameters

Supported S3 APIs	Request Parameters/Request Headers/Request Elements/Request Body
PUT Bucket	Nutanix Objects Common Headers on page 243 + CreateBucketConfiguration, LocationConstraint (By default, LocationConstraint “us-east-1” is only supported.)
PUT Bucket Lifecycle	Nutanix Objects Common Headers on page 243 + Content-MD5 + AbortIncompleteMultipartUpload, And, Date, Days, DaysAfterInitiation, Expiration, Filter, ID, Key, LifecycleConfiguration, NoncurrentDays, NoncurrentVersionExpiration, NoncurrentVersionTransition, Prefix, Rule, Status
PUT Bucket Policy	Nutanix Objects Common Headers on page 243 + JSON string containing the policy contents containing the policy statements
PUT Bucket versioning	Status, VersioningConfiguration
Complete Multipart Upload	Nutanix Objects Common Headers on page 243 + CompleteMultipartUpload, Part, PartNumber, ETag
PUT Object - Copy	Nutanix Objects Common Headers on page 243 + x-amz-copy-source, x-amz-metadata-directive, x-amz-copy-source-if-match, x-amz-copy-source-if-none-match, x-amz-copy-source-if-unmodified-since, x-amz-copy-source-if-modified-since, x-amz-tagging, x-amz-tagging-directive, x-amz-meta-, x-amz-object-lock-legal-hold, x-amz-object-lock-mode, x-amz-object-lock-retain-until-date
DELETE Nutanix Objects	Single query string parameter delete + Common Headers + Content-MD5, Content-Length, Delete, Quiet, Object, Key, VersionId

Supported S3 APIs	Request Parameters/Request Headers/Request Elements/Request Body
GET Object	Nutanix Objects Common Headers on page 243 + Range, If-Modified-Since, If-Unmodified-Since, If-Match, If-None-Match, x-amz-object-lock-legal-hold, x-amz-object-lock-mode, x-amz-object-lock-retain-until-date
HEAD Object	Nutanix Objects Common Headers on page 243 + Range, If-Modified-Since, If-Unmodified-Since, If-Match, If-None-Match, x-amz-object-lock-legal-hold, x-amz-object-lock-mode, x-amz-object-lock-retain-until-date
List Multipart Uploads	Nutanix Objects Common Headers on page 243 + delimiter, max-uploads, key-marker, prefix, upload-id-marker
List Object Versions	Nutanix Objects Common Headers on page 243 + delimiter, key-marker, max-keys, prefix, version-id-marker
GET Bucket (List Nutanix Objects) Version 1	Nutanix Objects Common Headers on page 243 + delimiter, marker, max-keys, prefix
GET Bucket (List Nutanix Objects) Version 2	Nutanix Objects Common Headers on page 243 + delimiter, max-keys, prefix, list-type, continuation-token, fetch-owner, start-after
List Parts	Nutanix Objects Common Headers on page 243 + uploadId, max-parts, part-number-marker
PUT Object	Nutanix Objects Common Headers on page 243 + Content-Length, Content-MD5, Expect, x-amz-tagging, x-amz-meta-, x-amz-object-lock-mode, x-amz-object-lock-retain-until-date, x-amz-object-lock-legal-hold
INITIATE Multipart Upload	Nutanix Objects Common Headers on page 243 + x-amz-meta-, x-amz-tagging, x-amz-object-lock-legal-hold, x-amz-object-lock-mode, x-amz-object-lock-retain-until-date
UPLOAD Part	Content-Length, Content-MD5, Expect
Upload Part - Copy	x-amz-copy-source, x-amz-copy-source-range, x-amz-copy-source-if-match, x-amz-copy-source-if-none-match, x-amz-copy-source-if-unmodified-since, x-amz-copy-source-if-modified-since
PUT Bucket Object Lock Configuration	Nutanix Objects Common Headers on page 243 + ObjectLockEnabled Rule
DELETE Bucket	Nutanix Objects Common Headers on page 243 + Bucket
Delete Multiple Nutanix Objects	Nutanix Objects Common Headers on page 243 + Bucket, Delete, Quiet, Nutanix Objects, Key, VersionId

Supported S3 APIs	Request Parameters/Request Headers/Request Elements/Request Body
GET Object Retention	Nutanix Objects Common Headers on page 243 + Bucket, Key, VersionId
PUT Object Retention	Nutanix Objects Common Headers on page 243 + Bucket, Key, VersionId, Retention, Mode, RetainUntilDate
PUT Object Legal Hold	Nutanix Objects Common Headers on page 243 + Bucket, Key, versionId, LegalHold, Status
GET Object Legal Hold	Nutanix Objects Common Headers on page 243 + Bucket, Key, versionId
Select Object Content	Expression, ExpressionType, InputSerialization, CSV, AllowQuotedRecordDelimiter, Comments, FieldDelimiter, FileHeaderInfo, QuoteCharacter, QuoteEscapeCharacter, RecordDelimiter, OutputSerialization, CSV

Note:

- User-provided metadata header names are stored in lowercase.
- User-defined object metadata is supported for the PUT Object and PUT Object -Copy APIs and is limited to 2 KB in size.
- For bucket tagging, tags that begin with the prefix `ntnx:` are reserved for internal services and cannot be used by external clients.
- The maximum size limit of a multipart object is 16 TiB. This limit is different from the standard S3 API limit of 5 TiB.

For more information on objects tagging API, see [Buckets and Nutanix Objects Tagging APIs Overview](#) on page 244.

For more information on S3 Select API, see [Nutanix Objects S3 Select API Overview](#) on page 246.

You can generate Presigned URLs for Nutanix Objects and restrict access by specifying an expiration time. The expiration time can exceed seven days. For more information, see *Signing and authenticating REST requests* in the *Amazon Simple Storage Service Developer Guide*.

Table 41: Supported APIs That Use Only Common Headers

The following table lists the supported APIs that use only [Nutanix Objects Common Headers](#) on page 243:

Supported APIs
GET Bucket lifecycle configuration
PUT Bucket lifecycle configuration
GET Bucket Object Lock Configuration
GET Bucket Location
GET Bucket Policy

Supported APIs

GET Bucket versioning
GET Bucket ACL
GET Bucket notification configuration
GET Bucket cors
GET Bucket website
GET Bucket replication
PUT Bucket cors
PUT Bucket notification configuration
PUT Bucket website
PUT Bucket replication
DELETE Bucket cors
DELETE Bucket website
DELETE Bucket Lifecycle
DELETE Bucket Policy

Note: DELETE Bucket Policy is only supported for regular buckets and not federated buckets.

DELETE Bucket replication
DELETE Object
HEAD Bucket
LIST Bucket
ABORT Multipart Upload

Nutanix Objects Common Headers

You can use the following headers while making requests:

Table 42: List of Common Headers

Authorization
Content-Length
Content-Type
Content-MD5
Date
Expect
Host
x-amz-content-sha256
x-amz-date
x-amz-security-token

For more information on common headers, see *Common Request Headers* in the *Amazon Simple Storage Service API Reference Guide*.

Nutanix Objects Unsupported APIs

The following table lists the unsupported S3 API methods:

Table 43: Unsupported S3 APIs

Unsupported S3 APIs

GET Bucket accelerate configuration
GET Bucket analytics configuration
GET Bucket encryption
GET Bucket inventory configuration
GET Bucket logging
GET Bucket metrics configuration
GET Bucket requestPayment
List Bucket Analytics Configurations
List Bucket Inventory Configurations
List Bucket Metrics Configurations
PUT Bucket accelerate configuration
PUT Bucket acl
PUT Bucket analytics configuration
PUT Bucket encryption
PUT Bucket inventory configuration
PUT Bucket logging
PUT Bucket metrics configuration
DELETE Bucket analytics configuration
DELETE Bucket inventory configuration
DELETE Bucket metrics configuration
DELETE Bucket encryption
Copy
GET Object ACL
PUT Object ACL
Restore Object

Buckets and Nutanix Objects Tagging APIs Overview

A tag is a label that you assign to a bucket or an object, and it consists of a key and a value pair that you can define.

The tagging APIs feature allows you to add, retrieve, and remove tags for your buckets and objects.

Buckets Tagging

- Retrieving bucket metadata also returns the number of tags associated with the bucket, if any.
- The maximum number of tags allowed per bucket is 50.
- Tag keys can be up to 128 Unicode characters in length, and tag values can be up to 256 Unicode characters in length.
- Tag keys and values are case-sensitive.

Nutanix Objects Tagging

- Retrieving objects metadata also returns the number of tags associated with the object, if any.
- Tagging is also supported with a few other Object APIs.
- The maximum number of tags allowed per object is 10.
- Tag keys can be up to 128 Unicode characters in length, and tag values can be up to 256 Unicode characters in length.
- Tag keys and values are case-sensitive.

API Operations Supported for Buckets and Nutanix Objects Tagging

Nutanix Objects supports APIs for tagging both buckets and objects. Tag-based bucket lifecycle policy management and object listing are not supported.

Bucket Tagging API Operations

Nutanix Objects supports the following bucket tagging API operations.

Note: For bucket tagging, tags starting with prefix `ntnx:` are reserved for internal services and not allowed by external clients.

- **PutBucketTagging:** Sets the tags for a bucket.
The following two scenarios are involved:
 - You can add tags to a bucket with no tags associated with it.
 - You can replace the existing tags associated with the bucket.
- **GetBucketTagging:** Returns the tag set associated with the bucket.
- **DeleteBucketTagging:** Deletes the tags associated with a bucket.

For more information on bucket tagging, see *PutBucketTagging* in the *Amazon Simple Storage Service API Reference* guide.

Object Tagging API Operations

Nutanix Objects supports the following object tagging API operations.

- **PUT Object tagging:** Replace the tags associated with an object. You can add the tags in the request body.
The following two scenarios are involved:
 - You can add tags to an object with no tags associated with it.
 - You can replace the existing tags associated with an object.

- **GET Object tagging:** Retrieves the tags associated with an object.
- **DELETE Object tagging:** Deletes the tags associated with an object.

The following Object APIs also support tagging:

- GET Object (returns tag count, if any)
- PUT Object
- PUT Object-Copy
- Initiate Multipart Upload

Nutanix Objects S3 Select API Overview

S3 Select is a feature to select the partial content of an object and returns the result.

With the S3 Select feature, you can only fetch the content of interest. This allows applications to speed up their queries drastically as the size of the sub-set of the object is smaller compared to the actual size of the entire object. S3 Select is supported on objects stored in CSV format and returns the result in CSV format.

Note: S3 select does not support compression.

The `SelectObjectContent` API filters the contents of an object located in Nutanix Objects using an SQL statement. In the request, you must specify the SQL expression and data serialization format (CSV) of the object. Nutanix Objects use this format to parse object data into records, and it returns only records that match the specified SQL expression. You must also specify the data serialization format (CSV) for the response.

Basic Command Support

Nutanix Objects S3 Select feature supports only the `SELECT` SQL command.

The S3 Select query operates on a single object. The primary commands involve data selection and filtering.

The following clauses are supported for the `SELECT` command.

- **SELECT list:** The `SELECT` list names the columns, functions, and expressions that you want the query to return.

The following are supported:

- Aggregate operators such as `AVG`, `COUNT`, `MAX`, `MIN`
- Arithmetic operations (For example, `+`, `/`)
- **FROM clause:** Using `FROM` clause, you can select from arrays or objects within a JSON object. You can use `S3Object` as a basic reference.
- **WHERE clause:** The `WHERE` clause filters rows based on the *condition*. A condition is an expression that has a Boolean result.

The following conditions are supported:

- `AND`, `NOT`, `OR` operators
- Comparison (For example, `<`, `>=`)
- Pattern matching (For example, `LIKE`, `%`)

- **LIMIT clause:** The `LIMIT` clause limits the number of records that you want the query to return based on *the number*.

The following is the general syntax of a `SELECT` query.

```
SELECT fields FROM S3Object WHERE condition LIMIT num
```

Note: The specific parameters differ based on the format and schema of the object.

For more information on the supported SQL functions, see [Supported SQL Functions by Nutanix Objects S3 Select](#) on page 247.

Requirements

The following is the requirement for using S3 SELECT:

- You must have `s3:GetObject` permission for the object you are querying.

Limitations

The following are limitations for using S3 SELECT:

- The maximum length of a SQL expression is 256 KB.
- The maximum length of a record in the input or result is 1 MB.
- Complex operations like sub-queries or joins are not supported.

Supported SQL Functions by Nutanix Objects S3 Select

S3 Select supports the following SQL functions.

Table 44: Type of Supported Functions

Type of Functions	Functions with Parameters
Aggregate	avg(X): Returns the average value of all X within a group. String and BLOB values that are not numbers are interpreted as 0. count(X) or count(*): Returns a count of the number of times X is in a group. The <code>count(*)</code> function (with no arguments) returns the total number of rows in the group. max(X): Returns the maximum value of all values in the group. It is the value that is returned last in an <code>ORDER BY</code> on the same column. min(X): Returns the minimum value of all values in the group. This is the first value that appears in an <code>ORDER BY</code> of the column. sum(X): Returns the sum of all values in the group.

Type of Functions	Functions with Parameters
Conditional	CASE: A CASE expression is similar to IF-THEN-ELSE. The optional expression that occurs in between the CASE keyword and the first WHEN keyword is called the <i>base</i> expression. There are two forms of the CASE expression: • A CASE with a base expression: The base expression is evaluated just once, and the result is compared against the evaluation of each WHEN expression from left to right. The result of the CASE expression is the evaluation of the THEN expression that corresponds to the first WHEN expression for which the comparison is true. Or, if none of the WHEN expressions evaluate to a value equal to the base expression, the result of evaluating the ELSE expression, if any. If there is no ELSE expression and none of the WHEN expressions produce a result equal to the base expression, the overall result is NULL. • A CASE without a base expression: Each WHEN expression is evaluated, and the result is treated as a boolean, starting with the leftmost and continuing to the right. The result of the CASE expression is the evaluation of the THEN expression that corresponds to the first WHEN expression that evaluates to true. Or, if none of the WHEN expressions are evaluated to be true, the result of evaluating the ELSE expression, if any. If there is no ELSE expression and none of the WHEN expressions are true, then the overall result is NULL. A NULL result is considered untrue when evaluating. coalesce(X,Y,...): Returns a copy of the first argument, or NULL if all arguments are NULL. Coalesce() must have at least two arguments. nullif(X,Y): Returns its first argument if the arguments are different and NULL if the arguments are the same. This function searches its arguments from left to right for an argument that defines a collating function and uses that collating function for all string comparisons.

Type of Functions	Functions with Parameters
Conversions	CAST: A CAST expression of the form <i>CAST(expr AS type-name)</i> is used to convert the value of <i>expr</i> to a different storage class specified by <i>type-name</i>. The following are the Conversion Processing ways for the different <i>type-name</i>: • Casting a value with no affinity converts into a BLOB. Casting to a BLOB first casts the value to TEXT, then interprets the resulting byte sequence as a BLOB instead of TEXT. • Casting a BLOB value to TEXT, the sequence of bytes that make up the BLOB is interpreted as text encoded using the database encoding. • Casting a BLOB value to a REAL, the value is first converted to TEXT. • Casting a BLOB value to INTEGER, the value is first converted to TEXT. • Casting a TEXT or BLOB value into NUMERIC yields either an INTEGER or a REAL result.

Type of Functions	Functions with Parameters
Date	date(time-value, modifier, modifier, ...) time(time-value, modifier, modifier, ...) datetime(time-value, modifier, modifier, ...) All the date and time functions take a time value as an argument. The time value is followed by zero or more modifiers. The date() function returns the date in the YYYY-MM-DD format. The time() function returns the time in the HH:MM:SS format. The datetime() function returns in the YYYY-MM-DD HH:MM:SS format. strftime(format, time-value, modifier, modifier, ...) The strftime() function takes a format string as its first argument. It returns the formatted date according to the format string specified as the first argument. The following is a list of valid strftime() substitutions: • %d - day of month: 00 • %f - fractional seconds: SS.SSS • %H - hour: 00-24 • %j - day of year: 001-366 • %J - Julian day number • %m - month: 01-12 • %M - minute: 00-59 • %s - seconds since 1970-01-01 • %S - seconds: 00-59 • %w - day of week 0-6 with Sunday==0 • %W - week of year: 00-53 • %Y - year: 0000-9999 • %% - %

Type of Functions	Functions with Parameters
String	length(X): Returns the following based on the value of X: • For a string value, returns the number of characters (not bytes) in X prior to the first NULL character, • For a blob value, returns the number of bytes in the blob. • For a NULL value, returns NULL. • For a numeric value, returns the length of a string representation of X. lower(X): Returns a copy of string X with all ASCII characters converted to lower case. Note: The default lower() function works for ASCII characters only. substring(X,Y,Z): Returns a substring of input string X that begins with the Y-th character and which is Z characters long. Returns the following based on different scenarios: • If Z is omitted, returns all characters through the end of the string X beginning with the Y-th. • If Y is negative, then the first character of the substring is found by counting from the right rather than the left. • If Z is negative, then the Z characters preceding the Y-th character are returned. • If X is a string, then the indices refer to actual UTF-8 characters. • If X is a BLOB, then the indices refer to bytes. ltrim(X,Y): Returns a string formed by removing any and all characters that appear in Y from the left side of X. If the Y argument is omitted, ltrim(X) removes spaces from the left side of X. rtrim(X,Y): Returns a string formed by removing any and all characters that appear in Y from the right side of X. If the Y argument is omitted, rtrim(X) removes spaces from the right side of X. trim(X,Y): Returns a string formed by removing any and all characters that appear in Y from both ends of X. If the Y argument is omitted, trim(X) removes spaces from both ends of X. upper(X): Returns a copy of input string X in which all lower-case ASCII characters are converted to their upper-case equivalent.

Server-Side Encryption with Customer-Provided Keys Supported APIs

Server-side encryption with customer-provided keys (SSE-C) encrypts objects using Advanced Encryption Standard with Galois Counter Mode (AES-256-GCM), ensuring data integrity and enabling efficient range reads.

Each object is encrypted using AES-256-GCM, which generates an authentication tag to ensure data integrity. Decryption requires the same encryption key used during encryption. Nutanix Objects are divided into blocks, and block-level encryption enables efficient range reads without requiring decryption of the entire object.

The following table lists the S3 API methods that require additional parameters to support this feature.

Table 45: S3 API Methods Requiring Additional Parameters for Feature Support

Supported S3 APIs	Response Headers/Request Headers
Head Object	x-amz-server-side-encryption-customer-algorithm: AES256,x-amz-server-side-encryption-customer-key: <Base64-encoded key>, x-amz-server-side-encryption-customer-key-MD5: <Base64-encoded MD5 checksum of Key>
Get Object	x-amz-server-side-encryption-customer-algorithm: AES256,x-amz-server-side-encryption-customer-key: <Base64-encoded key>, x-amz-server-side-encryption-customer-key-MD5: <Base64-encoded MD5 checksum of Key>
Put Object	x-amz-server-side-encryption-customer-algorithm: AES256,x-amz-server-side-encryption-customer-key: <Base64-encoded key>, x-amz-server-side-encryption-customer-key-MD5: <Base64-encoded MD5 checksum of Key>
Copy Object	x-amz-server-side-encryption-customer-algorithm: AES256,x-amz-server-side-encryption-customer-key: <Base64-encoded key>, x-amz-server-side-encryption-customer-key-MD5: <Base64-encoded MD5 checksum of Key>
Create Multipart Upload	x-amz-server-side-encryption-customer-algorithm: AES256,x-amz-server-side-encryption-customer-key: <Base64-encoded key>, x-amz-server-side-encryption-customer-key-MD5: <Base64-encoded MD5 checksum of Key>
Upload Part	x-amz-server-side-encryption-customer-algorithm: AES256,x-amz-server-side-encryption-customer-key: <Base64-encoded key>, x-amz-server-side-encryption-customer-key-MD5: <Base64-encoded MD5 checksum of Key>
Upload Part Copy	x-amz-server-side-encryption-customer-algorithm: AES256,x-amz-server-side-encryption-customer-key: <Base64-encoded key>, x-amz-server-side-encryption-customer-key-MD5: <Base64-encoded MD5 checksum of Key>

Note: SSE-C is incompatible with buckets with Network File System access, buckets in the Federated namespace, and buckets with replication configured. An object which is stored with SSE-C, cannot be accessed through S3 Select API.

The following table lists the compatibility of SSE-C with other Nutanix Objects features and the expected behavior when used together.

Table 46: SSE-C Compatibility with other Nutanix Objects Features

Feature	Compatibility Details
Network File System	Network File System can only be enabled at bucket creation. If Network File System access is enabled, an Object Put or Multipart Upload with SSE-C keys fails.
Federation	An Object PUT or MULTIPART UPLOAD on a bucket in the federated namespace with SSE-C keys fails.
Replication	If replication is configured on a bucket, an Object PUT or MULTIPART UPLOAD with SSE-C keys fails. Additionally, if a bucket contains encrypted objects, configuring replication rules fails.
S3 Select	Running S3 Select on an encrypted object results in failure.

The following table lists alerts related to SSE-C, along with their descriptions.

Table 47: Alerts

Alerts	Description
DecryptionAuthFailure	Decryption failed due to an AuthTag mismatch
EncryptionKeyMismatch	An encryption key mismatch detected.
EncryptionParamsNeeded	Missing encryption keys for request.
InsecureConnectionWithEncryptionKey	Encryption parameters sent over an insecure connection.

NUTANIX OBJECTS REST ERROR RESPONSES

This section provides reference information about Nutanix Objects error responses and codes.

When Nutanix Objects request returns an error, the client receives an error response. The format of the error response is API specific; however, all the error responses have common elements.

When an error occurs, the header information includes the following:

- Content-Type: application/xml
- An appropriate status code

The following elements are included in REST error responses:

- Code. For more information, see [List of Error Codes for Nutanix Objects](#) on page 254.
- Error
- Message
- RequestId
- Resource

For more information on the error responses, see *REST Error Responses* in the *Amazon Simple Storage Service API Reference Guide*.

List of Error Codes for Nutanix Objects

The following table lists the error codes:

Table 48: Error Codes for Nutanix Objects

Error Code	HTTP Status Code
AccessDenied	403 Forbidden
AuthorizationHeaderMalformed	400 Bad Request
BadDigest	400 Bad Request
BucketAlreadyExists	409 Conflict
BucketAlreadyOwnedByYou	409 Conflict
BucketNotEmpty	409 Conflict
EntityTooSmall	400 Bad Request
EntityTooLarge	400 Bad Request
IncompleteBody	400 Bad Request
InlineDataTooLarge	400 Bad Request
InternalServerError	500 Internal Server Error
InvalidAccessKeyId	403 Forbidden
InvalidArgument	400 Bad Request

Error Code	HTTP Status Code
InvalidBucketName	400 Bad Request
InvalidBucketState	409 Conflict
InvalidDigest	400 Bad Request
InvalidLocationConstraint	400 Bad Request
InvalidObjectState	403 Forbidden
InvalidPart	400 Bad Request
InvalidPartOrder	400 Bad Request
InvalidPolicyDocument	400 Bad Request
InvalidRange	416 Requested Range Not Satisfied
InvalidRequest	400 Bad Request
InvalidURI	400 Bad Request
KeyTooLongError	400 Bad Request
MalformedACLError	400 Bad Request
MalformedPOSTRequest	400 Bad Request
MalformedXML	400 Bad Request
MaxMessageLengthExceeded	400 Bad Request
MaxPostPreDataLengthExceededError	400 Bad Request
MetadataTooLarge	400 Bad Request
MethodNotAllowed	405 Method Not Allowed
MissingContentLength	411 Length Required
MissingRequestBodyError	400 Bad Request
NoSuchBucket	404 Not Found
NoSuchBucketPolicy	404 Not Found
NoSuchKey	404 Not Found
NoSuchLifecycleConfiguration	404 Not Found
NoSuchUpload	404 Not Found
InvalidVersion	404 Not Found
NotImplemented	501 Not Implemented
OperationAborted	409 Conflict
PermanentRedirect	301 Moved Permanent
PreconditionFailed	412 Precondition Failed
Redirect	307 Moved Temporarily
RequestIsNotMultiPartContent	400 Bad Request
RequestTimeout	400 Bad Request
RequestTimeTooSkewed	403 Forbidden
SignatureDoesNotMatch	403 Forbidden

Error Code	HTTP Status Code
ServiceUnavailable	503 Service Unavailable
SlowDown	503 Slow Down
TemporaryRedirect	307 Moved Temporarily
UnexpectedContent	400 Bad Request

For more information on the error codes, see *List of Error Codes* in the *Amazon Simple Storage Service API Reference Guide*.

NUTANIX OBJECTS INTEGRATION WITH BACKUP APPLICATIONS

Nutanix Objects is ideal for cost-effective, scale-out storage. It provides a fully distributed, API-accessible storage platform that integrates directly into applications or is used for backup, archiving, and data retention. Nutanix Objects offers you a seamless way to switch from your traditional backup to object store backup. You can also perform multipart uploads. Nutanix Objects supports integration with backup applications such as Commvault, HYCU, Veeam and Veritas.

Note: Some backup vendors have configurable size limits if larger than 5 TiB VM images need to be backed up (For example, HYCU). The backup appliance configuration needs to be changed in order to take advantage of the larger size limit in Nutanix Objects.

For more information on Commvault Integration, see *Commvault with Nutanix* guide on the Nutanix Support Portal.

TROUBLESHOOTING NUTANIX OBJECTS

This section explains how to troubleshoot issues that you might encounter while using Nutanix Objects.

Handling Deployment Failure

In case you encounter any deployment failure, contact Nutanix Support at <https://portal.nutanix.com/>.

Shutting Down Nutanix Objects VMs

This section describes the steps to perform a graceful shutdown of Nutanix Objects VMs on AHV or ESXi.

Before you begin

Make sure all read and write operations are stopped from the application before shutting down the VM.


About this task

Warning: Ensure that you only perform shutdown operations as described in this procedure. Do not perform any destructive actions such as deleting a VM.

To shut down Nutanix Objects VMs, follow these steps:

Procedure

1. Log on to the Prism Central web console, and select **Objects** in the Application Switcher.
2. Copy the Object Stores cluster name.



☐	Name	Version	Domain	Nodes	Usage (Logical)	Buckets
☐	OSS-161836167	3.11		1	797.09 KiB / 310 TiB	1
☐	OSS-161	3.11		1	797.09 KiB / 200 TiB	1
☐	OSS-16	3.11		2	797.09 KiB / 670 TiB	1

Figure 65: Object Store Cluster

3. SSH into any CVM on the Prism Element cluster where the object store is deployed.

4. To view the list of VMs, do any one of the following:

- » **AHV:** To view the list of VMs, run the following command.

```
nutanix@NTNX-18FM6F370127-A-CVM:xx.xx.xxx.xx:~$ acli vm.list |grep '<objectstore-name>' -i
```

For example, `nutanix@NTNX-18FM6F370127-A-CVM:xx.xx.xxx.xx:~$ acli vm.list |grep 'OSS-1611836167' -i`, where `OSS-1611836167` is the object store name.

```
nutanix@NTNX-18FM6F370127-A-CVM:XX.XX.XXX.XX:~$ acli vm.list |grep 'OSS-1611836167' -i
```

```
oss-1611836167-898db8-default-0 6596b630-cacd-48bd-90b5-b4f1698e260a
```

```
oss-1611836167-898db8-ijpsvgbict-envoy-0 c25b5fd8-2787-46b6-b3b4-c982f6ddf1b9
```

- » **ESXi:** To view the list of VMs in vCenter, click **Hosts and Clusters** tab, expand the ESXi cluster.

The list of VMs appear.

To find the primary and secondary MSP clusters, run the command `mspctl cluster list`. If the `cluster_type` value is `primary_msp`, then the cluster is a primary MSP cluster. Rest clusters are secondary.

Note: It is recommended to shut down the envoy VMs first, and then proceed to shut down the worker VMs. Also, ensure that the primary cluster is shut down at last as it hosts all the common services such as IAM.

5. To shut down the VMs, do any one of the following:

- » **AHV:** Run the following command:

```
acli vm.shutdown <vm name>
```

- » **ESXi:** Right click the VM, and then click **Power > Power Off**.

Note: First shut down the envoy VMs.

```
nutanix@NTNX-18FM6F370127-A-CVM:~$ acli vm.list |grep 'OSS-1611836167' -i
oss-1611836167-898db8-default-0 6596b630-cacd-48bd-90b5-b4f1698e260a
oss-1611836167-898db8-ijpsvgbict-envoy-0 c25b5fd8-2787-46b6-b3b4-c982f6ddf1b9
nutanix@NTNX-18FM6F370127-A-CVM:~$ acli vm.shutdown oss-1611836167-898db8-ijpsvgbict-envoy-0
oss-1611836167-898db8-ijpsvgbict-envoy-0: pending
oss-1611836167-898db8-ijpsvgbict-envoy-0: complete
nutanix@NTNX-18FM6F370127-A-CVM:~$
```

Figure 66: Shutting Down the VMs: AHV

For AHV, you can check the status of the VM from the Prism Central web console. Click the **Hamburger icon > Virtual Infrastructure > VMs** and confirm that the status of the VM is shown as Off in the **Power State** column.

For ESXi, you can check the status of the VM from the vCenter. Click the **Hosts and Clusters** tab, and expand the cluster where the VM is listed, and confirm that the status of the VM is shown as Off.

6. Perform Step 5 for all the VMs within the cluster.

What to do next

You can also start the Nutanix Objects VMs. See [Starting the Nutanix Objects VMs](#) on page 260.

Starting the Nutanix Objects VMs

This section describes the steps to start the Nutanix Objects VM after performing a graceful shutdown of Nutanix Objects VMs on AHV or ESXi.


About this task

Warning: Ensure that you only perform power on operations as described in this procedure. Do not perform any destructive actions such as deleting a VM.

To start the Nutanix Objects VMs, follow these steps:

Procedure

1. Log on to the Prism Central web console, and select **Objects** in the Application Switcher.
2. Copy the Object Stores cluster name.



☐	Name	Version	Domain	Nodes	Usage (Logical)	Buckets
☐	OSS-1611836167	3.1.1		1	797.09 KiB / 310 TiB	1
☐	OSS-161	3.1.1		1	797.09 KiB / 200 TiB	1
☐	OSS-16	3.1.1		2	797.09 KiB / 670 TiB	1

Figure 67: Object Store Cluster

3. SSH into any CVM on the Prism Element cluster where the object store is deployed.
4. To view the list of VMs, do any one of the following:

» **AHV:** To view the list of VMs, run the following command.

```
nutanix@NTNX-18FM6F370127-A-CVM:xx.xx.xxx.xx:~$ acli vm.list |grep '<objectstore-name>' -i
```

For example, `nutanix@NTNX-18FM6F370127-A-CVM:xx.xx.xxx.xx:~$ acli vm.list |grep 'OSS-1611836167' -i`, where `OSS-1611836167` is the Object Store name.

```
nutanix@NTNX-18FM6F370127-A-CVM:XX.XX.XXX.XX:~$ acli vm.list |grep 'OSS-1611836167' -i
```

```
oss-1611836167-898db8-default-0 6596b630-cacd-48bd-90b5-b4f1698e260a
```

```
oss-1611836167-898db8-ijpsvgbict-envoy-0 c25b5fd8-2787-46b6-b3b4-c982f6dddf1b9
```

» **ESXi:** To view the list of VMs in vCenter, click **Hosts and Clusters** tab, expand the ESXi cluster.

The list of VMs appear.

To find the primary and secondary MSP clusters, run the command `mspctl cluster list`. If the `cluster_type` value is `primary_msp`, then the cluster is a primary MSP cluster. Rest clusters are secondary.

5. To start the VMs, do any one of the following:

- » **AHV:** Run the following command:

```
accli vm.on <vm name>
```

- » **ESXi:** Right click the VM, and then click **Power > Power On**.

Note: First start the worker VMs.

6. Check the following to ensure that the VMs are powered on.

- For AHV, in the Prism Central web console, go to the **VMs** page. Select the VM that you powered on using the accli and perform the **Launch Console** action. You will view the login prompt if the VM is powered on.

For ESXi, you can check the status of the VM from the vCenter. Click **Hosts and Clusters** tab, and expand the cluster where the VM is listed, and confirm that the status of the VM is shown as On.
- In the Prism Central web console, go to **Entity Menu > Services > Objects**. Check for the following points for your Nutanix Objects cluster:
 - Statistics are visible in the Buckets and Nutanix Objects columns.
 - Click on the corresponding Nutanix Objects cluster to check if it is reachable. After the Object Store page opens, check if the various statistics are visible. For example, Performance, Usage Summary, and so on.
 - In the Alerts page, ensure that there are no active alerts. If there are active alerts, wait for a few minutes. If the alerts persist, contact Nutanix Support.
 - Perform read and write operations to ensure that the object cluster is running.

Detection of Slow Connections by Nutanix Objects

Nutanix Objects can detect slow-performing client connections to Nutanix Objects.

Any client connection must be able to read or write data at the rate of 2 MiB in every 600 seconds window (translates to approximately 4 KiB/s throughput). A client connection that is not able to transfer 2 MiB data in 600 seconds window will be treated as a slow connection.

Any slow connection will be terminated immediately to avoid Denial of Service (DoS) attacks and better manage the resources on the server.

Note: Up to 1000 active connections are accepted by each Nutanix Objects endpoint. The slow-performing client connections can potentially consume all slots and cause Denial of Service (DoS).

Prism Element Clusters or Subnets Not Listed in the User Interface

If clusters or subnets are missing in the UI, it may be due to synchronization issues, lack of user access, or ongoing upgrades in Prism Central.

If the **Cluster** or **Subnets** dropdown list appears empty or shows an error message, it indicates that no Prism Element clusters or subnets are currently available for deployment. This issue might also occur when adding a multicluster to the object store.

The following are the possible causes for the Prism Element clusters or subnets to not appear in the user interface:

- Prism Central has no other registered Prism Element clusters.

- A newly registered Prism Element has not completed synchronization with Prism Central.
- The logged-in user lacks access to any Prism Element cluster.
- The logged in user does not have access to any subnets on the allowed Prism Element cluster.
- An admin recently created an access policy granting user access to a Prism Element cluster.
- An admin recently created an access policy granting user access to a subnet.
- The Prism Element cluster is undergoing an upgrade or the upgrade status is not marked as Succeeded.

You can verify the cluster status in the Prism Central UI. For more information, see [Clusters Summary View](#) in the *Prism Central Infrastructure Guide*.

Note: If you just register to a Prism Element or receive access, wait five minutes and try again.

Object Stores Not Listed While Creating a Replication Rule

You might see that while creating a replication rule for a bucket, the **Object Store** drop-down under the **Destination** section shows an empty list.

The following are the possible causes for the object stores to not appear in the user interface:

- The logged in user does not have permission to view the object store.
- The deployment of the object store is not in a **Complete** state.
- The object store version is lower than Nutanix Objects 3.0 version.

Manually Updating the DNS Server for Nutanix Objects Cluster

When the DNS servers are updated in the managed network that is used by Nutanix Objects deployments, they are not automatically updated in the MSP VMs for these Nutanix Objects deployments. For updating them, you have to either restart the network service or restart the VMs. So, to overcome this issue, you can manually update the DNS servers for each of the Nutanix Objects clusters using the managed network.

About this task

To manually update the DNS servers for Nutanix Objects cluster, follow these steps:

Procedure

1. From the Prism Element CLI, update the name servers in the AHV managed networks.

```
nutanix@cvm$ acli net.update_dhcp_dns network_name servers=new server list
```

For example,

```
nutanix@cvm$ acli net.update_dhcp_dns OssNet servers=10.40.64.16
```

2. Update `etc/resolv.conf` manually on the Nutanix Objects cluster with the new name servers.

```
nutanix@cvm$ mspctl cluster ssh cluster --all --cmd "sudo chattr -i /etc/resolv.conf"
```


For example,

```
nutanix@cvm$ mspctl cluster ssh dns-test-store --all --cmd "sudo chatter -i /etc/resolv.conf"
```

- a. Remove each old name server by running the following command.

Caution: Do not remove 127.0.0.1 entry.

```
nutanix@cvm$ mspctl cluster ssh cluster --all --cmd "sudo sed -i '/nameserver old_ip_address/d' /etc/resolv.conf"
```

For example,

```
nutanix@cvm$ mspctl cluster ssh dns-test-store --all --cmd "sudo sed -i '/search /i nameserver 10.40.64.16' /etc/resolv.conf"
```

- b. For each name server, add entry in all nodes of the Nutanix Objects clusters.

```
nutanix@cvm$ mspctl cluster ssh cluster --all --cmd "sudo sed -i '/search /i nameserver new_ip_address' /etc/resolv.conf"
```

For example,

```
nutanix@cvm$ mspctl cluster ssh dns-test-store --all --cmd "sudo sed -i '/search /i nameserver 10.40.64.16' /etc/resolv.conf"
```

```
nutanix@cvm$ mspctl cluster ssh cluster --all --cmd "cat /etc/resolv.conf"
```

For example,

```
nutanix@cvm$ mspctl cluster ssh dns-test-store --all --cmd "cat /etc/resolv.conf"
```

3. Identify all the DNS pods.

```
nutanix@cvm$ mspctl cluster ssh cluster_name  
sudo kubectl get pods -A | grep -i dns
```

For example,

```
nutanix@cvm$ mspctl cluster ssh dns-test-store  
===== 10.44.90.142 =====  
[nutanix@dns-test-store-16dcfd-default-0 ~]$ sudo kubectl get pods -A | grep -i dns
```

4. Restart all the DNS pods. For each pod, run the following command:

```
nutanix@cvm$ mspctl cluster ssh <cluster_name>  
kubectl delete pod -n namespace pod
```

For example, for pod *coredns-d9bfcd7f-5xtc7*, run the following command:

```
nutanix@cvm$ mspctl cluster ssh dns-test-store  
===== 10.44.90.142 =====  
[nutanix@dns-test-store-16dcfd-default-0 ~]$ kubectl delete pod -n kube-system coredns-d9bfcd7f-5xtc7
```

5. Restart the *dnsmasq* *systemd* service.

```
nutanix@cvm$ mspctl cluster ssh cluster --all --cmd "sudo systemctl restart dnsmasq"
```

For example,

```
nutanix@cvm$ mspctl cluster ssh dns-test-store --all --cmd "sudo systemctl restart dnsmasq"
```

COPYRIGHT

Copyright 2025 Nutanix, Inc.

Nutanix, Inc.

1740 Technology Drive, Suite 150

San Jose, CA 95110

All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. Nutanix and the Nutanix logo are registered trademarks of Nutanix, Inc. in the United States and/or other jurisdictions. All other brand and product names mentioned herein are for identification purposes only and may be trademarks of their respective holders.

SOLUTIONS DOCUMENT

Nutanix AHV Virtualization

Legal

© 2025 Nutanix, Inc. All rights reserved. Nutanix, the Nutanix logo and all Nutanix product and service names mentioned are registered trademarks or trademarks of Nutanix, Inc. in the United States and other countries. All other brand names mentioned are for identification purposes only and may be the trademarks of their respective holder(s).

Certain information contained in this content may link or refer to, or be based on, studies, publications, surveys, and other data obtained from third-party sources and our own internal estimates and research. While we believe these third-party studies, publications, surveys, and other data are reliable as of the date of publication, they have not independently verified unless specifically stated, and we make no representation as to the adequacy, fairness, accuracy, or completeness of any information obtained from a third-party. Our decision to publish, link to or reference third-party data should not be considered an endorsement of any such content.

Nutanix, Inc.
1740 Technology Drive, Suite 150
San Jose, CA 95110

Contents

- 1. Executive Summary..... 4**
- 2. Nutanix AHV Components..... 7**
- 3. AHV Integrated Management Capabilities..... 10**
 - AHV Virtual Networking Management..... 11
 - Nutanix Flow Virtual Networking and Flow Network Security.....11
 - Tools and Processes for Managing Upgrades, Maintenance, and Scaling..... 13
 - AHV Virtual Machine Management.....14
- 4. AHV Performance and Scale..... 21**
 - AHV Memory Improvements for Performance.....23
- 5. AHV Graphics Processing Unit Support..... 26**
 - vGPU..... 26
- 6. AHV Migration with Nutanix Move and Cluster Conversion..... 29**
- 7. AHV Security..... 30**
- About Nutanix.....32**
- List of Figures.....33**

1. Executive Summary

Nutanix natively converges compute and storage into a single appliance that you can deploy in minutes to run any application out of the box. The Nutanix solution offers powerful virtualization capabilities—including core virtual machine (VM) operations, live migration, VM high availability, and virtual network management—as fully integrated features of the infrastructure stack rather than standalone products that require separate deployment and management.

The native Nutanix hypervisor, [AHV](#), represents a fresh approach to virtualization that brings substantial benefits to enterprise IT administrators by simplifying every step of the infrastructure life cycle, from buying and deploying to managing, scaling, and supporting.

For more information, see these documents:

- [AHV Administration Guide](#)
- [Nutanix AHV Best Practices](#)
- [Nutanix AHV Networking Best Practices](#)
- [Data Protection and Disaster Recovery Best Practices](#)

Table: Document Version History

Version Number	Published	Notes
1.0	January 2016	Original publication.
1.1	July 2016	Updated platform information.
1.2	December 2016	Updated for AOS 5.0.
1.3	May 2017	Updated for AOS 5.1.
2.0	December 2017	Updated for AOS 5.5.
2.1	March 2019	Updated for AOS 5.10 and updated the Nutanix overview.
2.2	January 2020	Updated the Nutanix overview and the AHV section.

Version Number	Published	Notes
3.0	June 2020	Updated terminology throughout.
4.0	January 2022	Updated for AOS 6.1 in the Live Migration, Backup APIs, and vGPU sections, and added the VM Templates, OVA Import and Export, Metro Availability, Memory Overcommit, vGPU Live Migration, Credential Guard, Flow Network Visualization, and Flow Microsegmentation sections.
4.1	September 2022	Updated for AOS 6.5.1 by adding the Virtual Trusted Platform Module section.
4.2	October 2023	Updated for AOS 6.7 in the Live Migration, Generation ID, Compute-Only Nodes, iSCSI Extensions for RDMA (iSER), and Migration sections.
4.3	July 2024	Updated the Graphics Processing Unit Support section.
4.4	September 2024	Updated the live migration guidance in the Graphics Processing Unit Support section.
4.5	October 2024	Added the Advanced Processor Compatibility and Automatic Cluster Selection sections.
4.6	February 2025	Updated the Security Baseline and Self-Healing and References sections.

Version Number	Published	Notes
4.7	March 2025	Updated and added references to related product documentation.
4.8	June 2025	Updated the Metro Availability section.
4.9	August 2025	Updated document structure.

2. Nutanix AHV Components

Nutanix built AHV on a proven open-source Kernel-based Virtual Machine (KVM) foundation and extended KVM's base functionality to include features such as high availability and live migration. AHV comes preinstalled on Nutanix appliances and you can configure it in minutes to deploy applications.

AHV has three main components:

KVM-kmod

KVM kernel module

Libvirt

An API, daemon, and management tool for managing KVM and the QEMU emulator. Communication between AHV, KVM, and QEMU occurs through libvirt

Qemu-KVM

A machine emulator and virtualization tool that runs in user space for every VM (domain). AHV uses it for hardware-assisted virtualization




Figure 1: AHV Components

Nutanix uses hypervisor-agnostic distributed storage to deliver data services such as storage provisioning, snapshots, clones, and data protection to VMs directly, rather than using the hypervisor's storage stack. On each AHV host, an iSCSI redirector service establishes a highly resilient storage path from each VM to storage across the Nutanix cluster. QEMU is configured with the iSCSI redirector as the iSCSI target portal. When the redirector receives a sign-in request, it performs an iSCSI sign-in redirect to a healthy Controller VM (CVM) (preferably the local one). For more information on AHV storage, see [Storage Overview](#) in the AHV Administration Guide.




Figure 2: Storage Configuration

3. AHV Integrated Management Capabilities

You can manage the Nutanix platform from a single pane of glass with Nutanix Prism. Prism provides integrated capabilities for cluster management and VM management that are available from the Prism web console, command line interface (CLI), PowerShell, and REST API. Managing clusters on AHV involves creating, updating, deleting, and monitoring cluster resources. These resources include hosts, storage, and networks.

Prism provides a central location for administrators to update host settings like virtual networking and high availability across all nodes in an AHV cluster. Controlling configuration at the cluster level eliminates the need for manual compliance checks and reduces the risk of having a cluster that isn't uniformly configured. For more information, see the [Settings Profile for Clusters](#) section of the Prism Central Infrastructure Guide.

Nutanix Prism offers in-depth analytics for every element in the infrastructure stack, including hardware, storage, and VMs. Administrators can use Prism views to monitor these infrastructure stack components, and they can use the Analysis view to see an integrated assessment of cluster resources or specific metrics for each element.

Prism makes detailed VM data available, grouping it into the following categories:

- VM Performance: Multiple charts with CPU- and storage-based reports around resource usage and performance
- Virtual Disks: In-depth data points that focus on I/O types, I/O metrics, read source, cache hits, working set size, and latency on a per-virtual disk level
- VM NICs: vNIC configuration summary for a VM
- VM Snapshots: A list of all snapshots for a VM with the ability to clone or restore from the snapshot or to delete the snapshot
- VM Tasks: A time-based list of all operational actions performed against the selected VM, including task summary, percent complete, start time, duration, and status
- Console: A pop-up console session or an inline console session for a VM

The Storage tab provides a direct view of the storage fabric running on an AHV cluster. Administrators can look at detailed storage configurations, capacity usage over time, space efficiency, and performance, as well as a list of alerts and events related to storage.

The Hardware tab gives you a direct view of the Nutanix blocks and nodes that make up a cluster. These reports are available in diagram and table format for easy consumption.

The Prism Analysis tab allows administrators to explore and understand what is going on in their clusters quickly and to identify steps for remediation as required. You can create custom interactive charts using hundreds of metrics available for elements such as hosts, disks, storage pools, storage containers, VMs, protection domains, remote sites, replication links, clusters, and virtual disks, then correlate trends in the charts with alerts and events in the system. You can also choose specific metrics and elements and set a desired time frame when building reports so that you can focus precisely on the data you're looking for.

AHV Virtual Networking Management

AHV uses Open vSwitch (OVS) for all VM networking. When you create a new AHV cluster, the system configures the CVM and management networking paths automatically. Administrators can easily create new VLAN-backed layer 2 networks through Prism. After the administrator creates a network, they can assign it to existing and new VMs. For more information on the AHV implementation of OVS, see the [AHV Administration Guide](#). For more information on AHV networking, see the [AHV Networking best practice guide](#).

Along with streamlining VM network creation, AHV can manage DHCP addresses for each network that administrators create. This functionality allows administrators to configure address pools for each network that they can automatically assign to VMs.

Nutanix Flow Virtual Networking and Flow Network Security

Nutanix Flow Virtual Networking is a software-defined network virtualization solution that provides overlay capabilities for on-premises AHV clusters. It uses networking features like Virtual Private Cloud (VPC) and Virtual Private Network (VPN) to support flexible,

app-driven networking that focuses on VMs and applications instead of virtual LANs and network addresses. Flow Virtual Networking provides the following benefits:

- A simplified, Prism Central–based workflow that deploys the application-driven network virtualization features
- A secure multitenancy solution that enables per-tenant isolation using VPC-based network segmentation and namespace isolation
- A secure, VPN-based connectivity solution for multiple sites, with automated VPN bundle upgrades
- NAT-based secure egress to external networks, with IP address retention and policy-based routing
- Self-service networking using REST APIs or Nutanix Cloud Manager Self-Service blueprints and projects
- Stretched layer 2 networking for more effective disaster recovery

Nutanix Flow Network Security protects against new threats like ransomware designed to spread laterally from one system to another in the same protected datacenter. Because perimeter-based firewalls traditionally only protect the environment from external threats, it can be difficult to repurpose them to protect internal traffic. The problem is compounded by virtualized workloads changing their network configurations and hosts as they start, stop, and migrate frequently. Flow Network Security applies security rules between all applications and VMs in the datacenter, adding internal protection behind your perimeter firewall.

Flow Network Security includes a policy-driven security framework that inspects traffic in the datacenter. Security policies are applied to categories (a logical grouping of VMs) and not to the VMs themselves. Therefore, it doesn't matter how many VMs you have in a category; traffic associated with that category is secured without administrative intervention even at scale.

The framework uses a workload-centric approach instead of a network-centric approach, so it can scrutinize traffic to and from VMs in a datacenter regardless of changes to their network configurations or their locations in the datacenter. This approach also enables the virtualization team to implement these security policies without having to rely on network security teams.

Prism Central works with Flow Network Security to offer a visualization-based approach to configuring security policies and monitoring the traffic a policy applies to.

Flow Network Security is complemented by the SaaS-based Nutanix Security Central to provide threat detection, compliance, and security planning using machine learning analysis of on-prem configuration and network traffic metadata.

Tools and Processes for Managing Upgrades, Maintenance, and Scaling

AHV supports and offers several tools and processes that make it easier to upgrade, maintain, and scale AHV clusters, including the Life Cycle Manager (LCM) one-click upgrade process, Advanced Processor Compatibility (APC), maintenance mode, and scale-out architecture.

Nutanix delivers a one-click upgrade process through LCM for all software in the Nutanix platform, including AOS, AHV, firmware, and Nutanix Cluster Check (NCC). Upgrades are nondisruptive and allow the cluster to run continuously during software maintenance while nodes upgrade on a rolling basis in the background. LCM manages the dependencies between products and versions and only shows the supported upgrade options for your environment. Nutanix qualifies firmware updates from the manufacturers of the hard or solid-state disk drives in the cluster and makes them available through the same upgrade process. For more information, see the [AHV Upgrades](#) section of Nutanix Upgrades: Life Cycle Manager.

Advanced Processor Compatibility (APC) can streamline upgrades and transition clusters to the latest generation of CPUs. With APC, clusters aren't limited by the slowest CPU. VMs can use APC to accurately identify and take advantage of the CPU features available to them, enabling improved migrations. APC offers flexibility in multiple ways. Clusters can be leveled for consistent CPU presentation to guest VMs, which is the default behavior. You can also set a baseline CPU model on a per-VM basis for on-demand, cross-cluster live migrations to different hardware generations. For more information, see the [Advanced Processor Compatibility in AHV](#) section of the AHV Administration Guide.

Administrators can place AHV hosts in maintenance mode during upgrades and maintenance-related operations. Maintenance mode live-migrates all VMs running on the node to other nodes in the AHV cluster, and the CVM can safely shut down if required. After the maintenance process finishes all the steps for the node, it returns the CVM to

service and synchronizes with other CVMs in the cluster. Maintenance mode enables graceful host suspension for routine cluster maintenance. For more information, see the [Node Maintenance Mode](#) section of the AHV Administration Guide.

The Nutanix solution's scale-out architecture enables incremental, predictable capacity and performance scaling in a Nutanix cluster running any hypervisor, including AHV. Administrators can start with as few as three nodes and scale out by easily adding nodes until they reach the targeted cluster size or cluster maximum, adding more clusters as needed. The system automatically discovers new nodes and makes them available for use. To expand a cluster, select the discovered nodes you want to add and provide network configuration details. Use the same process to deploy new clusters in the same location or at an edge site. With Prism, administrators can image or update new nodes to match the AHV version of their preexisting nodes for seamless node integration, no matter what version they installed originally.

AHV Virtual Machine Management

VM management on AHV involves creating, updating, deleting, protecting, and monitoring VMs and their resources. These cluster services and features are all available through the highly available Prism interface. For more information, see the [Virtual Machine Management](#) section of the AHV Administration Guide.

Prism displays a list of all VMs in an AHV cluster along with configuration, resource usage, and performance details on a per-VM basis. Administrators can create VMs and perform many operations on selected VMs, including power state change, reset, shut down, restart, snapshot and clone, migrate, category assignment, policy assignment, and launch a remote console. You can manage the entire life cycle of a VM and its policies that govern security and data protection from Prism.

The image management service in AHV is a centralized repository that provides access to virtual media and disk images and the ability to import from external sources. With AHV's image management service, you can store VMs or virtual disks as images, which you can then use to create new VMs quickly from a known-good base image. The image management service can store the virtual disk files you used to create the VMs or OS installation media as an .iso file that you can mount to provide a fresh OS install experience. Incorporated into Prism, the image management service can import and convert existing virtual disk formats, including .raw, .vhd, .vmdk, .vdi, and .qcow2.

The previous virtual hardware settings don't constrain an imported virtual disk, so administrators have the flexibility to fully configure CPU, memory, virtual disks, and network settings when they provision VMs.

Tools and Processes for Cloning, Importing, and Exporting AHV VMs

AHV supports several tools and processes that make it easier to manage cloning, importing, and exporting AHV VMs, including the Generation ID, the AHV image library, VM templates, and Open Virtual Appliance (OVA) files:

Generation ID

Multiple applications use the Generation ID to access a VM identifier to validate if the VM was cloned or duplicated for licensing or functionality verification. Starting in AOS version 6.7, AHV creates a Generation ID for each VM you create, which applications running inside that VM can access. Applications can then make decisions on how they want to behave based on whether the correct Generation ID is present. This feature ensures that safeguards exist for cloning and rolling back to protect applications such as Windows domain controllers from the issues that accidental rollbacks or cloning can cause. Generation ID provides this information to a VM, so that the app can test and enforce these restrictions. For more information, see the [VM Generation UUID and BIOS UUID Support in AHV](#) section of the AHV Administration Guide.

AHV image library

The AHV image library captures data in a single vDisk so that you can clone it easily, requiring input from the admin to declare the CPU, memory, and network details. VM templates simplify this concept further. You can create AHV VM templates from existing VMs, so that they inherit the attributes of the defining VM, such as the CPU, memory, and networking details. You can then configure the template to customize the guest OS on deployment and provide a Windows license key if needed.

VM templates

VM templates support maintenance of multiple versions of a template, making it easy to apply updates like operating system and application patches without creating new templates. With the template versioning function, you can declare which version of the template is active, so that you can roll back to a previous version if needed or stage updates before making them available. For more

information, see the [VM Template Management](#) section of the AHV Administration Guide.

Open Virtual Appliance files

Industry-standard OVA files are useful in many situations, but the most common are when you deploy virtual appliances (typically from a non-Nutanix vendor) and when you move VMs to a different hypervisor. To help you accomplish these tasks, the OVA import-and-export process is available in Prism Central. You can export VMs on AHV to the OVA format to keep them on external storage or import standard OVAs to launch third-party VMs quickly. For more information, see the [OVAs](#) section of the AHV Administration Guide.

Tools and Processes for Intelligent VM Placement

AHV uses the Acropolis Dynamic Scheduler (ADS), an automatic function enabled on every AHV cluster, to avoid hot spots in cluster nodes. ADS continually monitors CPU, memory, and storage data points to make migration and initial placement decisions for VMs and Nutanix Volumes. Starting with existing statistical data for the cluster, ADS watches for anomalies, honors affinity controls, and only makes move decisions to avoid hot spots. Using machine learning, ADS can adjust move thresholds over time from their initial fixed values to achieve the greatest efficiency without sacrificing performance.

ADS tracks each node's CPU and memory usage. When a node's CPU allocation breaches its threshold (currently 85 percent), Nutanix migrates VMs or Nutanix Volumes off that host as needed to rebalance the workload. ADS only uses migration when the system experiences contention. If the nodes experience skewed usage (for example, three nodes at 10 percent and one at 50 percent), ADS doesn't migrate the nodes, because migration offers minimal benefit if the system has enough resources. For more information, see the [Acropolis Dynamic Scheduling in AHV](#) section of the AHV Administration Guide.

When you create, restore, or recover VMs, the system assigns them to an AHV host in the cluster based on a recommendation from ADS. This VM placement process also takes into account the AHV cluster's high-availability configuration, so it doesn't violate any failover host or segment reservations. We explain these high-availability constructs in the Automated High Availability section.

You can use affinity controls with ADS to govern where VMs run. AHV has two types of affinity controls:

- VM-host affinity strictly ties a VM to a host or group of hosts, so the VM only runs on that host or group. Use cases that involve software licensing or VM appliances often require affinity. In such cases, you typically need to limit the number of hosts an application can run on or bind a VM appliance to a single host. For more information, see the [VM-Host Affinity Policies](#) section of the AHV Administration Guide.
- Antiaffinity lets you declare that a given list of VMs can't run on the same hosts. Antiaffinity allows clustered VMs or VMs running a distributed application to run on different hosts, increasing the application's availability and resilience. To prefer VM availability over VM separation, the system overrides this type of rule when a cluster becomes constrained.

In general-purpose environments, where a single Prism Central instance manages multiple clusters and an administrator manages cluster resource use, you can use the Automatic Cluster Selection feature to determine optimal cluster placement. The Automatic Cluster Selection feature intelligently determines the optimal cluster for VM deployment based on resource availability, host affinities, project specifications, and image distribution policies. Automatic Cluster Selection helps you deliver a cloud-like experience to your users with minimal management overhead. It ensures that resources are automatically balanced across clusters as new VMs are deployed. It also saves valuable time and ensures that workloads are evenly distributed for maximum performance and utilization across the infrastructure.

Live and Cross-Hypervisor VM Migration

Live migration allows the system to move VMs from one host to another while the VM is turned on without workload interruption, regardless of whether the administrator or an automatic process initiates the movement. Live migrations occur regularly in the cluster nodes, triggered by maintenance operations, ADS workload balancing, node expansion, or administrator-driven requests.

You can also use live migration to migrate VMs to another physical cluster in the same location or in a different location to rebalance workloads, run maintenance operations, or avoid planned interruptions. For cross-cluster live migrations, we recommend a network with 5 ms of network latency and can support up to 40 ms of network latency.

For more information on VM migration on AHV, see the [VM Migration Specifications](#) section of the AHV Administration Guide.

Nutanix simplifies the process of migrating existing VMs between an ESXi cluster and an AHV cluster using built-in data protection capabilities. You can create one or more protection domains on the source ESXi cluster and set the AHV cluster as the target remote cluster. Then, snapshot VMs on the source ESXi cluster and replicate them to the AHV cluster, where you can restore them and bring them online as AHV VMs. For more information, see the [VM Migration Specifications](#) section of the AHV Administration Guide.

Storage- and Compute-Only Nodes

With AHV, you can declare a node as never schedulable when joining it to a cluster. Commonly referred to as storage-only nodes, these never-schedulable nodes enable you to scale the storage performance and capacity of a cluster without expanding the compute resources. Because you configure this setting when you join a node to a cluster, you can't easily change it without removing the node from the cluster and rejoining it as a regular node. This setting is most helpful when you need to scale storage resources for workloads licensed by compute resources, such as popular database services. For more information, see the [Adding an AHV Storage-Only Node](#) section of the AHV Administration Guide.

Compute-only nodes provide an additional option for controlling and realizing value from software licenses. Compute-only nodes don't have CVMs, so VMs use the CVMs on the hyperconverged infrastructure (HCI) or storage-only nodes to access disks. As a result, licensed cores on the compute-only nodes are only used for the application VMs. As with storage-only nodes, you must configure the compute-only status when you image the node with Foundation before adding it to the cluster. For more information, see the [AHV Compute-Only Node Configuration](#) section of the AHV Administration Guide.

Applications or databases that are licensed on a per-CPU core basis require you to license the entire node, which includes the cores that the CVM runs on. With compute-only nodes, you might see a much higher return on investment on your database licenses because the CVM doesn't consume any compute resources from the node where the application or database runs.

AHV compute-only nodes support the following two types of deployments:

- AHV compute-only node + AHV HCI nodes

- AHV compute-only node + AHV storage-only nodes

This deployment type is also called the Nutanix Optimized Database Solution.

AHV Backup and Disaster Recovery

Nutanix clusters running any hypervisor have access to Nutanix data protection services, which safeguard VMs both locally and remotely for use cases ranging from basic file protection to recovery from a complete site outage. For more information on the built-in backup and disaster recovery capabilities in the Nutanix platform, see [Data Protection and Disaster Recovery Best Practices](#).

AHV publishes a rich set of APIs to support external backup vendors. The AHV backup APIs use changed region tracking to allow backup vendors to back up only the data that changed since the last backup job for each VM. Changed region tracking also allows backup jobs to skip reading zeros, further reducing backup times and bandwidth consumed. Nutanix backup APIs allow backup vendors that build integration to perform full, incremental, and differential backups. Changed region tracking is always on in AHV clusters; you don't need to enable it on each VM. Backups can be either crash-consistent or application-consistent.

Another feature available in the backup APIs is instant recovery, which enables supported backup vendors to instantly recover a VM or vDisk typically on their backup appliance and allow immediate access to its contents. You can use and delete the VM or vDisk or migrate it to a Nutanix volume for permanent storage.

AOS offers VM high availability to ensure VM availability in the event of a host or block outage. If a host fails, the VMs previously running on that host restart on healthy nodes throughout the cluster. Multiple high-availability configuration options are available to account for different cluster scenarios:

- By default, all AHV clusters provide best-effort high availability, even when you didn't configure the cluster for high availability. Best-effort high availability works without reserving any resources and doesn't enforce admission control, so the capacity might not be sufficient to start all the VMs from the failed host.
- You can also configure an AHV cluster for high availability with resource reservation to guarantee that the resources required to restart VMs are always available. To enable guaranteed high-availability reservations, select the **Enable HA Reservation** checkbox.

The Acropolis leader CVM restarts the VMs on the healthy hosts and tracks host health by monitoring connections to the libvirt on all cluster hosts. If the Acropolis leader becomes partitioned or isolated, or if it fails, the healthy portion of the cluster elects a new Acropolis leader from the Acropolis workers in the CVMs.




Figure 3: Acropolis Leader Monitoring

AHV Metro Availability combines the Nutanix hyperconverged infrastructure with a continuous-availability solution that limits downtime and preserves all data, even during a complete site failure. Metro Availability also enables workload mobility for disaster avoidance and planned maintenance scenarios. Administrators can use Metro Availability to extend hypervisor clustering technologies across datacenters. We call this type of configuration a stretched cluster, and it helps to minimize downtime during unplanned outages.

Metro Availability also supports VM migration across sites using technologies like live migration, which means that you have zero downtime while transitioning workloads between datacenters. Nutanix built safeguards into the platform for everything from minor events, such as individual VM deletion, to major ones, including unplanned datacenter failure.

4. AHV Performance and Scale

The Nutanix platform optimizes performance at both the AOS and hypervisor levels. The CVMs that represent the control and data planes contain the AOS optimizations that benefit all supported hypervisors. Although we built AHV on an open-source KVM foundation, we added a significant amount of innovation to make AHV a unique Nutanix offering, including virtual Trusted Platform Module (vTPM) and secure boot, AHV Turbo, and several memory features.

Nutanix AHV supports host secure boot and provides vTPM for VMs. A TPM is hardware component in a server or computer that provides secure storage for keys or credentials. A vTPM is an integrated feature in the hypervisor that provides the same functionality but is implemented as software for VMs to use. Windows 11 is the primary reason most organizations use a vTPM in their environments, as running that operating system virtually requires a vTPM. Other Windows features such as BitLocker and Credential Guard also use vTPM. For more information, see the [Securing AHV VMs with Virtual Trusted Platform Module \(aCLI\)](#) section of the AHV Administration Guide.

To maximize storage performance, Nutanix developed AHV Turbo technology, which benefits the following entities:

- VMs, by presenting a multiqueue virtio-scsi PCI controller
- The hypervisor, by processing different queues in parallel using multiple threads
- The interface between the hypervisor and the storage layer, by increasing efficiency when receiving the requests from the hypervisor and passing them to the CVM

The following figure explains the logical implementation of AHV Turbo.




Figure 4: AHV Turbo Technology Logical Implementation

AHV Turbo technology is transparent to VMs, but you can make the following configuration changes to achieve greater performance:

- Enable multiqueue on the VM. Consult your Linux distribution documentation to ensure that the guest operating system fully supports multiqueue before you enable it and for instructions on how to enable it. One common method for enabling multiqueue is to add one of the following lines to the kernel command line:

```
scsi_mod.use_blk_mq=y  
scsi_mod.use_blk_mq=1
```

- Install virtIO for Windows-based VMs. No additional configuration is required.
- Provide more than 1 vCPU to the VM.
- Ensure that the workloads are multithreaded.

AHV Turbo represents significant advances to the data path in AHV over the core KVM source code foundation. In the core KVM code, all I/O from a VM flows through the hosted VM monitor, QEMU. While this architecture can achieve impressive performance, some application workloads require more. AHV Turbo provides an I/O path that bypasses QEMU and services storage I/O requests, which decreases CPU usage and increases the amount of storage I/O available to VMs. When using QEMU, all I/O travels through a single queue, which can affect performance. The storage queues scale out automatically to match the number of vCPUs configured for a given VM, making even better performance possible as the workload scales up. While these improvements demonstrate immediate benefits, they also prepare AHV for future technologies such as NVMe and persistent memory advances that offer dramatically increased I/O capabilities with lower latencies. For more information, see the [AHV Turbo](#) section of the AHV Administration Guide.

AHV Memory Improvements for Performance

AHV brings several memory-related innovations for performance and scalability, including vNUMA, RDMA, and memory overcommit. Modern Intel server architectures assign memory banks to specific CPU sockets. In this design, one of the memory banks in a server is local to each CPU, so you see the highest level of performance when accessing memory locally, as opposed to accessing it remotely from a different memory bank. Each CPU and memory pair is a NUMA node. vNUMA allows a VM's architecture to mirror the NUMA architecture of the underlying physical host.

vNUMA doesn't affect most workloads, but it can be beneficial to large VMs configured with more vCPUs than available physical cores in a single CPU socket. In these scenarios, configure vNUMA nodes to use local memory access efficiently for each CPU to achieve the highest performance results. For more information, see the [Virtual Machine Memory Management \(vNUMA\)](#) section of the AHV Administration Guide.

Remote direct memory access (RDMA) enables a node to write to the memory of a remote node by granting a VM running in the user space access to a NIC directly. This approach avoids TCP and kernel overhead, resulting in CPU savings and performance gains. AOS RDMA support is reserved for inter-CVM communications and uses the standard RDMA over Converged Ethernet (RoCEv2) protocol on systems configured with RoCE-capable NICs.

iSCSI Extensions for RDMA (iSER) is an extension of iSCSI that uses RDMA direct data placement technology to eliminate copies in the data path. The system uses the CPU resources efficiently by taking advantage of Mellanox NICs for offloading functions and reduces latency between the AHV host and CVM I/O transfer. iSER provides the following key benefits:

- Enhances the I/O flow performance between AHV host and CVM: iSER bypasses both the AHV host and CVM kernel spaces and enables you to achieve low latency in the I/O data transfer as the data is sent and received between AHV host and CVM without involving the TCP software stack. iSER enables you to achieve a high bandwidth in the I/O data transfer as it removes the unwanted context switches and system calls between the AHV host and CVM.
- Saves platform resources: iSER reduces CPU utilization because an application can read remote memory without any intervention from the remote processor.

One of the central benefits of virtualization is the ability to overcommit compute resources, making it possible to provision VMs with more CPUs than are physically present on the server host. Most workloads don't need 100 percent of their assigned CPU all the time, and the hypervisor can dynamically allocate CPU cycles to workloads that need them.

You can also overcommit memory. The VMs on the host might not use all their allocated memory, and the hypervisor can share that unused memory with other workloads. Memory overcommit enables administrators to provision more VMs per host by combining the unused memory and allocating it to VMs that need it. AOS 6.1 and later

versions offer memory overcommit to AHV for environments like test and development that need additional memory and VM density. You define overcommit on a per-VM basis, so you control whether to share memory between all VMs on a cluster or just a subset. For more information, see the [Memory Overcommit](#) section of the AHV Administration Guide.

5. AHV Graphics Processing Unit Support

A graphics processing unit (GPU) is the hardware or software that displays graphical content to end users. In laptops and desktops, a GPU is either a physical card or built directly into the CPU hardware. Historically, virtualized GPU functions were software-driven and consumed additional CPU cycles. With modern operating systems, applications, and 3D tools, more organizations find themselves needing a hardware GPU. You can install physical GPU cards in qualified hosts and present them to guest VMs using passthrough or vGPU mode. Nutanix AHV allows you to assign up to 16 vGPU profiles to a single VM or a single GPU using passthrough. For more information, see the [GPU and vGPU Support](#) section of the AHV Administration Guide.

The GPU cards deployed in server nodes for virtualized use cases typically combine multiple GPUs in a single Peripheral Component Interconnect (PCI) card. AHV can pass a GPU through to a VM, allowing the VM to own that physical device in a 1:1 relationship. When you configure nodes with one or more GPU cards that attach multiple GPUs to a larger number of VMs, you can consolidate applications and users on each node. With passthrough, you can also use GPUs for offloading computational workloads—a more specialized situation than the typical graphical use cases.

AHV currently supports NVIDIA GRID cards for GPU passthrough. For the current list of supported devices, see the [AHV Administration Guide](#).

Nutanix LCM doesn't support live migration for VMs configured with GPU passthrough. For more information, see the limitations in the [GPU Passthrough for Guest VMs](#) section of the AHV Administration Guide.

vGPU

While passthrough is a method that works well for fewer VMs that require more GPU resources, workloads such as VDI often have different requirements. VDI workloads typically have more VMs that need varying GPU resources based on application types and usage.

Today's NVIDIA GRID GPU cards contain between one and four GPUs on each physical PCI card, and each physical host can support installing one or two cards. This capability allows up to eight GPUs in each node to meet the density requirements of VDI workloads. For maximum flexibility, vGPU mode allows you to slice each GPU into smaller segments that you can virtually assign to VMs.




Figure 5: HCI Node with One PCI Card and Three GPUs with vGPU Enabled

vGPU profiles allow you to assign different levels of resources to VMs, so that each VM can use a defined maximum number of displays and quality of resolution. Working within these parameters lets you choose the correct GPU profile to meet your application requirements, and when you also know the type and number of GPU cards in your deployment, you can accurately describe the maximum density possible per host.

In some rare use cases, you might need additional GPU resources for demanding workloads. With AHV, you can assign multiple vGPU profiles to a single VM to fulfill the requirements of these workloads.

AHV currently supports NVIDIA GRID cards for vGPU. For the current list of supported devices, see the [NVIDIA Drivers](#) section of the Compatibility and Interoperability Matrix.

To make maintenance operations like upgrades and failure handling as efficient as possible, you need the ability to migrate your VMs. With AHV, you can live-migrate VMs with vGPU assignments. For more information on live migration with your combination of GPU and NVIDIA host driver, see the [NVIDIA GRID Host Driver for AHV Installation Guide](#) or contact Nutanix support.

6. AHV Migration with Nutanix Move and Cluster Conversion

The most common migration method in virtualized environments involves moving an existing operating system and its applications without modification. AHV supports a wide range of operating systems that you can migrate from either physical or virtual environments. You can use Nutanix Move and cluster conversion to migrate workloads to AHV. For more information, see [Migrating VMs to Nutanix AHV](#).

Nutanix Move quickly and easily transfers VMs from ESXi, Hyper-V, and public cloud environments (including AWS and Azure) to Nutanix AHV, keeping VM data unmodified and intact. Nutanix Move is a VM appliance, typically hosted on the target AHV cluster. After you deploy this free product as a VM in your AHV environment, you can use its HTML5-based interface to manage it. With Nutanix Move, you can create migration groups that consist of a single VM or large groups. This migration group construct allows you to schedule and control data replication and to cut over the entire group, specific VMs, or a single VM to the target cluster. The source data remains intact, but the VMs are shut down and disconnected until it's safe to clean up.

Sometimes, you might need to convert a cluster rather than migrate to a new cluster; for example, you might have a single Nutanix cluster that is deployed at a site that you need to convert to AHV. The cluster conversion process allows you to convert a Nutanix cluster running VMware ESXi to AHV in place with minimal downtime. The process is a rolling cluster conversion similar to the upgrade process except that the node's hypervisor under maintenance is converted to AHV. The automated process places a node into maintenance mode and live-migrates all VMs on that node to other nodes. After the hypervisor is fully evacuated and in maintenance mode, the process converts it to AHV. Once AHV is running on that node, the process shuts down the VMs that were running on it, converts them into AHV format, and then turns them on in the AHV node. Because you changed the VM format, each VM experiences a brief downtime. After the process converts the first node, it moves to the next node and rolls through the cluster until each node is converted. When the process finishes, the cluster is fully converted to AHV, and you can roll it back to ESXi if you need to.

7. AHV Security

The fully integrated Nutanix infrastructure stack reduces security risks associated with legacy solutions that involve many vendors with a narrow, fragmented view of security. Nutanix designed AHV to be an integral component of the converged infrastructure stack rather than a general-purpose hypervisor. Consequently, to reduce the security surface area, we turned off many of the services that we don't use or that AHV obviates. For example, with AHV, you don't need to use Virtual Network Computing to console VMs, because Prism offers a more robust and secure service. The audits log in Prism provides a comprehensive list of all actions taken by administrators and users against AHV resources. You can easily locate details on who took an action (such as a VM creation, deletion, or update) and when.

To maintain agile and comprehensive security, Nutanix developed its own Security Development Life Cycle (SecDL) that addresses security at every step of the development process instead of applying it at the end. SecDL integrates security features into the software development process, including automated security testing during development and threat modeling to assess and mitigate customer risk from code changes.

Nutanix also developed custom Security Technical Implementation Guides (STIGs)—security tools based on well-established National Institute of Standards and Technology (NIST) standards—that administrators can apply to multiple baseline requirements for the United States Department of Defense (DoD) and PCI-DSS. Unlike general-purpose STIGs that make blanket security recommendations, Nutanix STIGs are specific to the Nutanix platform and therefore more effective. Encoded in a machine-readable format, Nutanix STIGs enable automated validation, ongoing monitoring, and self-remediation, reducing the time required to verify security compliance from weeks or months to days. For more information, see the [Nutanix security page](#).

If you host unstructured data on AHV using Nutanix Unified Storage (NUS)—including Nutanix Files and Nutanix Objects—Nutanix Data Lens offers ransomware resilience and data analytics for unstructured data on the Nutanix Cloud Platform. This service offers global data visibility to proactively assess and mitigate data security risks by identifying anomalous activity, auditing user behavior, and adhering to compliance requirements

while enabling efficient data life cycle management. Nutanix Data Lens centralizes information from all your clusters connected to Nutanix Pulse across various datacenter locations. Using cloud resources eliminates scaling constraints, allowing you to analyze an unlimited number of files and objects and unlimited amounts of data.

In addition to Nutanix features, AHV supports Credential Guard through the Nutanix vTPM. Credential Guard is an additional virtualization-based security layer for the Local Security Authority Subsystem Service (LSASS) process in Windows that stores credentials for New Technology LAN Manager (NTLM) and Kerberos. Credential Guard isolates the LSASS process in a virtual container that can't be accessed by users and creates a proxy process to communicate with it, protecting the system from further lateral incursion.

About Nutanix

Nutanix offers a single platform to run all your apps and data across multiple clouds while simplifying operations and reducing complexity. Trusted by companies worldwide, Nutanix powers hybrid multicloud environments efficiently and cost effectively. This enables companies to focus on successful business outcomes and new innovations. Learn more at [Nutanix.com](https://www.nutanix.com).

List of Figures

Figure 1: AHV Components..... 8

Figure 2: Storage Configuration..... 9

Figure 3: Acropolis Leader Monitoring.....20

Figure 4: AHV Turbo Technology Logical Implementation.....22

Figure 5: HCI Node with One PCI Card and Three GPUs with vGPU Enabled..... 27



NVIDIA L40S

Unparalleled AI and graphics performance for the data center.




Generative AI is fueling transformative change, unlocking a new frontier of opportunities for enterprises across every industry. To transform with AI, enterprises need more compute resources, greater scale, and a broad set of capabilities to meet the demands of an ever-increasing set of diverse and complex workloads.

The NVIDIA L40S GPU is the most powerful universal GPU for the data center, delivering end-to-end acceleration for the next generation of AI-enabled applications—from **gen AI**, LLM inference, small-model training and fine-tuning to 3D graphics, rendering, and video applications.

Accelerate Next-Generation Workloads


- > Generative AI
- > LLM inference
- > LLM fine-tuning and small-model training
- > NVIDIA Omniverse™ Enterprise
- > Rendering and 3D graphics
- > Streaming and video content

Generative AI Image Generation



Measured performance; NVIDIA L40S
Stable Diffusion v2.1, TRT 8.6.1, BS:1, FP16 |
Stable Diffusion XL 1.0, TRT 8.6.1, BS:1, FP16

Large Language Model (LLM) Inference



Measured performance; NVIDIA L40S
Llama 2-7B/13B/70B, ISL=2048, OSL=128,
BS=1: FP8.

Powered by the NVIDIA Ada Lovelace Architecture

Fourth-Generation Tensor Cores

Hardware support for structural sparsity and optimized TF32 format provides out-of-the-box performance gains for faster AI and data science model training. Accelerate AI-enhanced graphics capabilities with **DLSS** to upscale resolution with better performance in select applications.

Third-Generation RT Cores

Enhanced throughput and concurrent ray-tracing and shading capabilities improve ray-tracing performance, accelerating renders for product design and architecture, engineering, and construction workflows. See lifelike designs in action with hardware-accelerated motion blur and stunning real-time animations.

Transformer Engine

Transformer Engine dramatically accelerates AI performance and improves memory utilization for both training and inference. Harnessing the power of the **Ada Lovelace fourth-generation Tensor Cores**, Transformer Engine intelligently scans the layers of transformer architecture neural networks and automatically recasts between FP8 and FP16 precisions to deliver faster AI performance and accelerate training and inference.

Data Center Ready

The L40S GPU is optimized for 24/7 enterprise data center operations and designed, built, tested, and supported by NVIDIA to ensure maximum performance, durability, and uptime. The L40S GPU meets the latest data center standards, is Network Equipment-Building System (NEBS) Level 3 ready, and features secure boot with root of trust technology, providing an additional layer of security for data centers.

Technical Specifications	
GPU Architecture	NVIDIA Ada Lovelace Architecture
GPU Memory	48GB GDDR6 with ECC
Memory Bandwidth	864GB/s
Interconnect Interface	PCIe Gen4 x16: 64GB/s bidirectional
NVIDIA Ada Lovelace Architecture-Based CUDA® Cores	18,176
NVIDIA Third-Generation RT Cores	142
NVIDIA Fourth-Generation Tensor Cores	568
RT Core Performance TFLOPS	209
FP32 TFLOPS	91.6
TF32 Tensor Core TFLOPS	183 366*
BFLOAT16 Tensor Core TFLOPS	362.05 733*
FP16 Tensor Core	362.05 733*
FP8 Tensor Core	733 1,466*
Peak INT8 Tensor TOPS	733 1,466*
Peak INT4 Tensor TOPS	733 1,466*
Form Factor	4.4" (H) x 10.5" (L), dual slot
Display Ports	4x DisplayPort 1.4a
Max Power Consumption	350W
Power Connector	16-pin

Thermal	Passive
Virtual GPU (vGPU) Software Support	Yes
vGPU Profiles Supported	See the virtual GPU licensing guide
NVENC NVDEC	3x 3x (includes AV1 encode and decode)
Secure Boot With Root of Trust	Yes
NEBS Ready	Level 3
MIG Support	No
NVIDIA® NVLink® Support	No

* With sparsity

Ready to Get Started?

To learn more about the NVIDIA L40S, visit
www.nvidia.com/l40s

© 2024 NVIDIA Corporation and affiliates. All rights reserved. NVIDIA, the NVIDIA logo, CUDA, HGX, NVLink, and Omniverse are trademarks and/or registered trademarks of NVIDIA Corporation and affiliates in the U.S. and other countries. Other company and product names may be trademarks of the respective owners with which they are associated. 3110647. FEB24



Intelligent Operations Guide

Intelligent Operations pc.7.3

October 14, 2025

Contents

About This Publication.....	5
NCM Licensing.....	6
Intelligent Operations Overview.....	7
Enabling Intelligent Operations.....	7
Accessing Intelligent Operations.....	9
Intelligent Operations Dashboard.....	9
Disabling Intelligent Operations.....	10
Performance Monitoring.....	12
Analysis Dashboard.....	12
Modifying the Time Range.....	16
Viewing Alerts and Events.....	17
Generating Analysis Dashboard Summary Report.....	21
Chart Management.....	24
Sessions Management.....	41
Creating a New Session.....	43
Setting the Default Session View.....	44
Deleting a Session.....	44
Editing Session Details.....	45
Sharing a Session.....	45
Closing a Session.....	47
Switching Between Sessions.....	47
Application Discovery.....	49
Application Discovery Summary View.....	49
Enabling Application Discovery.....	60
Prerequisites for Enabling Application Discovery.....	64
Connecting to the Cloud.....	65
Adding or Removing Clusters to Monitor.....	66
Configuring vCenter Authentication.....	68
Discovering Applications.....	70
Publishing Applications.....	71
Creating Discovery Policies.....	71
Modifying or Viewing Discovery Policies.....	72
Application Monitoring.....	74
Application Monitoring Summary View.....	74
Application Instance Details View (SQL Server).....	77
Application Instance Details View (vCenter).....	84
Configuring Application Monitoring.....	86
Prerequisites for Configuring Application Monitoring.....	87
Configuring an SQL Server Application Monitoring.....	88

Configuring a vCenter Application Monitoring.....	89
Modifying Application Monitoring.....	91
Capacity Planning.....	92
Capacity Tab-storage View.....	94
Capacity Tab- CPU View.....	95
Capacity Tab-memory View.....	96
Resource Planning.....	98
Scenarios Summary View.....	98
Creating a Scenario.....	99
Analyzing a Scenario.....	110
Modifying a Scenario.....	113
Sharing a Scenario.....	114
Capacity Runway Summary View.....	116
Updating Capacity Configurations.....	118
Behavioral Learning Tools.....	119
Task Automation - Playbooks.....	123
Overview.....	123
Playbooks Summary View.....	123
Playbook Triggers.....	129
Playbook Actions.....	129
Prerequisites for VM Actions in Playbooks.....	144
User Impersonation to Run a Playbook.....	147
Create a Playbook.....	147
Creating Playbooks Using Alert.....	148
Creating Playbooks Using Alerts Matching Criteria.....	151
Creating Playbooks Using Event.....	152
Creating Playbooks Using Time.....	154
Creating Playbooks Using Manual Triggers.....	155
Creating Playbooks Using Webhook.....	156
Using Branch Action (Conditional Execution).....	158
Sharing a Playbook.....	160
Generating Playbooks Summary Report.....	161
Downloading a Report.....	162
Emailing a Report.....	163
Saving as a Report Configuration.....	163
Saving as a Report View.....	164
Running a Playbook (Manual Trigger).....	165
Configuring Manual Parameters.....	165
Parameters for Triggers and Actions.....	167
Exporting or Importing Playbooks.....	171
X-Play Integrations.....	172
Integrating Slack Action in Playbooks.....	172
Integrating MS Teams Action in Playbooks.....	173
Ansible Integration in Playbooks.....	174
ServiceNow Integration with Prism Central (X-Play Support).....	176
PagerDuty Integration with Prism.....	181
Reports Management.....	185
Reports Summary View.....	186

Managing a Predefined System Report Configuration.....	191
Running a Report Instance.....	194
Cloning a Predefined System Report Configuration.....	195
Scheduling a Predefined System Report Configuration.....	196
Sharing a Predefined System Report Configuration.....	197
Creating a Custom Report Configuration.....	199
Managing a Custom Report Configuration.....	200
Running a Report Instance.....	201
Editing a Custom Report Configuration.....	202
Cloning a Custom Report Configuration.....	203
Scheduling a Custom Report Configuration.....	203
Sharing a Custom Report Configuration.....	206
Exporting a Custom Report Configuration.....	208
Deleting a Custom Report Configuration.....	208
Importing a Custom Report Configuration.....	209
Managing a Report Instance from Configurations Tab.....	209
Viewing a Report Instance.....	209
Downloading a Report Instance.....	210
Resending a Report Instance.....	211
Deleting a Report Instance.....	212
Example Report - Environment Summary.....	212
Managing a Report Instance from Generated Reports Tab.....	221
Downloading a Report Instance.....	221
Resending a Report Instance.....	222
Deleting a Report Instance.....	223
Editing a Saved Report Configuration.....	223
Views in the Reports.....	224
Adding a Custom View.....	225
Adding a Predefined View.....	258
Adding a Saved View.....	261
Configuring Report Settings.....	262
Role Based Access Control.....	266
 Settings & Configurations.....	 268
Storing Non-Nutanix Credentials in Repository.....	268
Operations Policy Management.....	269
Operations Policies Summary View.....	269
Managing Operations Policies.....	271
 Copyright.....	 275

ABOUT THIS PUBLICATION

This document provides the information about Nutanix Cloud Manager Intelligent Operations.

Intended Audience

The intended audience for this document are site personnels, customer's operation teams and any tier or individual entity defined by the customer who is equipped to perform the site configurations and modifications. In addition to this, any individual with an intention to learn and develop understanding about Nutanix Cloud Manager Intelligent Operations usage, can refer this document.

NCM LICENSING

NCM Starter, NCM Pro and NCM Ultimate license tiers enable you to access Intelligent Operations. NCM Starter tier is ideal for infrastructure operations such as monitoring, planning, rightsizing and low code automation. NCM Pro and NCM Ultimate tiers provide comprehensive infrastructure operations features.

Nutanix provides a free trial period of 90 days to try out Intelligent Operations. The free trial starts on Prism Central deployment.

The Prism Element web console and Nutanix Support Portal provide the most current information on your licenses. For more information on licenses, see the [License Manager Guide](#).

INTELLIGENT OPERATIONS OVERVIEW

Nutanix Cloud Manager Intelligent Operations (formerly Prism Pro/Ultimate) is an end-to-end operations, monitoring and planning solution that combines aspects of administration, reporting, and intelligent automation for IT operations. Intelligent Operations empowers you to plan intelligently and optimize for capacity, proactively detect performance anomalies, and automate the IT operations. Intelligent Operations also includes capabilities such as advanced application insights and automation for troubleshooting applications.

Intelligent Operations comprises the following features:

- **Analysis:** Analyze cluster activity. For more information, see [Performance Monitoring](#) on page 12.
- **Discovery:** Discover applications running in a cluster. For more information, see [Application Discovery](#) on page 49.
- **Monitoring Configurations:** Monitor select applications in a cluster. For more information, see [Application Monitoring](#) on page 74.
- **Planning:** Plan for resource capacity and usage. For more information, see [Resource Planning](#) on page 98.
- **Playbooks:** Create playbooks for task automation. For more information, see [Task Automation - Playbooks](#) on page 123.
- **Reports:** Generate resource and activity reports for Nutanix and non-Nutanix environments. For more information, see [Reports Management](#) on page 185.

Access Intelligent Operations from Prism Central web console. Prism Central provides a unified interface for managing all your clusters from a single management interface.

Enabling Intelligent Operations

You can enable Intelligent Operations through Admin Center in Prism Central web console.

Before you begin

Ensure that you have four vCPUs and two GiB memory free in the PC VM that can be allocated to Intelligent Operations. If you do not have four vCPUs and two GiB memory free in the PC VM, you can add them while enabling Intelligent Operations. The features supported by Intelligent Operations might suffer performance degradation if the required number of vCPUs and memory are not allocated for Intelligent Operations.

About this task

To enable Intelligent Operations from Prism Central, do the following.

Procedure

1. Log in to Prism Central.
2. Select **Admin Center** in the Application Switcher.
For information on the application switcher functionality and Admin Center, see [Application Switcher Function](#) in the *Prism Central Infrastructure Guide*.
3. Click **Settings** in the navigation bar.

4. Navigate to **Setup > Intelligent Operations**. The **Intelligent Operations** page appears.

Enable Intelligent Operations

Run your IT ops efficiently with intelligent operational features. [Learn more](#)

Product Status: ⏻ Disabled ?

License Status: Active

Please follow these instructions to enable and experience Intelligent Operations.

1. Resource Requirements

Intelligent Operations requires additional resources on PC VM for better performance. Select a resource allocation option.

☒ **Add Resources**

Add 4 vCPUs and 2GiB memory to PC VM and allocate them to Intelligent Operations

☐ **Already Added?**

PC VM has additional 4 vCPUs and 2GiB memory. Allocate them to Intelligent Operations

[Allocate & Enable](#)

2. License Requirements

Please ensure Intelligent Operations is licensed. [Manage license](#)

Figure 1: Intelligent Operations

5. In the **Resource Requirements** section, do either of the following:

- » Select the **Add Resources** checkbox to add four new vCPUs to the PC VM, and allocate them for Intelligent Operations.

Note: Select this checkbox if you do not have four free vCPUs already added in your PC VM.

- » Select the **Already Added?** checkbox to allocate four existing vCPUs in the PC VM for Intelligent Operations.

Note: Select this checkbox if you already have four free vCPUs added in your PC VM.

6. Click **Allocate Resources**.

The **Product Status** toggle turns blue, and displays **Enabled** indicating that the vCPUs are allocated and Intelligent Operations is enabled in the Prism Central VM.

Accessing Intelligent Operations

About this task

Intelligent Operations is a part of Nutanix Cloud Manager offering. You can access Intelligent Operations through Prism Central web console. Prism Central provides a centralized environment to manage multiple Nutanix apps from a single workspace. The Application Switcher facilitates you to seamlessly switch between the Nutanix apps that you have enabled and have access to in Prism Central.

To get details about application switcher functionality, see [Application Switcher Function](#) in the *Prism Central Infrastructure Guide*.

Procedure

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher.




Figure 2: Accessing Intelligent Operations

The system displays the **Intelligent Operations** dashboard.

Intelligent Operations Dashboard

The **Operations Dashboard** allows you to view a summary of the following widgets:

- **Cluster Runway:** Displays a list of clusters with capacity runway. Click the name of the cluster to view the detailed capacity runway information.

- **Plays:** Displays the number of plays completed in the last 24 hours along with a breakup of failed and paused plays. Click the number to view the details.
- **Alerts:** Displays the summary of alerts raised during last 24 hours.
- **Reports:** Displays a summary of total reports and scheduled reports. Click on the number to view the details.
- **Impacted Clusters:** Displays a summary of impacted clusters with details such as alerts, anomalies, runway, inefficient VMs, and plays.
- **VM efficiency:** Displays the summary of VM efficiency with details such as over provisioned, inactive, constrained, and bully. Click **View All Inefficient VMs** to view the details of all the inefficient VMs. For more information about VM Right-sizing or VM Inefficiency Detection, see [Behavioral Learning Tools](#) on page 119.
- **Discovered Apps:** Displays a summary of discovered apps.

To access the **Operations Dashboard**:

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher.

The system displays the **Operations** dashboard with all the widgets.




Figure 3: Operations Dashboard

Disabling Intelligent Operations

About this task

To disable Intelligent Operations from Prism Central, do the following.

Procedure

1. Log in to Prism Central.
2. Select **Admin Center** in the Application Switcher.
3. Click **Settings** in the navigation bar.
4. Navigate to **Setup > Intelligent Operations**. The **Intelligent Operations** page appears.

5. Slide the **Product Status** toggle to the left.

The **Product Status** toggle is grayed out and displays **Disabled** indicating that Intelligent Operations is disabled in the PC VM.

The vCPUs will remain added to the PC VM even after you disable Intelligent Operations. However, the vCPUs will not be allocated for Intelligent Operations. You can use these free vCPUs to enable Intelligent Operations again or to enable other features such as NCM Self-Service (formerly known as Calm).

PERFORMANCE MONITORING

Prism Central allows you to monitor performance across the registered clusters.

- You can view performance graphs for selected elements from the **Analysis** dashboard. For more information, see [Analysis Dashboard](#) on page 12.
- You can create custom entity and metric performance charts. For more information, see [Chart Management](#) on page 24.

Note: These views reflect that Prism Central retains alerts and events for 90 days and hourly data for a year by default.

Analysis Dashboard

The **Analysis** dashboard allows you to create sessions with charts that can dynamically monitor a variety of performance measures.

To access the **Analysis** dashboard:

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher, and navigate to **Analysis** from the **Navigation Bar**.

The **Analysis** dashboard opens displaying the session that you were last working in as the default session.

When you access the **Analysis** dashboard for the first time after installing or upgrading Prism Central, the **Analysis** dashboard displays a session named **Analysis Session**. Prism Central displays the label **System** with the name of this session in the **Sessions** dashboard to indicate that this session is a system-generated session. For more information, see [Sessions Management](#) on page 41.

Note:

- the default system-generated Analysis session and the user-defined **Analysis** sessions are available only with a valid NCM license. For more information, see the [License Manager Guide](#).
- To optimize the display performance on the Analysis Dashboard, only the first 50 VMs are shown, so exporting this chart limits your view to 50 entities. To work around this limitation, go to the reports to generate the full data export.

The following figures are sample views of the **Analysis** dashboard. Some fields include a slide bar on the right to view additional information in that field. The displayed information is dynamically updated.




Figure 4: Analysis Dashboard - Prism-generated Session




Figure 5: Analysis Dashboard - User-defined Session

The **Analysis** dashboard elements are described in the table starting with the elements in the top left corner to the right side of the dashboard.

Table 1: Analysis Dashboard Elements

Field	Description
Session	The name of the session that is displayed in the Analysis dashboard.
Switch Session	By default, the Analysis dashboard displays the last session you viewed. You can change the default session using the Settings option in the Sessions dashboard.
	A dropdown menu of all the created and saved available sessions. Select the session you want to view and modify the details.
	Select Switch Sessions > View All Sessions to view the Sessions dashboard that lists all the created and saved sessions.
Add Chart	You can add a chart by clicking: - The Add Chart option at the top left of the dashboard. - The Add Chart option at the bottom center of the dashboard. You can define a chart to monitor many entities with the same single metric or several metrics for a single entity.
	Note: You cannot define a single chart with many entities and many metrics.

Field	Description
Alerts and Events monitor	Displays the alerts and events that occurred. The alerts and events occurring at any specific point of time are displayed as a stacked bar with three colored segments. The hover image of the stacked bar displays the following: • Time range for the stacked bar. • Red segment depicting Critical Alert with the number of alerts. • Yellow segment depicting Warning Alert with the number of alerts. • Grey segment depicting Events with the number of events. Click the stacked bar to open the Alerts and Events List panel on the right. The Alerts and Events List panel displays the details of the alerts and events represented by the stacked bar.
Alerts and Events List	Provides more information about the alerts and events. When you click: • The Alerts and Events List option at the top right of the Analysis dashboard, the Alerts and Events List panel provides information about all the alerts and events that have occurred during the selected time range. • A stacked bar in the Alerts and Events monitor, the Alerts and Events List panel provides information about all the alerts and events that occurred in the time range indicated by the stacked bar that you clicked on the timeline. See the <i>Alerts and Events List Panel Elements</i> table for descriptions of the tabs and fields in the Alerts and Events List panel.

Field	Description
Actions	Provides the following options for the current session: 1. Edit Session Details 2. Share 3. Download Report (csv, pdf) 4. Email Report (csv, pdf) 5. Save as Report Configuration 6. Delete (session) This action is not available and appears dimmed if the current session is the system-defined session, because the system-defined session cannot be deleted. 7. Close (session) When you close the current session, the Sessions dashboard is displayed.
(Defined chart monitors)	Displays the monitors for the charts that you defined for the session.

Table 2: Alerts and Events List Panel Elements

Field	Description
Alerts tab	Displays the list of alerts. The tab name includes colored icons with inscribed numbers. The color and number provide the following information: • Color depicts the severity type of the alert. Red indicates critical and yellow indicates warning. • Number indicates the number of respective alerts. In the example depicted in the figure, the red icon has 1 inscribed indicating that there is one critical alert in the stacked bar instance.
Group by field in Alerts tab	Provides the grouping criteria that you can use to group the alerts. You can group the alerts by: • Entity Type • Impact • Severity

Field	Description
Events tab	Displays the list of events. The tab name includes a grey-colored icon with an inscribed number. The inscribed number indicates the number of events. You can see the event details page by clicking the event in the list. The events in the list are grouped into two default groups: • System Action • User Action
(Time range for the stacked bar)	The time range considered for the stacked bar plotted on the Alerts and Events monitor. The time range is specifically displayed only when you open the Alerts and Events List panel by clicking a stacked bar in the Alerts and Events monitor.
Hide/Show toggle	Displayed for each group of alerts or events in the list. This toggle allows you to hide or show the alerts or events in a group.

Modifying the Time Range

You can select a time range for the session and the charts you choose in the **Analysis** dashboard.

About this task

The time range you select applies to the session as a whole including the charts you add to the session. The default time range is **1 Day**.

Note: The time range selection affects the whole session and all the charts added to the session.


You can select a predefined time range or customize a range. You can customize a range by moving the start and end time selectors on the time-line or you can set a range in the **Select a Time Range** dialog box.

Perform the following steps to select a predefined time range or customize a range for the session.

Procedure

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher, and navigate to **Analysis** from the **Navigation Bar**.
The **Analysis** dashboard opens displaying the session that you were last working on.
3. Click the **Range** dropdown menu.
4. Select one of the following:
 - » A predefined range from the list.
 - » **Custom Range**.

5. If you select **Custom Range**, perform the following in the **Select a Time Range** dialog box:
 - a. Select the date and time in the **From** section to define the start of the range.
 - b. Select the date and time in the **To** section to define the end of the range.
 - c. Click **Select**.



The image shows a dialog box titled "Select a Time Range" with a close button (X) in the top right corner. Inside the dialog, there are two rows of input fields. The first row is labeled "From" and the second row is labeled "To". Each row contains two input fields: one for the date (with a calendar icon) and one for the time (with a clock icon). At the bottom right of the dialog, there are two buttons: "Cancel" and "Select".

Figure 6:

If you select **Custom Range**, you can also select the range dynamically on the time-line by sliding the range start and end sliders. There is no minimum time range for **Custom Range**, and the maximum range is one year.




Figure 7: Time Range

Viewing Alerts and Events

About this task

You can view the alerts and events occurring during the selected time range in the following ways:

- From the **Alerts and Events** monitor available below the time-line of the **Analysis** dashboard.

- By clicking the **Alerts and Events List** option at the top right of the **Analysis** dashboard.

Procedure

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher, and navigate to **Analysis** from the **Navigation Bar**.
The **Analysis** dashboard opens displaying the session that you were last working on.
3. Hover over a stacked bar on the monitor to view more information about the alerts and events that occurred in the time range that the stacked bar depicts.
Click the stacked bar to view the detailed information of the alerts and events in the **Alerts and Events List** panel.



Figure 8: Stacked bar - Alerts and Events Information

Alternately, you can perform the following steps to view the detailed information of the alerts and events in the **Alerts and Events List** panel.

4. Click the **Alerts and Events List** option at the top right of the **Analysis** dashboard.

The **Alerts and Events List** panel displays the details of the alerts and events that occurred during the time range selected.

For more information, see the [Table 2: Alerts and Events List Panel Elements](#) on page 15 table.

- The **Alerts** tab provides the details of the alerts. You can filter the alerts.

Alerts and Events List

Alerts 1 2 Events 99+

Group by Entity Type :

Cluster

1 warning

Hide

• Flow Control Plane Fal... 08/17, 3:36 PM

VM

1 warning

Hide

• VM Recovery Point Cr... 08/17, 11:21 PM

None

1 critical

Hide

• Protection Domain Sn... 08/18, 8:33 AM

Figure 9: Alerts tab

Use the filters provided in the **Group by** dropdown list to filter the category of alerts.

- The **Events** tab provides the details of the events. The events cannot be filtered.




Figure 10: Events tab

Click **Show** to expand the category of alerts or events.

5. Click an alert or event to open the alert or event details page.

Generating Analysis Dashboard Summary Report

Use the Generate Report feature in Prism Central to export dashboard data. You can download, email, or schedule reports in PDF or CSV format.

In Prism Central, you can monitor infrastructure widgets by visualizing the data on the dashboards. However, the data is available only in the Prism Central interface. After generating a report, you can export the analysis dashboard data as a report for further analysis, archival, compliance, and other purposes.

Note: Generating a report requires a valid license - Nutanix Cloud Manager (NCM) Starter, Pro or Ultimate. For more information, see [Nutanix Cloud Manager Licensing](#). The **Generate Report** option does not appear if the appropriate license is not available.

You can generate a report in the following ways:

- Download the analysis dashboard data as a report in PDF format or CSV format.
- Email the analysis dashboard data as a report to the recipients in PDF format or CSV format, or both, as an attachment.
- Save the analysis dashboard data as a report configuration to schedule, share, and export reports.

You can access the report options from the Analysis Dashboard:




Figure 11: Analysis Dashboard - Report Options

Table 3: Report Options

Options	Description
Download Report (csv, pdf)	Click this action to download the report. For more information, see Downloading a Report on page 22.
Email Report (csv, pdf)	Click this action to email the report. For more information, see Emailing a Report on page 23.
Save as Report Configuration	Click this action to save the report as a configuration. For more information, see Saving as a Report Configuration on page 23.

Downloading a Report

This task describes how to download the **Analysis** dashboard data as a report, either in PDF format or CSV format.

About this task

Downloading a report has the following workflow:

- The report instance is generated in your preferred format and notified on the **Analysis** dashboard.
- The generated report instance is available for you to download under **Generated Reports** tab of the Reports page

Important:

- You can also resend the generated report instance as an email from the **Generated Reports** tab of the Reports page. For more information, see [Resending a Report Instance](#).
- You must download or resend the report within 24 hours from the time you generated the report instance. After 24 hours, Prism Central removes the report instance from the **Generated Reports** tab.

Procedure

1. Log in to Prism Central.
2. Select the **Intelligent Operations** application from **Application Switcher** function, and navigate to **Analysis** from the **Navigation Bar**.

The Analysis dashboard opens, displaying the session that you were last working on.

3. Click **Actions**.
4. Select **Download Report (csv, pdf)** and enter the following information:
 - a. **Report Name**: Enter the report name.
 - b. **Report Format**: Select PDF format or CSV format.
5. Click **Download**.

Prism Central displays a confirmation message on the dashboard screen that the report was generated.

What to do next

Select the **Intelligent Operations** application from **Application Switcher** function, and navigate to **Reports** from the **Navigation Bar**. Select **Generated Reports** tab to view and download the generated report instance. For more information, see [Downloading a Report Instance](#).

Emailing a Report

This task describes how to email the **Analysis** dashboard as a report.

About this task

You can email the **Analysis** dashboard reports as an attachment in PDF format, CSV format, or both. After you email the report, it is also available as a report instance under **Generated Reports** tab of the Reports page to resend or download the report.

Important: Ensure to resend or download the report within 24 hours from the time you generated the report instance. After 24 hours, Prism Central removes the report instance from the **Generated Reports** tab.

Procedure

1. Log in to Prism Central.
2. Select the **Intelligent Operations** application from **Application Switcher** function, and navigate to **Analysis** from the **Navigation Bar**.

The Analysis dashboard opens, displaying the session that you were last working on.
3. Click **Actions**.
4. Select **Email Report (csv, pdf)** and enter the following information:
 - a. **Recipients**: Enter the recipients (comma-separated) email address of the report.
 - b. **Subject**: Enter the subject of the email.
 - c. **Body**: Enter the desired text.
 - d. **Report Format**: Select the **PDF** checkbox or **CSV** checkbox or both to set a report format.
5. Click **Send**.

What to do next

Select the **Intelligent Operations** application from **Application Switcher** function, and navigate to **Reports** from the **Navigation Bar**. Select **Generated Reports** tab to view and manage the generated report instance. For more information, see [Managing a Report Instance from Generated Reports Tab](#).

Saving as a Report Configuration

This task describes how to save the **Analysis** dashboard data as a report configuration.

About this task

You can save the **Analysis** dashboard as a report configuration. The configuration of the report that you save is available under **Configurations** tab of the Reports page to perform report management actions.

Procedure

1. Log in to Prism Central.
2. Select the **Intelligent Operations** application from **Application Switcher** function, and navigate to **Analysis** from the **Navigation Bar**.

The Analysis dashboard opens, displaying the session that you were last working on.

3. Click **Actions**.
4. Select **Save as Report Configuration** and enter the following information:
 - **Report Configuration Name:** Enter the report configuration name.
5. Click **Save**.

What to do next

Select the **Intelligent Operations** application from **Application Switcher** function, and navigate to **Reports** from the **Navigation Bar**. Select **Configurations** tab to view and manage the saved report configuration. For more information, see [Managing a Custom Report Configuration](#).

Chart Management

You can perform the following tasks to manage charts from the **Analysis** dashboard.

- Add a chart (see [Adding a Chart](#) on page 25)
- Edit a chart (see [Editing a Chart](#) on page 27)
- Delete a chart (see [Deleting a Chart](#) on page 27)
- Clone a chart (see [Cloning a Chart](#) on page 28)
- Export a chart as a JSON or CSV file (see [Exporting a Chart](#) on page 28)
- Download a chart as a report (see [Downloading a Chart as a Report](#) on page 28)
- Email a chart as a report (see [Emailing a Chart as a Report](#) on page 29)
- Save a chart as a report configuration (see [Saving a Chart as a Report Configuration](#) on page 30)
- Save a chart as a report view (see [Saving a Chart as a Report View](#) on page 31)

The following sample figures show the options to manage charts from the **Analysis** dashboard.




Figure 12: Sample Chart - Chart Options




Figure 13: Sample Chart - Report Options

Adding a Chart

About this task

You can define a single chart to monitor a single metric for multiple entities of the same entity type. You can define a chart to monitor multiple metrics for a single entity.

Note:

You cannot define a single chart to monitor multiple metrics for multiple entities.

To add a chart in a session, perform the following steps:

Procedure

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher, and navigate to **Analysis** from the **Navigation Bar**.
The **Analysis** dashboard opens displaying the session that you were last working on.
3. Click one of the following:
 - » The **Add Chart** option at the top right of the dashboard.
 - » The **Add Chart** option at the bottom center of the dashboard.

4. In the **Add Chart** dialog box that opens, perform the following:

The screenshot shows the 'Add Chart' dialog box. It has a title bar with the text 'Add Chart' and a close button (X). Below the title bar, the instruction 'Select entitles and metrics to create a chart' is displayed. The main content area is divided into three sections: 'Entity', 'Metric', and 'Chart Name'. The 'Entity' section contains a 'Select entity type' dropdown menu and a 'Search entities' search bar. The 'Metric' section contains a 'Search metrics' search bar. The 'Chart Name' section contains a text input field with the placeholder text 'Enter chart name'. Below these sections is a red-outlined button labeled '+ Add Another Chart'. At the bottom right of the dialog are two buttons: 'Cancel' and 'Add'.

Figure 14: Add Chart dialog box

- Under **Entities**, select an entity type.
- In the **Search entities** bar, select the checkbox associated with the entities you want to choose.
- Under **Metric**, select the checkbox associated with the metrics you want to choose.

For the list of metrics you can select, see [Chart Metrics](#) on page 31.

Note:

In one chart, select metrics that use the same unit of measurement. Add a metric with a different unit of measurement to a separate chart.

See the [step to add another chart](#) below.

When adding a metric for multiple entities hosted on different hypervisors, select a metric only if it is supported on the hypervisors that are hosting the selected entities. Prism validates the list of available metrics provided for the **Metric** field. Prism removes from the list any metric that is not supported by all the hosting hypervisors for the selected entities. For example, for two entities hosted

on ESXi and AHV hypervisors respectively, the list includes only the metrics that are supported on both ESXi and AHV.

- d. Provide an appropriate **Chart Name** with appropriate keywords.

Prism provides a default name in the **Chart Name** field based on the entities and metrics you select in the preceding steps.

- e. To add another chart, click **Add Another Chart**.

Follow the instructions mentioned in Steps a through d, to define the chart.

- f. Click **Add** after adding all the charts.

5. You can also add any chart or metric from anywhere in Prism Central. For example, you can add a chart from the **Metrics** tab of a **Disk** details page. On the chart or metric widget, do one of the following:

- » Click **Actions** > **Add to Analysis** to add the metric to the current session in the **Analysis** dashboard.
- » Click **Actions** > **Select Analysis Session** to add the metric to any available session or a new session.

Editing a Chart

About this task

You can edit a chart in the **Analysis** dashboard to modify the entity or the metric you are monitoring with the chart. You can also modify the name of the chart.

Perform the following steps to edit a chart:

Procedure

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher, and navigate to **Analysis** from the **Navigation Bar**.
The **Analysis** dashboard opens displaying the session that you were last working on.
3. Click the three vertical dots on the top right corner of the chart that you want to modify.
4. Click **Edit**.
5. In the **Edit Chart** dialog box, make the necessary changes in the fields.
6. Click **Update**.
The changes are displayed in the chart on the **Analysis** dashboard.

Deleting a Chart

About this task

Perform the following steps to delete a chart:

Procedure

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher, and navigate to **Analysis** from the **Navigation Bar**.
The **Analysis** dashboard opens displaying the session that you were last working on.
3. Click the three vertical dots on the top right corner of the chart that you want to delete.

4. Click **Delete**.
5. Click **Confirm** to delete the chart widget.
Click **Cancel** to return to the **Analysis** dashboard without deleting the chart.

Cloning a Chart

About this task

Perform the following steps to clone a chart:

Procedure

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher, and navigate to **Analysis** from the **Navigation Bar**.
The **Analysis** dashboard opens displaying the session that you were last working on.
3. Click the three vertical dots on the top right corner of the chart that you want to clone.
4. Click **Clone**.
The cloned chart is the last chart in the row of charts in the **Analysis** dashboard.

Exporting a Chart

You can export a chart as a CSV or a JSON file.

About this task

Perform the following steps to export a chart:

Procedure

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher, and navigate to **Analysis** from the **Navigation Bar**.
The **Analysis** dashboard opens displaying the session that you were last working on.
3. Click the three vertical dots on the top right corner of the chart that you want to export.
4. Click one of the following:
 - » **Export CSV** to export as a CSV file.
 - » **Export JSON** to export as a JSON file.
5. Select **Save File**.
6. Browse to the location you want to save the file in the local folder and click **Save**.

Downloading a Chart as a Report

This task describes how to download an individual chart data as a report, either in PDF format or CSV format.

Before you begin

You have added the chart to the **Analysis** dashboard. For more information about chart management, see [Chart Management](#) on page 24.

About this task

Downloading a report has the following workflow:

- The report instance is generated in your preferred format and notified on the **Analysis** dashboard.
- The generated report instance is available for you to download under **Generated Reports** tab of the Reports page.

Important:

- You can also resend the generated report instance as an email from the **Generated Reports** tab of the Reports page. For more information, see [Resending a Report Instance](#).
- You must download or resend the report within 24 hours from the time you generated the report instance. After 24 hours, Prism Central removes the report instance from the **Generated Reports** tab.

Procedure

1. Log in to Prism Central.
2. Select the **Intelligent Operations** application from **Application Switcher** function, and navigate to **Analysis** from the **Navigation Bar**.

The Analysis dashboard opens, displaying the session that you were last working on.

3. Click **clip-board** icon  on the top right corner of the individual chart.
4. Select **Download Report (csv, pdf)** and enter the following information:
 - a. **Report Name:** Enter the report name.
 - b. **Report Format:** Select PDF format or CSV format.
5. Click **Download**.

Prism Central displays a confirmation message on the dashboard screen that the report was generated.

What to do next

Select the **Intelligent Operations** application from **Application Switcher** function, and navigate to **Reports** from the **Navigation Bar**. Select **Generated Reports** tab to view and download the generated report instance. For more information, see [Downloading a Report Instance](#).

Emailing a Chart as a Report

This task describes how to email an individual chart as a report.

Before you begin

You have added the chart to the **Analysis** dashboard. For more information about chart management, see [Chart Management](#) on page 24.

About this task

You can email an individual chart report as an attachment in PDF format, CSV format, or both. After you email the report, it is also available as a report instance under **Generated Reports** tab of the Reports page to resend or download the report.

Important: Ensure to resend or download the report within 24 hours from the time you generated the report instance. After 24 hours, Prism Central removes the report instance from the **Generated Reports** tab.

Procedure

1. Log in to Prism Central.
2. Select the **Intelligent Operations** application from **Application Switcher** function, and navigate to **Analysis** from the **Navigation Bar**.

The Analysis dashboard opens, displaying the session that you were last working on.

3. Click **clip-board** icon  on the top right corner of the individual chart.
4. Select **Email Report (csv, pdf)** and enter the following information:
 - a. **Recipients:** Enter the recipients (comma-separated) email address of the report.
 - b. **Subject:** Enter the subject of the email.
 - c. **Body:** Enter the desired text.
 - d. **Report Format:** Select the **PDF** checkbox or **CSV** checkbox or both to set a report format.
5. Click **Send**.

What to do next

Select the **Intelligent Operations** application from **Application Switcher** function, and navigate to **Reports** from the **Navigation Bar**. Select **Generated Reports** tab to view and manage the generated report instance. For more information, see [Managing a Report Instance from Generated Reports Tab](#).

Saving a Chart as a Report Configuration

This task describes how to save the individual chart as a report configuration.

Before you begin

You have added the chart to the **Analysis** dashboard. For more information about chart management, see [Chart Management](#) on page 24.

About this task

You can save the individual chart as a report configuration. The configuration of the report that you save is available under **Configurations** tab of the Reports page to perform report management actions.

Procedure

1. Log in to Prism Central.
2. Select the **Intelligent Operations** application from **Application Switcher** function, and navigate to **Analysis** from the **Navigation Bar**.

The Analysis dashboard opens, displaying the session that you were last working on.

3. Click **clip-board** icon  on the top right corner of the individual chart.
4. Select **Save as Report Configuration** and enter the following information:
 - **Report Configuration Name:** Enter the report configuration name.

5. Click **Save**.

What to do next

Select the **Intelligent Operations** application from **Application Switcher** function, and navigate to **Reports** from the **Navigation Bar**. Select **Configurations** tab to view and manage the saved report configuration. For more information, see [Managing a Custom Report Configuration](#).

Saving a Chart as a Report View

This task describes how to save the individual chart as a report view.

Before you begin

You have added the chart to the **Analysis** dashboard. For more information about chart management, see [Chart Management](#) on page 24.

About this task

You can save the individual chart as a report view. The report view you save is available under the **Saved** tab of the **ADD VIEWS** pane. Use the saved views when adding views to a new report or an existing report.

Procedure

1. Log in to Prism Central.
2. Select the **Intelligent Operations** application from **Application Switcher** function, and navigate to **Analysis** from the **Navigation Bar**.

The Analysis dashboard opens, displaying the session that you were last working on.

3. Click **clip-board** icon  on the top right corner of the individual chart.
4. Select **Save as Report View** and enter the following information:
 - **Report View Name:** Enter the report view name.

5. Click **Save**.

What to do next

Use the saved views while creating a report. For more information, see [Adding a Saved View](#).

Chart Metrics

The following metrics can be added to charts.

Note: The mapping between a metric and an entity type is hypervisor dependent.

Memory Usage	Entity Type(s)	Description
Metric		
Content Cache Hit Rate (%)	• Host • Cluster	Content cache hits over all lookups. ID: CONTENT_CACHE_HIT_PPM

Memory Usage Metric	Entity Type(s)	Description
Content Cache Hits	- Host - Cluster	Number of hits on the content cache. ID: <code>CONTENT_CACHE_NUM_HITS</code>
Content Cache Logical Memory Usage	- Host - Cluster	Logical memory (in bytes) used to cache data without deduplication. ID: <code>CONTENT_CACHE_LOGICAL_MEMORY_USAGE_BYTES</code>
Content Cache Logical SSD Usage	- Host - Cluster	Logical SSD memory (in bytes) used to cache data without deduplication. ID: <code>CONTENT_CACHE_LOGICAL_SSD_USAGE_BYTES</code>
Content Cache Lookups	- Host - Cluster	Number of lookups on the content cache. ID: <code>CONTENT_CACHE_NUM_LOOKUPS</code>
Content Cache Physical Memory Usage	- Host - Cluster	Real memory (in bytes) used to cache data by the content cache. ID: <code>CONTENT_CACHE_PHYSICAL_MEMORY_USAGE_BYTES</code>
Content Cache Reference Count	- Host - Cluster	Average number of content cache references. ID: <code>CONTENT_CACHE_NUM_DEDUP_REF_COUNT_PPH</code>
Content Cache SSD Usage	- Host - Cluster	Real SSD usage (in bytes) used to cache data by the content cache. ID: <code>CONTENT_CACHE_PHYSICAL_SSD_USAGE_BYTES</code>
Deduplication Fingerprints Cleared	- Host - Cluster	Number of written bytes for which fingerprints have been cleared. ID: <code>DEDUP_FINGERPRINT_CLEARED_BYTES</code>
Deduplication Fingerprints Written	- Host - Cluster	Number of written bytes for which fingerprints have been added. ID: <code>DEDUP_FINGERPRINT_ADDED_BYTES</code>
Disk I/O Bandwidth	- Host - Cluster - Disk - Storage Pool	Data transferred per second in KB/second from disk. ID: <code>STATS_BANDWIDTH</code>

Memory Usage	Entity Type(s)	Description
Metric		
Disk I/O Bandwidth - Read	• Host • Cluster • Disk • Storage Pool	Read data transferred per second in KB/second from disk. ID: STATS_READ_BANDWIDTH
Disk I/O Bandwidth - Write	• Host • Cluster • Disk • Storage Pool	Write data transferred per second in KB/second from disk. ID: STATS_WRITE_BANDWIDTH
Disk I/O Latency	• Host • Cluster • Disk • Storage Pool	I/O latency in milliseconds from disk. ID: STATS_AVG_IO_LATENCY
Disk IOPS	• Host • Cluster • Disk • Storage Pool	Input/Output operations per second from disk. ID: STATS_NUM_IOPS
Disk IOPS - Read	• Host • Cluster • Disk • Storage Pool	Input/Output read operations per second from disk. ID: STATS_NUM_READ_IOPS
Disk IOPS - Write	• Host • Cluster • Disk • Storage Pool	Input/Output write operations per second from disk. ID: STATS_NUM_WRITE_IOPS
GPU Framebuffer Usage	VM	Framebuffer usage in percentage. ID: FRAMEBUFFER_USAGE_PPM
Note: The VM entity is applicable for AHV hypervisor.		

Memory Usage Metric	Entity Type(s)	Description
GPU Usage	VM	GPU compute usage in percentage. ID: GPU_USAGE_PPM Note: The VM entity is applicable for AHV hypervisor.
GPU video decoder Usage	VM	GPU video decoder usage in percentage. ID: DECODER_USAGE_PPM Note: The VM entity is applicable for AHV hypervisor.
GPU video encoder usage	VM	GPU video encoder usage in percentage ID: ENCODER_USAGE_PPM Note: The VM entity is applicable to AHV hypervisor.
Hypervisor CPU Ready Time (%)	VM	Percentage of time that the VM was ready, but was not scheduled to run. ID: STATS_HYP_CPU_READY_TIME
Hypervisor CPU Usage (%)	- Host - Cluster - VM	Percent of CPU used by the hypervisor. ID: STATS_HYP_CPU_USAGE
Hypervisor I/O Bandwidth	- Host - Cluster - VM	Data transferred per second in KB/second from Hypervisor. ID: STATS_HYP_BANDWIDTH Note: VM entity is applicable for ESXi hypervisor
Hypervisor I/O Bandwidth - Read	- Host - Cluster - VM	Read data transferred per second in KB/second from Hypervisor. ID: STATS_HYP_READ_BANDWIDTH Note: VM entity is applicable for ESXi hypervisor.
Hypervisor I/O Bandwidth - Write	- Host - Cluster - VM	Write data transferred per second in KB/second from Hypervisor. ID: STATS_HYP_WRITE_BANDWIDTH Note: VM entity is applicable for ESXi hypervisor.

Memory Usage	Entity Type(s)	Description
Metric		
Hypervisor I/O Latency	- Host - Cluster - VM	I/O latency in milliseconds from Hypervisor. ID: STATS_HYP_AVG_IO_LATENCY Note: VM entity is applicable for ESXi hypervisor.
Hypervisor I/O Latency - Read	- Host - Cluster	I/O read latency in milliseconds from Hypervisor. ID: STATS_HYP_AVG_READ_IO_LATENCY Note: Cluster entity is applicable only for AHV hypervisor.
Hypervisor I/O Latency - Write	- Host - Cluster	I/O write latency in milliseconds from Hypervisor. ID: STATS_HYP_AVG_WRITE_IO_LATENCY Note: Cluster entity is applicable only for AHV hypervisor.
Hypervisor IOPS	- Host - Cluster - VM	Input/Output operations per second from Hypervisor. ID: STATS_HYP_NUM_IOPS Note: - Cluster entity is applicable for AHV hypervisor. - VM entity is applicable for ESXi.
Hypervisor IOPS - Read	- Host - Cluster - VM	Input/Output read operations per second from Hypervisor. ID: STATS_HYP_NUM_READ_IOPS Note: - Cluster entity is applicable for AHV hypervisor. - VM entity is applicable for ESXi.
Hypervisor IOPS - Write	- Host - Cluster - VM	Input/Output write operations per second from Hypervisor. ID: STATS_HYP_NUM_WRITE_IOPS Note: - Cluster entity is applicable for AHV hypervisor. - VM entity is applicable for ESXi.

Memory Usage Metric	Entity Type(s)	Description
Hypervisor Memory Usage (%)	- Host - Cluster - VM	Percent of memory used by the hypervisor. ID: STATS_HYP_MEMORY_USAGE
Logical Usage	Storage Container	Logical usage of storage (physical usage divided by replication factor). ID: STATS_UNTRANSFORMED_USAGE
Memory Usage (%)	- Host - Cluster - VM	Percentage of memory usage used by (any) hypervisor without high availability. ID: AGGREGATE_MEMORY_USAGE_PPM
Overall Memory Usage (%)	- Host - Cluster	Percentage of memory usage used by AHV with high availability. ID: OVERALL_MEMORY_USAGE_PPM
Network Rx Bytes	VM	Number of bytes received from the network that is reported by the hypervisor. ID: HYPERVISOR_NUM_RECEIVED_BYTES Note: VM entity is applicable for AHV hypervisor.
Network Tx Bytes	VM	Number of bytes transmitted to the network that is reported by the hypervisor. ID: HYPERVISOR_NUM_TRANSMITTED_BYTES Note: VM entity is applicable for AHV hypervisor.
Physical Usage	- Host - Cluster - Disk - Storage Pool - Storage Container	Actual usage of storage. ID: STATS_TRANSFORMED_USAGE
Read IOPS (%)	- Host - Cluster - Disk - Storage Pool	Percent of IOPS that are reads. ID: STATS_READ_IO_PPM

Memory Usage Metric	Entity Type(s)	Description
Replication Bandwidth - Received	- Cluster - Remote Site - Protection Domain - Replication Link	Replication data received per second in KB/second ID: STATS_REP_BW_RECEIVED
Replication Bandwidth - Transmitted	- Cluster - Remote Site - Protection Domain - Replication Link	Replication data transferred per second in KB/second ID: STATS_REP_BW_TRANSFERRED
Replication Bytes - Received	- Cluster - Remote Site - Protection Domain - Replication Link	Number of bytes received. ID: STATS_REP_NUM_RECEIVED_BYTES Note: Cluster entity is applicable for AHV hypervisor.
Replication Bytes - Total Received	Replication Link	Total number of bytes received. ID: STATS_REP_TOT_RECEIVED_BYTES
Replication Bytes - Total Transmitted	Replication Link	Total number of bytes transmitted. ID: STATS_REP_TOT_TRANSMITTED_BYTES
Replication Bytes - Transmitted	- Cluster - Remote Site - Protection Domain - Replication Link	Number of bytes transmitted. ID: STATS_REP_NUM_TRANSMITTED_BYTES Note: Cluster entity is applicable for AHV hypervisor.
Storage Controller Bandwidth	- Cluster - Storage Container - VM - Volume Group - Virtual Disk	Data transferred in KB/second from the Storage Controller. ID: STATS_CONTROLLER_BANDWIDTH

Memory Usage Metric	Entity Type(s)	Description
Storage Controller Bandwidth - Read	- Cluster - Storage Container - VM - Volume Group - Virtual Disk	Read data transferred in KB/second from the Storage Controller. ID: <code>STATS_CONTROLLER_READ_BANDWIDTH</code>
Storage Controller Bandwidth - Write	- Cluster - Storage Container - VM - Volume Group - Virtual Disk	Write data transferred in KB/second from the Storage Controller. ID: <code>STATS_CONTROLLER_WRITE_BANDWIDTH</code>
Storage Controller IOPS	- Host - Cluster - Storage Container - VM - Volume Group - Virtual Disk	Input/Output operations per second from the Storage Controller ID: <code>STATS_CONTROLLER_NUM_IOPS</code> Note: The Host entity is applicable for AHV hypervisor only.
Storage Controller IOPS - Read	- Host - Cluster - Storage Container - VM - Volume Group - Virtual Disk	Input/Output read operations per second from the Storage Controller ID: <code>STATS_CONTROLLER_NUM_READ_IOPS</code> Note: The Host entity is applicable for AHV hypervisor.
Storage Controller IOPS - Read (%)	- Cluster - Storage Container - VM - Volume Group - Virtual Disk	Percent of Storage Controller IOPS that are reads. ID: <code>STATS_CONTROLLER_READ_IO_PPM</code> Note: The Host entity is applicable for AHV hypervisor.

Memory Usage	Entity Type(s)	Description
Metric		
Storage Controller IOPS - Write	- Host - Cluster - Storage Container - VM - Volume Group - Virtual Disk	Input/Output write operations per second from the Storage Controller ID: STATS_CONTROLLER_NUM_WRITE_IOPS
Storage Controller IOPS - Write (%)	- Cluster - Storage Container - VM - Volume Group - Virtual Disk	Percent of Storage Controller IOPS that are writes. ID: STATS_CONTROLLER_WRITE_IO_PPM
Storage Controller Latency	- Cluster - Storage Container - VM - Volume Group - Virtual Disk	I/O latency in milliseconds from the Storage Controller. ID: STATS_CONTROLLER_AVG_IO_LATENCY Note: - Cluster entity is applicable only for AHV hypervisor. - VM entity is applicable for AHV and Hyper-V hypervisors.
Storage Controller Latency - Read	- Cluster - Storage Container - VM - Volume Group - Virtual Disk	Storage Controller read latency in milliseconds. ID: STATS_CONTROLLER_AVG_READ_IO_LATENCY Note: Cluster entity is applicable for AHV hypervisor.
Storage Controller Latency - Write	- Cluster - Storage Container - VM - Volume Group - Virtual Disk	Storage Controller write latency in milliseconds. ID: STATS_CONTROLLER_AVG_WRITE_IO_LATENCY Note: Cluster entity is applicable for AHV hypervisor.

Memory Usage Metric	Entity Type(s)	Description
Storage container own usage	Storage Container	Storage container's own usage + Reserved (not used). ID: <code>NEW_CONTAINER_OWN_USAGE_LOGICAL</code>
Swap In Rate	VM	Rate of data being swapped in. ID: <code>STATS_HYP_SWAP_IN_RATE</code> Note: VM entity is applicable for ESXi and Hyper-V hypervisors.
Swap Out Rate	VM	Rate of data being swapped out. ID: <code>STATS_HYP_SWAP_OUT_RATE</code> Note: VM entity is applicable for ESXi and Hyper-V hypervisors.
Virtual NIC bytes received packets with error.	VM	Virtual NIC bytes received packets with error. <code>STATS_NETWORK_ERROR_RECEIVED_PACKETS</code> Note: VM entity is applicable for ESXi hypervisor.
Virtual NIC bytes received rate.	VM	Virtual NIC bytes received rate in kbps. <code>STATS_NETWORK_RECEIVED_RATE</code> Note: VM entity is applicable only for ESXi hypervisor.
Virtual NIC bytes transmitted rate.	VM	Virtual NIC bytes transmitted rate in kbps. <code>STATS_NETWORK_TRANSMITTED_RATE</code> Note: VM entity is applicable for ESXi hypervisor.
Virtual NIC dropped transmitted packets.	VM	Number of dropped transmitted packets by the Virtual NIC. <code>STATS_NETWORK_DROPPED_TRANSMITTED_PACKETS</code> Note: VM entity is applicable for ESXi hypervisor.
Virtual NIC receive packets dropped.	VM	Number of receive packets dropped by the Virtual NIC. <code>STATS_NETWORK_DROPPED_RECEIVED_PACKETS</code> Note: VM entity is applicable for ESXi hypervisor.

Memory Usage Metric	Entity Type(s)	Description
Write IOPS (%)	- Host - Cluster - Disk - Storage Pool	Percent of IOPS that are writes. ID: STATS_WRITE_IO_PPM

Sessions Management

A session helps you correlate the metrics with the alerts and events for troubleshooting. You can create new sessions to troubleshoot new situations and preserve the old sessions created for old issues.

The **Sessions** page provides a list of sessions including the system-generated session. After an upgrade, the old sessions are also available in the **Sessions** page.

To access the **Sessions** page:

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher, and navigate to **Analysis** from the **Navigation Bar**.

The **Analysis** dashboard opens displaying the session that you were last working in as the default session.

3. Navigate to **Switch Sessions** > **View All Sessions** from the **Analysis** dashboard.

The **Sessions** page opens displaying a list of all the sessions including the system-generated session.




Figure 15: Sessions Page

Note: Prism Central automatically saves newly created sessions and all the changes in a session, except changes made using specific options that opens respective dialog boxes to facilitate the changes. For example, when you modify the session name or description in the **Sessions** dashboard using the **Edit Session Details** option, you must click **Update** to save the changes made in the dialog box. However, when you change the name of the session in the **Analysis** dashboard, then the change is automatically saved.

Note:

- the default system-generated Analysis session and the user-defined **Analysis** sessions are available only with a valid NCM license. For more information, see the [License Manager Guide](#).

- To optimize the display performance on the Analysis Dashboard, only the first 50 VMs are shown, so exporting this chart limits your view to 50 entities. To work around this limitation, go to the reports to generate the full data export.

The following table describes the fields that appear in the **Sessions** page.

Note: The fields vary based on the **View by** option. A dash (-) is displayed in a field when a value is not available or applicable.

Table 4: Sessions – Field Description

Field	Description
Session Name	Displays the name of the session.
Owner	Displays the name of the user who created the session.
Description	Displays a description of the session.
Created Time	Displays the date and time when the session was created.
Last Updated Time	Displays the date and time when the session was last updated.

You can perform the following actions for the sessions from the **Sessions** page:

- Click the name of a session to open the **Analysis** dashboard. For more information, see [Analysis Dashboard](#) on page 12.
- Create a session by clicking **Create Session**. For more information, see [Creating a New Session](#) on page 43.
- Set the default session view for the **Analysis** dashboard by clicking **Settings**. For more information, see [Setting the Default Session View](#) on page 44.
- Filter the sessions by name using the **Type text to filter by** Search bar. Type a word used in the session names that you want to filter in this search bar. Prism Central filters the list of all the sessions to provide a list of the sessions which have the word in their names.
- View subnets based on pre-defined criteria or create a custom view.
- Filter the sessions list based on a variety of parameter values using the **Filters** pane. For more information, see [Filters Pane - Sessions page](#).
- Perform the following session-specific actions on a session using the **Actions** dropdown menu. The **Actions** dropdown menu appears when a session is selected.

Table 5: Sessions Actions

Action	Description
Delete	Click this action to delete the session. For more information, see Deleting a Session on page 44.

Action	Description
Share	Click this action to share the session. For more information, see Sharing a Session on page 45.
Edit Sessions Details	Click this action to edit the name and description of the session. For more information, see Editing Session Details on page 45.

Filters Pane - Sessions page

You can filter the information in the **Sessions** page based on the following fields that are available in the **Filters** pane.

Table 6: Filter Pane Field Description - Sessions page

Field	Description	Values
Session Name	Filters based on the session name. It returns a list of sessions that satisfy the name condition/string.	(Session name string)
Description	Filters based on the description of the session. It returns a list of sessions that satisfy the description condition/string.	(Session description string)

Creating a New Session

You can create a new session from the **Analysis** dashboard or the **Sessions** page.

About this task

To create a new session, perform the following steps:

Procedure

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher, and navigate to **Analysis** from the **Navigation Bar**.
The **Analysis** dashboard opens displaying the session that you were last working on.
3. Navigate to **Switch Session > Create New Session**.

Note: To create a new session from the **Sessions** page, click **Create Session** in the **Sessions** page.

4. In the new session page displayed in the **Analysis** dashboard, edit the **Session Name** and add a **Description** help quickly identify the session.

Note:

- Adding a name for a session is mandatory. The default name is Untitled Session. Change the default name to a suitable name with appropriate keywords to facilitate filtering the sessions by session name.
- Adding a description for a session is not mandatory. However, to facilitate filtering the sessions by description, enter a description with appropriate keywords for the sessions you create.

To edit the **Session Name** and add a **Description**, see [Editing Session Details](#) on page 45.

What to do next

You can add charts to the new session.

Setting the Default Session View

You can set the default session view for the **Analysis** dashboard in the **Sessions** page.

About this task

To set the default session view, perform the following steps:

Procedure

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher, and navigate to **Analysis** from the **Navigation Bar**. The **Analysis** dashboard opens displaying the session that you were last working on.
3. Navigate to **Switch Sessions** > **View All Sessions** from the **Analysis** dashboard. The **Sessions** page opens displaying a list of all the sessions including the system-generated session.
4. Click **Settings**.
5. Select one of the following options:
 - » **Last working session**: This session is displayed by default in the **Analysis** dashboard. This is the last session that you worked on before signing out of Prism. This option is pre-selected.
 - » **System defined session**: This session is named **Analysis Session**. Prism creates and opens this session after installation of or upgrade to Prism Central 5.18 or later. Prism migrates and persists all the charts that you created before upgrading, to **Analysis Session** in Prism Central 5.18.

Prism also opens this session in the **Analysis** dashboard after Prism services are restarted.
6. Click **Save**.

Deleting a Session

You can delete a session from the **Sessions** page.

About this task

To delete a session, perform the following steps:

Procedure

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher, and navigate to **Analysis** from the **Navigation Bar**. The **Analysis** dashboard opens displaying the session that you were last working on.
3. Navigate to **Switch Sessions** > **View All Sessions** from the **Analysis** dashboard. The **Sessions** page opens displaying a list of all the sessions including the system-generated session.
4. Select the session you want to delete.
5. Click **Actions** > **Delete**.
6. In the confirmation dialog box that opens, click **Confirm** to delete the session.
Click **Cancel** to return to the **Sessions** dashboard without deleting the session.

Editing Session Details

You can edit the name and description of the session from the **Sessions** page.

About this task

To edit the session details, perform the following steps:

Procedure

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher, and navigate to **Analysis** from the **Navigation Bar**.
The **Analysis** dashboard opens displaying the session that you were last working on.
3. Navigate to **Switch Sessions** > **View All Sessions** from the **Analysis** dashboard.
The **Sessions** page opens displaying a list of all the sessions including the system-generated session.
4. Click **Actions** > **Edit Session Details**.
5. Enter an appropriate **Session Name** and **Description** to help quickly identify the session.
6. Click **Update**.

Sharing a Session

Any user with edit permission can share a session with other users and provide view or edit access to them.

About this task

To share a session, perform the following steps:

Procedure

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher, and navigate to **Analysis** from the **Navigation Bar**.
The **Analysis** dashboard opens displaying the session that you were last working on.
3. Navigate to **Switch Sessions** > **View All Sessions** from the **Analysis** dashboard.
The **Sessions** page opens displaying a list of all the sessions including the system-generated session.

4. Select a session, and click **Actions** > **Share**.




Figure 16: Sessions Dashboard

The **Sharing Session** window opens.




Figure 17: Sharing Session

5. In the **Add Users & Groups** tab, do the following:

- a. **Share with:** In the Search bar, type the name of the users or groups to share the session. The Search field displays a list of matching user or groups whose names match the letters that you type. Click the user name or group name to add. You can add more than one user or a group.
- b. **Permission Type:** You can provide view only or edit permissions to a user or group by selecting **View only** or **Editor** from the dropdown menu.

A user with **Editor** permission can perform these tasks:

- View a session
- Edit a session
- Share a session

- c. In the **Manage Users & Groups** tab, you can change the existing permission type for users or groups. To change the permission type, select the type from the dropdown menu associated with the user name or group name.

You can also remove the access for a user or group.

6. Click **Save**.

A message that the session is shared successfully appears.

After sharing the session, wait for few minutes for changes to take effect.

Users will not be able to view any information about entities in the analysis chart without entity permissions.

Closing a Session

About this task

To close a session, perform the following steps:

Procedure

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher, and navigate to **Analysis** from the **Navigation Bar**.
The **Analysis** dashboard opens displaying the session that you were last working on.
3. Navigate to **Switch Sessions** > **View All Sessions** from the **Analysis** dashboard.
The **Sessions** page opens displaying a list of all the sessions including the system-generated session.
4. Click **Actions** > **Close**.

Switching Between Sessions

About this task

To switch between sessions, perform the following steps:

Procedure

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher, and navigate to **Analysis** from the **Navigation Bar**.
The **Analysis** dashboard opens displaying the session that you were last working on.

3. Click **Switch Session**.
4. From the dropdown menu, select the session you want to view.
The selected session opens in the **Analysis** dashboard.

If you do not find the session you want to switch to in the dropdown menu, perform the following steps:

5. Click **Switch Session > View All Sessions** to go to the **Sessions** dashboard.
6. Select the session you want to open in the **Analysis** dashboard.
The selected session opens in the **Analysis** dashboard.

APPLICATION DISCOVERY

Application discovery allows you to monitor applications in Nutanix, non-Nutanix (vCenter managed), or mixed (cluster with mixed hypervisors) environment. It provides the following capabilities:

- Discover applications in a specified set of clusters.
- Identify the VM in which each application runs and what ports it uses.
- Publish the data to an internal database which is accessible through an external API call.

Note: When you add a mixed cluster for application discovery, only ESXi hypervisor nodes are enabled for discovery.

Application discovery sends data about the VM network traffic stats to a SaaS-like service in the cloud for analysis. The information sent to the cloud includes network flow data (source IP, destination IP, destination port, protocol, packet count, byte count) and VM data (cluster UUID, source VM name, UUID, VLAN ID, and IP addresses).

- To configure application discovery:
 - enable application discovery (see [Enabling Application Discovery](#) on page 60)
 - connect to the cloud-based discovery service (see [Connecting to the Cloud](#) on page 65)
 - add clusters to monitor (see [Adding or Removing Clusters to Monitor](#) on page 66)
 - authenticate with vCenter for ESXi cluster access (see [Configuring vCenter Authentication](#) on page 68)
- To generate current findings at any time:
 - start an application discovery (see [Discovering Applications](#) on page 70)
 - publish data for identified applications (see [Publishing Applications](#) on page 71)
 - create policies for unidentified applications (see [Creating Discovery Policies](#) on page 71)
- To view discovery findings, see the application discovery dashboard (see [Application Discovery Summary View](#) on page 49).

Application Discovery Summary View

The application discovery dashboard allows you to view information about discovered applications and associated policies.

To access the **Discovery** dashboard:

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher. The system displays the **Intelligent Operations** dashboard.
3. Navigate to **Discovery** from the **Navigation Bar**.

The **Discovery** dashboard opens displaying the **Discovered Apps** tab.

The **Discovery** dashboard has the **Summary**, **Discovered Apps**, **Published Apps**, **Policies**, **App Relationships**, **Monitored Clusters**, and **Alerts** tabs.

Note:

- This section describes the information and options that appear in the **Discovery** dashboard.

- If application discovery is not enabled, an enablement page appears instead of the dashboard. For more information, see [Enabling Application Discovery](#) on page 60.

Summary Tab

The **Summary** tab displays the following four widgets:

Widgets	Description
Discovery Highlights	Displays the total number of application instances, the number of identified (known application to port mapping) and unidentified instances, the number of application types, and the number of ports.
Status & Settings	Displays the number of monitored clusters, the number of discovery policies, and the status (active or inactive) of the discovery process, vCenter authentication, vCenter monitoring, and cloud connection. Click View Details to see the vCenter authentication and monitoring details (see Configuring vCenter Authentication on page 68) or click Manage to see the cloud connection details (see Connecting to the Cloud on page 65). If vCenter monitoring is not configured, an Add vCenter link appears (instead of the View Details link); click the link to add clusters from any configured vCenter instances. For more information, see Adding or Removing Clusters to Monitor on page 66.
Top Discovered Apps	Lists the applications with the most instances.
Top Discovered Ports	Lists the ports with the most instances.




Figure 18: App Discovery Summary Tab

Discovered Apps Tab

The **Discovered Apps** tab displays a list of the discovered applications.




Figure 19: Discovered Apps Tab

The following table describes the fields that appear in the **Discovered Apps** page.

Note: The fields vary based on the **View by** and **Group by** options. A dash (-) is displayed in a field when a value is not available or applicable.

Table 7: Discovered Apps – Field Description

Field	Description	Values
App Instance	Displays the application name for the instance. Click the name to display the Application Info window. The window displays the application name, VM name, TCP ports, and UDP ports.	(app name)
VM	Displays the VM name. Click the name to display the details page for the VM. For more information, see VM Details View in <i>Prism Central Infrastructure Guide</i> .	(VM name)
TCP Ports	Displays one or more TCP ports used by the application.	(port numbers)
UDP Ports	Displays one or more UDP ports used by the application.	(port numbers)

Field	Description	Values
Cluster	Displays the name of the cluster in which the application runs. Click the name to view the details for that cluster. For more information, see Cluster Details View topic in <i>Prism Central Infrastructure Guide</i> . For information about authenticating application discovery on the cluster with vCenter, see Configuring vCenter Authentication on page 68.	(cluster name)
Environment	Displays whether the cluster resides in a Nutanix or non-Nutanix (vCenter managed) environment.	Nutanix, vCenter
Status	Indicates whether the instance is an identified application type and whether it is newly discovered (<i>new app</i> meaning discovered during the most recent pass) or discovered previously (<i>existing app</i> meaning discovered in an earlier pass). An <i>identified</i> application is one that is a known type. The discovery service can classify many applications, but for those it cannot identify, you can assign them to a type through a custom discovery policy. For more information, see Creating Discovery Policies on page 71.	Identified, Unidentified

You can perform the following actions from the **Discovered Apps** page:

- Click the name of a discovered application to view the detailed information about the application.
- Start a discovery search by clicking **Discover**. For more information, see [Discovering Applications](#) on page 70.
- Publish the application information by clicking **Publish**. For more information, see [Publishing Applications](#) on page 71.
- Download the contents of the fields that appear in the **Discovered Apps** page in CSV format by clicking the **Export** link.
- Group the application instances based on pre-defined criteria.
- View the application instances based on pre-defined criteria or create a custom view.
- Perform the following actions on a single or multiple application instances using the **Actions** dropdown menu. The **Actions** dropdown appears when one or more application instances are selected.

Table 8: Discovered Apps Actions

Action	Description
Identify	Click this action to set the selected instances to Identify status.
Unidentify	Click this action to set the selected instances to Unidentify status.
Ignore	Click this action to set the selected instances to Ignore status.
Unignore	Click this action to set the selected instances to Unignore status.

You can filter the information in the **Discovered Apps** page based on the following fields that are available in the **Filters** pane.

Table 9: Filter Pane Field Description - Discovered Apps Page

Field	Description	Values
Status	Filters by whether the instance has been identified as a known application type.	Identified, Unidentified
App State	Filters based on the state of the application.	New, Existing
Ignored	Filters for ignored applications.	Show Ignored Apps (checkbox)
App Instance	Filters based on the application instance name. Select a condition from the dropdown list (Contains , Doesn't contain , Starts with , Ends with , or Equal to) and enter a string in the field.	(app instance name string)
Cluster	Filters based on the cluster name. Select a condition from the dropdown list (Contains , Doesn't contain , Starts with , Ends with , or Equal to) and enter a string in the field.	(cluster name string)
Environment	Filters based on whether the applications are in a Nutanix or non-Nutanix (vCenter) managed environment.	Nutanix, vCenter
VM	Filters based on the VM name (same options as App Instance).	(VM name string)
TCP Ports	Filters based on the TCP port numbers. Enter a range in the From <low> to <high> field.	([xx] to [yy] range)
UDP Ports	Filters based on the UDP port numbers. Enter a range in the From <low> to <high> field.	([xx] to [yy] range)

Published Apps Tab

The **Published Apps** tab displays a list of the published applications. A *published* application is one that is made available to the API. Publishing involves copying the data for the identified (and not ignored) applications into an internal database that you can access through an external API call. An *identified* application is one where the VM in which it runs is known and the used ports are mapped.




Figure 20: Published Apps Tab

The fields that appear in the **Published Apps** page are the same as **Discovered Apps** page, except there is no *Status* field. For more information about the fields, see the *Discovered Apps – Field Description* table in [Discovered Apps Tab](#).

You can perform the following actions from the **Published Apps** page:

- Click the name of a published application to view the detailed information about the application.
- Copy the API link to your clipboard by clicking **Get API Link**. The link is the access point to consume the published applications through the API.
- Delete an application instance using the **Delete** option in the **Actions** dropdown menu. The **Actions** dropdown appears when one or more application instances are selected.
- Download the contents of the fields that appear in the **Published Apps** page in CSV format by clicking the **Export** link.
- Group the application instances based on pre-defined criteria.
- View the application instances based on pre-defined criteria or create a custom view.
- Filter the information in the **Published Apps** page based on the following fields that are available in the **Filters** pane.

Table 10: Filter Pane Field Description - Published Apps Page

Field	Description	Values
App Instance	Filters based on the application instance name. Select a condition from the dropdown list (Contains , Doesn't contain , Starts with , Ends with , or Equal to) and enter a string in the field.	(app instance name string)
Cluster	Filters based on the cluster name. Select a condition from the dropdown list (Contains , Doesn't contain , Starts with , Ends with , or Equal to) and enter a string in the field.	(cluster name string)
Environment	Filters based on whether the applications are in a Nutanix or non-Nutanix (vCenter) managed environment.	Nutanix, vCenter

Field	Description	Values
VM	Filters based on the VM name (same options as App Instance).	(VM name string)

Policies Tab

The **Policies** tab displays a list of application discovery policies. The policies define the application types.

Note: Existing discovered applications are remapped whenever there is a policy action change, based on the existing VM traffic statistics.




Figure 21: Policies Tab

The following table describes the fields that appear in the **Policies** page.

Note: The fields vary based on the **View by** option. A dash (-) is displayed in a field when a value is not available or applicable.

Table 11: Policies – Field Description

Field	Description	Values
Policy Name	Displays the name of the policy. Click the name to display the Edit Discovery Policy window for user-defined policies or the View Discovery Policy window for system-defined policies. For more information, see Modifying or Viewing Discovery Policies on page 72.	(policy name)

Field	Description	Values
Application	Displays the application type. Application discovery can identify various common applications by default (referred to as system-defined applications) such as DNS, NTP, SSH, Web Server, MS Exchange Server, MS SQL Server, Oracle DB, SAP HANA. For user-defined applications, the user provides the name. See Creating Discovery Policies on page 71.	(application type)
TCP Ports	Displays one or more TCP ports used by the application.	(port numbers)
UDP Ports	Displays one or more UDP ports used by the application.	(port numbers)
VMs	Displays the scope of the policy, that is to which VMs the policy applies. A value of "All" indicates it applies to all VMs. One or more names indicates that the policy applies to the named VMs only.	All, (VM names)
Status	Displays the current operational status of the policy.	Enabled, Disabled
Description	Displays the description of the policy.	(policy description)
Rule Type	Displays the rule type that defines the policy.	(System, User)

You can perform the following actions from the **Policies** page:

- Click the name of a policy to view the detailed information about the policy.
- Create a custom discovery policy by clicking **New Policy**. For more information, see [Creating Discovery Policies](#) on page 71.
- Download the contents of the fields that appear in the **Policies** page in CSV format by clicking the **Export** link.
- View the policies based on pre-defined criteria or create a custom view.
- Perform the following actions on a single or multiple policies using the **Actions** dropdown menu. The **Actions** dropdown appears when one or more policies are selected.

Table 12: Policies Actions

Action	Description
Enable	Click this action to enable the selected policy.
Disable	Click this action to disable the selected policy.
Delete	Click this action to delete the selected policy.

Note: You can delete user-defined policies (not system policies) only.

You can filter the information in the **Policies** page based on the following fields that are available in the **Filters** pane.

Table 13: Filter Pane Field Description - Policies Page

Field	Description	Values
Policy Name	Filters based on the policy name. Select the checkbox, choose a condition from the dropdown list (Contains , Doesn't contain , Starts with , Ends with , or Equal to), and enter a string in the field.	(policy name string)
Status	Filters based on whether the policy is enabled or disabled.	(Enabled, Disabled)

App Relationships Tab

The **App Relationships** tab displays the top 20 applications based on the number of entities that are communicating with these applications. The displayed information includes the VMs where these applications are running, along with the number of entities that are communicating with these VMs.




Figure 22: App Relationships Tab

The following table describes the fields that appear in the **App Relationships** page.

Note: The fields vary based on the **View by** option. A dash (-) is displayed in a field when a value is not available or applicable.

Table 14: App Relationships – Field Description

Field	Description	Values
App Name	Displays the application name.	(app name)

Field	Description	Values
VM	Displays the VM name. Click the name to display the details page for the VM. For more information, see VM Details View in <i>Prism Central Infrastructure Guide</i> .	(VM name)
IP Address	Displays the IP Address of the VM.	(IP Address)
Ports	Displays the port used by the application.	(port numbers)
Dependent Entities	Displays the number of entities that are communicating with the application.	(number of dependent entities)
Cluster	Displays the name of the cluster in which the application runs. Click the name to view the details for that cluster. For more information, see Cluster Details View topic in <i>Prism Central Infrastructure Guide</i> . For information about authenticating application discovery on the clusters with vCenter, see Configuring vCenter Authentication on page 68.	(cluster name)
Identification Status	Indicates whether the application is an identified application type. An identified application is one that is a known type per application discovery service. You can use custom discovery policy to assign a type to the unidentified applications listed under App Relationships tab. If an application is displayed as unidentified regardless of it belonging to the known types, try refreshing the list of discovered applications. You can refresh the list by running the Discover operation from Discovered Apps tab.	Identified, Unidentified

You can filter the information in the **App Relationships** page based on the following fields that are available in the **Filters** pane.

Note: When you use filters for information displayed on the App Relationships tab, the filters fetch the latest discovery and communication data and accordingly display the result set.

Table 15: Filter Pane Field Description - App Relationships Page

Field	Description	Values
App Name	Filters based on the application name. Select a condition from the dropdown list (Contains, Equal to, Not equal to, Doesn't contain, Starts with, Ends with) and enter a string in the field.	(app name string)
VM	Filters based on the VM name. (same options as App Name).	(VM name string)
Cluster	Filters based on the cluster name (same options as App Name).	(cluster name string)

Field	Description	Values
IP Address	Filters based on the IP address (same options as App Name).	(IP address string)
TCP Port	Filters based on the TCP port number.	(TCP port number)
UDP Port	Filters based on the UDP port number.	(UDP port number)
Time Range	Filters based on the time range.	Last 30 Mins, Last 1 Hour, Last 12 Hours, and Last 24 Hours

To download the contents of the fields that appear in the **App Relationships** page in CSV format, click **Export**.

Monitored Clusters Tab

The **Monitored Clusters** tab displays the list of clusters that application discovery monitors when it does a discovery search.




Figure 23: Monitored Clusters Tab

The following table describes the fields that appear in the **Monitored Clusters** page.

Note: The fields vary based on the **View by** option. A dash (-) is displayed in a field when a value is not available or applicable.

Table 16: Monitored Clusters – Field Description

Field	Description	Values
Cluster Name	Displays the cluster name. Click the name to displays the details page for the cluster. See Cluster Details View topic in <i>Prism Central Infrastructure Guide</i> .	(cluster name)
# of Nodes	Displays the number of nodes in the cluster.	(integer)
# of VMs	Displays the number of VMs in the cluster.	(integer)

Field	Description	Values
vCenter IP	Displays the IP address of the vCenter for this cluster.	(IP address)
Environment	Displays whether the cluster resides in a Nutanix or non-Nutanix (vCenter managed) environment.	Nutanix, vCenter
Monitored Hypervisor	Displays the monitored hypervisor in the cluster.	ESXi, AHV, vCenter
Health Status	Displays the current VM network traffic monitoring status.	Active, Inactive
License	Displays the type of the Prism license in the cluster.	(Prism license name)

You can perform the following actions from the **Monitored Clusters** page:

- Click the name of a cluster to open the cluster details page, which displays the detailed information about the cluster. [Cluster Details View](#) topic in *Prism Central Infrastructure Guide*.
- Add a cluster for monitoring by clicking **Add Cluster**. For more information, see [Adding or Removing Clusters to Monitor](#) on page 66.
- Authenticate application discovery on the monitored clusters with vCenter by clicking **vCenter Authentication**. For more information, see [Configuring vCenter Authentication](#) on page 68.
- Remove a cluster from application discovery monitoring by selecting the cluster and clicking **Remove**.
- Download the contents of the fields that appear in the **Monitored Clusters** page in CSV format by clicking the **Export** link.
- View the clusters based on pre-defined criteria or create a custom view.

Alerts Tab

The **Alerts** tab displays a list of alerts. This tab provides the same features and options as the *Alerts dashboard (List tab)*, except it is filtered to display just application discovery-related alerts. For more information, see the *List Tab* section in [Alerts Summary View \(Prism Central\)](#) of the *Prism Central Alerts and Events Reference Guide*.

Enabling Application Discovery

Enabling application discovery (disabled by default) involves connecting to the discovery service in the cloud, specifying clusters to monitor, and authenticating to the vCenter instances for the specified clusters.

About this task

When you select application discovery for the first time, a wizard appears to guide you through the setup process. To configure application discovery for the first time, do the following:

Before you begin

Ensure that all the prerequisites specified in [Prerequisites for Enabling Application Discovery](#) on page 64 are met before you enable application discovery.

Procedure

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher. The system displays the **Intelligent Operations** dashboard.

3. Navigate to **Discovery** from the **Navigation Bar**.
If application discovery is not enabled, a **Welcome to App Discovery** page appears.
4. In the **Welcome to App Discovery** page, select the **I have increased the PC VM's memory and vCPUs** checkbox (as specified in the prerequisites) and click **Enable App Discovery**.
The **App Discovery Setup** window opens.
5. In the **Cloud Connection** page, enter your API key and key authentication values in the indicated fields and click **Next**.

App Discovery Setup

1 Cloud Connection 2 Select Clusters

✓ Enablement successful! Now, complete the setup to discover apps.

The discovered network flow data and VM metadata will be sent to the cloud for analysis. [Learn more](#)

Provide your API Key

An API Key is required to connect to the cloud. Your keys can be managed from the API Keys page on the Nutanix Support Portal. [Get API Key and Key ID](#)

API Key

Key ID

Cancel Next

Figure 24: App Discovery Setup: Cloud Connection screen

6. In the **Select Clusters** page, perform the following:
- » If the clusters to be selected are in a Nutanix environment, continue with this step.
 - » If the clusters to be selected are in a non-Nutanix (vCenter managed) or mixed environment, click the **Configure vCenter Monitoring** link and then skip to Step 7.




Figure 25: App Discovery Setup: Select Clusters (initial)

- a. Select the checkboxes associated with the target clusters from the dropdown list.
- b. Enter the vCenter administrator user name and password in the respective fields.

Note: Ensure that the administrator has write permission to configure Netflow collector settings.

An entry for each registered vCenter appears here. For more information, see [Registering vCenter Server \(Prism Central\)](#) in the *Prism Central Infrastructure Guide*.

- c. Click **Discover Apps**.

This starts a discovery search. When the search is complete, the **App Discovery** window and the dashboard appears (see [Application Discovery Summary View](#) on page 49). This completes the setup for a Nutanix environment.




Figure 26: App Discovery Setup: Select Clusters

7. [vCenter only] Select **vCenter** in the **External Entity Type** field and then click **Enable**.

8. [vCenter only] Do the following in the indicated fields to configure a vCenter instance:
 - a. **External Entity Type:** Select **vCenter** from the dropdown menu.
 - b. **vCenter Host :** Enter the IP address of the vCenter host VM.

If the IP address does not resolve automatically to the proper port number, enter the address in the form `ip_address:port#`. The standard vCenter Server port is 443.
 - c. **Username:** Enter the vCenter administrator user name.
 - d. **Password:** Enter the vCenter administrator password.
 - e. **SSL Verify:** Select the SSL mode to use (**Verify** or **Don't Verify**) when connecting to the target.

This setting determines whether to verify (or not verify) the target's SSL certificate as part of the connection process.

Configure external entity for monitoring

Configure a vCenter for monitoring in Prism. Then, discover apps on ESXi clusters not running Nutanix AOS.

External Entity Type
vCenter

vCenter Host

User Name

Password

SSL Verify (i)
Don't Verify

Cancel Test Connection Next

Figure 27: Configure vCenter Instance

9. [vCenter only] Click **Test Connection** to verify (test) the connection.

A *Test connection success* message appears after a successful test. After a successful test, click **Next**.
10. [vCenter only] Select the clusters to monitor.

A list of discovered clusters for the vCenter instance appears. Select (check the boxes for) the clusters you want to monitor from the list and then click **Save**. The instance configuration and data collection begins. When it completes (which can take a few minutes), a *vCenter Configuration successful* message appears.
11. [vCenter only] After the vCenter configuration is complete, click **Back to App Discovery**.

The **Back to App Discovery** takes you to the monitored clusters page and opens a cluster selection window. The vCenter clusters specified in step 8 appear in the list for you to select.

Prerequisites for Enabling Application Discovery

Ensure that the following prerequisites are met so that you can successfully enable application discovery.

- **Sufficient Prism Central VM resources:** Application discovery is one of the NCM Pro license features that requires additional Prism Central VM resources (more memory and vCPUs). The requirements are listed in [KB 8932](#). If your Prism Central VM resources are below the recommended amount, manually increase the memory and vCPUs to the recommended amount.
- **Internet connection:** An internet connection is required to connect to the Nutanix cloud.
- Up and running firewall in Prism Central.
- **VDS enabled on vCenter:** This requires a vSphere version (Enterprise Plus license) that supports vSphere Distributed Switch (VDS) and netflow version 10.
- My Nutanix account that is either admin role or account admin.
- **API key and key ID:** Log in to the Nutanix support portal and get your API key information. For information about how to generate a new key or view information about an existing key, see the [License Manager Guide](#).

Each Prism Central instance requires a unique API key/ID pair. If you have multiple Prism Central instances (or a key was revoked for any reason), create additional API key/ID pairs as needed so you have a unique and valid pair for each Prism Central instance.




Figure 28: Create API Key Flow (example)

Connecting to the Cloud

Using application discovery requires connecting to a cloud-based discovery service.

About this task

Nutanix provides a SaaS-like service in the cloud that analyzes application data from the monitored clusters. The connection is established during initial setup (see [Enabling Application Discovery](#) on page 60), but to update the connection after setup, do the following:

Before you begin

Log in to the Nutanix support portal and get your API key information. For information about how to generate a new key or view information about an existing key, see the [License Manager Guide](#).

Procedure

1. Log in to Prism Central.

2. Select **Intelligent Operations** from application switcher. The system displays the **Intelligent Operations** dashboard.
3. Navigate to **Discovery** from the **Navigation Bar**.

The **Discovery** dashboard opens displaying the **Discovered Apps** tab.

4. Click the **Summary** tab, and click **Manage** associated with the **Cloud Connection** field of the **Status & Settings** widget.
5. In the **Cloud Connection** window, perform the following in the indicated fields:
 - a. **API Key**: Enter the API key value.
 - b. **Key ID**: Enter the key identification number.
 - c. Click **Save**.

Only one API key/ID pair is allowed per Prism Central instance. If you have multiple Prism Central instances (or your initial key was revoked for any reason), create additional API key/ID pairs as needed (through the Nutanix support portal) so you have a unique and valid pair for each Prism Central instance.

Adding or Removing Clusters to Monitor

Using application discovery requires first specifying the clusters to monitor.

About this task

Specifying clusters to monitor is done during initial setup (see [Enabling Application Discovery](#) on page 60), but you can add or remove clusters after setup at any time.

Procedure

- To add a cluster to the monitored list, do the following:
 - a. Log in to Prism Central.
 - b. Select **Intelligent Operations** from application switcher. The system displays the **Intelligent Operations** dashboard.
 - c. Navigate to **Discovery** from the **Navigation Bar**.

The **Discovery** dashboard opens displaying the **Discovered Apps** tab.
 - d. Click the **Monitored Clusters** tab and then click the **Add Cluster**.
 - e. In the **Add Cluster to App Discovery** window, select the checkboxes associated with the target clusters from the dropdown list.

If you add a cluster from a vCenter server, you must enter the vCenter administrator user name and password in the respective fields. The administrator must have write permission.

Click **Add**. Network data collection begins immediately and may take some time to complete when multiple clusters are added.

Add Clusters to App Discovery
?
X

For clusters with mixed hypervisors, only ESXi hypervisors are monitored for application discovery.
X

Select from 2 applicable clusters

Note: App discovery is supported for AHV clusters with AHV 7.3 and above with AOS 6.1 and above

Clusters : 1 Selected

auto_cluster_prod
X

You have selected clusters from a vCenter that requires authentication. Please provide write permissions. This will be updated centrally for all features connecting to this vCenter.
X

vCenter 10.46.1.215

User Name

Password
Show

★
Want to discover apps on ESXi clusters that are not running Nutanix AOS?

You'll need visibility into vCenter infra first.

Try "Monitoring Integrations", a Prism Pro feature.

[Configure vCenter Monitoring](#)

Close
Add

Figure 29: Add Clusters to App Discovery Window

- f. Run discovery to identify applications on the new clusters. For more information, see [Discovering Applications](#) on page 70.
- To remove a cluster from the monitored list, select the target cluster(s) in the list and then click **Remove**.

Note: Removing a cluster stops application monitoring for that cluster. You can effectively disable application monitoring by removing all clusters, which means no data is sent between Prism Central and the Nutanix cloud.

Configuring vCenter Authentication

Using application discovery with ESXi clusters requires first authenticating with the associated vCenter.

About this task

Authenticating with vCenter is done during the initial setup (see [Enabling Application Discovery](#) on page 60), but to update authentication after setup, do the following:

Procedure

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher. The system displays the **Intelligent Operations** dashboard.
3. Navigate to **Discovery** from the **Navigation Bar**.
The **Discovery** dashboard opens displaying the **Discovered Apps** tab.
4. Click the **Monitored Clusters** tab, and click **vCenter Authentication**.
5. In the **vCenter Authentication** window, perform the following in the indicated fields:
 - a. **User Name:** Enter the user name for a vCenter administrator or any user with write permission to the clusters.
 - b. **Password:** Enter the user password.
 - c. Click **Save**.

An entry for each registered vCenter appears here. For more information, see [Registering vCenter Server \(Prism Central\)](#) in *Prism Central Infrastructure Guide*. The authentication and monitoring status for each

vCenter is displayed. If a vCenter is not monitored currently, you can click **Configure** to enable monitoring. For more information, see [Configuring Application Monitoring](#) on page 86.

The screenshot shows a window titled "vCenter Authentication". It contains two sections for different vCenter instances. The first section is for vCenter 10.46.1.213, showing "Authentication: Active" and "Monitoring: Active". The second section is for vCenter 10.46.1.214, showing "Authentication: Active" and "Monitoring: Not Configured" with a blue "Configure" button. Both sections have input fields for "User Name" (containing "administrator@vsphere.local") and "Password" (masked with dots). At the bottom right are "Close" and "Save" buttons.

Figure 30: vCenter Authentication Window

Discovering Applications

You can check for the current applications on the monitored clusters at any time.

About this task

To start an application discovery search, do the following:

Procedure

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher. The system displays the **Intelligent Operations** dashboard.
3. Navigate to **Discovery** from the **Navigation Bar**.

The **Discovery** dashboard opens displaying the **Discovered Apps** tab.

4. Click **Discover**.

The search begins immediately and may take a few minutes to complete depending on the number and size of the monitored clusters. When the search is complete, the **App Discovery** window appears. The window displays the search status (successful, unsuccessful); number of discovered, identified, and unidentified applications; number of application types; and number of new application instances.

5. Click **View App Instances** to redisplay the tab.

If some applications are misclassified, you can modify the results by manually identifying, unidentifying, or ignoring (hiding) selected applications. For more information, see the *Discovered Apps Actions* table in [Application Discovery Summary View](#) on page 49.

Publishing Applications

You can publish the application discovery findings to a database.

About this task

Applications are discovered based on discovery policies that map applications and ports. You may validate, fix, and sanitize the list of discovered applications, and then publish the final list of discovered application. Publishing means copying the VM and port-mapping data for the identified applications to an internal database, which is accessible through an external API call. To publish data for the current identified applications, do the following:

Procedure

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher. The system displays the **Intelligent Operations** dashboard.
3. Navigate to **Discovery** from the **Navigation Bar**.

The **Discovery** dashboard opens displaying the **Discovered Apps** tab.

4. Click **Publish**.
Information for the identified applications is published immediately.
5. Go to the **Published Apps** tab to view the list of published applications.

You can delete applications from the published list that you do not want in the database. For more information, see [Published Apps Tab](#) information.

Creating Discovery Policies

You can define custom discovery policies for unidentified applications.

About this task

Application discovery can identify a variety of common (system defined) application types such as DNS, NTP, SSH, Web Server, MS Exchange Server, MS SQL Server, Oracle DB, and SAP HANA. However, it cannot determine the type for some applications. To create a custom (user defined) discovery policy for an unidentified application, do the following:

Procedure

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher. The system displays the **Intelligent Operations** dashboard.
3. Navigate to **Discovery** from the **Navigation Bar**.

The **Discovery** dashboard opens displaying the **Discovered Apps** tab.

4. Click the **Policies** tab and click **New Policy**.

5. In the **New Discovery Policy** window, do the following in the indicated fields:

- a. **Application:** Enter the application name.
- b. **Policy Name:** Enter a name for the policy.
- c. **Description:** Enter a description for the policy (optional).
- d. **VMS:** Choose either of the following:
 - » If you want the policy to apply to all VMs, select **All VMs**.
 - » If you want the policy to apply to a limited set of VMs, select **Specific VMs** and then select the target VMs from the dropdown list.
- e. **TCP Ports:** Enter the TCP port number(s) used by the application. Enter a comma separated list for multiple port numbers.
- f. **UDP Ports:** Enter the UDP port number(s) used by the application. Enter a comma separated list for multiple port numbers.
- g. Click **Save and Apply**.

Information about existing discovered applications is refreshed, reflecting any changes from the new policy.

Modifying or Viewing Discovery Policies

You can view any discovery policy, enable or disable any policy, or modify a user-defined policy at any time.

About this task

To view or modify a policy, do the following:

Procedure

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher. The system displays the **Intelligent Operations** dashboard.
3. Navigate to **Discovery** from the **Navigation Bar**.

The **Discovery** dashboard opens displaying the **Discovered Apps** tab.

4. Click the **Policies** tab, and click the policy you want to view or modify from the list.

If it is a system policy, the **View Discovery Policy** window appears. If it is a user-defined policy, the **Edit Discovery Policy** window appears. The displayed information is the same, but the **View Discovery Policy** window is read-only.

5. To modify one or more parameters in a user-defined policy, change the entry in the appropriate field(s) and then click **Save and Apply**.

Application

APP UDP 5353

Policy Name

APP UDP 5353_SB1_MyWS_CL_01

Description

VMs

☒ All VMs ☐ Specific VMs

TCP Ports

Note: Input specific ports separated by commas, ie. *3333, 3340*

UDP Ports

5353 x

Note: Input specific ports separated by commas, ie. *3333, 3340*

Close Save and Apply

Figure 31: Edit Discovery Policy Window

6. To enable, disable, or delete a policy, select the policy from the list and then select the desired option (**Enable**, **Disable**, or **Delete**) from the **Actions** dropdown menu. You can only delete user-defined policies (not system policies).

APPLICATION MONITORING

Application monitoring provides visibility into integrated applications by collecting application metrics using Nutanix and third-party collectors, providing a single pane of glass for both application and infrastructure data, correlating application instances with virtual infrastructure, and providing deep insights into applications performance metrics.

Application monitoring provides visibility into the following application:

- **VMware vCenter Server:** When you configure a vCenter Server for monitoring, the **Monitoring Configurations** page lists the vCenter Server. You can also view the vCenter Server listed on the **Non-Nutanix** page (**Compute > VMs > Non-Nutanix**). The vCenter details are also available on the **Hardware > Clusters > Non-Nutanix** page.

See [Application Monitoring Summary View](#) on page 74 and [Application Instance Details View \(vCenter\)](#) on page 84.

- **Microsoft SQL Server:** You can monitor a Microsoft SQL Server running on a Nutanix or non-Nutanix cluster. When you configure a Microsoft SQL Server for monitoring, only the **Monitoring Configurations** page lists the SQL Server.

See [Application Monitoring Summary View](#) on page 74 and [Application Instance Details View \(SQL Server\)](#) on page 77.

Application Monitoring Summary View

The **Monitoring Configurations** page allows you to view information about select applications, such as SQL Server or vCenter instances, running in the cluster.

To access the **Monitoring Configurations** page:

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher. The system displays the **Intelligent Operations** dashboard.
3. Navigate to **Monitoring Configurations** from the **Navigation Bar**.

The **Monitoring Configurations** page opens displaying the SQL Server or vCenter instances running in the cluster. This page allows you to view summary information about application instances and access detailed information about each instance.

Note: If monitoring integrations is not enabled, a **Welcome to Monitoring Configurations** page appears instead of the **Monitoring Configurations** page. Click **Get Started**. This enables monitoring integrations, after which the **Monitoring Configurations** page appears. However, the **Monitoring Configurations** page is blank until you configure monitoring. For more information, see [Configuring Application Monitoring](#) on page 86.




Figure 32: Monitoring Configurations Page

Note: If you delete any entities in the source application being monitored, the entities keep getting listed in the monitoring view. The deleted entities are removed from the monitoring view after 12 hours.

The following table describes the fields that appear in the **Monitoring Configurations** page.

Note: The fields vary based on the **View by** option. A dash (-) is displayed in a field when a value is not available or applicable.

Table 17: Monitoring Configurations – Field Description

Field	Description	Values
Name	Displays the name of the application instance. Click the name to Monitoring Configurations display the details for that instance (see Application Instance Details View (SQL Server) on page 77 or Application Instance Details View (vCenter) on page 84).	(instance name)
IP address	Displays the IP address of the VM where the application instance is hosted.	(IP address)
Type	Displays the type of application instance.	Microsoft SQL Server, vCenter
Data Provider	Displays the data provider (agentless collector) that supplies the application metrics.	Nutanix, (third-party name)
Historical Interval	Historical data collection interval for the vCenter instance. No data is displayed for SQL server instance.	(interval in minutes)
Health Status	Displays the current health status of the instance.	Active, Failed, Stopped, Partial

You can perform the following actions from the page:

- Click the name of an application instance to open the application instance details page, that displays the detailed information about the instance. For more information, see [Application Instance Details View \(SQL Server\)](#) on page 77 or [Application Instance Details View \(vCenter\)](#) on page 84.

- Add an application instance for monitoring by clicking **Configure New Instance**. For more information, see [Configuring Application Monitoring](#) on page 86.
- Download the contents of the fields that appear in the page in CSV format by clicking the **Export** link.
- View the application instances based on pre-defined criteria or create a custom view.
- Filter the application instances list based on a variety of parameter values using the **Filters** pane. For more information, see [Filters Pane - Monitoring Configurations page](#).
- Perform the following actions on a single or multiple application instances using the **Actions** dropdown menu. The **Actions** dropdown appears when one or more application instances are selected.

Table 18: Monitoring Configurations Actions

Action	Description
Update	Click this action to update the instance configuration. For more information, see Modifying Application Monitoring on page 91.
Delete	Click this action to stop monitoring the instance. For more information, see Modifying Application Monitoring on page 91.

Filters Pane - Monitoring Configurations page

You can filter the information in the **Monitoring Configurations** page based on the following fields that are available in the **Filters** pane.

Table 19: Filter Pane Field Description - Monitoring Configurations page

Field	Description	Values
Name	Filters based on the application instance name. Select the checkbox, and choose a condition from the dropdown menu (Contains , Doesn't contain , Starts with , Ends with , or Equal to) and enter a string in the field.	(name string)
	Filters based on the IP address of the host VM. Select the checkbox, and choose a condition from the dropdown menu (Contains , Doesn't contain , Starts with , Ends with , or Equal to) and enter a string in the field.	(IP address)
Type	Filters based on the application instance type.	(Microsoft SQL Server, vCenter)
Health Status	Filters based on the status of the application instance.	(Active, Stopped, Disabled)
Last Updated Time	Filters based on the time of last update of the application instance.	(from, to)
Created Time	Filters based on the time of the application instance creation.	(from, to)

Application Instance Details View (SQL Server)

The details page of an SQL Server application instance provides a range of data and performance metrics for the SQL Server instance.

The details page has the **Summary**, **Alerts**, **Events**, **Databases**, **Queries**, and **Metrics** tabs.

To access the details page of an SQL Server application instance:

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher. The system displays the **Intelligent Operations** dashboard.
3. Navigate to **Monitoring Configurations** from the **Navigation Bar**.

The **Monitoring Configurations** page opens displaying the SQL Server or vCenter instances running in the cluster.

4. Click an SQL Server application instance to open the details page of the instance.

The **Summary** tab opens displaying the detailed information about the instance in widgets.

Click the Back icon to return to the **Monitoring Configurations** page.

Summary Tab

The **Summary** tab provides detailed information about the instance in widgets.




Figure 33: Application Instance - Summary Tab

The **Summary** tab has the following widgets. A dash (-) is displayed in a field when a value is not available or applicable.

Table 20: Summary Tab - Widgets

Widget Name	Information provided
Properties	Provides the following: • Host VM — Displays the host VM IP address. Click the IP address to displays the details page for that VM. See VM Details View in <i>Prism Central Infrastructure Guide</i>. • Type — Displays the type of SQL server. • SQL Version — Displays the SQL server version number for the instance. • Database Count — Displays the number of databases managed by the SQL server instance. • Host — Displays the SQL server host name. • Instance — Displays the instance type. • VM Name — Displays the host VM name. • VM Alert Count — Displays the number of alerts associated with the instance.
Metrics	Provides the following: • CPU Usage — Displays the percentage of CPU capacity currently used by the instance. • Total Server Memory (MB) — Displays the amount of memory currently assigned to SQL Server instance. • Target Server Memory (MB) — Displays the amount of memory required for best performance of the SQL Server instance. • Available Physical Memory — Displays the amount of available memory. • Buffer Cache Hit Ratio % — Displays the buffer cache hit rate for the past 15 minutes. • Number of Active Connections — Displays the number of active connections.
Alerts	Displays a graph of related alerts that occurred during the specified interval. Select either Last 24 hours (default) or Last week from the dropdown menu. When an alert appears, you can click on the graph, which then displays a list of those alerts. Clicking an alert displays the details page for that alert.
Anomalies	Displays a graph of memory, I/O, CPU, networking, disk or other anomalies that occurred during the specified interval. Select either Last 24 hours (default) or Last week from the dropdown menu. When an anomaly appears, you can click on the graph, which then displays a list of those anomalies. Clicking an anomaly displays the event page for that anomaly.

Alerts Tab

The **Alerts** tab displays a list of alerts for the SQL Server application instance. This tab provides the same features and options as the *Alerts dashboard (List tab)*, except it displays just SQL Server-related alerts. For more information, see the *List Tab* section in [Alerts Summary View \(Prism Central\)](#) of the *Prism Central Alerts and Events Reference Guide*.

Events Tab

The **Events** tab displays a list of events for the SQL Server application instance. This tab provides the same features and options as the *Events dashboard*, except it displays just SQL Server-related events. For more information, see the *List Tab* section in [Alerts Summary View \(Prism Central\)](#) of the *Prism Central Alerts and Events Reference Guide*.

Databases Tab

The **Databases** tab displays a list of databases supported by the SQL Server instance.




Figure 34: Application Instance - Databases Tab

The following table describes the fields that appear in the **Databases** tab.

Note: The fields vary based on the **View by** option. A dash (-) is displayed in a field when a value is not available or applicable.

Table 21: Databases Tab - Field Description

Field	Description	Values
Name	Displays the database name.	(database name)
Reads	Displays the total number of read operations.	(number of reads)
Writes	Displays the total number of write operations.	(number of writes)
Read Bytes	Displays the total number of bytes read.	(number of bytes read)
Write Bytes	Displays the total number of bytes written.	(number of bytes written)
Write Latency ms	Displays the total time taken to complete the write operations (in milliseconds).	(number of write latency in milliseconds)
Read Latency ms	Displays the total time taken to complete the read operations (in milliseconds).	(number of read latency in milliseconds)

You can perform the following actions from the **Databases** tab:

- Download the contents of the fields that appear in the **Databases** tab in CSV format by clicking the **Export** link.

- Filter the database list based on the database name using the **Filters** pane. To do this:
 1. Click **Modify Filters**.
 2. In the **Filters** pane that opens, select the checkbox for Name, choose a condition from the dropdown menu (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field.
- View the databases based on pre-defined criteria or create a custom view.
- Order the rows by clicking a column header, which toggles the list alphabetically or numerically (as appropriate) based on that parameter.
- Click the Back icon to return to the **Monitoring Configurations** page.
- Click the name of a database to open the database details page, which displays six graphs. Each graph is a rolling time interval performance or usage monitor for a metric. The metrics are **Read Operations**, **Write Operations**, **Write Latency**, **Read Latency**, **Bytes Read**, and **Bytes Written**.




Figure 35: Database Details Page

Perform the following actions on each graph using the **Actions** dropdown menu.

Table 22: Database Details Page - Actions

Action	Description
Add to Analysis	Click this action to add the metric to the analysis page. For more information, see Analysis Dashboard on page 12.
Select Analysis Session	Click this action to add the metric to an existing or new analysis session. Perform the following steps to add the metric to an existing or new analysis session: 1. Click Select Analysis Session. 2. In the Select A Session window that opens, select a session from the Session dropdown menu. 3. Click Add to Session.

Action	Description
Alert Settings	Click this action to create a custom alert policy for the metric. For more information, see Creating Custom Alert Policies topic in <i>Prism Central Alerts and Events Reference Guide</i> .

You can perform the following actions from the database details page:

- View a subset of the graphs by clicking the **Filters** option and selecting the desired graphs.
- Specify a tile view format (default) or stack view format for the graphs by clicking the tile or stack icon.
- Specify a time interval for the graphs by selecting **Last 1 hour**, **Last 24 hours** (default), or **Last Week** from the dropdown menu.
- Click the Back icon to return to the **Databases** tab.

Queries Tab

The **Queries** tab displays a list of SQL query statements.




Figure 36: Application Instance - Queries Tab

The following table describes the fields that appear in the **Queries** tab.

Note: The fields vary based on the **View by** option. A dash (-) is displayed in a field when a value is not available or applicable.

Table 23: Query Tab - Field Description

Parameter	Description	Values
Query Text	Displays the query statement text.	(query text)
Avg Execution Time	Displays the average amount of time it took to run the query.	xxx secs
Execution Count	Displays the number of times the query statement has been executed.	(integer)
Last Execution Duration	Displays the amount of time it took to run the last query.	xxx secs
Last Execution Time	Displays the date and time when the query was last executed.	(date and time)

You can perform the following actions from the **Queries** tab:

- Order the rows by clicking a column header, which toggles the list alphabetically, numerically, or chronologically (as appropriate) based on that parameter.
- Download the contents of the fields that appear in the **Queries** tab in CSV format by clicking the **Export** link.
- Filter the queries list based on the query name using the **Filters** pane. To do this:
 1. Click **Modify Filters**.
 2. In the **Filters** pane that opens, select the checkbox, choose a condition from the dropdown menu (**Contains**, **Doesn't contain**, **Starts with**, **Ends with**, or **Equal to**) and enter a string in the field.
- View the queries based on a pre-defined criteria or create a custom view.
- Click the Back icon to return to the **Monitoring Configurations** page.
- Click a query text to open the details page for that query.




Figure 37: Query Details Page

The query details page has a **Properties** widget. The following table describes the fields that appear in the widget.

Table 24: Properties Widget - Field Description

Parameter	Description	Values
Execution Count	Displays the number of times the query statement has been executed.	(integer)

Parameter	Description	Values
Last Execution Duration	Displays the amount of time it took to run the last query.	xxx secs
Query Text	Displays the query statement text.	(query text)

You can perform the following actions from the query details page:

- View the full query statement by clicking the **View** link.
- Copy the query statement by clicking the **Copy** link.
- Click the Back icon to return to the **Queries** tab.

Metrics Tab

The **Metrics** tab displays a set of graphs for various metrics. Each graph is a rolling time interval performance or usage monitor for that metric.

The following table describes the graphs that appear in the **Metrics** tab.

Table 25: Metric Graphs - Description

Metric	Description
CPU Utilization	Displays the CPU bandwidth usage by all requests in all workload groups belonging to all related resource pools.
Active Connections	Displays the number of active connections to this instance.
Buffer Cache Hit Ratio	Displays the percentage of pages found in the buffer cache without having to read from disk. The ratio is the total number of cache hits divided by the total number of cache lookups over the last few thousand page accesses. Because reading from the cache is much less expensive than reading from disk, you want this ratio to be high. Generally, you can increase the buffer cache hit ratio by increasing the amount of memory available to SQL Server or by using the buffer pool extension feature.
Batch Request Rate	Displays the rate of transact-SQL command batches received. This statistic is affected by all constraints (such as I/O, number of users, cache size, complexity of requests, and so on). High batch requests mean good throughput. This is a calculated metric displaying the batch requests that occurred per second since the previous collection.
SQL Compilations Rate	Displays the rate of SQL compilations of the instance. This metric displays the SQL compilations that occurred per second since the previous collection.
SQL Recompilations Rate	Displays the rate of statement recompiles. Counts the number of times statement recompiles are triggered. Generally, you want the recompiles to be low. This metric displays the SQL Recompilations that occurred per second since the previous collection.
Buffer Page Life Expectancy	Displays the number of seconds a page will stay in the buffer pool without references.

Metric	Description
Checkpoint Pages Rate	Displays the number of pages flushed to disk by a checkpoint or other operation that requires all dirty pages to be flushed. This metric displays the checkpoint pages that occurred per second since the previous collection.
Available Physical Memory	Displays the amount of physical memory available.
Lock Wait Time	Displays the total wait time for locks in the last second.
User Connections	Displays the number of connections currently used for user sessions.

Perform the following actions on each graph using the **Actions** dropdown menu.

Table 26: Metrics Tab - Actions

Action	Description
Add to Analysis	Click this action to add the metric to the analysis page. For more information, see Analysis Dashboard on page 12.
Select Analysis Session	Click this action to add the metric to an existing or new analysis session. Perform the following steps to add the metric to an existing or new analysis session: 1. Click Select Analysis Session. 2. In the Select A Session window that opens, select a session from the Session dropdown menu. 3. Click Add to Session.
Alert Settings	Click this action to create a custom alert policy for the metric. For more information, see Creating Custom Alert Policies topic in <i>Prism Central Alerts and Events Reference Guide</i> .

You can perform the following actions from the database details page:

- View a subset of the graphs by clicking the **Filters** option and selecting the desired graphs.
- Specify a tile view format (default) or stack view format for the graphs by clicking the tile or stack icon.
- Specify a time interval for the graphs by selecting **Last 1 hour**, **Last 24 hours** (default), or **Last Week** from the dropdown menu.
- Click the Back icon to return to the **Databases** tab.

Application Instance Details View (vCenter)

The details page of a vCenter instance provides a range of data and performance metrics for the vCenter instance.

The details page has the **Summary**, **Alerts**, and **Monitored Clusters** tabs.

To access the details page of a vCenter instance:

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher. The system displays the **Intelligent Operations** dashboard.

3. Navigate to **Monitoring Configurations** from the **Navigation Bar**.

The **Monitoring Configurations** page opens displaying the SQL Server or vCenter instances running in the cluster.

4. Click a vCenter instance to open the details page of the instance.

The **Summary** tab opens displaying the detailed information about the instance in widgets.

Click the Back icon to return to the **Monitoring Configurations** page.

Summary Tab

The **Summary** tab shows summary information about the instance in widgets.




Figure 38: Application Instance - Summary Tab

The following table describes the fields in the **Properties** widget. A dash (-) is displayed in a field when a value is not available or applicable.

Table 27: Summary Tab - Widgets

Widget Name	Information provided
Properties	Provides the following: • VMs — Displays the number of VMs managed by the vCenter instance. • Clusters — Displays the number of clusters managed by the vCenter instance. • Nodes — Displays the number of clusters managed by the vCenter instance. • Critical Alerts — Displays the current number of critical alerts associated with the vCenter instance. • Warning Alerts — Displays the current number of warning alerts associated with the vCenter instance. • Name — Displays the name of the vCenter instance. • Monitoring Status — Displays the status of monitoring the vCenter instance. • vCenter Authentication — Displays the status of vCenter authentication. • Configured Alerts Count — Displays the number of alerts configured to monitor the vCenter instance.

Alerts Tab

The **Alerts** tab displays a list of alerts. This tab provides the same features and options as the *Alerts dashboard (List tab)*, except it displays just vCenter-related alerts. For more information, see the *List Tab* section in [Alerts Summary View \(Prism Central\)](#) of the *Prism Central Alerts and Events Reference Guide*.

Monitored Clusters Tab

The **Monitored Clusters** tab displays the list of vCenter clusters being monitored. This tab provides the same features and options as the *Clusters (List tab)* for *Non-Nutanix* clusters. For more information, see the *List Tab* section in [Clusters Summary View](#) of the *Prism Central Infrastructure Guide*.

Configuring Application Monitoring

Monitoring an application instance requires configuring an instance profile.

Before you begin

Ensure that all the prerequisites specified in [Prerequisites for Configuring Application Monitoring](#) on page 87 are met before you configure application monitoring.

Procedure

To configure an instance profile, do the following:

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher. The system displays the **Intelligent Operations** dashboard.
3. Navigate to **Monitoring Configurations** from the **Navigation Bar**.
The **Monitoring Configurations** page opens displaying the SQL Server or vCenter instances running in the cluster.
4. Click **Configure New Instance**.

5. In the **New Monitoring Configuration** window, perform the following in the indicated fields:

- a. **Name:** Enter a name for the application monitoring configuration.
- b. **External Entity Type:** Select the target type from the dropdown menu.

The options are **Microsoft SQL Server** or **vCenter**. If you are configuring an instance for the first time, an **Enable** option appears. Click the **Enable** option to provision instance monitoring. You must log in as a user with super admin role to perform this operation.




Figure 39: Configuration Page (enable)

When the provisioning is complete (which can take a few minutes), the remaining fields appear based on your selection in the **External Entity Type** field.

- c. Do one of the following:
 - For **Microsoft SQL Server** monitoring configuration, see [Configuring an SQL Server Application Monitoring](#) on page 88.
 - For **vCenter** monitoring configuration, see [Configuring a vCenter Application Monitoring](#) on page 89.

Note: If you want to monitor an SQL Server hosted on a non-Nutanix (vCenter Server) cluster, follow the steps mentioned in [Configuring a vCenter Application Monitoring](#) on page 89.

Prerequisites for Configuring Application Monitoring

Ensure that the following prerequisites are met so that you can successfully configure application monitoring.

- **Sufficient Prism Central VM resources:** Application monitoring is a feature that requires additional Prism Central VM resources (more memory and vCPUs). The requirements are listed in [KB 8932](#) (see instructions for Prism Ultimate features). If your Prism Central VM resources are below the recommended amount, manually increase the memory and vCPUs to the recommended amount.
- TLS version 1.2.

- (Only for SQL Server): Ensure the following:
 - Each SQL Server instance port must be accessible through your network, but no other special network configuration is required.
 - Create a read-only login to connect to your server:

```
CREATE LOGIN <username> WITH PASSWORD = '<PASSWORD>';
CREATE USER <username> FOR LOGIN <username>;
GRANT SELECT on sys.dm_os_performance_counters to <username>;
GRANT VIEW SERVER STATE to <username>;
GRANT VIEW ANY DEFINITION to <username>;
```

After creating the login make sure that your SQL Server instance supports SQL Server authentication by enabling **SQL Server and Windows Authentication mode** in the server properties (**Server Properties** > **Security** > **SQL Server and Windows Authentication mode**).

- Active Directory users and Windows authentication are not supported for SQL Server monitoring.
- (Only for vCenter Server): Telegraf collector (used for application monitoring) supports only vSphere versions 6.5, 6.7, 7.0, and 8.0.
- (Only for vCenter Server): vCenter user must have minimum permission as read-only, with **Propagate to children** option enabled.

Configuring an SQL Server Application Monitoring

Procedure

To configure an SQL Server Application Monitoring instance, do the following:

1. Select one of the following options:
 - **Registered SQL Server:** If you want to monitor an SQL Server application instance that is already registered.
 - **New SQL Server:** If you want to monitor a new SQL Server application.
2. (Registered SQL Server only):
 - Select one of the registered SQL Server instances, and click **Test Connection**.
 - If the connection testing is successful, click **Save** to save the configuration instance.

3. (New SQL Server only):

- Enter the following SQL Server configuration parameters:
 - **Microsoft SQL Server Host:** Host name or IP address of the Microsoft SQL Server.
If you want to monitor a SQL Server hosted on a vCenter Server, enter the IP address of the vCenter Server host, provided the host is being monitored through vCenter Server monitoring. Configuration for monitoring a SQL Server hosted on a vCenter Server might take 15 minutes to 60 minutes.
 - **Microsoft SQL Server Port:** The default port for Microsoft SQL Server is 1433. You can change the port, if required.
 - **User Name:** Username for SQL Server.
 - **Password:** Password for SQL Server.

Note: You must use the same username and password that were used while configuring monitoring for the same SQL server for the first time.

- Click **Test Connection**.
- If the connection testing is successful, click **Save** to save the configuration instance.




Figure 40: New SQL Server Monitoring Configuration

Once you save the configuration, the data collection begins. The monitoring configuration and data collection processes can take a few minutes. When the process completes, an SQL Server Configuration Successful screen appears.

4. Click **Close** to close the page or click the App Discovery link (if displayed) to enable application discovery for the selected clusters. For more information, see [Enabling Application Discovery](#) on page 60.

Configuring a vCenter Application Monitoring

Procedure

To configure an vCenter Server Application Monitoring instance, do the following:

1. Select one of the following options:
 - **Registered vCenter:** If you want to monitor a vCenter Server application instance that is already registered.
 - **New vCenter:** If you want to monitor a new vCenter Server application.

2. (Registered vCenter only):

- Select one of the registered vCenter instances, historical statistics interval for the instance, and click **Test Connection**.
- If the connection testing is successful, click **Save** to save the configuration instance.

3. (New vCenter only):

- Enter the following vCenter configuration parameters:
 - **vCenter Host:** Host name or IP address of the vCenter.
 - **User Name:** Username for vCenter.
 - **Password:** Password for vCenter.
 - **Historical Interval:** Select the historical data collection interval for the vCenter instance. Historical interval must be the smallest interval that is enabled in vCenter.
- Click **Test Connection**.
- If the connection testing is successful, click **Next**.

Figure 41: New vCenter Monitoring Configuration

4. In the **Configure New Monitoring Configuration** dialog box, select the clusters to monitor.

A list of discovered clusters for the vCenter instance appears. Select the checkbox associated with the clusters that you want to monitor from the list, and then click **Save**.

Note:

Nutanix clusters monitored by vCenter that are also registered with this Prism Central instance do not appear in the discovered clusters list. Conversely, non-registered Nutanix clusters selected here but later registered with Prism Central appear in both the Nutanix and non-Nutanix tabs of the clusters dashboard.

5. Once you save the configuration, the data collection begins. The monitoring configuration and data collection processes can take a few minutes. When the process completes, a **vCenter Configuration Successful** screen appears.

Once you save the configuration, the data collection begins. The monitoring configuration and data collection processes can take a few minutes. When the process completes, a vCenter Configuration Successful screen appears.

Note:

If vCenter is slow or unresponsive, the relevant API might time out and fail to fetch the clusters. Even after a successful configuration, a slow vCenter might impact the details displayed for a cluster or other entities.

6. Click **Close** to close the page or click the App Discovery link (if displayed) to enable application discovery for the selected clusters. For more information, see [Enabling Application Discovery](#) on page 60.

Modifying Application Monitoring

You can view, update, or remove the application configuration for an instance.

About this task

To update application monitoring for an instance, do the following:

Procedure

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher. The system displays the **Intelligent Operations** dashboard.
3. Navigate to **Monitoring Configurations** from the **Navigation Bar**.

The **Monitoring Configurations** page opens displaying the SQL Server or vCenter instances running in the cluster.

4. Select the instance that you want to modify, and then select **Update** from the **Actions** dropdown menu.

The **Configure new monitoring configuration** window opens.

5. To modify one or more parameters, change the entry in the appropriate field(s), click **Test Connection** to test the connection, and then click **Save**.

For more information, see [Configuring Application Monitoring](#) on page 86.

Note: [vCenter] Modifying the cluster selection occurs after testing the connection.

Perform the following to remove application monitoring for an instance:

6. Select the instance that you want to remove, and then select **Delete** from the **Actions** dropdown menu.

The entry for that instance disappears from the list.

CAPACITY PLANNING

You can review and analyze current and historical usage information and use the resource planning tools in the **Capacity** tab under Cluster details view.

Note:

- Capacity planning requires a minimum of 21 days of data from a cluster to calculate the baseline runway estimates. No estimates appear when insufficient data is available. In addition, it takes a day after registering a cluster for the data to appear in Prism Central.
- Capacity planning checks the past 90 days of data for monthly seasonality, it also uses up to 365 days of past data (when available) for calculating and adjusting the baseline accordingly if it finds a seasonality pattern.
- For a Non-Nutanix cluster, the vCenter datastores dashboard appears only when at least one vCenter instance has been registered. For more information, see [Application Monitoring](#) on page 74.

To access the **Capacity** tab:

1. Log in to Prism Central.
2. Select the **Infrastructure** application from application switcher, and navigate to **Hardware** > **Clusters** from the **Navigation Bar**.

The system displays the **List** tab by default with all the registered clusters in **Nutanix** environment.

3. Click the target `<cluster_name>`, and navigate to **Capacity** tab of an individual cluster. Replace `<cluster_name>` with the actual cluster name at your site.

The **Capacity** tab of an individual cluster includes the following sections:

- A **<cluster_name> Runway** area (upper left) that displays the current cluster runway overall and broken down by CPU, memory, and storage. *Runway* refers to how long the cluster can continue to run normally, based on the current consumption rate, before the existing resources are used to capacity. The overall runway value is the same as the lowest value for any single resource (storage, CPU, or memory). For example, if the storage runway is 14 days while the CPU and memory runways are 45 days, the overall runway is 14 days (the storage value).

A plus sign in any runway value, for example the 89+ in the Cluster Capacity Tab figure, indicates that the predicted runway is more than that number of days, but the estimation stopped at that number.

Note: Prism Central version 2023.3 onwards, Storage Runway feature is not available for vCenter clusters.

- Clicking **Storage Runway** displays a storage usage graph and table (to the right).
- Clicking **CPU Runway** displays a CPU usage graph and table.
- Clicking **Memory Runway** displays a memory graph and table.

- Click **Optimize Resource** (middle left) to open the **Optimize Resources** window, which identifies areas to improve cluster resource allocation or capacity. Recommendations might include
 - Table of *inefficient* (over-provisioned, inactive, constrained, and bully) VMs based on the VM behavioral learning engine. Click the appropriate link to see a list of the VMs in that state. For more information, see [Behavioral Learning Tools](#) in *Intelligent Operations Guide*.
 - An add node recommendation. The recommendation is based on the current (historical) CPU, memory, and storage usage demand across the cluster.

Click the **Get Report** link to run the Cluster Efficiency report. For more information, see [Reports Management](#) in the *Intelligent Operation Guide*.

- Get Started** (lower left). Click **Get Started** to open the resource planning scenario page. From this page you can create *what if* scenarios to test future workloads against current or hypothetical storage, CPU, and memory resource capacity. For more information, see [Creating a Scenario](#) in *Intelligent Operations Guide*.
- Ignore Time Window** (upper right). Click **Ignore Time Window** to open the **Ignore Time** window. On this page, you can specify the time window you want to ignore while calculating the capacity runway. You should ignore any one-time intentional cluster usage that can adversely affect the cluster runway calculation, such as when a workload is transferred to another cluster due to planned maintenance. You should set the Ignore Time window slightly wider than the usage spike you want to ignore. For example, if the spike is from March 5th to March 7th, you should set a window from March 4th to March 8th.

Once defined, the Ignore Time window is disregarded from the time series data for CPU, memory, and storage usage for a Nutanix or vCenter cluster, thus making the runway prediction more accurate. It takes a couple of hours for the runway to get updated after you define an Ignore Time window. You can add multiple Ignore Time windows for each cluster.

- A usage history graph (middle right). The graph displays storage, CPU, or memory usage depending on which is selected in the left column. Placing the cursor anywhere on the horizontal axis displays the value(s) at that time. A (solid red) maximum capacity line appears near the top of the graph, and an additional (dotted red) effective capacity line appears on the storage graph. A (vertical) Today line is displayed to indicate where the current date is located on the graph. When there is an alert or event, a bell icon appears either on the Today line or at the starting point of the alert. Clicking the bell icon (red for critical alert, orange for warning alert, or gray for event) displays more information about that alert or event.
- A table (lower right) that displays usage information. The information in the table varies depending on which runway (storage, CPU, or memory) is selected.




Figure 42: Cluster Capacity Tab (Nutanix Cluster)




Figure 43: Cluster Capacity Tab (Non-Nutanix Cluster)

Capacity Tab-storage View

When **Storage Runway** is selected, the graph and table display storage usage information. The maximum and effective storage capacity values (in TiBs) for the cluster appear above the graph. Effective capacity of a cluster is computed based on the configured reservations (from the capacity configurations settings) and the replication factor for all the containers in the cluster.

The storage view includes two tabs:

- Click the **By Usage** button to see cluster-wide storage information (see previous figure).
- Click **By Storage Container** to see storage information per storage container. You can select all storage containers from the dropdown list (upper left of graph) or a single storage container from the dropdown list or by clicking the storage container name in the table.




Figure 44: Storage Container View (Nutanix Cluster)




Figure 45: Storage Container View (Non-Nutanix Cluster)

The following table describes the fields in the storage table.

Table 28: Capacity Tab: Storage Fields

Parameter	Description	Values
By Usage tab (cluster-wide)		
Name	Displays a descriptive name for a type of storage. There are four types: - Live Usage: The amount of live storage. - Snapshot Usage: The amount of storage used for snapshots. - System Usage: The amount of storage for everything else (total usage - (reserved + snapshot + live)). System usage includes garbage, parity, and other miscellaneous uses. - Reserved Usage: The amount of storage reserved for thick provisioned VMs (ESXi).	Live Usage, Reserved Usage, Snapshot Usage, System Usage
Current Storage Usage	Displays the amount of used storage space in the cluster or storage container.	xxx [GB TB]
By Storage Container tab: All Storage Containers		
Storage Container Name	Displays the name of the storage container.	(storage container name)
Current Usage	Displays the amount of storage space used by the storage container currently.	xxx [GB TB]
Capacity	Displays the total amount of storage capacity available to the storage container.	xxx [TB]
By Storage Container tab: <storage container name>		
Storage Container Name	Displays one of the following: - If All Storage Containers is selected, the storage container names are displayed. Click a name to see the details for that storage container. - If a single storage container is selected, it displays the four storage types described for the Name field in the By Usage tab (Live Usage, Reserved Usage, Snapshot Usage, and System Usage).	(storage container name) OR Live Usage, Reserved Usage, Snapshot Usage, System Usage
Current Storage Usage	Displays the amount of used storage space in the cluster or storage container.	xxx [GB TB]

Capacity Tab- CPU View

When **CPU Runway** is selected, the graph and table display CPU usage information. The maximum and effective CPU capacity values (in GHz) for the cluster appear above the graph. The CPU view includes two tabs:

- Click **Overall** to see cluster-wide CPU information.

- Click **By Host** to see CPU information per host. You can select all hosts from the dropdown list (upper left of graph) or a single host from the dropdown list or by clicking the host name in the table.

The following table describes the fields in the CPU table. Entries for (up to) the top 10 VMs (overall tab) or hosts appear in the CPU table.

Table 29: Capacity Tab: CPU Fields

Parameter	Description	Values
Overall tab		
VM Name	Displays the VM name. There is a line for each VM in the cluster.	(VM name)
CPU Usage (%)	Displays the percent of CPU capacity used by the VM currently.	(0–100%)
Num vCPUs	Displays the number of vCPUs allocated to the VM.	(1-max)
By Host tab: All Hosts		
Host Name	Displays the name of the host.	(host name)
Average CPU Usage over last 1 hour	Displays the average host CPU usage during the last hour.	xxx [GHz]
Current CPU Capacity	Displays the CPU capacity of the host.	xxx [GHz]
Runway	Displays the current runway (estimated time interval that current CPU capacity is sufficient for the workload).	(number of days)
By Host tab: <host name>		
VM Name	Displays the VM name. There is a line for each VM in the cluster.	(VM name)
CPU Usage (%)	Displays the percent of CPU capacity used by the VM.	(0–100%)
Num vCPUs	Displays the number of vCPUs allocated to the VM.	(1-max)

Capacity Tab-memory View

When **Memory Runway** is selected, the graph and table display memory usage information. The maximum and effective memory capacity values (in GiBs) for the cluster appear above the graph. The memory view includes two tabs:

- Click **Overall** to see cluster-wide memory information.
- Click **By Host** to see memory information per host. You can select all hosts from the dropdown list (upper left of graph) or a single host from the dropdown list or by clicking the host name in the table.

The following table describes the fields in the memory table. Entries for (up to) the top 10 VMs (overall tab) or hosts appear in the memory table.

Table 30: Capacity Tab: Memory Fields

Parameter	Description	Values
Overall tab		
VM Name	Displays the VM name. There is a line for each VM in the cluster.	(VM name)
Memory Usage	Displays the amount of memory used by the VM.	xxx [GiB]
Memory Usage (%)	Displays the percentage of memory capacity used by the VM.	(0–100%)
By Host tab: All Hosts		
Host Name	Displays the name of the host.	(host name)
Average Memory Usage over last 1 hour	Displays the average host memory usage during the last hour.	xxx [GiB]
Runway	Displays the current runway (estimated time interval that current memory capacity is sufficient for the workload).	(number of days)
Current Memory Capacity	Displays the memory capacity of the host.	xxx [GiB]
By Host tab: <host name>		
VM Name	Displays the VM name. There is a line for each VM in the cluster.	(VM name)
Memory Usage	Displays the amount of memory used by the VM.	xxx [GiB]
Memory Usage (%)	Displays the percentage of memory capacity used by the VM currently.	(0–100%)

RESOURCE PLANNING

You can review and analyze current and potential resource needs through the Planning page. This page has the **Scenarios**, and **Capacity Runway** tabs.

To access the Planning page, perform the following:

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher. The system displays the **Intelligent Operations** dashboard.
3. Navigate to **Planning** from the **Navigation Bar**.

The Planning page opens displaying the **Scenarios** tab. This tab displays a list of scenarios created to assess the future resource requirements.

Prism Central includes machine-learning capabilities that analyze resource usage over time and provide tools to monitor resource consumption, identify abnormal behavior, and guide resource planning. For more information, see [Behavioral Learning Tools](#) in *Intelligent Operations Guide*.

Scenarios Summary View

The **Scenarios** page displays the list of scenarios created to assess the future resource requirements for a cluster.

To access the **Scenarios** page:

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher. The system displays the **Intelligent Operations** dashboard.
3. Navigate to **Planning** from the **Navigation Bar**.

The **Scenarios** page opens displaying a list of scenarios created to assess the future resource requirements for a cluster.

If there are no saved scenarios, a **Welcome to One-click Planning** message appears with a **Get Started** option. Create a scenario by clicking **Get Started** or **New Scenario**.




Figure 46: Scenarios Page

The following table describes the fields that appear in the **Scenarios** page. A dash (-) is displayed in a field when a value is not available or applicable.

Table 31: Scenarios – Field Description

Field	Description	Values
Scenario	Displays the scenario name.	(scenario name)
Last Cluster	Displays the name of the cluster to which the scenario applies.	(cluster name)
	Note: The column title is Last Cluster because you can change the target cluster when modifying a scenario.	
Workloads	Displays the workloads added to the scenario. For example, in the figure above, the Tech Pubs scenario lists sql in this field indicating an SQL workload was added to the scenario.	(workload type)
Owner	Displays the name of the user who created the scenario.	(user name)
Target Runway	Displays the target runway duration for this scenario, that is how many days are planned for.	xxx [days]
Last Saved	Displays the time and date for the last (most recent) time that the scenario was saved.	(time and date)
Actions	- To edit a scenario, click the Edit icon. - To share a scenario, click the Share icon. - To delete a scenario, click the x icon.	(n/a)

You can perform the following actions for the scenarios from the **Scenarios** page:

- Click the name of a scenario to view the details of the scenario.
- Create a scenario by clicking **New Scenario**. For more information, see [Creating a Scenario](#) on page 99.
- Enter a string in the search field to filter the list for scenarios with that string in the name.

Creating a Scenario

About this task

Anticipating future resource needs can be a challenging task. To address this task, Prism Central provides an option to create *what if* scenarios that assess the resource requirements for possible future work loads. This allows you to evaluate questions like

- How many new VMs can the current cluster support?
- If I need a new database server in a month, does the cluster have sufficient resources to handle that increased load?
- If I create a new cluster for a given set of workloads, what kind of cluster do I need?
- If I remove a set of VMs or nodes, how will my cluster look?

You can create various *what if* scenarios to answer these and other questions. The answers are derived by applying industry standard consumption patterns to the hypothetical workloads and current consumption patterns for existing workloads. To create a *what if* scenario, do the following:

Procedure

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher. The system displays the **Intelligent Operations** dashboard.
3. Navigate to **Planning** from the **Navigation Bar**.

The **Scenarios** page opens displaying a list of scenarios created to assess the future resource requirements for a cluster.

4. Click **New Scenario**.
The Scenario configuration page appears.




Figure 47: Scenario Configuration Page (example)

5. In the top left field, select **New Scenario** (to start with a blank scenario configuration) or the name of an existing scenario (to start with that scenario configuration) from the dropdown menu and do the following in the indicated fields:

If there are no existing scenarios (see [Scenarios Summary View](#) on page 98), Prism Central displays only **New Scenario**.

- a. **Cluster:** Select either **New Cluster** or an existing cluster from the dropdown menu.

Selecting an existing cluster means you start with the resources and workloads currently in that cluster; selecting **New Cluster** means you start with a blank slate (no resources or workloads).

Note: You can select a cluster at any time to see the runway for that cluster using this scenario.

- b. **Target:** Select a runway duration (1 to 12 months) from the dropdown menu.
- c. **Vendor Type** (new cluster only): For a new cluster, select the hardware vendor of choice (Nutanix, Dell, Lenovo, Cisco, IBM) from the dropdown menu.
- d. **Capacity configuration:** Select this checkbox to reserve some capacity prior to analyzing the scenario.

All the storage, CPU, and memory capacity in the cluster is used by default. However, you have the option to reserve (withhold) some capacity before doing the analysis. To specify the capacity to reserve, click the Edit icon and do the following in the **Edit cluster configuration** window that opens:

1. **Replication Factor:** Select the desired replication factor (2 or 3) from the dropdown menu. This dropdown appears only for a new cluster configuration. You cannot modify the replication factor for an existing cluster.
2. **CPU Reservation:** Specify the percentage of CPU capacity to keep in reserve. (The default is 0.)
3. **Memory Reservation:** Specify the percentage of memory capacity to keep in reserve. (The default is 0.)
4. **Storage Reservation:** Specify the percentage of storage capacity to keep in reserve. (The default is 0.)
5. **CPU Overcommit Ratio for New VM Workloads:** Specify the ratio of virtual CPU to physical CPU. (The default is 1.)
6. **Memory Overcommit Ratio for New VM Workloads:** Specify the ratio of virtual memory to physical memory. (The default is 1.)
7. Click **Save Configuration** to save the configuration and close the window. Select the **Capacity configuration** checkbox to apply this configuration.

Note: Continue to the next step (e) to add a workload to the scenario. To remove the impact of selected VMs in an existing workload from the scenario, skip to the following step (f).

- e. To add a workload, click the **Add/Adjust** link in the **Workload** section. In the **Add/Adjust Workload** window that opens, click the **New Workload** radio button.

This displays the fields for adding a workload. In the **Workload** field, select the desired workload from the dropdown menu, enter the appropriate information in the remaining fields, and click the **Save** button to add

that workload to the scenario. You can choose from a variety of workloads. The configuration fields vary depending on the selected workload. The following table describes the configuration fields for each workload.

Note: If you want to include multiple workloads in the scenario, repeat this step as many times as needed.

The screenshot shows the 'Add/Adjust Workload' window. At the top, there are two radio buttons: 'New Workload' (selected) and 'Disregard Existing Workloads'. Below this is a 'Workload' dropdown menu currently showing 'SQL Server'. A red arrow points to this dropdown with the text 'select workload'. A large red rectangle encloses the 'workload-specific' configuration fields. These include: 'Number Of DB VMs (1 DB Instance Per VM)' with a value of '1'; 'Profile Type' with radio buttons for 'Small', 'Medium' (selected), and 'Large'; 'Database Type' with radio buttons for 'OLAP' (selected) and 'OLTP'; 'Business Critical' with radio buttons for 'Yes' and 'No' (selected); and an 'ON' date field with the value '06/27/2021'. At the bottom right of the dialog are 'Cancel' and 'Save' buttons.

Figure 48: New Workload Fields (SQL Server example)

The following table describes the fields that appear in the **Add/Adjust Workload** window.

Table 32: Workload - Field Description

Workload	Field	Description	Values
(all)	Workload	Select the desired workload type from the dropdown menu.	SQL Server, VMs, VDI, Splunk, Xen App, Change in Demand
SQL Server (This emulates the workload for a specified set of SQL databases.)	Number of Databases	Enter the number of databases to support. The default is 1.	(number)
	Profile Type	Select (click the radio button for) the database size. The default is medium.	Small, Medium, Large
	Database Type	Select the database type. There are two options: - on-line transaction processing (OLTP) is characterized by a large number of short on-line transactions. - on-line analytical processing (OLAP) characterized by a relatively low volume of transactions. This is the default.	OLAP, OLTP
	On	Enter the starting date for this workload. Clicking in the field displays a calendar from which you can select the date.	(date)
VMs (This emulates the workload for a specified set of VMs.)	Based On	Select (click the radio button for) the source for the VM profile: Note: If the list of available profiles or VMs is long, you can filter the list by entering a string in the "filter by" field. VM Profiles: Select this option to use an existing profile. Select (click the line for) the desired profile from the provided list in the following table. Existing VMs: Select this option to use an existing VM as the profile. Select the desired VM from the provided list. New Profile: Select this option to create a profile from scratch. For a new profile, enter a name, the number of vCPUs (4 by default), and the amount of memory (4 GiB by default) and HDD (20 GiB by default) capacity in the displayed fields.	VM Profiles, Existing VMs, New Profile

Workload	Field	Description	Values
	On	Enter the starting date for this workload. Clicking in the field displays a calendar from which you can select the date.	(date)
	Number of VMs	Enter the number of VMs to add.	(number)
VDI (This emulates a virtual desktop infrastructure workload for a specified number of users.)	Vendor	Select (click the radio button for) the vendor product, either View from VMWare (default) or XenDesktop from Citrix.	View, XenDesktop
	User Type	Select the target user type from the following options: • Task worker (default): The Task Worker profile is designed to run fewer applications than the other workloads. • Knowledge worker: The Knowledge Worker profile is designed to run the standard office and email applications with browsers. • Power user: The Power user profile is the most intensive of the standard workloads. • Developer: The Developer profile is for super heavy user demands. See the table Table 33: Resources for VDI User Types on page 106 for resources for each of the profiles.	task worker, knowledge worker, power user, developer
	Provision Type	Select how to provision the VMs. The selection of provisioning types vary depending on the selected vendor: • View: full clones, V2V/P2V, linked clones • XenDesktop: full clones, V2V/P2V, provisioning services, machine creation services	full clones, V2V/P2V, linked clones, provisioning services, machine creation services
	Number of users	Enter the number of users.	(number)
Splunk (This emulates a Splunk workload for a specified configuration.)	On	Enter the starting date for this workload. Clicking in the field displays a calendar from which you can select the date.	(date)
	Daily Average Index Size	Enter a value (number of events) for the daily average index size. The default is 500.	(number)
	Hot Retention Days	Enter the number of days to retain hot data. The default is 7.	(number)

Workload	Field	Description	Values
	Cold Retention Days	Enter the number of days to retain cold data. The default is 60.	(number)
	Search Users	Enter the number of search users. The default is 5.	(number)
	On	Enter the starting date for this workload. Clicking in the field displays a calendar from which you can select the date.	(date)
Xen App (This emulates a XenApp workload for a specified configuration.)	Vendor	Select (click the radio button for) the vendor, either Citrix (default) or Microsoft	Citrix (XenApp), Microsoft
	Operating System	Select the operating system to run. The default is Windows 2012R2.	Windows 2012R2, Windows 2008R2
	Provision Type	Select how to provision XenApp. The choice of provisioning types vary depending on the selected vendor: - Citrix: PVS, MCS - Microsoft: VM clone	PVS, MCS, VM clone
	System Data	Enter the amount of system data (in GB), that is the space consumed by each RDSH Server image. The default is 100.	xxx [GB]
	Number of Concurrent Users	Enter the number of concurrent users. The default is 1000.	(number)
	MCS Diff Per VM	Enter a size (in GB) for the MCS difference disk per VM. The default is 20.	xxx [GB]
	User Profile Data	Enter a size per user (in MB) for the user profile data, which includes the Windows profile and registry settings. The default is 20. The size is typically in the 20-100 MB range.	xxx [MB]
	PVS Write Cache Size Per VM	Enter a size (in GB) for the PVS write cache per VM. The default is 15.	xxx [GB]
	On	Enter the starting date for this workload. Clicking in the field displays a calendar from which you can select the date.	(date)
Change in Demand	Increase/Decrease	Specify how to change the current workload, that is select (from the dropdown menu) whether to increase or decrease the workload.	Increase, Decrease
	By Percentage	Enter the amount you want to change the workload as a percentage of the current workload. Enter the value as an integer (no percent mark).	(1-unlimited)

Workload	Field	Description	Values
	On	Enter a date when the workload change should take effect by clicking in the field and then selecting the target date from the displayed calendar. You can make the change effective immediately or at any date in the future.	(date)

Table 33: Resources for VDI User Types

User Type	vCPU	vCPU: pCore Ratio	RAM	User Data Capacity	Gold Image Capacity
Task worker (default)	2	12	2 GiB	5 GiB	20 GiB
Knowledge worker	2	10	4 GiB	10 GiB	20 GiB
Power user	2	6	8 GiB	25 GiB	20 GiB
Developer	2	2	16 GiB	50 GiB	20 GiB

Adding a workload creates demand reflected in the Runway chart. Red indicates one or more of the resources are insufficient to support the scenario; blue indicates current resources are sufficient. In the case of a new cluster, the chart is red initially because no resources have been added yet. A scenario based on an existing cluster may or not be red for the applied workload(s) running on the current resources. For more information, see [Analyzing a Scenario](#) on page 110.

- f. To remove the impact of selected VMs from an existing workload, click the **Add/Adjust** link in the **Workload** section and then click the **Disregard Existing Workloads** radio button (in the **Add/Adjust Workload** window).

Select the category of VMs you want excluded from the workload in the **VMs To Disregard** field, enter the starting date in the **Disregard From** field, and then click the **Save** button. Note the following:

- If the selected category contains VMs from multiple clusters, only those VMs from the current cluster are considered for removal.
- If a VM is a member of multiple categories, it is removed just once even when both categories are selected. For example, if VM1 is a member of both Category1 and Category2 and both categories are marked for removal, VM1 is considered only for the first workload listed for removal, which in this case applies to Category1 (and ignored for Category2).

Note: You must first create a category to use this option. For more information, see [Category Management](#) in *Prism Central Infrastructure Guide*.

6. To add a resource, do one of the following:

- Click **Recommend** in the **Resources** section of the screen. The system automatically analyzes the scenario and then adds one or more models to the **Resources** list. The **Overall Runway** chart turns all blue indicating sufficient resources are now available. (The **Recommend** is active only when at least one resource is red in the **Overall Runway** chart.)

- Click **Add/Adjust** in the **Resources** section and then click the **New Nodes** radio button in the **Add/Adjust Resource** window to explicitly add a resource. Enter the appropriate information in the following fields, and then click **Save**.

- **Model:** Select the model type from the dropdown menu. The models match the selected vendor.

Note: The options for the following fields are adjusted to match the selected model type.

- **CPU:** Select the amount of CPU (in GHz) from the dropdown menu.
- **Memory:** Select the amount of memory (in GB) from the dropdown menu.
- **HDD:** Select the amount of HDD capacity (in GB) from the dropdown menu.
- **SSD:** Select the amount of SSD capacity (in GB) from the dropdown menu.
- **# of Nodes:** Enter the number of nodes to include (1-4 depending on the model type).
- **On:** Enter the starting date for this resource. Clicking in the field displays a calendar from which you can select the date.

Add/Adjust Resource

Resource Type

☒ New Nodes ☐ Disregard Existing Nodes

Model

NX-1050

CPU

24.0 GHz

Memory

64.0 GB

HDD

5120.0 GB

SSD

400.0 GB

Of Nodes

2

ON

11/18/2018

Cancel

Save

Figure 49: New Nodes Fields

The **Overall Runway** chart is updated accordingly. However, unlike with the **Recommend** button which adds a resource(s) guaranteed to satisfy the target runway, the resource you add explicitly may or may not satisfy the target runway completely. If it does not, one or more of the resources will remain red in the chart.

Note: You can make changes at any time to the allocated resources and any of the configuration fields to try various scenarios.

7. To remove one or more nodes from a resource, click the **Add/Adjust** button in the **Resources** section, click the **Disregard Existing Nodes** radio button in the **Add/Adjust Resource** window, select the node(s) to disregard from the displayed list, and then click the **Save** button.

The total, disregarded, and effective CPU, memory, and storage are displayed at the bottom. These values are adjusted as you select (unselect) each node to be removed from the scenario. In addition, various stability checks are evaluated when a node is removed including whether this results in too many lost nodes on the same day, system capacity falls below system usages, resources are insufficient to maintain the RF level, and the erasure coding ratio can no longer be maintained (when erasure coding is enabled). When one of the checks fails, an appropriate warning or error message appears.

Add/Adjust Resource

Resource Type

☐ New Nodes

☒ Disregard Existing Nodes

Select nodes to disregard

	MODEL	DISREGARD DATE	CPU	MEMORY	STORAGE
☐	NX-3060	11/18/2018	56 GHz	125.83 GiB	3.7 TiB
☒	NX-3060	11/18/2018	56 GHz	125.83 GiB	3.7 TiB
☐	NX-3060	11/18/2018	56 GHz	125.83 GiB	3.7 TiB
☐	NX-3060	11/18/2018	56 GHz	125.83 GiB	3.7 TiB

Total Existing Resources

224 GHz503.32 GiB14.78 TiB

Disregarded

56 GHz125.83 GiB3.7 TiB

Effective Available Resources

168 GHz377.49 GiB11.09 TiB

Cancel

Save

Figure 50: Disregard Existing Nodes Fields

8. To save a scenario, click **Save Scenario**. The saved scenario appears in the **Scenarios** tab. For more information, see [Scenarios Summary View](#) on page 98.
 - To download a PDF version of the scenario, click **Generate PDF**.
If the target runway in the scenario is not satisfied (one or more resources are red), the PDF includes a list of node recommendations that satisfy the runway target.
 - To share a scenario, click **Share**.
 - To access the end user help documentation, click the Help icon.
9. To close the scenario screen, click the X icon.

Analyzing a Scenario

The scenario screen provides a simple interface to configure and analyze a variety of scenarios.

Example Analysis

To illustrate this process, consider the following example. In this case an existing cluster consists of a NX8150-G4 with four nodes. It has total CPU, memory, and storage capacity that easily handles the current workload.

For this example let's make the target runway six months and add a VDI workload for 100 full clone users that will start in two months. When that workload is added, the runway (previously blue for all) shows that while the current CPU capacity is adequate for the added workload (blue), the memory and storage capacity (red) will not be when the workload starts in two months.




Figure 51: Scenario Screen (with added workload)

In order to meet the VDI workload demand, more resource capacity is required. You can experiment with how to address this deficit by clicking the **+ Add Resource** and adding a model of your choice, but the simplest approach is to let the planning algorithm determine what to add by clicking **Recommend**. The result in this example is a recommendation to add a two node NX-1065-G5. Adding it turns the overall runway blue for all resources.




Figure 52: Scenario Screen (with added workload and recommended resource)

You can evaluate a variety of scenarios without leaving the configuration screen. For example, you can change the target runway duration, change the demand by adding or deleting workloads (check/uncheck the box for a workload or resource to add/remove from a scenario without deleting it), modify the workload parameters, or change the allocated resources. In each case the projected runway is updated immediately to reflect the changes you make. This provides a quick and robust method for testing out a variety of scenarios.

Viewing Options

The runway view, which you display by clicking the runway view icon, provides a view of the projected runway overall and by each resource (CPU, memory, and storage). You can also view the usage of each resource by clicking the usage view icon. The following three figures display the storage, CPU, and memory usage views for the example before adding the recommended resource.

- The first two months of each graph (from the left) shows the current workload usage.
- The red dot indicates when the VDI workload begins and the corresponding impact on the resource usage.
- The dotted line indicates the current capacity for each resource. In this case the graphs show that the added workload does not exceed the current CPU capacity but does exceed the storage capacity (by a little) and the memory capacity (by a lot).
- You can use the cursor to move the select point line anywhere on the graph to display the projected usage at that point in time.




Figure 53: Storage Usage View




Figure 54: CPU Usage View




Figure 55: Memory Usage View

Modifying a Scenario

About this task

To modify an existing planning scenario, do the following:

Procedure

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher. The system displays the **Intelligent Operations** dashboard.
3. Navigate to **Planning** from the **Navigation Bar**.
The **Scenarios** page opens displaying a list of scenarios created to assess the future resource requirements for a cluster.
4. Click the Edit icon associated with a scenario to display the screen for the scenario.

Note: The following steps are brief. For more information about details and screen shots, see [Creating a Scenario](#) on page 99.
5. To change a target cluster (or start with a new cluster), select the cluster from the **Cluster** field dropdown menu.
6. To change the runway duration, select the desired length (1-12 months) from the **Target** field dropdown menu.
7. To change the reserved capacity, select the **Capacity configuration** checkbox and click the associated Edit icon to modify the parameters as desired.
8. To change the workload, do one or more of the following:
 - To add a new workload, click the **Add/Adjust** link. In the **Add/Adjust Workload** window that opens, click the **New Workload** radio button, select the desired workload from the **Workload** dropdown menu, enter appropriate information in the remaining fields, and then click **Save**.
 - To remove the impact of selected VMs from an existing workload, click the **Add/Adjust** link. In the **Add/Adjust Workload** window that opens, click the **Disregard Existing Workloads** radio button, select the

category of VMs you want excluded from the workload in the **VMs To Disregard** field, enter the starting date in the **Disregard From** field, and then click **Save**.

- To modify an existing workload, click the associated pencil icon to open the **Add/Adjust Workload** window to that workload and make the desired changes.
- To delete an existing workload, click the associated X icon.

9. To change the resources, do one or more of the following:

- To add a new (recommended) resource automatically, click **Recommend**.

Note: The **Recommend** button is active only when at least one resource is red in the runway chart.

- To add a new (specified) resource manually, click **Add/Adjust**, click the **New Nodes** radio button in the Add/Adjust Resource window, enter appropriate information in the displayed fields, and then click **Save**.
- To remove a resource, click **Add/Adjust**, click the **Disregard Existing Nodes** radio button in the Add/Adjust Resource window, select the node(s) to remove from the displayed list, and then click **Save**.
- To modify an existing resource, click the associated pencil icon to open the Add Resource window to that resource and make the desired changes.
- To delete an existing resource, click the associated X icon.

10. To save the changed scenario, click **Save Scenario**.

Sharing a Scenario

About this task

Any user with edit permission can share a scenario with other users and provide view or edit access to them.

Note: Enable CMSP to use this feature.

Procedure

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher. The system displays the **Intelligent Operations** dashboard.
3. Navigate to **Planning** from the **Navigation Bar**.

The **Scenarios** page opens displaying a list of scenarios created to assess the future resource requirements for a cluster.

4. Select the scenario that you want to share and click the Share icon under **Actions**.




Figure 56: Scenarios Tab

The **Sharing Scenario** window opens.




Figure 57: Sharing Scenario

5. In the **Add Users & Groups** tab, do the following:

- a. **Share with:** In the Search field, enter the name of the users or groups to share the scenario. The Search field displays a list of users or groups whose names match the letters that you type. Click the username or group name to add that user or group. You can add more than one user or a group.
- b. **Permission Type:** You can provide view only or edit permissions to a user or group by selecting **View only** or **Editor** from the dropdown menu.

A user with **Editor** permission can perform these tasks:

- View a session
- Edit a session
- Share a session

- c. In the **Manage Users & Groups** tab, you can change the existing permission type for users or groups. To change the permission type, select the type from the dropdown menu associated with the user name or group name.

You can also remove the access for a user or group.

6. Click **Save**.

A message that the scenario is shared successfully appears.

After sharing the scenario, wait for few minutes for changes to take effect.

Note: Users will not be able to view any information about a cluster with permissions related only to scenarios and not a cluster.

Capacity Runway Summary View

The **Capacity Runway** page displays a list of registered Nutanix and Non-Nutanix clusters and a summary view of resource runway information for these clusters.

To access the **Capacity Runway** page:

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher. The system displays the **Intelligent Operations** dashboard.
3. Navigate to **Planning** from the **Navigation Bar**.

The **Scenarios** page opens displaying a list of scenarios.

4. Click the **Capacity Runway** tab.

The **Capacity Runway** page opens displaying a list of registered Nutanix and Non-Nutanix clusters and a summary view of resource runway information for these clusters.

Note:

- Capacity planning requires a minimum of 21 days of data from a cluster to calculate the baseline runway estimates. No estimates appear when insufficient data is available. In addition, it takes a day after registering a cluster for the data to appear in Prism Central.
- Capacity planning checks the past 90 days of data for monthly seasonality. It also uses up to 365 days of past data (when available) for calculating and adjusting the baseline accordingly if it finds a seasonality pattern.
- Capacity planning supports rebuild capacity reservation of the cluster.

- For a Non-Nutanix cluster, the vCenter datastores dashboard appears only when at least one vCenter instance is registered. For more information, see [Application Monitoring](#) on page 74.




Figure 58: Capacity Runway Page (Nutanix Cluster)



Figure 59: Capacity Runway Page (Non-Nutanix Cluster)

The following table describes the fields that appear in the **Capacity Runway** page. A dash (-) is displayed in a field when a value is not available or applicable.

Table 34: Capacity Runway Page - Field Description

Field	Description	Values
Name	Displays the cluster name.	(cluster name)
Cluster Runway	Displays the overall runway for the cluster. The <i>cluster</i> runway is the shortest CPU, memory, or storage runway; that is, how long will it be before one of these resources runs out of capacity? Note: <i>Runway</i> is the number of days the existing CPU, memory, or storage capacity can handle the workload based on the current workload consumption model.	xxx [days]
CPU Runway	Displays the cluster runway for CPU usage.	xxx [days]
Memory Runway	Displays the cluster runway for memory usage.	xxx [days]
Storage Runway	Displays the cluster runway for storage usage.	xxx [days]

Clicking the cluster name displays the **Capacity** tab of the details page for that cluster. This tab displays current, historical, and projected usage information. For more information, see [Cluster details View](#) in *Prism Central Infrastructure Guide*.

Updating Capacity Configurations

About this task

The Prism Central planning tools (see [Resource Planning](#) on page 98) assume by default that all the storage, CPU, and memory capacity in the cluster should be considered when estimating the runway. However, you have the option to reserve some capacity before doing the runway analysis. To reserve storage, CPU, or memory capacity, do the following:

Procedure

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher. The system displays the **Intelligent Operations** dashboard.
3. Navigate to **Operations Settings** from the **Navigation Bar**. The **Settings** page opens.
4. Click **Capacity Configurations**. The **Update Capacity Configurations** window appears.




Figure 60: Update Capacity Configurations Window

5. Perform the following in the indicated fields:
 - a. **Reserve Capacity For Failure:** Click the appropriate radio button to specify whether runway analysis should account for the impact of a node failure.
 - Click **Auto Detect** to account for a node failure in the runway estimates. In this case the reserved capacity percentages (following three fields) refer to the amount of CPU, memory, or storage resources in the cluster minus the largest single node. This means the runway values are computed with enough cushion to account for a single node failure.
 - Click **None** to ignore failure scenarios in the runway estimates. In this case no capacity is reserved for a potential node failure.
 - b. **Reserve CPU Capacity (%):** Enter the percentage of cluster CPU to reserve (0-100).

Note: The default is 0% of the cluster total for CPU, memory, and storage capacity. Setting a higher value in one or more of these three fields means leaving out that amount of the total cluster CPU, memory, or storage

capacity from the runway analysis. For example, setting this field to "10" means 10% of the total available CPU capacity is subtracted before doing the runway analysis.

- c. **Reserve Memory Capacity (%)**: Enter the percentage of cluster memory to reserve (0-100).
 - d. **Reserve Storage Capacity (%)**: Enter the percentage of cluster storage to reserve (0-100).
6. Verify the entries and click **Save** to apply the values and close the window.

Behavioral Learning Tools

Prism Central uses machine learning to analyze resource usage, monitor consumption, detect abnormal behavior, and guide effective resource planning.

Prism Central includes machine-learning capabilities that analyze resource usage over time and provide tools to monitor resource consumption, identify abnormal behavior, and guide resource planning. These tools include

- *VM right sizing* where VMs are analyzed and those that exhibit inefficient profiles are identified.
- Anomaly detection to record when performance or resource usage is outside an expected range based on learned VM baseline behavior.
- *Smart* alerts that trigger when specified anomalies are recorded.
- Reports that summarize cluster efficiency.

VM Right Sizing

It is useful to look at the profile of your VMs when analyzing problems in a cluster or assessing future resource needs. This can help you identify VMs that are not optimally configured such as ones that consume too many resources, are constrained, are over provisioned, or are inactive. *VM right sizing* is supported for vCenter-monitored VMs as well.

The right sizing feature identifies inefficient VMs that fit one of the profiles described in the following table.

Table 35: Inefficient VM Profiles

Type	Description
Bully VM	A <i>bully</i> VM is one that consumes too many resources and causes other VMs to starve. A VM is considered a bully when it exhibits one or more of the following conditions for over an hour: • CPU ready time > 5% • Memory swap rate > 0 Kbps • Host I/O Stargate CPU usage > 85% One or more bully VMs might cause cluster performance to degrade. Identifying bully VMs can help in analyzing whether one or more of those VMs are misbehaving or need additional resources.

Type	Description
Constrained VM	A <i>constrained</i> VM is one that does not have enough resources for the demand and can lead to performance bottlenecks. A VM is considered constrained when it exhibits one or more of the following baseline values, based on the past 21 days: • CPU usage > 90% (moderate), 95% (severe) • CPU ready time > 5% , 10% • Memory usage > 90%, 95% • Memory swap rate > 0 Kbps (no moderate value) One or more constrained VMs might cause a performance bottleneck. To provide adequate host resources, resize (increase) the constrained VMs. While comparing with constrained thresholds, the 95th percentile of hourly averages of metrics is considered.
Over-provisioned VM	An <i>over-provisioned</i> VM is the opposite of a constrained VM, meaning it is a VM that is over-sized and wasting resources which are not needed. A VM is considered over-provisioned when it exhibits one or more of the following baseline values, based on the past 21 days: • CPU usage < 50% (moderate) or < 20% (severe) and CPU ready time < 5% • Memory usage < 50% (moderate) or < 20% (severe) and memory swap rate = 0 Kbps To prevent host resource wastage, resize (decrease) the over-provisioned VMs. While comparing with over-provisioned thresholds, the 99.5th percentile of hourly averages of metrics is considered.
Inactive VM	A VM is <i>inactive</i> in either of the following states: • <i>Dead</i> VM: A VM is considered dead when it has been powered off for at least 30 days. • <i>Zombie</i> VM: A VM is considered zombie and marked inactive if it has very few read/write IOs and minimal network traffic every day for the past 21 days. To prevent host resource wastage, delete the dead and zombie VMs.

Note: If a VM is already classified as inactive, it is not marked as over-provisioned. Inactive category takes precedence over over-provisioned category.

You can view the information about inefficient VMs from the **VM Efficiency** and **Impacted Cluster** widgets on the home or a custom dashboard (see [Main Dashboard - Infrastructure](#) and [Adding Dashboard Widgets](#) in *Prism Central Infrastructure Guide*) and from the **Capacity** tab of a cluster details page (see [Cluster details View](#) in *Prism Central Infrastructure Guide*).




Figure 61: Impacted Cluster and VM Efficiency Widgets

You can also use a search filter from the VMs Summary View to identify inefficient VMs.



Figure 62: Filtering VMs

Anomaly Detection

The system predicts a normal behavior band for various metrics based on historical data. The anomaly detection module monitors a predefined set of metrics on a daily basis and publishes baseline values for each of the metrics.

- Twenty-seven metrics are monitored for VMs, hosts, and clusters.
- Data for each metric from the past 21 days is recorded and analyzed, a normal behavior band is established, and predictions for the next 2 days are formulated.
- The behavior bands and predictions are adjusted accordingly when time period or trend patterns are observed, for example low CPU on weekends or increasing CPU usage.

The anomaly detection module measures usage every five minutes and compares that usage with the predicted values. If the observed value is outside the band, it flags that value as an anomaly. Each anomaly is recorded as an event

(see [Event Details View](#) in *Prism Central Alerts and Events Reference Guide*). Anomalies appears as outliers in the behavioral anomaly event details screen and the VM details Metrics tab.




Figure 63: Anomaly Event Details Screen

Smart Alerts

You can create custom policies to generate alerts when behavioral anomalies occur. You can generate a critical or warning alert when a behavioral anomaly occurs for the following conditions:

- a specified VM, host, or cluster
- a specified metric
- every time or only when the anomaly resides within a certain range

For instructions on how to create a custom anomaly-based alert, see [Creating Custom Alert Policies](#) in *Prism Central Alerts and Reference Guide*.

Cluster Efficiency Reports

One of the default reports (see [Reports Summary View](#) on page 186) is the **Cluster Efficiency Summary** report. This report provides cluster statistics (host and VM counts), performance and usage statistics (average and peak CPU and memory usage), runway metrics, and information about inefficient VMs (bully, constrained, over-provisioned, and inactive). You can run this report to get a quick view of how efficiently the cluster is performing.

TASK AUTOMATION - PLAYBOOKS

You can automate routine administrative tasks through Prism Central by using the X-Play feature.

Overview

The X-Play feature allows you to automate certain administrative tasks.

Note: X-Play is a Nutanix Cloud Manager (NCM) feature. If you disable the NCM features, you can still create playbooks and actions, but you cannot enable or run the playbooks.

X-Play is an easy to use automation tool that helps you to automate routine administrative tasks, and auto-remediate issues that may occur in your system. You can achieve this automation by creating **Playbooks**.

Playbook allows you to define a trigger that results in the execution of an action or a series of actions. A trigger may be an event that occurs in the system, such as an alert or a request made by you. The resultant actions that you configure can be VM actions, communication actions, alert actions, or report actions.

To get started, you can clone one of the predefined playbooks or create a playbook based on your requirement. You can create playbooks based on predefined alert triggers or you can create a manual trigger. For more information, see [Playbook Triggers](#) on page 129.

You can view the instances of playbook execution from the **Plays** tab in the **Playbooks** page. You can see a snapshot view of the executed plays from the **Plays** widget in the **Main Dashboard**. Additionally, you can view the plays executed in the last 24 hours in the **Impacted Cluster** widget.

Roles

You can perform the following X-Play operations if you have the Prism Admin or the Super Admin role.

- Create, read, clone, export, update, delete, enable, disable and run playbooks.
- Clone, read, update, and delete action template in the action gallery.
- View executed plays, and stop or resume a playbook that is in running state.

If you have the Prism Viewer role, you can only view the playbooks, plays, and action gallery. All other roles do not have permissions to the playbook, play, or action gallery.

Limitations

- When Prism Central Disaster Recovery is enabled in your setup, playbooks with application passwords, authentication tokens, or secrets fail on a recovered Prism Central. You must manually enter the credentials again in playbooks.
- Playbooks currently support only standard time zones. Daylight Saving Time (DST) adjustments are not considered. Ensure that all time-based configurations use standard time zones to maintain consistency and accuracy.

Playbooks Summary View

The **Playbooks** page displays the list of playbooks across all the registered clusters. It has the **List**, **Plays**, and **Action Gallery** tabs.

To access the **Playbooks** page:

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher. The system displays the **Intelligent Operations** dashboard.

3. Navigate to **Playbooks** from the **Navigation Bar**.
The **Playbooks** page opens displaying the **List** tab.




Figure 64: Playbooks Page

List Tab

The **List** tab provides information about all the playbooks configured for the registered clusters. The following table describes the fields that appear in the **List** tab.

Note: The fields vary based on the **View by** and **Group by** options. A dash (-) is displayed in a field when a value is not available or applicable.

Table 36: List Tab – Field Description

Field	Description	Values
Name	Displays the playbook name.	(playbook name)
Autopilot Status	Displays the status of the autopilot feature.	(enabled, disabled)
Created Time	Displays the date and time when the playbook was created.	(date and time)
Description	Displays the user-provided description of the playbook.	(description)
Last Updated By	Displays the user who has last updated the playbook.	(user name)
Last Updated Time	Displays the date and time of last update to the playbook.	(date and time)
Owner	Displays the name of the user who created the playbook.	(user name)
Run on behalf of	Displays the name of the impersonated user assigned by the owner to run the playbook. For more information, see User Impersonation to Run a Playbook on page 147.	(user name)
Status	Displays the status of the playbook.	(enabled, disabled)

Field	Description	Values
Trigger	Displays the type of trigger configured for the playbook.	(trigger name)
Type	Displays whether the playbook is system generated or user created.	(system, custom)
Validated	Displays whether the playbook is validated according to the defined schema.	(true, false)

You can perform the following actions from the **List** tab:

- Click the name of a playbook to open the playbook details page, which displays the detailed information about the playbook.
- Create a playbook by clicking **Create Playbook**. For more information, see [Playbook Triggers](#) on page 129.
- Import any previously exported playbook by clicking **Import**. For more information, see [Exporting or Importing Playbooks](#) on page 171.
- Generate the playbook list as a report by clicking **Generate Reports**. For more information, see [Generating Playbooks Summary Report](#) on page 161.
- Group the playbooks based on pre-defined criteria.
- View playbooks based on a pre-defined criteria or create a custom view.
- Filter the playbooks list based on a variety of parameter values using the **Filters** pane.
- Perform the following actions on a playbook using the **Actions** dropdown menu. The **Actions** dropdown appears when a playbook is selected.

Table 37: Playbook Actions

Action	Description
Clone	Click this action to clone the playbook.
Update	Click this action to update the playbook.
Disable	Click this action to disable the playbook.
Share	Click this action to share the playbook. For more information, see Sharing a Playbook on page 160.
Delete	Click this action to delete the playbook.
Export	Click this action to export the playbook. For more information, see Exporting or Importing Playbooks on page 171.

Filters Pane - List Tab

You can filter the information in the **List** tab based on the following fields that are available in the **Filters** pane.

Table 38: Filter Pane Field Description - List Tab

Field	Description	Values
Enabled	Filters based on the status of the playbook.	(Enabled/Disabled)
Type	Filters based on the playbook type.	(System/Custom)
Trigger	Filters based on the trigger type.	(Alerts/Alerts Matching Criteria/Event/Manual/Time/Webhook)
Autopilot Status	Filters based on the status of the autopilot feature.	(Enabled/Disabled)
Owner	Filters based on the name of the user who created the playbook.	(Name of owner)
Run on behalf of	Filters based on the name of the impersonated user assigned by the owner to run the playbook.	(Name of user)

Plays Tab

Plays are the execution instances of playbooks that you can use to monitor the status of all the playbooks that you have executed.

The following table describes the fields that appear in the **Plays** tab.

Note: The fields vary based on the **View by** option. A dash (-) is displayed in a field when a value is not available or applicable.

Table 39: Plays Tab – Field Description

Field	Description	Values
Playbook Name	Displays the playbook name.	(playbook name)
Autopilot Status	Displays the status of the autopilot feature.	(enabled, disabled)
Cluster	Displays the name of the cluster for which this playbook is configured.	(cluster name)
Last Updated Time	Displays the time when the play finished executing.	(date and time)
Occurrence Time	Displays the date and time when the play occurred.	(date and time)
Run on behalf of	Displays the name of the impersonated user assigned by the owner to run the playbook. For more information, see User Impersonation to Run a Playbook on page 147.	(user name)
Source Entity	Displays the entity (such as a cluster, host, or VM name) to which the alert applies.	(entity name)
Status	Displays the status of the play.	(succeeded, failed)
Trigger	Displays the type of trigger configured for the playbook.	(trigger name)

You can perform the following actions from the **Plays** tab:

- Click the name of a play to open the details page of the play, which displays the detailed information about the play.
- Generate the plays list as a report by clicking **Generate Reports**. For more information, see [Generating Playbooks Summary Report](#) on page 161.
- View plays based on a pre-defined criteria or create a custom view.
- Filter the plays list based on a variety of parameter values using the **Filters** pane.
- Perform the following actions on a play using the **Actions** dropdown menu. The **Actions** dropdown appears when a play is selected.

Table 40: Plays Tab - Actions

Action	Description
Resume	Click this action to resume any play that is paused.
Abort	Click this action to prevent a paused play from resuming.

Filters Pane - Plays Tab

You can filter the information in the **Plays** tab based on the following fields that are available in the **Filters** pane.

Table 41: Filter Pane Field Description - Plays Tab

Field	Description	Values
Enabled	Filters based on the status of the playbook.	(Enabled/Disabled)
Type	Filters based on the playbook type.	(System/Custom)
Trigger	Filters based on the trigger type.	(Alerts/Alerts Matching Criteria/Event/Manual/Time/Webhook)
Autopilot Status	Filters based on the status of the autopilot feature.	(Enabled/Disabled)
Owner	Filters based on the name of the user who created the playbook.	(Name of owner)
Run on behalf of	Filters based on the name of the impersonated user assigned by the owner to run the playbook.	(Name of user)

Action Gallery Tab

The **Action Gallery** tab displays the actions that can be used in Playbooks. The actions available within the action gallery serve the most common IT automation use cases. For details on the available system actions, see the [Playbook Actions](#) on page 129 table.

You can clone system actions from the action gallery, to create custom actions to use in your Playbooks. You can also update or delete any cloned actions.

Note: You cannot clone actions with any password filled in the action parameter.




Figure 65: Action Gallery Tab

The following table describes the fields that appear in the **Action Gallery** tab.

Note: The fields vary based on the **View by** option. A dash (-) is displayed in a field when a value is not available or applicable.

Table 42: Action Gallery Tab – Field Description

Field	Description	Values
Name	Displays the action name.	(action name)
Action Type	Displays the type of action.	(name of action type)
Created Time	Displays the date and time when the action was created.	(date and time)
Description	Displays the user-provided description of the action.	(description)
Last Updated Time	Displays the date and time of last update to the action.	(date and time)
Owner	Displays the name of the user who created the action.	(user name)
Template Type	Displays the type of the template associated with the action.	(template name)

You can perform the following playbook actions from the **Action Gallery** tab:

- Click the name of a playbook action to open the action details page, which displays the detailed information about the action.
- Generate the playbook actions as a report by clicking **Generate Reports**. For more information, see [Generating Playbooks Summary Report](#) on page 161.
- View playbook actions based on a pre-defined criteria or create a custom view.
- Filter the playbook actions list based on a variety of parameter values using the **Filters** pane.
- Perform the following actions on a playbook action using the **Actions** dropdown menu. The **Actions** dropdown appears when an playbook action is selected.

Table 43: Action Gallery Tab - Actions

Action	Description
Clone	Click this option to clone the playbook action.
Update	Click this option to update the playbook action.
Delete	Click this option to delete the playbook action.

Filters Pane - Action Gallery Tab

You can filter the information in the **Action Gallery** tab based on the following fields that are available in the **Filters** pane.

Table 44: Filter Pane Field Description - Action Gallery Tab

Field	Description	Values
Name	Filters based on the name of the action.	(action name)
Action Type	Filters based on the action type.	(action types)
Template Type	Filters based on the template type.	(System/Custom)
Owner	Filters based on the name of the user who created the action.	(Name of owner)

Playbook Triggers

You can create playbooks based on the following triggers.

- [Alert Trigger](#) - The **Alert** trigger allows you to select an alert policy as a trigger. When an alert is generated from this policy, the action configured in the playbook is triggered.
- [Alerts Matching Criteria](#) - The **Alerts Matching Criteria** trigger allows you to trigger playbook on any alert matching specified criteria.
- [Event](#) - The **Event** trigger allows you to trigger playbook when an event of a certain type is generated.
- [Manual Trigger](#) - The **Manual** trigger allows you to select an entity type (VM, cluster, or host) to manually trigger the playbook from the **Explore** view for the selected entity type.
- [Time](#) - The **Time** trigger allows you to trigger playbook at the specified time.
- [Webhook](#) - The **Webhook** trigger allows triggering the playbook from a POST to the webhook URL.

Playbook Actions

Playbook actions are user configured tasks to be performed by a playbook.

To access the **Actions** menu:

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher. The system displays the **Intelligent Operations** dashboard.
3. Navigate to **Playbooks** from the **Navigation Bar**.

The **Playbooks** page opens displaying the **List** tab.

4. Click **Create Playbook**.

The **Create Playbook** window opens.

5. Select a Trigger of your choice.

6. Click **Add Action**.

Prism Central displays the Actions page. You can switch between the **Tiles** or **List** views by clicking the respective icons.

Note: The VM actions configured in playbooks have certain prerequisites based on the hypervisor (AHV or ESXi) and the AOS version. For more information, see [Prerequisites for VM Actions in Playbooks](#) on page 144.

You can integrate the following applications for Playbook actions:

- Slack - see **Slack** under *Communication Actions* in the table.
- MS Teams - see **MS Teams** under *Communication Actions* in the table.
- Ansible Tower - see **Ansible** under *Execution Actions* in the table.
- ServiceNow - see **Send Alert to ServiceNow** under *Alert Actions* in the table.
- PagerDuty - see **Send Alert to PagerDuty** under *Alert Actions* in the table.

You can manage vCenter-monitored entities through Playbook actions.

Table 45: Playbook Actions

Action	Action Description	Configurable attributes
VM Actions		
Eject CDROM	Ejects CDROM for a VM. Note: This action is supported only on VMs hosted on Nutanix.	You can eject only one CDROM or eject all CDROMs at a time. 1. Select the VM from the Target VM dropdown. 2. Choose the Disk to Eject. If you select Specific CDROM, mention the Bus Type and Index of the CDROM disk. If you select All CDROMs, system ejects all CDROMs from the target VM. 3. Select the Bus Type from the dropdown. PC supports ejecting only IDE or SATA bus types. 4. Type the Index of the CDROM. The index detail is present in the disk address, such as ide.1, and sata.3. For more information, see VM Details View section in the <i>Prism Central Infrastructure Guide</i>.

Action	Action Description	Configurable attributes
Add to Category	Add an entity like a VM, Host or Cluster to one or more Categories.	1. Select an Entity Type such as VM, Host or Cluster. 2. Select the target entity in the Target <entity> field. 3. Select the Category or categories you want to add the entity to.
Power Off VM	Power off a VM.	1. Select the VM from the Target VM dropdown menu. 2. Select the Type of Power Off Action (Guest Shutdown, ACPI, or Power Off). Note: Guest shutdown requires NGT to be installed on the VM.
Power On VM	Power on a VM.	Select the VM from the Target VM dropdown menu.
Remove from Category	Remove an entity like a VM, Host or Cluster from one or more Categories.	1. Select an Entity Type such as VM, Host or Cluster. 2. Select the target entity in the Target <entity> field. 3. Select the Category or categories you want to remove the entity from.
VM Add CPU	Add CPU for a VM.	1. Select the VM from the Target VM dropdown menu. 2. Enter the vCPUs to Add and Maximum Number of vCPUs. For more information, see Prerequisites for VM Actions in Playbooks on page 144.

Action	Action Description	Configurable attributes
VM Restore	Restores a VM from a Recovery Point.	- Select the VM from the Target VM dropdown menu. - Select Type. Clone or Revert Note: PC does not support reverting ESX VMs hosted on Nutanix infra. - Select Recovery Point Time. Latest or Latest before Date. Note: The revert action creates a new recovery point.
VM Add Disk	Add vDisk to a VM.	- Select the VM from the Target VM dropdown menu. - Select the Bus Type and enter the Disk Size in GiB.
VM Add Memory	Add memory for a VM.	- Select the VM from the Target VM dropdown menu. - Enter the Add Absolute Memory and Absolute Maximum.

Action	Action Description	Configurable attributes
VM Clone	Clone a VM.	1. Select a VM to clone from the Source VM dropdown menu. 2. Type a new name for the clone VM. 3. Choose one of the following: • One Clone • Multiple Clones Number of clones: Enter the number of VMs to clone from the source VM. You can clone up to 10 VMs. Starting Index Number: Enter the starting index number of the cloned VM. For example, if you enter 4 as starting index, and you are cloning 3 VMs, then the name of the cloned VMs are <cloned_VM_name>-4, <cloned_VM_name>-5, <cloned_VM_name>-6 4. (Optional) Enable Override CPU config. 5. (Optional) Enable Override Memory config. 6. Wait for job completion toggle - Prism Central waits for the current action to complete before beginning the next action. By default, the Wait for job completion toggle is turned on and the default value is 10 minutes. If you turn off the Wait for job completion toggle, then Prism Central does not wait for the current action to complete and begins the next action.

Action	Action Description	Configurable attributes
VM Expand Disk	Expand vDisk for a VM.	1. Select the VM from the Target VM dropdown menu. 2. Enter the Expand Disk By in GiB. 3. Enter the Maximum Disk Size in GiB. 4. Choose from: • Specific Disk - Select the Bus Type and enter the Index value for the selected disk type. • Most Used Disk
VM Reboot	Reboot a VM.	1. Select the VM from the Target VM dropdown menu. 2. Optionally, enable the Use Guest reboot if available option. This option is available for AHV hypervisors with PE running AOS 6.1 or greater. Also, NGT should be installed on the VM.
VM Recovery Point	Captures the state of a VM and saves it.	1. Select the VM from the Target VM dropdown menu. 2. Enter the Time To Live in number of days.
VM Reduce CPU	Reduce CPU for a VM.	1. Select the VM from the Target VM dropdown menu. 2. Enter the vCPUs to Remove and Minimum Number of vCPUs.
VM Reduce Memory	Reduce memory for a VM.	1. Select the VM from the Target VM dropdown menu. 2. Enter the value for Reduce Memory by and Minimum Limit.
Communication Actions		

Action	Action Description	Configurable attributes
Email	Send an email.	Enter the Recipient, Subject, Message, and Attachment. Note: - The Attachment field is applicable only if you have selected Generate Forecast Report as a previous action. - This action requires the SMTP server to be configured. For more information, see Configuring an SMTP Server in the <i>Prism Element Web Console Guide</i>.
Microsoft Teams	Send a message to Microsoft Teams	Enter the Webhook URL that you configured in Microsoft Teams and the Message . For more information, see Integrating MS Teams Action in Playbooks on page 173.
Slack	Send a Slack message.	Enter the Slack Token , Channel Name , and Message . For more information, see Integrating Slack Action in Playbooks on page 172.
Alert Actions		
Acknowledge Alert	Mark an Alert as acknowledged.	Select the target Alert from the Target Alert dropdown menu.
Resolve Alert	Mark an Alert as resolved.	Select the alert from the Target Alert dropdown menu.
Send Alert to PagerDuty	Send an Alert to PagerDuty.	Enter the Routing Key. The Routing Key is an integration key that you find in your PagerDuty instance.
Send Alert to ServiceNow	Send an Alert to one of the following ServiceNow target.	Enter the ServiceNow Instance Name , Username , and Password for the ServiceNow instance.
	- ServiceNow Instance - Mid Server	Enter the Mid Server Address , Mid Server Port , Username , and Password .

Action	Action Description	Configurable attributes
Send Alert to Webhook	Send a JSON Alert Object to a REST API endpoint.	Enter the Alert , Method , and URL for the outgoing webhook. Optionally, select the Username , Password , and Request Headers .
Report Actions		
Generate Forecast Report	Generate a just in time forecast report.	Select the cluster for which you want to generate the report and enter the days for the runway period.
Generate Report	Generate a custom report.	1. Select the Report Type as New Report. You can also select any of previously created reports to be generated. 2. Enter the Report Instance Name. 3. Optionally, enter a Description. 4. Select the Time Period for Report.
Wait Actions		
Wait for Some Time	Wait for a specified number of minutes.	Enter the time period in Minutes in the Wait for field. Select the action to be taken at the end of the entered time period. • Resume - automatically resume the play once the wait time is up. • Stop - requires the play to be manually resumed. Note: If a play is not resumed before the wait time is up, then it will be aborted. Optionally, check to enable Check trigger validity when the Playbook is resumed. If this option is enabled, the system checks if the alert that caused it to run is still valid or if it has been resolved already.

Action	Action Description	Configurable attributes
Wait until Day of Month	Wait until the next occurrence of a specified day and time of the month.	Select the date in the Wait until field and set the Time. Select the action to be taken at the end of the entered time period. • Resume - automatically resume the play once the wait time is up. • Stop - requires the play to be manually resumed. If a play is not resumed before the wait time is up, then it will be aborted. Optionally, check to enable Check trigger validity when the Playbook is resumed. If this option is enabled, the system checks if the alert that caused it to run is still valid or if it has been resolved already.
Wait until Day Of Week	Wait until the next occurrence of a specified day and time of the week.	Select the day in the Wait until field and set the Time. Select the action to be taken at the end of the entered time period. • Resume - automatically resume the play once the wait time is up. • Stop - requires the play to be manually resumed. Note: If a play is not resumed before the wait time is up, then it will be aborted. Optionally, check to enable Check trigger validity when the Playbook is resumed. If this option is enabled, the system checks if the alert that caused it to run is still valid or if it has been resolved already.
Execution Actions		

Action	Action Description	Configurable attributes
Execute a Self-Service Runbook	Executes a runbook that you created in Self-Service. Note: To use this action in the playbook, the NCM Self-Service must be enabled with appropriate license in your Prism Central instance.	1. Select the project to which the runbook belongs from the Project dropdown menu. 2. Select the Marketplace or Non-Marketplace radio button for Runbook Type. Select Marketplace if the runbook you want to associate with the playbook is published in the Marketplace. Select Non-Marketplace if the runbook you want to associate with the playbook is not published in the Marketplace. 3. Select the runbook from the Runbook dropdown menu. The runbooks appear in the Runbook dropdown menu based on the options you select for Project and Runbook Type. 4. (Optional) Select an endpoint from the Default Endpoint dropdown menu. The runbook uses the default endpoint only when you have not configured any endpoint at the task level in the runbook. 5. Specify the values for the variables defined in the runbook. The Variable fields appear when you select the runbook. The number and type of variables vary depending on the runbook you select. 6. (Optional) Turn on the Wait for job completion toggle if you want Prism Central to wait for the current action to complete before beginning the next action. If you turn on the Wait for job completion toggle, then specify the wait time for the job completion in the Wait Time field. The default value is 10 minutes.

Action	Action Description	Configurable attributes
Ansible	Run an Ansible Workflow or Job template.	Select the Template Type , and enter the Ansible Tower/ AWX Hostname and OAuth authentication Token to log in to the Tower or AWX host. For more information, see Integrating Ansible Action in Playbooks on page 175.
IP Address PowerShell	Remotely run a PowerShell script for a given VM IP address.	1. Enter the VM IP Address, Username, Password, and the Path to Script. 2. Enable or disable HTTPS.
IP Address SSH	Remotely run an ssh command for a given VM IP address.	Enter the VM IP Address, Username, Password, and the Command to Run. Note: SSH action has a timeout of 5 minutes. The execution of ssh command should complete within 5 minutes or it should be executed as a daemon. Otherwise, the action will be marked as failed even if command successfully executes after the timeout.

Action	Action Description	Configurable attributes
Puppet	Run a Puppet Task or Plan.	1. Enter Puppet Enterprise HostName. Fully qualified domain name of a host where the Puppet server is running. 2. The default Puppet Enterprise Port is 8143. The port number in which Puppet accepts connections. You can change the default port number. 3. Enter Authorization Token. The authorization token provided by the Puppet server. 4. Select Command Type Task or Plan. 5. Enter Task or Plan name depending on the Command Type. 6. Enter Scope. This is the JSON payload that determines set of nodes on which this task runs. Optionally, add other parameters shown in the screen.
REST API	Make a call to a REST API endpoint.	1. Select the Method. 2. Enter the URL. 3. Enable Use Stored Credentials. This uses the non-Nutanix system user access credentials stored in a PC. 4. Select Authorization Type and enter the Username, and Password. 5. Enter the Request Body, and Request Headers
VM Powershell	Remotely execute PowerShell script for a given VM.	1. Select the VM from the Target VM dropdown menu. 2. Enter the Username, Password, and the Path to Script. 3. Enable or disable HTTPS.

Action	Action Description	Configurable attributes
VM SSH	Remotely execute an ssh command for a given VM	1. Select the VM from the Target VM dropdown menu. 2. Enter the Username, Password, and the Command to Run.
Utility Actions		
Run Health Check	Runs health check and collects cluster report	1. Select the health check from the dropdown. 2. Select the cluster.
Run SQL Query	Run queries on SQL server VM	VM Type: Existing SQL Server VM or New SQL Server VM. Existing SQL Server VM • Select SQL VM. A list of Microsoft SQL servers that you have configured in the Monitoring Configurations page appears. • Select Database. A list of databases available in all SQL server VMs appears. • Type or paste SQL Query. PC supports only SQL SELECT operation. New SQL Server VM • Enter the IP Address or Hostname of the SQL server VM • Type Username and Password • Enter the Database name. • Type or paste SQL Query. PC supports only SQL SELECT operation.

Action	Action Description	Configurable attributes
Math Operation	Evaluates a math formula and return a result	- Enter Formula. You can use any math formula or function. - Type a Default Value. If the above formula fails to return a valid result, then the default value will be used.
String Patch	Update JSON, Regular Expression or XML string	Type of patch: JSON, Regular Expression or XML JSON: - Enter String to Patch. - Select Operation. - Enter Destination Path. Follow the format defined in the RFC-6902 to add the destination JSON path. - Enter the Value. Regular Expression: - Enter String to Patch. - Enter Regular Expression in RE2 format. - Enter Replacement String. XML - Enter String to Patch. - Select Operation. - Select Value. - Select Position. - Enter Destination Path. Follow the format defined in the RFC-5261 to add the destination JSON path. - Enter the Value.

Action	Action Description	Configurable attributes
Metric Action	Captures details of a given metric over a period of time	1. Select the Entity Type. 2. Select the Entity. 3. Select the Metric Name. The options displayed depends on the Entity Type selected earlier. 4. Mention the Monitor Duration in minutes. Monitoring duration is of minimum 5 minutes and maximum 60 minutes. 5. Select the Aggregate Type.
Lookup Cluster Details	Captures cluster details	Select the cluster from the Target Cluster dropdown menu.
Lookup Host Details	Captures host details	Select the host from the Target Host dropdown menu.
Lookup VM Details	Captures VM details	Select the VM from the Target VM dropdown menu.
String Parser	Parse data from a string. You can use this action to parse response from any other actions like REST API or SSH Action .	• String to Parse - Enter the string. • Format - Select one of the following options. • JSON - Enter the field to retrieve as JSONPath. • XML - Enter the node/tag/text to retrieve as XPath in the XPath field. Optionally, enter the comma separated list of namespaces as key=value in the XML Namespaces field. • Regular Expression - Enter the string to parse as regular expression
Conditional Actions		
Branch	Switch execution based on a matching criteria.	For more information, see Using Branch Action (Conditional Execution) on page 158.

Action	Action Description	Configurable attributes
Post Actions Flow	Executes an action at the end of a workflow.	You add Post Actions Flow action to the main branch only once. Post Actions Flow should be the last action in the main branch and it can have upto five actions in it. After you have created the playbook and added main actions, do the following: 1. Select Post Actions Flow action Note: The Post Actions Flow is not available as the first action. It should be added as the last action. 2. Choose post action flow conditions: • Run on playbook execution failure. • Run on playbook execution success. • Run when playbook is aborted. You can choose more than one conditions. 3. Click Add Action. You can add upto five actions under Post Actions Flow.

Prerequisites for VM Actions in Playbooks

The VM actions configured in playbooks have certain prerequisites based on the hypervisor (AHV or ESXi) and the AOS version. Refer to the following table to understand the prerequisites and other dependencies for VM actions defined in playbooks.

Note: VM actions are not supported for Hyper-V.

Table 46: VM Action Prerequisites

Action	Hypervisor	Prerequisites	Notes
VM Add CPU	AHV	None	- The VM has a guest operating system that supports CPU hot add functionality. - Number of cores per socket cannot be changed while VM is powered on.
	ESXi	None	The VM has a guest operating system that supports CPU hot add functionality.
VM Reduce CPU	AHV	VM must be powered off	
	ESXi	VM must be powered off	
VM Add Memory	AHV	None	The VM has a guest operating system that supports Memory hot add functionality.
	ESXi	None	- The VM has a guest operating system that supports Memory hot add functionality. - You cannot increase the memory above 3 GB due to vSphere restrictions.
VM Reduce Memory	AHV	VM must be powered off	
	ESXi	VM must be powered off	
Power On VM	AHV	None	
	ESXi	None	
Power Off VM (Guest shutdown or ACPI shutdown)	AHV	Nutanix Guest Tools (optional power functions)	
	ESXi	VMware guest tools (optional power functions)	

Action	Hypervisor	Prerequisites	Notes
VM Recovery Point	AHV	None	- Maximum Time to Live value allowed is 365 days. - Snapshot can not be created for VMs which are protected by protection domain of legacy DR.
	ESXi	None	- The VM snapshot action creates recovery points for all ESXi VMs. - Maximum Time to Live value allowed is 365 days. - Snapshot can not be created for VMs which are protected by protection domain of legacy DR.
VM Add Disk	AHV	AHV does not support hot add of SATA disks.	
	ESXi	PCI disk type is not supported.	
VM Expand Disk	AHV	None	
	ESXi	PCI disk type is not supported.	
Lookup VM Details	AHV	None	
	ESXi	None	
VM Clone	AHV	None	
	ESXi	None	
VM Powershell	AHV	Powershell remoting must be enabled.	
	ESXi	Powershell remoting must be enabled.	
VM Reboot	AHV	NGT and the Prism Element web console must support AOS 6.1 or later versions.	Default mode is ACPI reboot for which no prerequisite is required. Optionally user can select guest reboot for which the prerequisite is required.

Action	Hypervisor	Prerequisites	Notes
	ESXi	VMware guest tools	Guest reboot is not supported.
VM SSH	AHV	SSH must be enabled.	
	ESXi	SSH must be enabled.	

User Impersonation to Run a Playbook

The user impersonation feature allows Prism to run a playbook on behalf of another user (executor) other than the creator of the playbook. The creator of the playbook assigns another user as an executor while creating the playbook. When Prism runs a playbook, it executes the actions in the playbook with the permission set of the impersonated user.

To run a playbook successfully, the impersonated user must have all the required permissions to view alerts and run actions in the playbook.

An admin user creates a custom role with user impersonation settings and permissions to allow a playbook creator to impersonate a playbook executor.

A local user accounts and Active Directory or LDAP user accounts supports user impersonation.

Note:

Before impersonating a local user, ensure that the user has logged into the PC once before they are listed to be impersonated.

Currently playbooks are only available to an admin user. When Prism executes playbooks, it runs all actions in the playbook using administrator level permissions. With the user impersonation feature, that allows the system to impersonate a particular user when running playbooks, the actions in the playbook are executed with the permission set of the given user.

Here is the work flow of the user impersonation:

1. An admin creates a custom role with user impersonation permission and playbook permissions; assigns the role to a user A with permission on all playbooks and access to impersonate User B.
2. User A can now create a playbook and assign User B as executor of the playbook.

Here, the executor is the impersonated user.

3. Prism governs the impersonation operation by suitable access control policies at runtime.
4. When the trigger defined in the playbook occurs, the system runs all actions defined in the playbook on behalf of the execution user (impersonated user). All privileges are authorized against this runtime execution user. At this point, the playbook creator and his accesses are not relevant. All audit logs indicate that the impersonated user performed these tasks.

Note: For Webhook and Manual triggers, the system runs the playbook on behalf of the triggering user.

Playbook Running on non-Nutanix Systems

For playbooks actions running on non-Nutanix systems, the playbook fetches the credentials from the credential repository. A user saves the access credentials in the credential repository before creating a playbook.

For information about how to save the credentials, see [Storing Non-Nutanix Credentials in Repository](#) on page 268.

Create a Playbook

Creating a playbook involves selecting a trigger for the playbook and configuring the action that the playbook has to perform.

For more information on playbook triggers, see [Playbook Triggers](#) on page 129.

For more information on playbook actions, see [Playbook Actions](#) on page 129.

Creating Playbooks Using Alert

About this task

To create a playbook using a predefined alert, perform the following steps.

Note: Alert as trigger allows you to select only a specific alert that can be used in a playbook. To select multiple alerts based on a matching criteria, see [Creating Playbooks Using Alerts Matching Criteria](#) on page 151.

Procedure

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher. The system displays the **Intelligent Operations** dashboard.
3. Navigate to **Playbooks** from the **Navigation Bar**.
The **Playbooks** page opens displaying the **List** tab.

4. Click **Create Playbook**.

Note: If you are creating a playbook for the first time, a **Get Started** option appears. Click **Get Started**, and then click **Create Playbook**.

5. Click **Select** within the **Alert** (option) card.
For information on alert messages, see [Alerts Summary View \(Prism Central\)](#) in *Prism Central Infrastructure Guide*.
6. In the Search field, enter the name of an existing alert policy. The Search field displays a list of matching policy names as you type the characters. Select the policy from the list.
7. Configure the filter criteria to target the most relevant alerts and the entities associated with the alert.
8. Select the target entity type.
 - » **All Applicable Clusters** - The alerts from all the applicable clusters and matching the alert criteria are used as trigger.
 - » **Clusters in Categories** - The alerts from only the clusters (or entity type) within the specified categories and matching the alert criteria are used as trigger.
 - » **Specific Clusters** - The alerts from only the specified clusters and matching the alert criteria are used as trigger.

9. Click **Add Action** to view the list of all the available actions on the deck. Click **All Actions** to filter from a list of following action categories.

- **VM Actions**
- **Communication Actions**
- **Alert Actions**
- **Report Actions**
- **Wait Actions**
- **Execution Actions**
- **Utility Actions**
- **Branch Actions**

The **Parameters** option allows you to select certain inputs based on the **Trigger** that you have selected. Parameters also allows you to select any available input based on the previous actions in the series of actions that you have configured.




Figure 66: Parameters

To configure an action, set the action attributes based on your requirement. Refer to the [Playbook Actions](#) on page 129 table for action attributes and instructions on how to set the attributes. For each action, you must define the one of the following setting that you want to do in the event of action failure.

- **Stop** - Stop any further execution of the playbook when the action fails.
- **Continue** - Continue the execution of playbook when the action fails.

Note: You can select multiple actions for a trigger. To add another action, click **Add Action** and select the action. You can also add actions before or after a particular action.

10. (Optional) In the **Execute behalf of** field, select **Other**. Search the user and click **OK**.




Figure 67: User Impersonation

This step enables user impersonation. Prism runs the playbook on behalf of the user selected in this step. For more information, see [User Impersonation to Run a Playbook](#) on page 147.

11. (Optional) Configure Autopilot for playbook.
The Autopilot feature allows you to control the execution playbooks and automate the optimization of plays by adding specific metrics like key performance indicator (KPI) and optimal threshold for your automation tasks, see [Configuring Autopilot for Playbook](#) on page 150.
12. To save the playbook, do the following.
 - a. Click **Save and Close**.
 - b. Specify the **Name** and **Description**.
 - c. Select the **Self** option to execute the playbook by yourself or the **Other** option to enable user impersonation. Select a user and click **OK**. For more information, see [User Impersonation to Run a Playbook](#) on page 147.
 - d. Click the slider to enable the playbook. Only playbooks in the enabled state are executed.
 - e. Click **Save**.

Configuring Autopilot for Playbook

The Autopilot feature allows you to control the execution playbooks and automate the optimization of plays by adding specific metrics like key performance indicator (KPI) and optimal threshold for your automation tasks. You can use the autopilot to runs playbooks in the autopilot mode until the specified KPI is met.

About this task

To configure autopilot for new or existing playbooks, perform the following steps.

Before you begin

Note: The autopilot feature is only applicable for playbooks that have metric-based alerts as triggers.

Procedure

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher.
The system displays the **Intelligent Operations** dashboard.

3. Navigate to **Playbooks** from the **Navigation Bar**.
The **Playbooks** page opens displaying the **List** tab.
4. Click **Create Playbook**.
The **Create Playbook** window opens.
5. In the **Create Playbook** page, click **Set Autopilot**.
6. Select the **KPI Metric** from the dropdown menu. For example, for a **VM Memory Constrained** alert, select **Memory Usage (%)** as the KPI metric.
7. In **KPI Target**, select the Range option and specify the KPI target range value. You can also select **<** (less than) or **>** (greater than) options and specify a fixed value for the KPI target.
8. Enter the **Monitoring Duration** in minutes.
9. In the **Max # Attempts** field, enter the number of attempts that the playbook must run in the autopilot mode to achieve the KPI target.
10. (Optional) Select the **Wait for approval for every attempt** checkbox to get an alert (INFO level) for every execution attempt of the playbook.

If you select this option, you will have to resume the playbook after each run from the **Plays** page. The playbook remains in a suspended state until it is manually resumed after each run.
11. Click the **Autopilot Status** toggle button to enable or disable the Autopilot.
12. Click **Save**.

What to do next

You can view the details of the autopilot enabled playbooks from the **Plays** tab in the [Playbooks](#) page.

Creating Playbooks Using Alerts Matching Criteria

About this task

To create a playbook using alerts matching criteria as trigger, perform the following steps.

Procedure

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher.
The system displays the **Intelligent Operations** dashboard.
3. Navigate to **Playbooks** from the **Navigation Bar**.
The **Playbooks** page opens displaying the **List** tab.
4. Click **Create Playbook**.

Note: If you are creating a playbook for the first time, a **Get Started** option appears. Click **Get Started**, and then click **Create Playbook**.

5. Click **Select** within the **Alerts Matching Criteria** (option) card.

6. Select the following options for either **All Alert Policies** or **Specific Alert Policies**:

- **Impact Type** - Choose **Any** to select all impact types; or choose a single value or a combination of values from **Capacity**, **Performance**, **Configuration**, **Availability**, and **System Indicator** impact types.
- **Severity** - Choose **Any** to select alerts having any severity levels; or choose a single value or combination of values from **Critical**, **Warning**, and **Info** severity levels.
- **Cluster** - Choose **All Clusters** for alerts from all clusters or **Specific Clusters** for alerts from only the specified clusters.

For **Specific Alert Policies**, the available options are **Policies**, **Severity**, and **Cluster**.

The actions defined in this playbook will be available for the selected entity type in the entities view.

7. Click **Add Action** to view the list of all the available actions on the deck.

8. Click **All Actions** to filter from a list of following action categories.

- **VM Actions**
- **Communication Actions**
- **Alert Actions**
- **Report Actions**
- **Wait Actions**
- **Execution Actions**
- **Utility Actions**
- **Branch Actions**

To configure an action, set the action attributes based on your requirement. Refer to the [Playbook Actions](#) table for action attributes and instructions on how to set the attributes. For each action, you must define the required setting that you want to do in the event of action failure (Stop/Continue).

Note: You can select multiple actions for a trigger. To add another action, click **Add Action** and select the action.

9. To save the playbook, do the following.

- a. Click **Save and Close**.
- b. Specify the **Name** and **Description**.
- c. Select the **Self** option to execute the playbook by yourself or the **Other** option to enable user impersonation. Select a user and click **OK**. For more information, see [User Impersonation to Run a Playbook](#) on page 147.
- d. Click the slider to enable the playbook. Only playbooks in the enabled state are executed.
- e. Click **Save**.

Creating Playbooks Using Event

About this task

To create a playbook using event trigger, perform the following steps.

Procedure

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher.
The system displays the **Intelligent Operations** dashboard.
3. Navigate to **Playbooks** from the **Navigation Bar**.
The **Playbooks** page opens displaying the **List** tab.
4. Click **Create Playbook**.

Note: If you are creating a playbook for the first time, a **Get Started** option appears. Click **Get Started**, and then click **Create Playbook**.

5. Click **Select** within the **Event** (option) card.

Note: Event triggers (except anomalous events) are supported only for AOS (AHV or ESXi) entities. However, you can use anomalous events as triggers for vCenter-monitored entities as well. Except for anomaly detection, no other events are detected for vCenter-monitored entities.

6. Select an event type from the events dropdown menu.
For details on event messages, see [Events Summary View \(Prism Central\)](#) in *Prism Central Infrastructure Guide*.
7. Select the target entity type by choosing one of the following options.
 - **All Applicable VMs** - The selected event type for all the applicable VMs is used as the trigger.
 - **VMs in Categories** - The selected event type for the VMs in the specified categories is used as the trigger.
 - **Specific VMs** - The selected event type for the specified VMs is used as the trigger.
8. Click **Add Action** to view the list of all the available actions on the deck.
9. Click **All Actions** to filter from a list of following action categories.
 - **VM Actions**
 - **Communication Actions**
 - **Alert Actions**
 - **Report Actions**
 - **Wait Actions**
 - **Execution Actions**
 - **Utility Actions**
 - **Branch Actions**

To configure an action, set the action attributes based on your requirement. Refer to the [Playbook Actions](#) table for action attributes and instructions on how to set the attributes. For each action, you must define the required setting that you want to do in the event of action failure (Stop/Continue).

Note: You can select multiple actions for a trigger. To add another action, click **Add Action** and select the action.

10. To save the playbook, do the following.
 - a. Click **Save and Close**.
 - b. Specify the **Name** and **Description**.
 - c. Select the **Self** option to execute the playbook by yourself or the **Other** option to enable user impersonation. Select a user and click **OK**. For more information, see [User Impersonation to Run a Playbook](#) on page 147.
 - d. Click the slider to enable the playbook. Only playbooks in the enabled state are executed.
 - e. Click **Save**.

Creating Playbooks Using Time

Create a playbook with a time trigger to automate tasks. Set a one time or recurring schedule, add actions, and then save and enable the playbook.

About this task

To create a playbook using a time trigger, perform the following steps.

Procedure

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher.
The system displays the **Intelligent Operations** dashboard.
3. Navigate to **Playbooks** from the **Navigation Bar**.
The **Playbooks** page opens displaying the **List** tab.
4. Click **Create Playbook**.

Note: If you are creating a playbook for the first time, a **Get Started** option appears. Click **Get Started**, and then click **Create Playbook**.

5. Click **Select** within the **Time** (option) card.
6. Select the **Trigger Action(s)** from the following occurrence type.
 - **Once** - Specify the day, time, and timezone.
 - **At a recurring interval** - Specify the **Recurrence** interval from daily, weekly, or monthly along with the required interval details.

Note: Playbooks currently support only standard time zones. Daylight Saving Time (DST) adjustments are not considered. Ensure that all time-based configurations use standard time zones to maintain consistency and accuracy.

7. Click **Add Action** to view the list of all the available actions on the deck.

8. Click **All Actions** to filter from a list of following action categories.

- **VM Actions**
- **Communication Actions**
- **Alert Actions**
- **Report Actions**
- **Wait Actions**
- **Execution Actions**
- **Utility Actions**
- **Branch Actions**

To configure an action, set the action attributes based on your requirement. Refer to the [Playbook Actions](#) table for action attributes and instructions on how to set the attributes. For each action, you must define the required setting in the event of action failure (Stop/Continue).

Note: You can select multiple actions for a trigger. To add another action, click **Add Action** and select the action.

9. To save the playbook, follow these steps.

- a. Click **Save and Close**.
- b. Specify the **Name** and **Description**.
- c. Select the **Self** option to execute the playbook by yourself or the **Other** option to enable user impersonation. Select a user and click **OK**. For more information, see [User Impersonation to Run a Playbook](#) on page 147.
- d. Click the slider to enable the playbook. Only playbooks in the enabled state are executed.
- e. Click **Save**.

What to do next

Once you create and enable a playbook with time trigger, you can go to the list view and click on the playbook to see the next time it will be executed.

Creating Playbooks Using Manual Triggers

About this task

To create a playbook using manual trigger for a VM, host, or cluster, perform the following steps.

Procedure

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher.
The system displays the **Intelligent Operations** dashboard.
3. Navigate to **Playbooks** from the **Navigation Bar**.
The **Playbooks** page opens displaying the **List** tab.

4. Click **Create Playbook**.

Note: If you are creating a playbook for the first time, a **Get Started** option appears. Click **Get Started**, and then click **Create Playbook**.

5. Click **Select** within the **Manual** (option) card.

6. Select the **Entity Type** from the following options.

- VM
- Host
- Cluster

The actions defined in this Playbook will be available for the selected entity type in the entities view.

7. Click **Add Action** to view the list of all the available actions on the deck.

8. Click **All Actions** to filter from a list of following action categories.

- **VM Actions**
- **Communication Actions**
- **Alert Actions**
- **Report Actions**
- **Wait Actions**
- **Execution Actions**
- **Utility Actions**
- **Branch Actions**

To configure an action, set the action attributes based on your requirement. Refer to the [Playbook Actions](#) table for action attributes and instructions on how to set the attributes. For each action, you must define the required setting that you want to do in the event of action failure (Stop/Continue).

Note: You can select multiple actions for a trigger. To add another action, click **Add Action** and select the action.

9. To save the playbook, do the following.

- a. Click **Save and Close**.
- b. Specify the **Name** and **Description**.
- c. Select the **Self** option to execute the playbook by yourself or the **Other** option to enable user impersonation. Select a user and click **OK**. For more information, see [User Impersonation to Run a Playbook](#) on page 147.
- d. Click the slider to enable the playbook. Only playbooks in the enabled state are executed.
- e. Click **Save**.

Creating Playbooks Using Webhook

About this task

To create a playbook using a webhook as trigger, perform the following steps.

Procedure

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher.
The system displays the **Intelligent Operations** dashboard.
3. Navigate to **Playbooks** from the **Navigation Bar**.
The **Playbooks** page opens displaying the **List** tab.
4. Click **Create Playbook**.

Note: If you are creating a playbook for the first time, a **Get Started** option appears. Click **Get Started**, and then click **Create Playbook**.

5. Click **Select** within the **Webhook** (option) card.
6. Click **Add Action** to view the list of all the available actions on the deck.
7. Click **All Actions** to filter from a list of following action categories.

- **VM Actions**
- **Communication Actions**
- **Alert Actions**
- **Report Actions**
- **Wait Actions**
- **Execution Actions**
- **Utility Actions**
- **Branch Actions**

To configure an action, set the action attributes based on your requirement. Refer to the [Playbook Actions](#) table for action attributes and instructions on how to set the attributes. For each action, you must define the required setting that you want to do in the event of action failure (Stop/Continue).

Note: You can select multiple actions for a trigger. To add another action, click **Add Action** and select the action.

8. To save the playbook, do the following.
 - a. Click **Save and Close**.
 - b. Specify the **Name** and **Description**.
 - c. Select the **Self** option to execute the playbook by yourself or the **Other** option to enable user impersonation. Select a user and click **OK**. For more information, see [User Impersonation to Run a Playbook](#) on page 147.
 - d. Click the slider to enable the playbook. Only playbooks in the enabled state are executed.
 - e. Click **Save**.

After saving, open this playbook to view the instructions for using this trigger.

9. Once saved, open the playbook to view and copy the webhook parameters.
Use the parameters (**API Method**, **URL**, **ID**, and **Body**) and values shown to create an incoming webhook for the external application to trigger playbook.

Using Branch Action (Conditional Execution)

You can create an X-Play playbook with conditional branch action for a trigger. Branch action allows you to switch execution based on a matching criteria.

About this task

To use a branch action, perform the following steps.

Procedure

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher.
The system displays the **Intelligent Operations** dashboard.
3. Navigate to **Playbooks** from the **Navigation Bar**.
The **Playbooks** page opens displaying the **List** tab.
4. Click **Create Playbook**.
5. Select and configure one of the following trigger.
 - [Alert Trigger](#)
 - [Alerts Matching Criteria](#)
 - [Event](#)
 - [Manual Trigger](#)
 - [Time](#)
 - [Webhook](#)
6. Click **Add Action** and select **Branch**.

7. Define condition for the branch.

- **Branch** - Click the Edit icon to add a description for the conditional branch.
- **Condition** - Select either **If** or **Else** depending on the condition context that you want to set.

Note:

- You can create multiple branch actions for a trigger.
- Once the playbook is run, you can view the **Execution Flow** of your playbook. To view the execution flow of any previously run playbook, go to **Operations > Plays** and click on the playbook name. To view the complete execution flow, click **View Full Playbook**.

- **Operand** - Click **Parameters** to select a **Global** or **Trigger** entity.
- **Operator** - Select the required operator from the dropdown menu. You can choose **regex** for pattern matching.
- **Value** - Enter the value for the operator, or click **Parameters** to select a **Global** or **Trigger** entity.




Figure 68: Branch Action

8. Click **Add Action** to view the list of all the available actions on the deck.

To configure an action, set the action attributes based on your requirement. Refer to the [Playbook Actions](#) on page 129 table for action attributes and instructions on how to set the attributes.

Note: For each action branch condition, you must add at least one action.

9. Create the execution flow using more branch actions (**Else** option) and adding required actions for each condition.

You can also configure nested branch actions as required. To delete any added branch action, click the three-dots icon select **Delete**.

10. To save the playbook, do the following.
 - a. Click **Save and Close**.
 - b. Specify the **Name** and **Description**.
 - c. Select the **Self** option to execute the playbook by yourself or the **Other** option to enable user impersonation. For more information, see [User Impersonation to Run a Playbook](#) on page 147.
 - d. Click the slider to enable the playbook. Only playbooks in the enabled state are executed.
 - e. Click **Save**.

Sharing a Playbook

About this task

If you are associated with a role that has edit permissions, then you can share playbooks with other users and provide view or edit access to them.

Perform the following steps to share a user-defined playbook:

Procedure

1. Log in to Prism Central with a role that has edit permissions.
2. Select **Intelligent Operations** from application switcher.
The system displays the **Intelligent Operations** dashboard.
3. Navigate to **Playbooks** from the **Navigation Bar**.
The **Playbooks** page opens displaying the **List** tab.
4. Select the user-defined playbook that you want to share.
5. Click **Actions > Share**.
The **Sharing Playbook** window opens.
6. In the **Add Users & Groups** tab, do the following:
 - a. **Share with:** In the Search field, enter the name of the users or groups to share the playbook. The Search field displays a list of users or groups whose names match the letters that you type. Click the username or group name to add that user or group. You can add more than one user or a group.
 - b. **Permission Type:** You can provide view only or edit permissions to a user or group by selecting **View only** or **Editor** from the dropdown menu.
A user with **Editor** permission can perform these tasks:
 - Run a playbook
 - Edit a playbook
 - Share a playbook
 - c. In the **Manage Users & Groups** tab, you can change the existing permission type for users or groups. To change the permission type, select the type from the dropdown menu associated with the user name or group name.
You can also remove the access for a user or group.

7. Click **Save**.

A message that the playbook is shared successfully with users or groups appears.

Generating Playbooks Summary Report

Use the Generate Report feature in Prism Central to export dashboard data. You can download, email, or schedule reports in PDF or CSV format.

In Prism Central, you can monitor infrastructure entities by visualizing the data on the Entities page. However, the data is available only in the Prism Central interface. Generating a Report feature allows you to export the playbooks data as a report for further analysis, archival, compliance, and other purposes.

Note: Generating a report requires a valid license - Nutanix Cloud Manager (NCM) Starter, Pro or Ultimate. For more information, see [Nutanix Cloud Manager Licensing](#). The **Generate Report** option does not appear if the appropriate license is not available.

You can generate a report in the following ways:

- Download the playbooks data as a report in PDF format or CSV format.
- Email the playbooks data as a report to the recipients in PDF format or CSV format, or both, as an attachment.
- Save the playbooks data as a report configuration to schedule, share, and export reports.
- Save the playbooks data as a report view to use later while adding views to a report.

You can access the dropdown menu of the **Generate Report** option from the Playbooks page:




Figure 69: Playbooks - Generate Report Option

The following table describes the dropdown menu list of the **Generate Report** option:

Table 47: Generate Report Dropdown List

Options	Description
Download Report (csv, pdf)	Click this action to download the report. For more information, see Downloading a Report on page 162.
Email Report (csv, pdf)	Click this action to email the report. For more information, see Emailing a Report on page 163.

Options	Description
Save as Report Configuration	Click this action to save the report as a configuration. For more information, see Saving as a Report Configuration on page 163.
Save as Report View	Click this action to save the report as a view. For more information, see Saving as a Report View on page 164.

Downloading a Report

This task describes how to download the playbook list, plays, and actions page as a report, either in PDF format or CSV format.

About this task

Downloading a report has the following workflow:

- The report instance is generated in your preferred format and notified on the playbook page.
- The generated report instance is available for you to download under **Generated Reports** tab of the Reports page.

Important:

- You can also resend the generated report instance as an email from the **Generated Reports** tab of the Reports page. For more information, see [Resending a Report Instance](#).
- You must download or resend the report within 24 hours from the time you generated the report instance. After 24 hours, Prism Central removes the report instance from the **Generated Reports** tab.

Procedure

1. Log in to Prism Central.
2. Select the **Intelligent Operations** application from **Application Switcher** function, and navigate to **Playbooks** from the **Navigation Bar**.

The Playbooks page displays the **List** tab by default. The **List** tab provides information about all the playbooks configured for the registered clusters.
3. Select **List**, **Plays**, or **Action Gallery** tab.
4. Click **Generate Report**.
5. Select **Download Report (csv, pdf)** and enter the following information:
 - a. **Report Name:** Enter the report name.
 - b. **Report Format:** Select PDF format or CSV format.
6. Click **Download**.
Prism Central displays a confirmation message on the dashboard screen that the report has been successfully generated.

What to do next

Select the **Intelligent Operations** application from **Application Switcher** function, and navigate to **Reports** from the **Navigation Bar**. Select **Generated Reports** tab to view and download the generated report instance. For more information, see [Downloading a Report Instance](#).

Emailing a Report

This task describes how to email the playbook list, plays, and actions page as a report.

About this task

You can email the playbook page reports as an attachment in PDF format, CSV format, or both. After you email the report, it is also available as a report instance under **Generated Reports** tab of the Reports page to resend or download the report.

Important: Ensure to resend or download the report within 24 hours from the time you generated the report instance. After 24 hours, Prism Central removes the report instance from the **Generated Reports** tab.

Procedure

1. Log in to Prism Central.
2. Select the **Intelligent Operations** application from **Application Switcher** function, and navigate to **Playbooks** from the **Navigation Bar**.

The Playbooks page displays the **List** tab by default. The **List** tab provides information about all the playbooks configured for the registered clusters.
3. Select **List**, **Plays**, or **Action Gallery** tab.
4. Click **Generate Report**.
5. Select **Email Report (csv, pdf)** and enter the following information:
 - a. **Recipients:** Enter the recipients (comma-separated) email address of the report.
 - b. **Subject:** Enter the subject of the email.
 - c. **Body:** Enter the desired text.
 - d. **Report Format:** Select the **PDF** checkbox or **CSV** checkbox or both to set a report format.
6. Click **Send**.

What to do next

Select the **Intelligent Operations** application from **Application Switcher** function, and navigate to **Reports** from the **Navigation Bar**. Select **Generated Reports** tab to view and manage the generated report instance. For more information, see [Managing a Report Instance from Generated Reports Tab](#).

Saving as a Report Configuration

This task describes how to save the playbook list, plays, and actions page as a report configuration.

About this task

You can save the playbook page as a report configuration. The configuration of the report that you save is available under **Configurations** tab of the Reports page to perform report management actions.

Procedure

1. Log in to Prism Central.
2. Select the **Intelligent Operations** application from **Application Switcher** function, and navigate to **Playbooks** from the **Navigation Bar**.

The Playbooks page displays the **List** tab by default. The **List** tab provides information about all the playbooks configured for the registered clusters.
3. Select **List**, **Plays**, or **Action Gallery** tab.
4. Click **Generate Report**.
5. Select **Save as Report Configuration** and enter the following information:
 - **Report Configuration Name:** Enter the report configuration name.
6. Click **Save**.

What to do next

Select the **Intelligent Operations** application from **Application Switcher** function, and navigate to **Reports** from the **Navigation Bar**. Select **Configurations** tab to view and manage the saved report configuration. For more information, see [Managing a Custom Report Configuration](#).

Saving as a Report View

This task describes how to save the playbooks list, plays, and actions page as a report view.

About this task

You can save the playbooks page as a report view. The report view you save is available under the **Saved** tab of the **ADD VIEWS** pane. Use the saved views when adding views to a new report or an existing report.

About this task

Procedure

1. Log in to Prism Central.
2. Select the **Intelligent Operations** application from **Application Switcher** function, and navigate to **Playbooks** from the **Navigation Bar**.

The Playbooks page displays the **List** tab by default. The **List** tab provides information about all the playbooks configured for the registered clusters.
3. Select **List**, **Plays**, or **Action Gallery** tab.
4. Click **Generate Report**.
5. Select **Save as Report View** and enter the following information:
 - **Report View Name:** Enter the report view name.
6. Click **Save**.

What to do next

Use the saved views while creating a report. For more information, see [Adding a Saved View](#).

Running a Playbook (Manual Trigger)

About this task

You can run a playbook that has a manual trigger associated with the playbook.

Procedure

1. Log in to Prism Central.
2. Select **Infrastructure** from application switcher, and navigate to **Compute > VMs** from the **Navigation Bar**.
The **VMs** page opens displaying the **List** tab.
To run playbooks with manual trigger for other entities like **Clusters** or **Hosts**, go to **Hardware > Clusters** or **Hardware > Hosts**.
3. Select a VM, host, or cluster from the list and click **Actions > Run Playbook**.
4. In the **Run Playbook** window that opens, select the playbook that you want to run for the selected entity, and click **Run**.

Configuring Manual Parameters

About this task

Parameters are the values that are dynamically filled in at **Playbook** runtime. You can either choose the parameters using the **Parameters** option at the time of creating **Playbooks**, or enter the parameters manually. Configuring manual parameters is useful if you want to use a parameter while cloning an action in the **Action Gallery**.

Note: This task demonstrates the usage of configuring parameters manually with a scenario on cloning the default **Email** action.

Procedure

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher.
The system displays the **Intelligent Operations** dashboard.
3. Navigate to **Playbooks** from the **Navigation Bar**.
The **Playbooks** page opens displaying the **List** tab.
4. Click **Action Gallery**.
5. Select **Email** from the Actions list.
6. Click **Actions > Clone**.
7. In the window that opens, edit the **Name** of the **Action**.
8. In the **Recipient** field, enter the recipient email address.
9. In the **Subject** field, enter Playbook `{{playbook.playbook_name}}` addressed Alert `{{trigger[0].alert_entity_info.name}}`.
Here, `{{playbook.playbook_name}}` and `{{trigger[0].alert_entity_info.name}}` are the available parameters. For the complete list of parameters, see [Parameters for Triggers and Actions](#) on page 167.

10. In the **Message** field, enter Alert, `{{trigger[0].alert_entity_info.name}}`, has been addressed. Action has been taken on `{{trigger[0].source_entity_info.name}}`.
11. Click **Copy** to complete.

Copy of Alert Email ? | X

Name

Copy of Alert Email

Description

Send an email after acting on an alert.

This action requires the SMTP server to be configured.

Recipient

Subject

Playbook `{{playbook.playbook_name}}` addressed Alert

Message

Alert, `{{trigger[0].alert_entity_info.name}}`, has been addressed. Action has been taken on `{{trigger[0].source_entity_info.name}}`.

Copy

Figure 70: Configuring parameters manually during Clone operation

What to do next

The parameters are dynamically recognized when the cloned action is selected while configuring a **Playbook**.

Note:

- Not all fields accept parameters. The fields with a **Parameters** hyperlink in the Playbook creation workflow accept parameters as inputs.
- If you enter a parameter string into a field that does not accept a parameter, it is read as a string and no value is filled in dynamically.
- Validation errors arise if a parameter of incorrect data type is supplied to a field that is expecting another data type. For the complete list of supported parameters, accepted data type and parameter format, see [Parameters for Triggers and Actions](#) on page 167.

Parameters for Triggers and Actions

Refer to the following table for the complete list of output parameters available for Playbook triggers.

Table 48: Triggers and Parameters

Trigger	Output Parameter	API Parameter
Alert	Alert (object)	{{trigger[0].alert_entity_info}}
	Source Entity (object)	{{trigger[0].source_entity_info}}
	Severity (string)	{{trigger[0].severity}}
	Creation Time (long)	{{trigger[0].creation_time}}
	Description (string)	{{trigger[0].description}}
Alerts Matching Criteria	Alert (object)	{{trigger[0].alert_entity_info}}
	Source Entity (object)	{{trigger[0].source_entity_info}}
	Severity (string)	{{trigger[0].severity}}
	Source Cluster (object)	{{trigger[0].cluster_entity_info}}
	Impact Type (string)	{{trigger[0].impact_type}}
	Creation Time (long)	{{trigger[0].creation_time}}
	Description (string)	{{trigger[0].description}}
Event	Event (object)	{{trigger[0].audit_entity_info}}
	Source Entity (object)	{{trigger[0].source_entity_info}}
	Creation Time (long)	{{trigger[0].creation_time}}
	Username (string)	{{trigger[0].user_name}}
	Client IP (string)	{{trigger[0].client_ip}}
	Source Cluster (object)	{{trigger[0].cluster_name}}

Trigger	Output Parameter	API Parameter
Manual	Source Entity (object)	{{trigger[0].source_entity_info}}
	Playbook UUID (string)	{{trigger[0].action_rule_uuid}}
Time	Trigger Time (string)	{{trigger[0].timestamp}}
	Playbook (object)	{{trigger[0].rule_entity}}
Webhook	Webhook Id (string)	{{trigger[0].webhook_id}}

Refer to the following table for the complete list of output parameters available for Playbook actions.

Table 49: Actions and Parameters

Action	Output Parameter	API Parameter
Branch Actions		
Branch	Selected Branch (string)	{{action[index].selected_branch}}
VM Actions		
Add to Category	Task (object)	{{action[index].task_entity}}
Power Off VM	Task (object)	{{action[index].task_entity}}
Power On VM	Task (object)	{{action[index].task_entity}}
Remove from Category	Task (object)	{{action[index].task_entity}}
VM Add CPU	Task (object)	{{action[index].task_entity}}
	Resultant vCPU count (long)	{{action[index].result_vcpu_count}}
VM Add Disk	Added Disk (string)	{{action[index].disk_name}}
	Task (object)	{{action[index].task_entity}}
VM Add Memory	Task (object)	{{action[index].task_entity}}
	Resultant Memory (size)	{{action[index].total_memory}}
VM Clone	Task (object)	{{action[index].task_entity}}
VM Expand Disk	Expanded Disk (object)	{{action[index].disk}}
	Resultant Disk Size (size)	{{action[index].resultant_disk_size}}
VM Reboot	Task (object)	{{action[index].task_entity}}
VM Recovery Point	Recovery Point (object)	{{action[index].vm_recovery_point_entity}}
	Task (object)	{{action[index].task_entity}}
VM Reduce CPU	Task (object)	{{action[index].task_entity}}
	Resultant vCPU count (long)	{{action[index].result_vcpu_count}}
VM Reduce Memory	Task (object)	{{action[index].task_entity}}
	Resultant Memory (size)	{{action[index].total_memory}}

Action	Output Parameter	API Parameter
Communication Actions		
Email	None	None
Microsoft Teams	None	None
Slack	Channel Id (string)	{{action[index].channel_id}}
	Message Timestamp (string)	{{action[index].timestamp}}
Alert Actions		
Acknowledge Alert	Task (object)	{{action[index].task_entity}}
Resolve Alert	Task (object)	{{action[index].task_entity}}
Send Alert to PagerDuty	None	None
Send Alert to ServiceNow	None	None
Send Alert to Webhook	Response Status Code (long)	{{action[index].status_code}}
	Response Headers (string)	{{action[index].response_headers}}
	Response Body (string)	{{action[index].response_body}}
Report Actions		
Generate Forecast Report	Report PDF (file)	{{action[index].report}}
Generate Report	Report (file)	{{action[index].report}}
	Task (object)	{{action[index].task_entity}}
Wait Actions		
Wait for Some Time	Resume Time (string)	{{action[index].resume_time}}
	Planned Resume Time (string)	{{action[index].planned_resume_time}}
Wait until Day of Month	Resume Time (string)	{{action[index].resume_time}}
	Planned Resume Time (string)	{{action[index].planned_resume_time}}
Wait Until Day of Week	Resume Time (string)	{{action[index].resume_time}}
	Planned Resume Time (string)	{{action[index].planned_resume_time}}
Execution Actions		
Ansible	Job ID (string)	{{action[index].job_id}}
IP Address Powershell	Standard Output (string)	{{action[index].stdout}}
	Standard Error (string)	{{action[index].stderr}}
	Return Code (long)	{{action[index].ret_code}}
IP Address SSH	Standard Output (string)	{{action[index].stdout}}
	Standard Error (string)	{{action[index].stderr}}
	Return Code (long)	{{action[index].ret_code}}

Action	Output Parameter	API Parameter
Puppet	Job Name (string)	{{action[index].job_name}}
REST API	Response Status Code (long)	{{action[index].status_code}}
	Response Headers (string)	{{action[index].response_headers}}
	Response Body (string)	{{action[index].response_body}}
VM Powershell	Standard Output (string)	{{action[index].stdout}}
	Standard Error (string)	{{action[index].stderr}}
	Return Code (long)	{{action[index].ret_code}}
VM SSH	Standard Output (string)	{{action[index].stdout}}
	Standard Error (string)	{{action[index].stderr}}
	Return Code (long)	{{action[index].ret_code}}
Utility Actions		
Lookup Cluster Details	Cluster Name (string)	{{action[index].cluster_name}}
	Cluster UUID (string)	{{action[index].cluster_uuid}}
	Cluster VIP (string)	{{action[index].cluster_vip}}
	CVM IP List (string)	{{action[index].cvm_ip_list}}
	Storage Capacity (size)	{{action[index].storage_capacity}}
	CPU Count (long)	{{action[index].cpu_counts}}
	Memory Size (size)	{{action[index].memory}}
	Categories (object)	{{action[index].categories}}
Lookup Host Details	Host Name (string)	{{action[index].host_name}}
	Host UUID (string)	{{action[index].host_uuid}}
	Host IP (string)	{{action[index].host_ip}}
	VM Count (long)	{{action[index].vm_count}}
	CPU Count (long)	{{action[index].cpu_count}}
	Memory Size (size)	{{action[index].memory}}
	Disk Capacity (size)	{{action[index].disk_capacity}}
	Hosting Cluster VIP (string)	{{action[index].hosting_cluster_vip}}
	Hosting CVM IP (string)	{{action[index].hosting_cvm_ip}}
	Hosting Maintenance Mode Status (string)	{{action[index].host_maintenance_mode_status}}
	Categories (object)	{{action[index].categories}}
Lookup VM Details	VM Name (string)	{{action[index].vm_name}}
	VM UUID (string)	{{action[index].vm_uuid}}

Action	Output Parameter	API Parameter
	VM IP Address (string)	{{action[index].vm_ip}}
	Host IP (string)	{{action[index].host_ip}}
	VM Hypervisor Type (string)	{{action[index].hypervisor_type}}
	Cores per vCPU (long)	{{action[index].vm_cores_per_vcpu}}
	vCPU Count (long)	{{action[index].vcpus}}
	Memory Size (size)	{{action[index].memory_size}}
	Disk Capacity (size)	{{action[index].disk_capacity}}
	Power State (string)	{{action[index].power_state}}
	Hosting Cluster VIP (string)	{{action[index].hosting_cluster_vip}}
	Hosting CVM IP (string)	{{action[index].hosting_cvm_ip}}
	VM Owner Name (string)	{{action[index].owner_name}}
	Categories (object)	{{action[index].categories}}
	Number of Cores (long)	{{action[index].total_cores}}
String Parser	Parsed String (string)	{{action[index].parsed_data}}

Exporting or Importing Playbooks

You can export a playbook for backup reasons or import any previously exported playbook. If you export a playbook using a Prism Central instance, you can import that playbook to another Prism Central instance that is running the same or newer version only.

About this task

Perform the following steps to export or import a playbook.

Procedure

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher.
The system displays the **Intelligent Operations** dashboard.
3. Navigate to **Playbooks** from the **Navigation Bar**.
The **Playbooks** page opens displaying the **List** tab.
4. Select the playbook that you want to export.
5. Click **Actions > Export**.
The **Export Playbooks** window opens.
6. (Optional) Edit the default **Filename**.
7. Click **Export**.
The playbook backup is downloaded in the .pbk format.

To import a playbook, perform the following steps.

8. Click **Import** from the **List** tab.
The **Import Playbooks** window opens.
9. Click **Browse** and select the previously exported playbook backup file.
10. Click **Import**.
The playbook is imported to Prism Central.

X-Play Integrations

This section provides information about integrating the various tools such as MS Teams, Ansible and Slack.

Integrating Slack Action in Playbooks

You can create a Playbook to send a Slack message. The Playbook sends the message from Prism Central to the defined channels in a Slack instance.

Before you begin

You must integrate Slack to automate receiving messages in Slack. To do this, you must create a new app in Slack, set the necessary app permissions, install the app in your workspace, and retrieve the Slack authorization token. For more information, see *Basic app setup* in [Slack documentation](#).

About this task

Perform the following steps to create a Playbook to send a Slack message.

Procedure

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher.
The system displays the **Intelligent Operations** dashboard.
3. Navigate to **Playbooks** from the **Navigation Bar**.
The **Playbooks** page opens displaying the **List** tab.
4. Click **Create Playbook**.

Note: If you are creating a playbook for the first time, a **Get Started** option appears. Click **Get Started**, and then click **Create Playbook**.

The **Select a Trigger** page appears.

5. Select a trigger of your choice.
6. Click **Add Actions**.
A list of all the available actions on the deck appears.
7. Click **Select** within the **Slack** (option) card.
8. Enter the **Slack Token**.

The Slack Token is an OAuth access token that you can find in your Slack instance. For more information, see [Slack documentation](#).

9. Enter the **Channel Name**.

Type the channel name that should receive the slack message that you describe in Step 8.

10. Type a **Message**.

In the message, you can mention users and channels. The mentioned users and members of the channels get slack notifications.

To mention a user, enter @ symbol and enter the email username of the user. For example: if the email ID of the user is [emailusername@company.com](#) then to mention a user, type: @emailusername.

Do not use a display name to mention a user in the message.

To mention a channel, enter the # symbol and enter the name of the channel.

11. (Optional) Select **Parameters**.

The Parameters option allows you to select certain inputs based on the trigger that you have selected. Parameters also allows you to select any available input based on the previous actions in the series of actions that you have configured.

12. To save the playbook, do the following.

- a. Click **Save and Close**.
- b. Specify the **Name** and **Description**.
- c. Select the **Self** option to execute the playbook by yourself or the **Other** option to enable user impersonation. Select a user and click **OK**. For more information, see [User Impersonation to Run a Playbook](#) on page 147.
- d. Click the slider to enable the playbook. Only playbooks in the enabled state are executed.
- e. Click **Save**.

Integrating MS Teams Action in Playbooks

This action is used to send a message (without attachments) to a Microsoft Teams channel using a configured webhook URL.

Before you begin

You must configure a webhook for your Microsoft Teams channel to allow Prism Central to share content in the Microsoft Teams channel. For more information, see *Create Incoming Webhooks* in [Microsoft documentation](#).

About this task

Perform the following steps to create a Playbook to send a message to a Microsoft Teams channel.

Procedure

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher.
The system displays the **Intelligent Operations** dashboard.
3. Navigate to **Playbooks** from the **Navigation Bar**.
The **Playbooks** page opens displaying the **List** tab.
4. Click **Create Playbook**.

Note: If you are creating a playbook for the first time, a **Get Started** option appears. Click **Get Started**, and then click **Create Playbook**.

The **Select a Trigger** page appears.

5. Select a trigger of your choice.
6. Click **Add Actions**.
A list of all the available actions on the deck appears.
7. Click **Select** within the **Microsoft Teams** (option) card.
8. Enter the **Webhook URL**.

The Webhook URL is a unique URL found in your MS Teams instance. You can use it to send information to Microsoft Teams. For more information, see *Create Incoming Webhooks* in [Microsoft documentation](#).
9. Type a **Message**.
10. (Optional) Select **Parameters**.

The Parameters option allows you to select certain inputs based on the trigger that you have selected. Parameters also allows you to select any available input based on the previous actions in the series of actions that you have configured.
11. To save the playbook, do the following.
 - a. Click **Save and Close**.
 - b. Specify the **Name** and **Description**.
 - c. Select the **Self** option to execute the playbook by yourself or the **Other** option to enable user impersonation. Select a user and click **OK**. For more information, see [User Impersonation to Run a Playbook](#) on page 147.
 - d. Click the slider to enable the playbook. Only playbooks in the enabled state are executed.
 - e. Click **Save**.

Ansible Integration in Playbooks

The Ansible action in X-play can be used to trigger job or workflow templates in your Ansible Tower or AWX instance. Refer to the following sections to understand the prerequisites and other dependencies for Ansible actions defined in playbooks.

- For information about job templates and how to create them, see [Ansible documentation](#).
- For information about workflow templates and how to create them, see [Ansible documentation](#).
- For information about how to create an application in Ansible Tower, see *Applications* information in [Ansible documentation](#).
- For information about how to create application token in Ansible Tower, see *Users - Tokens* information in [Ansible documentation](#).

Note:

- The X-Play Ansible action internally uses the following Tower APIs to launch the job and workflow templates respectively.
- If you need to pass extra variables to your job or workflow template, ensure that you select **PROMPT ON LAUNCH** for **EXTRA VARIABLES** while creating your template from the Tower UI.
 - [/api/v2/job_templates/{id}/launch/](/api/v2/job_templates/{id}/launch/)
 - [/api/v2/workflow_job_templates/{id}/launch/](/api/v2/workflow_job_templates/{id}/launch/)

Integrating Ansible Action in Playbooks

The Ansible action uses a token-based authentication to integrate with Ansible Tower.

Before you begin

You must create a job template, a workflow template, an application, and an application token in Ansible Tower to integrate Prism Central with Ansible Tower. For more information, see [Ansible documentation](#).

About this task

Perform the following steps to create a Playbook to trigger job or workflow templates in your Ansible Tower or AWX instance.

Procedure

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher.
The system displays the **Intelligent Operations** dashboard.
3. Navigate to **Playbooks** from the **Navigation Bar**.
The **Playbooks** page opens displaying the **List** tab.
4. Click **Create Playbook**.

Note: If you are creating a playbook for the first time, a **Get Started** option appears. Click **Get Started**, and then click **Create Playbook**.

The **Select a Trigger** page appears.

5. Select a trigger of your choice.
6. Click **Add Actions**.
A list of all the available actions on the deck appears.
7. Click **Select** within the **Ansible** (option) card.
8. Select **Template Type**.
9. Enter the appropriate details into the following fields:
 - **Workflow Template Name** or **Job Template Name** (as per the selected template type): Enter the name for the workflow template.
 - **Tower/AWX Hostname:** Enter the name of the Ansible tower or the AWX host.
 - **Token:** Enter the OAuth authentication token created in the Ansible Tower.
This token is required to log in to the Tower or AWX host. For information, see *Users - Tokens* information in [Ansible documentation](#).
 - (Optional) **Request Body JSON (Optional):** Enter any other parameters that you want to use in JSON format.
10. (Optional) Select **Parameters**.
The Parameters option allows you to select certain inputs based on the trigger that you have selected. Parameters also allows you to select any available input based on the previous actions in the series of actions that you have configured.
11. (Optional) Turn on the toggle for **Wait for job completion?** if you want Prism Central to wait for the current action to complete before beginning the next action.

12. To save the playbook, do the following.
 - a. Click **Save and Close**.
 - b. Specify the **Name** and **Description**.
 - c. Select the **Self** option to execute the playbook by yourself or the **Other** option to enable user impersonation. Select a user and click **OK**. For more information, see [User Impersonation to Run a Playbook](#) on page 147.
 - d. Click the slider to enable the playbook. Only playbooks in the enabled state are executed.
 - e. Click **Save**.

ServiceNow Integration with Prism Central (X-Play Support)

ServiceNow Integration Overview

ServiceNow provides IT services management (ITSM), and IT operations management (ITOM) as a software service.

ServiceNow integration with Prism Central allows seamless incident management for Prism Central alerts using ServiceNow. Export of alerts from Prism Central is achieved by creating an X-Play playbook in Prism Central. The X-Play Action (Send to ServiceNow) pushes the Prism Central alerts to ServiceNow to be processed by the ServiceNow event management pipeline.

Note: To integrate Prism Central with ServiceNow, enable the **Event Management** plug-in on your ServiceNow instance.

Enabling Event Management Plug-in on ServiceNow (Production Instance)

The ServiceNow **Event Management** plugin requires a separate subscription and can be activated by installing it from the ServiceNow store. For information about how to enable the plugin for your production instance, see [ServiceNow documentation](#).

Enabling Event Management Plug-in on ServiceNow (Developer Instance)

For information about how to enable X-Play support on the ServiceNow developer instance, see *Activate a plugin on a personal developer instance* in [ServiceNow documentation](#).

Configuring Alert Playbook for ServiceNow Alert Management

About this task

You can create a **Playbook** to send Prism Central alerts to ServiceNow.

Procedure

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher.
The system displays the **Intelligent Operations** dashboard.
3. Navigate to **Playbooks** from the **Navigation Bar**.
The **Playbooks** page opens displaying the **List** tab.
4. Click **Create Playbook**.

Note: If you are creating a playbook for the first time, a **Get Started** option appears. Click **Get Started**, and then click **Create Playbook**.

5. Click **Select** within the **Alerts Matching Criteria** (option) card.

This option allows you to trigger playbook on any alert matching specified criteria such as severity levels, impact type, and so on. For more information, see [Creating Playbooks Using Alerts Matching Criteria](#) on page 151.

6. Click **Add Action**.

7. Click **Select** within the **Send Alert to ServiceNow** (option) card.

8. Configure the ServiceNow action.

Note: This action requires the *Event Management Plug-in* to be enabled on the ServiceNow instance. For more information, see [Enabling Event Management Plug-in on ServiceNow \(Production Instance\)](#) on page 176.

- a. Click **Parameters** and select Alert Matching Criteria: Alert in the **Alert** dropdown menu.
- b. Enter any additional information in the **Additional Info (Optional)** field.
- c. Select the appropriate ServiceNow instance based on the ServiceNow deployment for the **Send to** parameter. Select **Cloud Instance**, **Private Instance**, or **MID Server**.

If you select **MID Server**, you need to provide the MID Server IP address and Port details in next sections

Note: The option to send the alert to ServiceNow is a default action in X-Play that calls the [ServiceNow Event Collection API](#) when a Nutanix alert matches the criteria specified.

- d. Enter the **ServiceNow Instance Name** in case of **Cloud Instance** or **Private Instance**. Enter the **MID Server IP Address** and **MID Server Port** details in case you selected **MID Server** in previous step.

Note: Do not enter the complete URL of your ServiceNow instance in the **ServiceNow Instance Name** field. If your instance URL is `https://foo.servicenow.com`, the **ServiceNow Instance Name** you must enter is `foo`.

- e. Enter the **Username** and **Password** for the ServiceNow instance.

Figure 71: ServiceNow Action

9. To save the playbook, do the following.

- a. Click **Save & Close**.
- b. Specify the **Name** and provide a **Description**.
- c. Select the **Self** option to execute the playbook by yourself or the **Other** option to enable user impersonation. Select a user and click **OK**. For more information, see [User Impersonation to Run a Playbook](#) on page 147.
- d. Click the slider to enable the playbook. Prism Central executes only the playbooks that are enabled.
- e. Click **Save**.

All Prism Central alerts of the specified type are forwarded to ServiceNow with the execution of the playbook created above.

What to do next

Log on to ServiceNow to view Nutanix alerts in the **Alert Intelligence All Alerts** menu after completing [Importing Nutanix Event Rules to ServiceNow](#) on page 178 and [Activating the Create Incident Rule](#) on page 179.

You can also create playbooks for actions that you want ServiceNow to initiate on Prism Central for the alerts sent to Service now.

Importing Nutanix Event Rules to ServiceNow

Ensure that you have imported Nutanix event rules to ServiceNow after you setup the alert playbook.

Before you begin

In ServiceNow, go to **Event Management > Event Rules** and check if the Nutanix event rules are already available.

The Nutanix event rules are provided in the Nutanix Event Rule Creation XML file. Download the *Nutanix Event Rule Creation XML* file from [Nutanix XML download link](#).

About this task

To import event rules using the *Nutanix Event Rule Creation xml* file, do the following.

Procedure

1. In ServiceNow, go to **System Update Sets > Retrieved Update Sets**.

2. On the right pane, click **Import Update Set from XML**.

Import Update Set from XML is a clickable link on the right pane.

3. Browse and select the Nutanix Event Rule Creation xml file that you downloaded. Click **Upload**.

4. On the **Update Sources** page, click the uploaded Update Set to open update set details page.

5. Click **Preview Update Set**.

If you find errors after the preview run is complete, click **Accept Remote Update** for each error. This resolves the error. Resolve all the errors.

6. After resolving all the errors, click **Commit Update Set**.

What to do next

Go to **Event Management** > **Event Rules**. Select **Source** in the search dropdown and type *Nutanix* in the **Search** box. The Nutanix Event Rules are now displayed.

Activating the Create Incident Rule

You need to activate the **Create Incident** rule in ServiceNow.

About this task

The incidence creation does not occur in ServiceNow until you set the active status of the **Create Incident** rule in **Alert Management Rules** in **Event Management** of the ServiceNow instance to true (Active status being false by default).

To change the Active status of the **Create Incident** to true, do the following:

Procedure

1. Go to **Event Management** > **Alert Management**. A list of Alert Management Rules is displayed.
2. Click the rule named **Create Incident**. Click (to check) the **Active** status checkbox below the rule name to set the **Create Incident** rule status to active.

The screenshot shows the configuration page for an alert rule in ServiceNow. At the top, there are tabs for 'Alert Info', 'Alert Filter', and 'Actions'. Below these is a blue instruction bar: 'Specify the alert rule name, order and description. Specify whether to look for other alert rules that match the same alert. Ensure that at least one action is specified in the Actions tab.' The main form has fields for 'Name' (Create Incident), 'Order' (100), 'Active' (checked), and 'Description' (OOB rule. Create an incident for all alerts that are not in maintenance state.). There are also buttons for 'Update' and 'Delete'.

Figure 72: Activate the Create Incident Rule

Configuring a Webhook Playbook for ServiceNow Action communications to Prism

X-Play now supports a webhook feature to invoke a playbook through REST calls. You can create webhook based playbooks for all the actions that you want ServiceNow to initiate on Prism Central for the alerts sent to ServiceNow.

About this task

Create webhook based playbooks using the **Webhook** trigger.

Note: You can add one or more actions that you want SNOW to perform on Prism in one playbook.

To create a webhook based playbook, perform the following steps.

Procedure

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher.
The system displays the **Intelligent Operations** dashboard.
3. Navigate to **Playbooks** from the **Navigation Bar**.
The **Playbooks** page opens displaying the **List** tab.

4. Click **Create Playbook**.

Note: If you are creating a playbook for the first time, a **Get Started** option appears. Click **Get Started**, and then click **Create Playbook**.

5. Click **Select** within the **Webhook** (option) card.




Figure 73: Select a Trigger

The **Webhook** page displays a message informing you to return to this playbook and review the instructions configured for this trigger.




Figure 74: Webhook page

6. Click **Add Action**.

7. Click **Select** in the appropriate action card(s) in the X-Play actions deck.

You can select one or more actions for the playbook. After you complete configuring an X-Play action, click **Add Action** to add another action to the playbook.




Figure 75: Select an Action page

When you click **Select** on a X-Play action card, the action page opens. Provide all the necessary values for the parameters on the action page.

8. To save the playbook, do the following.
 - a. Click **Save & Close**.
 - b. Specify the **Name** and provide a **Description**.
 - c. Select the **Self** option to execute the playbook by yourself or the **Other** option to enable user impersonation. Select a user and click **OK**. For more information, see [User Impersonation to Run a Playbook](#) on page 147.
 - d. Click the slider to enable the playbook. Prism Central executes only the playbooks that are enabled.
 - e. Click **Save**.

What to do next

Open the playbook you just created. Notice that the **Webhook** section is populated with the URL details and the body of the webhook. The URL of the webhook points to Prism Central; each webhook has a unique ID. Call this URL with the requisite parameters to trigger the playbook.

PagerDuty Integration with Prism

PagerDuty is an incident management platform that provides reliable notifications, automatic escalations, on-call scheduling, and other functionality to help teams detect and fix infrastructure problems quickly.

The PagerDuty integration with Prism Central allows receiving Prism alerts as incidents in PagerDuty interface. You can create a Playbook using the new alert action called **Send Alert to PagerDuty**. This alert action sends Prism alerts (from Prism Central and Prism Element) to PagerDuty. The alerts appear as PagerDuty incidents in the PagerDuty **Incidents** page.

Contact Nutanix Support for any issues with PagerDuty integration with Prism.

Configuring a Playbook with the PagerDuty Integration

You can create a Playbook to send Prism alerts to PagerDuty. The Playbook sends the alerts from Prism Central and Prism Element to the PagerDuty instance.

Before you begin

You must set up the PagerDuty integration key. This key is used as the routing key to integrate PagerDuty with Playbooks in Prism Central. For information about how to generate an integration key, see *Create a Generic Events API Integration* in [PagerDuty documentation](#).

About this task

Perform the following steps to create a Playbook to send Prism alerts to PagerDuty.

Procedure

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher.
The system displays the **Intelligent Operations** dashboard.
3. Navigate to **Playbooks** from the **Navigation Bar**.
The **Playbooks** page opens displaying the **List** tab.
4. Click **Create Playbook**.

Note: If you are creating a playbook for the first time, a **Get Started** option appears. Click **Get Started**, and then click **Create Playbook**.

The **Select a Trigger** page appears.

5. Click **Select** within the **Alerts Matching Criteria** option.
The **Alerts Matching Criteria** allows you to choose all or specific alert policies.
6. In the **Alerts Matching Criteria** page, do the following:
 - a. Select **All Alerts Policies** to send alerts coming from all existing alert policies, or select **Specific Alert Policies** to send alerts coming from one or more specific alert policies. You can type in the alert policy name and search, or scroll the list of existing alert policies and select multiple alert policies.
 - b. Specify the **Impact Type**.
You can specify more than one Impact Type.
Impact type is the category in which the alerts are classified such as Availability, Configuration, Capacity, Performance, and System Indicator alerts.
 - c. Specify the **Severity** of the alerts.
 - d. Specify the cluster, either choose **All Clusters** or **Specific Clusters**.
If you choose a specific cluster, then only alerts coming from that cluster trigger the action specified in Step 6.
Every time there is an alert raised by the selected alert policy, the system will trigger actions defined in this Playbook.

7. Configure the filter criteria to target the most relevant alerts and the entities associated with the alert. The filter criteria options that appear depend on the type of alert policy that you have selected in the Step 4.

The filter options can be any of the following:

- Severity
 - Target VM
 - Target Cluster
 - Target Host
8. Click **Add Actions**.
A list of all the available actions on the deck appears.
 9. Click **Select** within the **Send Alert to PagerDuty** option.




Figure 76: Send Alert to PagerDuty

10. Enter the **Routing Key**.

The Routing Key is an integration key that you can find in your PagerDuty instance. For more information, see [PagerDuty documentation](#).

The image shows a configuration window titled "Send Alert to PagerDuty" with a subtitle "Add Description". At the top right are buttons for "Set Autopilot", "Save & Close", and a close icon. The main area has a "Change Action" link. Below it is an "Alert" dropdown menu currently showing "Alert: Alert". Underneath is a "Parameters" section containing a "Routing Key" text input field, which is highlighted with a red rectangular border. A grey callout bubble with the text "enter integration key" and an arrow points to this field. Below the "Routing Key" field is an "Additional Info (Optional)" text area. At the bottom right of the main area is another "Parameters" label. On the right side of the window, there is a "Settings" section with "On Action Failure:" set to "Stop".

Figure 77: Routing Key

11. Select **Parameters**.

The Parameters option allows you to select certain inputs based on the trigger that you have selected. Parameters also allows you to select any available input based on the previous actions in the series of actions that you have configured.

12. To save the playbook, do the following.

- Click **Save and Close**.
- Specify the **Name** and **Description**.
- Select the **Self** option to execute the playbook by yourself or the **Other** option to enable user impersonation. Select a user and click **OK**. For more information, see [User Impersonation to Run a Playbook](#) on page 147.
- Click the slider to enable the playbook. Only playbooks in the enabled state are executed.
- Click **Save**.

REPORTS MANAGEMENT

Use the report management feature to configure and deliver historical reports containing information about the infrastructure resources across Nutanix-managed and non-Nutanix-managed environments.

The report management feature provides operational insights about your infrastructure to your mailbox according to the schedule that you configured.

The following are the important report management considerations:

- All the built-in roles except the **Prism Viewer** role can manage the reports (create, update, share, and delete). Users with **Prism Viewer** privileges can only view the reports.
- When cluster-based licensing is in use, data from an unlicensed cluster (where cluster-based license is not applied) is filtered out from the report. For more information, see [License Manager Guide](#).

You can configure the following characteristics of reports:

- **Style:** Defines the report structure, such as background color, logos, and so on.
- **Representation of data:** Defines how to represent the data. You can include different views (pre-defined, customizable, and saved).
- **Data:** Defines the actual data depending on your selection, for example information about different metrics; CPU usage, memory usage, IOPS bandwidth, VM count, host count, cluster count, license summary, and so on.

Report management includes the following features:

- Generate and customize end-to-end reports
- Run predefined system reports
- Share the report configuration with other users
- Add different views to customize what data is displayed and how that data is represented
- Download the report in PDF and CSV format
- Create a report definition from another report definition. You can copy and reuse the read-only report definition created by another user and modify it as per your requirements.
- Support for internationalization in simplified Chinese and Japanese. The generated report can include English, Chinese, or Japanese characters.
- Customize the report definition.
- Add a logo and make changes to the appearance of the report
- Send the PDF or CSV version of the report through email
- Schedule reports. You can define only one schedule for a report definition.
- Retain the reports for the specified period of time
- Check the report log details for the status of the report and any error message (if the report generation fails)
- Create RBAC to define roles and assign permissions
- Export and import the report configurations

Limitations

The report management feature in Prism Central has the following limitations:

- Report generation in PDF format is supported for a maximum of 2,000 entities. Nutanix recommends generating reports in CSV format for more than 1,500 entities.
- Report generation in PDF and CSV formats supports repeating a maximum of 20 entities if the **Repeat views in this group** checkbox is selected.
- Report generation supports repeating a maximum of 10 entities if the vCenter application monitoring is configured.
- Report generation in CSV format supports a maximum of 25,000 entities.
- Report generation in PDF format can only consist of a maximum of 10 columns for a data table.
- When you select both the report formats (PDF and CSV) during the report generation, the entire report generation process fails if the report in either of the formats (PDF or CSV) is not generated.

Alerts Generation

Alerts are generated during the following failure scenarios:

- Generating a report
- Sending of a scheduled email message

Reports Summary View

The Reports page displays a list of report configurations created about the Nutanix and vCenter entities and predefined system report configurations.

The Reports page has the **Configurations** and **Generated Reports** tab.

To access the Reports page:

1. Log in to Prism Central.
2. Select the **Intelligent Operations** application from Application Switcher function, and navigate to **Reports** from the **Navigation Bar**.

The Reports page displays the **Configuration** tab by default. The **Configuration** tab provides information about custom and predefined system report configurations and saved report configurations about the Nutanix and vCenter entities.

Configurations Tab

The **Configurations** tab lists the following information:

- Report configurations saved from the infrastructure main and custom dashboards, Analysis dashboard, individual Charts, Alerts and Events page, and Playbooks page about the Nutanix and vCenter entities.
- Custom report configurations created about the Nutanix and vCenter entities.
- Predefined system report configurations.

Report Name	Report Description	Owner	Imported	Scheduled	Frequency	Last Updated Time
Report-Stage		admin	No	No		09/25/23, 03:11:35 PM
Playbook-List		admin	No	No		09/25/23, 03:11:03 PM
Main-Dashboard		admin	No	No		09/25/23, 03:09:46 PM
Cluster Configuration Summary System	Cluster Configuration Summar...		No	No		09/25/23, 12:30:22 PM
VM Snapshot Usage System	Snapshot Usage report shows...		No	No		09/25/23, 12:30:22 PM
Cluster Capacity and Reclaim Summary System	Cluster Capacity and Reclaim ...		No	No		09/25/23, 12:30:22 PM
Storage Container Overview System	Storage Container Overview r...		No	No		09/25/23, 12:30:22 PM
VM Overview System	VM Overview report shows th...		No	No		09/25/23, 12:30:22 PM
Cluster Efficiency Summary System	Cluster Efficiency Summary re...		No	No		09/25/23, 12:30:22 PM
Data Protection And Recovery Summary System	The Data Protection and Recoa...		No	No		09/25/23, 12:30:22 PM
Disk Usage Summary System	Disk Usage Summary report R...		No	No		09/25/23, 12:30:22 PM
VM Performance Summary System	VM Performance Summary re...		No	No		09/25/23, 12:30:22 PM

Figure 78: Reports Configurations Tab

The following table describes the fields that appear in **Configurations** tab.

Table 50: Configurations Tab - Field Description

Field	Description	Values
Report Name	Displays the name of the report.	(report name)
	Note: The SYSTEM label appears against the report name for pre-defined reports.	
Report Description	Displays the description of the report.	(report description)
Owner	Displays the name of the user who created the report.	(user name)
Scheduled	Displays whether the report is scheduled to run at specified times. - Yes means that the report is scheduled to run automatically at specified times (see Frequency). - No means that the report is not scheduled to run. To run the report you must either manually run it or update the report to specify a schedule.	(Yes, No)
Imported	Displays whether the report is imported or not.	(Yes, No)
Frequency	Displays how frequently the report is scheduled to run.	(Daily, Weekly, Monthly, Yearly)
Last Updated Time	Displays the last updated date and time of the report.	(mm/dd/yy, hh:mm:ss AM/PM)
Report Type	Displays the report type.	(User Report, System Report)

Field	Description	Values
Creation Timestamp	Displays when the report was created.	(mm/dd/yy, hh:mm:ss AM/PM)
Editable	Displays whether the report is editable or not.	(Yes, No)
Template (UUID)	Displays the unique identifier of the template that is associated with the report.	(UUID)
Pre-defined	Displays whether the report is predefined or not.	(Yes, No)
Schedule Multiple	Displays the number of times the report is scheduled to run.	(integer)
Start Time (Secs)	Displays when the scheduled run of the report begins.	(mm/dd/yy, hh:mm:ss AM/PM)
End Time (Secs)	Displays when the scheduled run of the report ends.	(mm/dd/yy, hh:mm:ss AM/PM)

Note: The fields vary based on the **View by** and **Group by** options. A dash (-) is displayed in a field when a value is not available or applicable.

You can perform the following actions from the **Configurations** tab:

- Create a custom report configuration about the Nutanix and vCenter entities and metrics of your choice by clicking **+ Report Configurations**. For more information, see [Creating a Custom Report Configuration](#) on page 199.
- Run predefined system report configurations. For more information, see [Managing a Predefined System Report Configuration](#) on page 191.
- Import report configurations by clicking **Import**. For more information, see [Importing a Custom Report Configuration](#) on page 209.
- Configure the appearance, e-mail, and retention settings of a report by clicking **Report Settings**. For more information, see [Configuring Report Settings](#) on page 262.
- Create a report instance from saved report configuration. For more information, see [Running a Report Instance](#) on page 201.
- Click the name of a report configuration to view the list of instances generated for the report. For more information, see [Viewing a Report Instance](#) on page 209.
- Manage report instances created from custom and predefined system report configurations and saved report configurations. For more information, see [Managing a Report Instance from Configurations Tab](#) on page 209.
- Edit the report configurations saved from dashboards and entities. For more information, see [Editing a Saved Report Configuration](#) on page 223.
- Group the report configurations based on pre-defined criteria. For information about how to group the report configurations, see [Group by](#).
- View the report configurations based on a pre-defined criteria or create a custom view. For information about available views and how to create a custom view, see [View by](#).
- Filter the report configurations list based on a variety of parameter values using the **Filters** pane. For more information about Filters pane, see [Filters Pane - Report Configurations Tab](#) on page 189.

- Perform the following actions on a report configuration using the **Actions** dropdown menu. The **Actions** dropdown menu appears when a report is selected.

Table 51: Report Configuration Actions

Action	Description
Run	Click this action to generate the report. For more information, see Running a Report Instance on page 201.
Edit	Click this action to edit the report. For more information, see Editing a Custom Report Configuration on page 202. Note: This action is unavailable for predefined system report configurations.
Clone	Click this action to clone the report. For more information, see Cloning a Custom Report Configuration on page 203.
Schedule	Click this action to schedule the report. For more information, see Scheduling a Custom Report Configuration on page 203.
Share	Click this action to share the report. For more information, see Sharing a Custom Report Configuration on page 206.
Export Configuration	Click this action to export the report. For more information, see Exporting a Custom Report Configuration on page 208. Note: This action is unavailable for predefined system report configurations.
Delete Configuration	Click this action to delete the report. For more information, see Deleting a Custom Report Configuration on page 208. Note: This action is unavailable for predefined system report configurations.

Filters Pane - Report Configurations Tab

You can filter the information in the Reports Configuration tab based on the following fields that are available in the **Filters** pane.

Table 52: Filter Pane Field Description - Configurations Tab

Field	Description	Values
Report Name	Filters based on the application instance name. Select the checkbox, choose a condition from the dropdown menu (Contains , Doesn't contain , Starts with , Ends with , or Equal to), and enter a string in the field.	(Report name string)
Owner	Filters based on the name of the user who created the report.	(user name)

Field	Description	Values
Imported	Filters based on the imported reports. Select the Yes checkbox to view only imported reports.	(Yes)
Report Type	Filters based on the report configuration type. Select the System Report checkbox to view only predefined system reports.	(System Report)

Generated Reports Tab

The generated reports tab lists the following information:

- Report instances generated from infrastructure main and custom dashboards, Analysis dashboard, individual Charts, Alerts and Events page, and Playbooks page.
- Report instances generated while creating a custom report or on an existing report.
- Report instances generated from predefined system report configurations.

Reports

Configurations

Generated Reports

21 Total Reports

Actions

View by my-view

Group by

View: my-view

Modify Filters

1 - 6 of 6

20 rows

Viewing all 4 Reports

Report Name	Report Type	Report Configuration	Generated on	Created By	Status	Download
Cluster Efficiency Summary	System Report	Cluster Efficiency Summary	09/25/23, 03:25:05 PM	admin	Successful	CSV
Playbook List	User Report	Playbook List	09/25/23, 03:18:40 PM	admin	Successful	PDF
Main-Dashboard-PDF	User Report	Main-Dashboard	09/25/23, 03:16:55 PM	admin	Successful	PDF
VM-Snapshot-Usage-CSV	System Report	VM Snapshot Usage	09/25/23, 03:16:33 PM	admin	Successful	CSV

Figure 79: Generated Reports Tab

The following table describes the fields that appear in the **Generated Reports** tab.

Table 53: Generated Reports Tab - Field Description

Field	Description	Values
Report Name	Displays the name of the report.	(report name)
Generated on	Displays when the report was created.	(mm/dd/yy, hh:mm:ss AM/PM)
Report Configuration	Displays the name of the report configuration from which the report instance was created.	(report configuration name)
Report Type	Displays the report type.	(User Report, System Report)
Created by	Displays the user (role) who generated the report.	(user name)
Status	Displays the status of the generated report.	(Successful, Failed)
Download	Displays the download button of the generated report.	(CSV, PDF)

Note: The fields vary based on the **View by** and **Group by** options. A dash (-) is displayed in a field when a value is not available or applicable.

You can perform the following actions from the **Generated Reports** tab:

- Download the reports by clicking **PDF** or **CSV** under the **Download** column next to the report instance. For more information, see [Downloading a Report Instance](#) on page 221.
- Group the generated reports based on the report name. For information about how to group the generated report, see [Group by](#).
- View the generated report based on predefined criteria or create a custom view. For information about available views and how to create a custom view, see [View by](#).
- Filter the generated reports list based on a variety of parameter values using the **Filters** pane. For more information, see [Filters Pane - Generated Reports Tab](#) on page 191.
- Perform the following actions on a report using the **Actions** dropdown menu. The **Actions** dropdown menu appears when a report instance is selected.

Table 54: Generated Reports Actions

Action	Description
Resend Report	Click this action to resend a report. For more information, see Resending a Report Instance on page 222.
Delete Report	Click this action to delete a report, see Deleting a Report Instance on page 223.

Filters Pane - Generated Reports Tab

You can filter the information in the Generated Reports tab based on the following fields that are available in the **Filters** pane.

Table 55: Filter Pane Field Description - Generated Reports Tab

Field	Description	Values
Report Type	Filters based on the report instance generated from the report configuration. Select System Report or User Report or both the checkboxes.	(System Report, User Report)

Managing a Predefined System Report Configuration

By default, Prism Central provides 18 predefined system report configurations that address common use cases and requirements of an organization's infrastructure management. You can use the predefined system report configurations in ready-to-consume format. The following actions are available on predefined system report configurations:

- Run

- Download
- Clone
- Schedule
- Share

Note: You cannot edit, export, or delete the predefined system report configurations.

The screenshot shows the 'Reports Configurations' tab in a web interface. It displays a table with 18 predefined system reports. The table columns are: Report Name, Report Description, Owner, Imported, Scheduled, Frequency, and Last Updated Time. The reports listed are: Cluster Configuration Summary, VM Snapshot Usage, Cluster Capacity and Reclaim Summary, Storage Container Overview, VM Overview, Cluster Efficiency Summary, Data Protection And Recovery Summary, Disk Usage Summary, VM Performance Summary, Infrastructure Performance Summary, VM Efficiency Summary, Main Dashboard, and Cluster Usage Summary. Each report has a 'System' tag and a 'Last Updated Time' of 09/25/23, 12:30:22 PM.

Report Name	Report Description	Owner	Imported	Scheduled	Frequency	Last Updated Time
Cluster Configuration Summary	Cluster Configuration Summary...	-	No	No	-	09/25/23, 12:30:22 PM
VM Snapshot Usage	Snapshot Usage report shows...	-	No	No	-	09/25/23, 12:30:22 PM
Cluster Capacity and Reclaim Summary	Cluster Capacity and Reclaim ...	-	No	No	-	09/25/23, 12:30:22 PM
Storage Container Overview	Storage Container Overview r...	-	No	No	-	09/25/23, 12:30:22 PM
VM Overview	VM Overview report shows th...	-	No	No	-	09/25/23, 12:30:22 PM
Cluster Efficiency Summary	Cluster Efficiency Summary re...	-	No	No	-	09/25/23, 12:30:22 PM
Data Protection And Recovery Summary	The Data Protection and Reco...	-	No	No	-	09/25/23, 12:30:22 PM
Disk Usage Summary	Disk Usage Summary report h...	-	No	No	-	09/25/23, 12:30:22 PM
VM Performance Summary	VM Performance Summary re...	-	No	No	-	09/25/23, 12:30:22 PM
Infrastructure Performance Summary	Infrastructure Performance Su...	-	No	No	-	09/25/23, 12:30:22 PM
VM Efficiency Summary	VM Efficiency Summary report...	-	No	No	-	09/25/23, 12:30:22 PM
Main Dashboard	Main dashboard report is the ...	-	No	No	-	09/25/23, 12:30:22 PM
Cluster Usage Summary	Cluster Usage Summary repor...	-	No	No	-	09/25/23, 12:30:22 PM

Figure 80: Reports Configurations Tab

The following table lists the predefined system report configurations available in report **Configurations** tab:

Table 56: List of Predefined System Report Configurations

Entity Type	System Report Name	Description
VM	VM Overview	Aggregated metrics of performance and usage for all VMs across the clusters.
	VM Snapshot Usage	VM by disk usage and snapshot usage along with snapshot size per VM.
	VM Performance Summary	Summarizes usage metrics of all VMs.
	VM Efficiency Summary	Identifies inefficient VMs (overprovisioned, constrained, inactive, bully) and corresponding usage metrics.
	VM Network Summary	Lists configuration details like IP addresses, virtual NICs and network statistics packets transmitted or dropped.

Entity Type	System Report Name	Description
Cluster	VM GPU Summary	Shows GPU configuration details as well as performance metrics for each VM.
	Cluster Usage Summary	Summarizes usage metrics for all clusters.
	Cluster Efficiency Summary	Lists inefficient VMs per every cluster and identifies ways to reclaim resources.
	Cluster Configuration Summary	Summarizes environment details like AOS version, hypervisor type, VM count of all the clusters.
	Cluster Capacity and Reclaim Summary	Shows cluster runway and details about ways to reclaim unused resources.
Dashboard	Main Dashboard	Summarizes PC's default infrastructure.
	Intelligent Operations Dashboard	Summarizes intelligent operations-related metrics.
Infrastructure Summary	Infrastructure Performance Summary	Summarizes infrastructure resources.
	Environment Summary	Summarizes infrastructure entities with key metrics configured in this environment.
	Host Performance Summary	Summarizes usage metrics of all hosts.
Storage	Disk Usage Summary	Summarizes overall performance of all disks.
	Storage Container Overview	Summarizes overall performance metrics of all storage containers.
Disaster Recovery	Data Protection And Recovery Summary	Shows backup and recovery plan details through corresponding replication metrics.

You can perform the following actions on a predefined system report configuration using the **Actions** dropdown menu. The **Actions** dropdown appears when a predefined system report configuration is selected.

Table 57: Predefined System Report Configuration Actions

Action	Description
Run	Click this action to generate the report instance. For more information, see Running a Report Instance on page 194.
Clone	Click this action to clone the report configuration. For more information, see Cloning a Predefined System Report Configuration on page 195.

Action	Description
Schedule	Click this action to schedule the report configuration. For more information, see Scheduling a Predefined System Report Configuration on page 196.
Share	Click this action to share the report configuration. For more information, see Sharing a Predefined System Report Configuration on page 197.

Running a Report Instance

This task describes how to run a report instance from the predefined system report configuration.

About this task

Perform the following steps to run a report instance from the predefined system report configuration.

Procedure

1. Log in to Prism Central.
2. Select the **Intelligent Operations** application from **Application Switcher** function and navigate to **Reports** from the **Navigation Bar**.

The Reports page displays the **Configuration** tab by default. The **Configuration** tab provides information about custom and predefined system report configurations and saved report configurations about the Nutanix and vCenter entities.
3. Select the checkbox associated with the predefined system report configuration to run.
4. From the **Actions** dropdown menu, select **Run**.
The **Run Report** window opens.
5. Enter the name of the report (new instance of the report) in the **Report Instance Name** field.
6. Enter the description of the report in the **Description** field.
7. Select **Nutanix Entities** and the entity type (**VM**, **Node**, **Cluster**, or **Category**) to include in the report.

Note: **Entities** option is not available for the following predefined system report configurations:

- Infrastructure Performance Summary
- Intelligent Operations Dashboard
- Data Protection And Recovery Summary
- Main Dashboard

8. Select the entity from the **Entity** dropdown menu.

Note: If you do not choose the entity, by default the report contains the first 20 entities sorted alphabetically by the entity name.

9. Select the duration of data that you want to display in the report from the **Time Period For Report** dropdown menu.

You can display the data for the last 24 hours, last week, or last month. You also have the option to select the date range.

Note: If you select **Custom Date Range**, you must specify the date and time range in the **From** and **To** fields. If you do not specify the date and time, the data is selected from the configured report.

10. Select your time zone from the **Time Zone** dropdown menu.

11. In the **Report Format** section, select a file format.

- Select the checkbox associated with **PDF** and **CSV** to generate the reports in both formats.
- Select the checkbox associated with **PDF** or **CSV** to generate the reports in either of the two formats.

Note:

- By default, the report is generated in PDF format if you do not select a format. Also, by default, the email recipient receives the report in PDF format.
- PDF format is not supported for a data table report with more than 10 columns.

12. In the **Email Report** section, in the **Additional Recipients** text box, enter recipients (comma-separated).

Note:

- Email is sent to the recipients specified in this field along with the recipients that were configured during the report configuration. For more information about configuring email settings, see [Configuring Report Settings](#) on page 262.
- Email recipients can download the report in the format you selected in Step 11.

13. Click **Run**.

A report instance is generated, and it is also available under **Generated Reports** tab.

What to do next

You can perform the following actions after a report instance is generated:

- Download a report instance. For more information, see [Download a Report Instance from Configurations Tab](#) and [Download a Report Instance from Generated Reports Tab](#).
- Resend a report instance. For more information, see [Resend a Report Instance from Configurations Tab](#) and [Resend a Report Instance from Generated Reports Tab](#).
- Delete a report instance. For more information, see [Delete a Report Instance from Configurations Tab](#) and [Delete a Report Instance from Generated Reports Tab](#).

Cloning a Predefined System Report Configuration

Cloning a predefined system report configuration creates a new copy of the report configuration. You can also add new views to the report when cloning the report.

About this task

Perform the following steps to clone a predefined system report configuration.

Procedure

1. Log in to Prism Central.
2. Select the **Intelligent Operations** application from **Application Switcher** function and navigate to **Reports** from the **Navigation Bar**.

The Reports page displays the **Configuration** tab by default. The **Configuration** tab provides information about custom and predefined system report configurations and saved report configurations about the Nutanix and vCenter entities.

3. Select the checkbox associated with the predefined system report configuration to clone.
4. From the **Actions** dropdown menu, select **Clone**.
The **Report Preview** window opens.
5. Update fields as needed.
You can also add new views to the report. For more information, see [Creating a Custom Report Configuration](#) on page 199.
The report configuration is cloned.

Scheduling a Predefined System Report Configuration

You can schedule the predefined system report configuration to run at a custom time period. The scheduling applies only to the specific predefined system report configuration for which you are configuring the schedule.

About this task

Perform the following steps to schedule a predefined system report configuration.

Procedure

1. Log in to Prism Central.
2. Select the **Intelligent Operations** application from **Application Switcher** function and navigate to **Reports** from the **Navigation Bar**.

The Reports page displays the **Configuration** tab by default. The **Configuration** tab provides information about custom and predefined system report configurations and saved report configurations about the Nutanix and vCenter entities.

3. Select the checkbox associated with the predefined system report configuration to schedule.
4. Click **Actions > Schedule**.
The **Schedule** window opens.
5. Select the frequency to run the report from the **Report Runs** dropdown menu.
You can run the report on daily, weekly, monthly, or yearly basis.
6. Select the time to repeat the report.
For example, if you select **Yearly** from the **Report Runs** dropdown menu, you can repeat the report by selecting the month when you want to run the report from the **Every** dropdown menu, date of the month from the **On** dropdown menu, time of the day when you want to run the report from the **Generated Time** dropdown menu.

7. Select the data range to include in the report from the **Report Time Period** dropdown menu.
You can display the data for last 24 hours, last week, or last month. You also have an option to select a custom data range.

- a. If you select a **Custom Date Range** from the **Report Time Period**, you need to select the date range from the **From** and **To** option.

If you select a date range from 10 a.m. (Monday) to 10 a.m. (Tuesday), the data for this time is displayed in the report. For example, in a line chart, the X-axis displays the time range from 10 a.m. (Monday) to 10 a.m. (Tuesday) and the Y-axis displays the variation in data.

8. Select **Nutanix Entities** and the entity type (**VM**, **Node**, **Cluster**, or **Category**) to include the report.

Note: **Entities** option is not available for the following predefined system report configurations:

- Infrastructure Performance Summary
- Intelligent Operations Dashboard
- Data Protection And Recovery Summary
- Main Dashboard

9. Select the entity from the **Entity** dropdown menu.

Note: If you do not choose the entity, by default the report contains the first 20 entities sorted alphabetically by the entity name.

10. Select the user impersonation to run the report from the **Run Report on Behalf of** option.

- If you are the creator of the report, select **Self** button.
- If you want to assign another user as an executor while creating the report, click the **Other** button and choose the impersonated user from the dropdown menu.

11. Click **Save**.

Sharing a Predefined System Report Configuration

Any user with edit permission can share a predefined system report configuration with other users and provide a view or edit access to these users.

About this task

Perform the following steps to share a predefined system report configuration.

Procedure

1. Log in to Prism Central.
2. Select the **Intelligent Operations** application from **Application Switcher** function and navigate to **Reports** from the **Navigation Bar**.

The Reports page displays the **Configuration** tab by default. The **Configuration** tab provides information about custom and predefined system report configurations and saved report configurations about the Nutanix and vCenter entities.

3. Select the checkbox associated with the predefined system report configuration to share.

4. From the **Actions** dropdown menu, select **Share**.
The **Sharing Report** window opens.

Sharing Report "██████████"

Add Users & Groups Manage Users & Groups 1

Share with

Search for the Users & Groups...

Permission Type

Please Select

Cancel Save

Figure 81: Sharing Report

5. In the **Add Users & Groups** tab, do the following:
- Share with:** In the **Search** field, enter the name of the users or groups to share the report with. The **Search** field displays a list of users or groups whose names match the letters you type. Click the username or group name to add that user or group. You can add more than one user or a group.
 - Permission Type:** You can provide view only or edit permissions to a user or group by selecting **View only** or **Edit** from the dropdown menu.
A user with **Edit** permission can perform these tasks:
 - View a report
 - Edit a report
 - Run a report
 - Share a report
 - In the **Manage Users & Groups** tab, you can change the existing permission type for users or groups with whom you have shared the report earlier. To change the permission type, select the type from the dropdown menu associated with the user name or group name.
You can also remove access for a user or group.

6. Click **Save**.

A message that the report is shared successfully appears.

After sharing the report, wait for few minutes for the changes to take effect.

Note: Users cannot view any information about entities in the reports without entity permissions.

Creating a Custom Report Configuration

This task describes how to create a custom report configuration from the Reports **Configuration** tab.

About this task

Prism Central allows you to create and customize the report by adding metrics of your choice for the Nutanix and vCenter entities.

Procedure

1. Log in to Prism Central.
2. Select the **Intelligent Operations** application from **Application Switcher** function and navigate to **Reports** from the **Navigation Bar**.

The Reports page displays the **Configuration** tab by default. The **Configuration** tab provides information about custom and predefined system report configurations and saved report configurations about the Nutanix and vCenter entities.
3. Click **+ Report Configurations**.
The **Report Preview** window opens. By default, The report configuration name is set to New Report.
4. To edit the name, click the pencil icon next to the default name, update the name, and then click **OK**.
5. Add the views that you want to add to the report. For information about different views, see [Adding a Custom View](#) on page 225.
6. Configure the settings specific for the report by clicking **Report Settings**. For more information, see [Configuring Report Settings](#) on page 262.
7. If you have scheduled the report or want to modify the report later, click **Save**.
For more information about editing the report, see [Editing a Custom Report Configuration](#) on page 202.

8. To save the report configuration and run immediately to generate a report instance, click **Save and Run Now** and enter the following information:

The **Run Report** window opens.

- Report Instance Name:** Enter the name of the new report instance.
- Description:** Enter the description of the new report instance.
- Time Period For Report:** Select the duration of data that you want to display in the report.
You can display the data for the last 24 hours, last week, or last month. You also have an option to select the date range.

Note: If you select **Custom Date Range**, you must specify the date and time range in the **From** and **To** fields. If you do not specify the date and time, the data is selected from the configured report.

- Time Zone:** Select your time zone from the dropdown list.
- Report Format:** Select the **PDF** checkbox or **CSV** checkbox or both to set a report format.
- Email Report:** Enter recipients (comma-separated) in the **Additional Recipients** text box.

Note:

- Email is sent to the recipients specified in this field along with the recipients that were configured during the report configuration. For more information about configuring email settings, see [Configuring Report Settings](#) on page 262.
- Email recipients can download the report in the format that you have selected in **Report Format**.

- Click **Run**.

A report instance is generated and it is also available under **Generated Reports** tab.

For example, in the following image, a new report is configured with different views (for example, bar chart, a metric summary, and a data table).



Figure 82: Create Report

Managing a Custom Report Configuration

You can perform the following actions to manage a custom report configuration from Prism Central.

- [Create a report](#)
- [Edit a report](#)

- [Clone a report](#)
- [Schedule a report](#)
- [Share a report](#)
- [Export a report](#)
- [Delete a report](#)
- [Import a report](#)

Running a Report Instance

This task describes how to run a report instance from the custom report configuration.

About this task

You can generate a report instance while creating a custom report configuration, using an existing report configuration, or from the saved report configuration of dashboards and entity pages.

Note:

- You can create multiple reports instances for a single report configuration.
- The report instance is not created by default when you save the report configuration from the dashboards and entities page. You must run the saved report configuration to create a report instance.

Procedure

1. Log in to Prism Central.
2. Select the **Intelligent Operations** application from **Application Switcher** function and navigate to **Reports** from the **Navigation Bar**.

The Reports page displays the **Configuration** tab by default. The **Configuration** tab provides information about custom and predefined system report configurations and saved report configurations about the Nutanix and vCenter entities.
3. Select the checkbox associated with the report configuration you want to run.
4. From the **Actions** dropdown menu, select **Run**.
The **Run Report** window opens.
5. Enter the name of the report (new instance of the report) in the **Report Instance Name** field.
6. Enter the description of the new report instance in the **Description** field.
7. Select the duration of data that you want to display in the report from the **Time Period For Report** dropdown list.
You can display the data for the last 24 hours, last week, or last month. You also have the option to select the date range.

Note: If you select **Custom Date Range**, you must specify the date and time range in the **From** and **To** fields. If you do not specify the date and time, the data is selected from the configured report.

8. Select your time zone from the **Time Zone** dropdown list.

9. In the **Report Format** section, select the format you want the report to be generated.
 - Select the checkbox associated with **PDF** and **CSV** to generate the reports in both formats.
 - Select the checkbox associated with **PDF** or **CSV** to generate the reports in either of the two formats.

Note:

- By default, the report is generated in PDF format if you do not select any format of the report. Also, by default, the email recipient receives the report in PDF format.
- PDF format is not supported for a data table report with more than 10 columns.

10. In the **Email Report** section, enter recipients (comma-separated) in the **Additional Recipients** text box.

Note:

- Email is sent to the recipients specified in this field along with the recipients that were configured during the report configuration. For more information about configuring email settings, see [Configuring Report Settings](#) on page 262.
- Email recipients can download the report in the format you selected in Step 9.

11. Click **Run**.

A report instance is generated, and it is also available under **Generated Reports** tab.

Editing a Custom Report Configuration

You can edit an existing custom report configuration. You can modify views and change other report settings in the Edit mode.

About this task

Perform the following steps to edit a custom report configuration.

Procedure

1. Log in to Prism Central.
2. Select the **Intelligent Operations** application from **Application Switcher** function and navigate to **Reports** from the **Navigation Bar**.

The Reports page displays the **Configuration** tab by default. The **Configuration** tab provides information about custom and predefined system report configurations and saved report configurations about the Nutanix and vCenter entities.

3. Select the checkbox associated with the report configuration you want to edit.
4. From the **Actions** dropdown menu, select **Edit**.
5. Edit the report by updating the desired fields as described in the [Creating a Custom Report Configuration](#) on page 199 topic.

Update the views that are added to the report by performing the following steps:

6. Select the checkbox associated with the view you want to update and perform the following actions using the **Actions** dropdown menu on the right.

You can edit or remove the view, and move or remove the view from a group (if the group already exists).

- a. To edit the view, click **Edit**.
- b. To remove the view, click **Remove**.
- c. To move the view inside a group view, click the name of the group. The view is moved into the group.

Note: This option only appears if you have a group view that is already configured and if you have adhered to the hierarchical nature of the group view. For more information, see [Adding a Group View](#) on page 253.

- d. To remove the view from a group, click **Remove from Group**.

7. After you edit the required fields, do either of the following.

- Click **Save and Run Now** to save and create the report instance.
- Click **Save** to save the report.

The custom report is edited and saved.

Cloning a Custom Report Configuration

Cloning a custom report configuration creates a new copy of the report configuration. You can also add new views to the report when cloning the report.

About this task

Perform the following steps to clone a custom report configuration.

Procedure

1. Log in to Prism Central.
2. Select the **Intelligent Operations** application from **Application Switcher** function and navigate to **Reports** from the **Navigation Bar**.

The Reports page displays the **Configuration** tab by default. The **Configuration** tab provides information about custom and predefined system report configurations and saved report configurations about the Nutanix and vCenter entities.

3. Select the checkbox associated with the report configuration you want to clone.
4. From the **Actions** dropdown menu, select **Clone**.

The **Report Preview** window opens.

5. Update the desired fields.

You can also add new views to the report. For more information, see [Creating a Custom Report Configuration](#) on page 199.

The report configuration is cloned.

Scheduling a Custom Report Configuration

After you add all the views to the reports, you have an option to run the report at a custom time period. This scheduling is applicable only for the specific report for which you are configuring the schedule.

About this task

Perform the following steps to schedule a custom report configuration.

Procedure

1. Log in to Prism Central.
2. Select the **Intelligent Operations** application from **Application Switcher** function and navigate to **Reports** from the **Navigation Bar**.

The Reports page displays the **Configuration** tab by default. The **Configuration** tab provides information about custom and predefined system report configurations and saved report configurations about the Nutanix and vCenter entities.

3. Select the checkbox associated with the custom report configuration that you want to schedule.
4. Click **Actions** > **Schedule**.
The **Schedule** window opens.

5. Select when you want to run the report from the **Report Runs** dropdown menu.
You can run the report on daily, weekly, monthly, or yearly basis.

Schedule

?

×

Report

main-dashboard

REPORT RUNS

Daily

GENERATE TIME

1:12 pm

Data duration for the report

REPORT TIME PERIOD

Last 24 Hours

RUN REPORT ON BEHALF OF

☒ Self ☐ Other

Cancel

Save

Figure 83: Report Scheduling

6. Select the time when you want to repeat the report.
For example, if you select **Yearly** from the **Report Runs** dropdown menu, you can repeat the report by selecting the month when you want to run the report from the **Every** dropdown menu, date of the month from the **On** dropdown menu, time of the day when you want to run the report from the **Generated Time** dropdown menu.

7. Select the duration for the data that you want to display in the report from the **Report Time Period** dropdown menu.

You can display the data for last 24 hours, last week, or last month. You also have an option to select a custom data range.

- a. If you select a **Custom Date Range** from the **Report Time Period**, you need to select the date range from the **From** and **To** option.

If you select a date range from 10 a.m. (Monday) to 10 a.m. (Tuesday), the data for this time is displayed in the report. For example, in a line chart, the X-axis will display the time range from 10 a.m. (Monday) to 10 a.m. (Tuesday) and the Y-axis will display the variation in data.

8. Select the user impersonation to run the report from the **Run Report on Behalf of** option.

- If you are the creator of the report, select **Self** button.
- If you want to assign another user as an executor while creating the report, select **Other** button and choose the impersonated user from the dropdown list.

9. Click **Save**.

Sharing a Custom Report Configuration

Any user with edit permission can share a custom report configuration with other users and provide a view or edit access to these users.

About this task

Perform the following steps to share a custom report configuration.

Procedure

1. Log in to Prism Central.
2. Select the **Intelligent Operations** application from **Application Switcher** function and navigate to **Reports** from the **Navigation Bar**.

The Reports page displays the **Configuration** tab by default. The **Configuration** tab provides information about custom and predefined system report configurations and saved report configurations about the Nutanix and vCenter entities.

3. Select the checkbox associated with the custom report configuration you want to share.

4. From the **Actions** dropdown menu, select **Share**.
The **Sharing Report** window opens.

Sharing Report " [blurred name] X

Add Users & Groups Manage Users & Groups 1

Share with

Search for the Users & Groups...

Permission Type

Please Select

Cancel Save

Figure 84: Sharing Report

5. In the **Add Users & Groups** tab, do the following:
- Share with:** In the Search field, enter the name of the users or groups to share the report. The Search field displays a list of users or groups whose names match the letters you type. Click the username or group name to add that user or group. You can add more than one user or a group.
 - Permission Type:** You can provide view only or edit permissions to a user or group by selecting **View only** or **Edit** from the dropdown menu.
A user with **Edit** permission can perform these tasks:
 - View a report
 - Edit a report
 - Run a report
 - Share a report
 - In the **Manage Users & Groups** tab, you can change the existing permission type for users or groups with whom you have shared the report earlier. To change the permission type, select the type from the dropdown menu associated with the user name or group name.
You can also remove the access for a user or group.

6. Click **Save**.

A message that the report is shared successfully appears.

After sharing the report, wait for few minutes for the changes to take effect.

Note: Users cannot view any information about entities in the reports without entity permissions.

Exporting a Custom Report Configuration

This task describes how to export a custom report configuration.

About this task

You can export a custom report configuration for backup reasons.

Procedure

1. Log in to Prism Central.
2. Select the **Intelligent Operations** application from **Application Switcher** function and navigate to **Reports** from the **Navigation Bar**.

The Reports page displays the **Configuration** tab by default. The **Configuration** tab provides information about custom and predefined system report configurations and saved report configurations about the Nutanix and vCenter entities.

3. Select the checkbox associated with the custom report configuration that you want to export.
4. Click **Actions > Export Configuration**.
The **Export Report Configurations** window opens.
5. (Optional) Edit the default **Filename**.
6. Click **Export**.
The custom report configuration is downloaded in the .rpt format.

Deleting a Custom Report Configuration

Deleting a custom report configuration removes all the generated report instances along with the selected report configuration.

About this task

Perform the following steps to delete a custom report configuration.

Procedure

1. Log in to Prism Central.
2. Select the **Intelligent Operations** application from **Application Switcher** function and navigate to **Reports** from the **Navigation Bar**.

The Reports page displays the **Configuration** tab by default. The **Configuration** tab provides information about custom and predefined system report configurations and saved report configurations about the Nutanix and vCenter entities.

3. Select the checkbox associated with the custom report configuration you want to delete.
4. From the **Actions** dropdown menu, select **Delete Configuration**.

5. From the confirmatory window that opens, select **Delete**.
The custom report configuration is deleted.

Importing a Custom Report Configuration

This task describes how to import a custom report configuration.

About this task

You can import any previously exported custom report configurations. If you export a custom report configuration using a Prism Central instance, you can import that custom report configuration to another Prism Central instance that is running the same or newer version only.

Procedure

1. Log in to Prism Central.
2. Select the **Intelligent Operations** application from **Application Switcher** function and navigate to **Reports** from the **Navigation Bar**.

The Reports page displays the **Configuration** tab by default. The **Configuration** tab provides information about custom and predefined system report configurations and saved report configurations about the Nutanix and vCenter entities.

3. Click **Import**.
The **Import Report Configurations** window opens.
4. Click **Browse** and select the previously exported custom report configuration file to upload.

Note: The custom report configuration file extension is in .rpt format.

5. Click **Import**.
The custom report configuration is imported to Prism Central.

Managing a Report Instance from Configurations Tab

You can perform the following actions to manage a report instance from the **Configurations** tab.

- [Viewing a Report Instance](#) on page 209
- [Downloading a Report Instance](#) on page 210
- [Resending a Report Instance](#) on page 211
- [Deleting a Report Instance](#) on page 212

Viewing a Report Instance

This task describes how to view a report instance from the **Configurations** tab.

Before you begin

Ensure that you have generated at least one instance of the report that you are viewing. Otherwise, Prism Central displays a no instance found message.

About this task

You can view the list of instances generated for the report configuration.

Procedure

1. Log in to Prism Central.
2. Select the **Intelligent Operations** application from **Application Switcher** function and navigate to **Reports** from the **Navigation Bar**.

The Reports page displays the **Configuration** tab by default. The **Configuration** tab provides information about custom and predefined system report configurations and saved report configurations about the Nutanix and vCenter entities.

3. Click the name of a report configuration to view the list of instances generated for the report.

Prism Central displays all the instances of the report with the timestamp information.

Note:

- You cannot view multiple reports at the same time.
- Prism Central displays an instance of the report only if you have manually generated the report or the report has been generated according to a configured schedule.
- If the report generation fails, the Status column displays the status as Failed. Hover over the field to view the reason for the failure.



Figure 85: Viewing Instance Information (Nutanix)



Figure 86: Viewing Instance Information (vCenter)

Downloading a Report Instance

You can download a report instance in PDF format, CSV format, or both.

About this task

- A PDF generated report displays a graphical representation of the infrastructure resources.
- A CSV generated report contains data of views separated by commas. You can use this CSV file to load CSV data into other systems. Data of views that you add to a report is populated to the CSV format report.
- The CSV report that you download is in the form of a .zip file. The .zip file contains CSV generated files for each view.
- For each group, a folder is created containing CSV files of each view of that group.

- If a view is repeated over a cluster entity or VM entity, the view has a CSV file for every repeating cluster. For example, a view named **View** that is repeating over clusters **Cluster_1**, **Cluster_2**, **Cluster_3**, the name of the CSV file is displayed as follows:
 - `1_1_View_Cluster_1.csv`
 - `1_2_View_Cluster_2.csv`
 - `1_3_View_Cluster_3.csv`
- If a group has a view repeating over a cluster or VM entity, the folder of the group consists of a folder for every cluster and that cluster contains a CSV file for the view. For example, a group named Group having view named **View** on clusters **Cluster_1**, **Cluster_2**, **Cluster_3**, the CSV folder structure is displayed as follows:
 - `Group / Cluster_1 / View.csv`
 - `Group / Cluster_2 / View.csv`
 - `Group / Cluster_3 / View.csv`

To download a report, do the following.

Procedure

1. Log in to Prism Central.
2. Select the **Intelligent Operations** application from **Application Switcher** function and navigate to **Reports** from the **Navigation Bar**.

The Reports page displays the **Configuration** tab by default. The **Configuration** tab provides information about custom and predefined system report configurations and saved report configurations about the Nutanix and vCenter entities.
3. Click the name of a report configuration to view the list of instances generated for the report.
4. In the window that opens, click **PDF** or **CSV** under the **Download** column next to the report instance you want to download.

Note: You can set the output format of the report, see [Creating a Custom Report Configuration](#) on page 199.

Resending a Report Instance

This task describes how to resend a report instance from the **Configurations** tab.

About this task

Prism Central allows you to resend a report instance created from custom and predefined system report configurations as an attachment either in PDF format or CSV format, or both through email.

Procedure

1. Log in to Prism Central.
2. Select the **Intelligent Operations** application from **Application Switcher** function and navigate to **Reports** from the **Navigation Bar**.

The Reports page displays the **Configuration** tab by default. The **Configuration** tab provides information about custom and predefined system report configurations and saved report configurations about the Nutanix and vCenter entities.

3. Click the name of a report configuration to view the list of instances generated for the report.
4. Select the checkbox associated with the report instance you want to resend.
5. Click **Actions** and select **Resend Report** from the dropdown menu.
6. In the **Resend Report** window that opens, enter the following information:
 - a. **Report Format:** Choose the format of the report (**PDF** or **CSV** or both) by selecting the checkboxes associated with it.
 - b. **Email Report:** Click **+ Add Additional Recipients** link and enter the recipients (comma-separated) of the report.

Note: If you do not add any recipients, the report is sent to the recipients already configured in the **Report Settings** Nutanix Flow.

7. Click **Send**.

Deleting a Report Instance

This task describes how to delete a report instance from the **Configurations** tab.

About this task

Prism Central allows you to delete a report instance created from a custom and predefined system report configurations.

Procedure

1. Log in to Prism Central.
2. Select the **Intelligent Operations** application from **Application Switcher** function and navigate to **Reports** from the **Navigation Bar**.

The Reports page displays the **Configuration** tab by default. The **Configuration** tab provides information about custom and predefined system report configurations and saved report configurations about the Nutanix and vCenter entities.
3. Click the name of a report configuration to view the list of instances generated for the report.
4. Select the checkbox associated with the report instance you want to delete.
5. Select **Actions** and click **Delete Report** from the dropdown menu.
6. From the confirmatory window that opens, click **Delete Instances**.

Example Report - Environment Summary

A generated environment summary report has the following information.

Sample Environment Summary Report

- Stylistic information: The first page of the report is used to represent this kind of information. The stylistic information includes the name of the report, time it was generated, any logo that you have configured, header and footer information (for example, report name and page number).



Figure 87: Stylistic Information

- Metrics information: Metrics summary displays the cluster details information (for example, cluster usage, performance details, licensing information, etc).

Following sample images displays the detailed information about the environment summary information for Nutanix and vCenter clusters.




Figure 88: Detailed Metrics Information (Nutanix: Image 1)




Figure 90: Detailed Metrics Information (Nutanix: Image 3)




Figure 91: Detailed Metrics Information (vCenter: Image 1)




Figure 92: Detailed Metrics Information (vCenter: Image 2)




Figure 93: Detailed Metrics Information (vCenter: Image 3)




Figure 94: Detailed Metrics Information (vCenter: Image 4)




Figure 95: Detailed Metrics Information (vCenter: Image 5)

Managing a Report Instance from Generated Reports Tab

You can perform the following actions to manage a report instance from the **Generated Reports** tab.

- [Downloading a Report Instance](#) on page 221
- [Resending a Report Instance](#) on page 222
- [Deleting a Report Instance](#) on page 223

Downloading a Report Instance

This task describes how to download a report instance from the **Generated Reports** tab.

About this task

You can download the report instance created from custom and predefined system report configurations and report instances generated from the dashboards and entity pages. The **Download** column next to the report instance shows the report format you have selected, either PDF format or CSV format, or both.

Procedure

1. Log in to Prism Central.
2. Select the **Intelligent Operations** application from **Application Switcher** function and navigate to **Reports** from the **Navigation Bar**.

The Reports page displays the **Configuration** tab by default. The **Configuration** tab provides information about custom and predefined system report configurations and saved report configurations about the Nutanix and vCenter entities.

3. Select **Generated Reports** tab.
4. Select the checkbox associated with the report instance you want to download.
5. Click **PDF** or **CSV** under the **Download** column next to the report instance you want to download.

Resending a Report Instance

This task describes how to resend a report instance from the **Generated Reports** tab.

About this task

Prism Central allows you to resend a report instance created from custom and predefined report configurations and report instances generated from dashboards and entity pages as an attachment either in PDF format or CSV format, or both through email.

Procedure

1. Log in to Prism Central.
2. Select the **Intelligent Operations** application from **Application Switcher** function and navigate to **Reports** from the **Navigation Bar**.

The Reports page displays the **Configuration** tab by default. The **Configuration** tab provides information about custom and predefined system report configurations and saved report configurations about the Nutanix and vCenter entities.

3. Select **Generated Reports** tab.
4. Select the checkbox associated with the report instance you want to resend.
5. Click **Actions** and select **Resend Report** from the dropdown menu.
6. In the **Resend Report** window that opens, enter the following information:
 - a. **Report Format:** Choose the format of the report (**PDF** or **CSV** or both) by selecting the checkboxes associated with it.
 - b. **Email Report:** Enter the recipients (comma-separated) of the report.

Note: If you do not add any recipients, the report is sent to the recipients already configured in the **Report Settings** Nutanix Flow.

7. Click **Send**.

Deleting a Report Instance

This task describes how to delete a report instance from the **Generated Reports** tab.

About this task

Prism Central allows you to delete a report instance created from custom and predefined system report configurations and report instances generated from dashboards and entity pages.

Procedure

1. Log in to Prism Central.
2. Select the **Intelligent Operations** application from **Application Switcher** function and navigate to **Reports** from the **Navigation Bar**.

The Reports page displays the **Configuration** tab by default. The **Configuration** tab provides information about custom and predefined system report configurations and saved report configurations about the Nutanix and vCenter entities.
3. Select **Generated Reports** tab.
4. Select the checkbox associated with the report instance you want to delete.
5. Click **Actions** and select **Delete Report** from the dropdown menu.
6. From the confirmatory window that opens, click **Delete Instances**.

Editing a Saved Report Configuration

This task describes how to edit the report configuration saved from dashboards and entities.

About this task

You can edit a saved report configuration by adding the views you want to add to the report.

Procedure

1. Log in to Prism Central.
2. Select the **Intelligent Operations** application from **Application Switcher** function and navigate to **Reports** from the **Navigation Bar**.

The Reports page displays the **Configuration** tab by default. The **Configuration** tab provides information about custom and predefined system report configurations and saved report configurations about the Nutanix and vCenter entities.
3. Select the checkbox associated with the report configuration you want to edit.
4. From the **Actions** dropdown menu, select **Edit**.
5. In the **Report Preview** window, add the views you want to add to the report. For information about different views, see [Adding a Custom View](#) on page 225.

Note: You cannot edit the widgets added by a report configuration. You can only view the widget details and remove the widget.

6. Configure the settings specific to the report by clicking **Report Settings**. For more information, see [Configuring Report Settings](#) on page 262.
7. If you have scheduled the report or want to modify the report later, click **Save**.

8. To save and run the report configuration to immediately generate a report instance, click **Save and Run Now**.

Views in the Reports

The report management feature provides the following views that you can add to a report.

- **Custom:** You can customize these views during the addition process itself. These views are displayed under the **Custom** tab of the **ADD VIEWS** pane. For more information, see [Adding a Custom View](#) on page 225.
- **Predefined:** You cannot customize these views, but they can be directly included in the reports. These views are displayed under the **Predefined** tab of the **ADD VIEWS** pane. For more information, see [Adding a Predefined View](#) on page 258.
- **Saved:** You can find the list of report views you saved from the Charts, Playbooks, Alerts, and Events page. These views are displayed under the **Saved** tab of the **ADD VIEWS** pane. For more information, see [Adding a Saved View](#) on page 261.



Figure 96: Custom Views



Figure 97: Predefined Views



Figure 98: Saved Views



Figure 99: All Views (Custom and Predefined Combined)

Adding a Custom View

Addition of Custom Views

You can add the following custom views to a report:

- To add a bar chart, line chart, or histogram view, see [Adding a Bar Chart or a Line Chart or a Histogram View to a Report \(Nutanix\)](#) on page 226 and [Adding a Bar Chart or a Line Chart or a Histogram View to a Report \(vCenter\)](#) on page 230.
- To add a data table view, see [Adding a Data Table View to a Report \(Nutanix\)](#) on page 233 and [Adding a Data Table View to a Report \(vCenter\)](#) on page 238.
- To add a configuration summary view, see [Adding Configuration Summary View](#) on page 241.
- To add a metric summary view, see [Adding Metric Summary View \(Nutanix\)](#) on page 243 and [Adding Metric Summary View \(vCenter\)](#) on page 246.

- To add an entity count view, see [Adding Entity Count View \(Nutanix\)](#) on page 249 and [Adding Entity Count View \(vCenter\)](#) on page 251.
- To add a title and description view, see [Adding Title and Description View](#) on page 253.
- To add a group view, see [Adding a Group View](#) on page 253.

Note:

- After you add the views to a report, the report displays the dummy data for the view that you have added. Only when you generate the report, the data for the view is updated. You can download a report to view its information.
- Collecting some VM data requires that NGT is installed. As a result, reports include such data only from NGT-enabled VMs.

Adding a Bar Chart or a Line Chart or a Histogram View to a Report (Nutanix)

You can use the bar chart, line chart, or a histogram view to represent the historical data in a graphical format. Perform the following procedure to add these views to the report.

Procedure

1. Log in to Prism Central.
2. Select the **Intelligent Operations** application from **Application Switcher** function and navigate to **Reports** from the **Navigation Bar**.

The Reports page displays the **Configuration** tab by default. The **Configuration** tab provides information about custom and predefined system report configurations and saved report configurations about the Nutanix and vCenter entities.

3. Click **+ Report Configurations**.
The **Report Preview** window opens.
4. Select the type of view (bar chart, line chart, or a histogram view) that you want to add from the **Custom** view pane.

5. Configure the bar chart, line chart, or a histogram view.
Following image displays the bar chart view.

Add Bar Chart?×

ENTITY TYPE

Nutanix Entities

VM

METRIC

Storage Controller IOPS

TITLE

VM Storage Controller IOPS

AGGREGATION

Average

NUMBER OF ENTITIES

-

10

+

SORT ORDER

Ascending

Entities

Select the entities that need to be reported in the view.

☐ All vms

☒ Specific vms

Rules

In case of multiple rules, a conjunction (AND operator) will be applied between them

+

Show Advanced Settings

Cancel

Add

Figure 100: Bar Chart View

- a. Select **Nutanix Entities** and the entity type (**VM, Disk, Storage Container, Host, Cluster, Virtual Disk, or Virtual Group**) from the **Entity Type** dropdown menu for which you want to configure the view.

Note: If a vCenter is enabled, you can select **vCenter Entities** in the entity type. This option does not appear in the dropdown menu if a vCenter is not configured. To configure the view for vCenter entities, see [Adding a Bar Chart or a Line Chart or a Histogram View to a Report \(vCenter\)](#) on page 230

- b. Select a specific metric of the entity type.

The metric selection depends on what you have selected in the **Entity Type** dropdown menu. For example, if you select **Cluster** as an entity type, you can select one of the following metric.

CPU Usage, Free Physical Storage, Controller Read I/O Bandwidth, Controller AVG Write I/O Bandwidth, Storage Savings Ratio, Controller I/O Bandwidth, Controller AVG Write I/O Latency, Controller IOPS, Controller Read IOPS, Controller Write I/O Bandwidth, Memory Usage (%), Physical Usage, Memory Usage Total, Storage Logical Usage, Controller Write IOPS, Controller AVG Read I/O Latency, Storage Savings

- c. Enter the title of the bar chart, line chart, or histogram view in the **Title** field.

- d. Select the aggregation from the **Aggregation** dropdown menu.

For the metrics that has time-series data, you can define the aggregation field. For example, if you want to include maximum CPU usage for a metric, you can use this field and select the **CPU Usage** as the metric and select aggregation as **Max**. Depending on the report schedule that you have defined, maximum value of the CPU usage is captured and included as part of the report.

You can define the aggregation value of **Sum, Max, Min, Average, Last, or Count**.

Note: Aggregation is not supported for Line Chart view.

- e. Enter the number of entities that you want to include in the report from the **Number of Entities** text box.

You can enter a maximum of 500 entities per Line Chart. Prism splits the line chart view into multiple views to display large number of entities in the generated report. Each view has a maximum of 10 entities.

Note: This field does not appear for the Histogram view.

- f. Select how you intend to represent the entities (**Ascending or Descending**) from the **Sort Order** dropdown menu.

- g. (Only for Histogram view) Enter number of buckets in the **Number of Objects** field.

You can define the frequency distribution for all the entities in the entity type for the selected metric by using this field. For example, for a selected metric if the minimum value is 10 and maximum value is 100 and you have entered number of buckets as 5 then you will have 18 buckets configured $(100-10/5)$. Depending on the value of the metric, different metrics will fall in the appropriate buckets.

- h. Select the entities that you want to add.

- » **All Entities:** Select this option if you want to include all the entities that you have selected in the view.
- » **Specific Entities:** Select this option to define rules on a particular entity that you have selected from the **Entity Type** dropdown menu. For example, if you select entity type as VM and desire to limit the view for

the 50 powered on VMs, you can select **Specific vms** option and define the rule **Power State > Equal to > On** . You can specify multiple such rules by clicking the **+** icon.

To generate a list of entities that are not associated with any category in a report, define the rule as **Categories > Equal to > <blank>** .

Important: Leaving the last field blank lists all the entities without a user-assigned category.

The screenshot shows a 'Rules' configuration interface. At the top, it says 'Rules' and 'In case of multiple rules, a conjunction (AND operator) will be applied between them'. Below this, there is a rule configuration row with three input fields: 'Categories', 'Equal to', and an empty field. Each field has a dropdown arrow on its right. To the right of these fields is a blue button labeled 'OR' and a red minus sign. Below the first field, there is a blue plus sign icon for adding more rules.

Figure 101: Rules Filter for Uncategorized Entities

- i. To configure advanced setting, click the **Show Advanced Settings** link. Select the **Group Data** checkbox to select the entities based on which the data in the view should be grouped and the views should be repeated.

6. Click Add.

The configured view gets added to the report.

Adding a Bar Chart or a Line Chart or a Histogram View to a Report (vCenter)

You can use the bar chart, line chart, or a histogram view to represent the historical data in a graphical format. Perform the following procedure to add these views to the report.

Procedure

1. Log in to Prism Central.
2. Select the **Intelligent Operations** application from **Application Switcher** function and navigate to **Reports** from the **Navigation Bar**.

The Reports page displays the **Configuration** tab by default. The **Configuration** tab provides information about custom and predefined system report configurations and saved report configurations about the Nutanix and vCenter entities.
3. Click **+ Report Configurations**.
The **Report Preview** window opens.
4. Select the type of view (bar chart, line chart, or a histogram view) that you want to add from the **Custom** view pane.

5. Configure the bar chart, line chart, or a histogram view.
Following image displays the bar chart view.

Add Bar Chart ? X

ENTITY TYPE

vCenter Entities External vCenter Cluster

METRIC

CPU Usage

TITLE

External vCenter Cluster CPU Usage

AGGREGATION

Average

NUMBER OF ENTITIES

- 10 +

SORT ORDER

Ascending

Entities

Select the entities that need to be reported in the view.

☒ All external vcenter clusters

☐ Specific external vcenter clusters

[Show Advanced Settings](#)

Cancel Add

Figure 102: Bar Chart View

- a. Select **vCenter Entities** and the entity type (**External vCenter Cluster**, **External vCenter Host**, or **External vCenter VM**) from the **Entity Type** dropdown menu for which you want to configure the view.

- b. Select a specific metric of the entity type.

The metric selection depends on what you have selected in the **Entity Type** dropdown menu. For example, if you select **External vCenter VM** as an entity type, you can select one of the following metric.

Ready, CPU Usage, Highest latency, Average read requests per second, Average read requests per second, Read rate, Read latency, Write latency, Controller Write I/O Bandwidth, Controller Read I/O Latency, Controller Write I/O Latency, Storage Logical Usage, Controller Read IOPS, Controller Write IOPS.

- c. Enter the title of the bar chart, line chart, or histogram view in the **Title** field.

- d. Select the aggregation from the **Aggregation** menu.

For the metrics that has time-series data, you can define the aggregation field. For example, if you want to include maximum CPU usage for a metric, you can use this field and select the **CPU Usage** as the metric and select aggregation as **Max**. Depending on the report schedule that you have defined, maximum value of the CPU usage is captured and included as part of the report.

You can define the aggregation value of **Sum, Max, Min, Average, Last, or Count**.

Note: Aggregation is not supported for Line Chart view.

- e. Enter the number of entities that you want to include in the report from the **Number of Entities** text box.

You can enter a maximum of 500 entities per Line Chart. Prism splits the line chart view into multiple views to display large number of entities in the generated report. Each view has a maximum of 10 entities.

Note: This field does not appear for the Histogram view.

- f. Select the **Limit entities in the chart** option to limit the number of entities that you want to include in the view and how you represent the entities (**Ascending** or **Descending**).

- g. (Only for Histogram view) Enter number of buckets in the **Number of Objects** field.

You can define the frequency distribution for all the entities in the entity type for the selected metric by using this field. For example, for a selected metric if the minimum value is 10 and maximum value is 100 and you have entered number of buckets as 5 then you will have 18 buckets configured $(100-10/5)$. Depending on the value of the metric, different metrics will fall in the appropriate buckets.

- h. Select the entities that you want to add.

- » **All Entities:** Select this option if you want to include all the entities that you have selected in the view
- » **Specific Entities:** Select this option to define rules on a particular entity that you have selected from the **Entity Type** dropdown menu. For example, if you select entity type as VM and desire to limit the view for

the 50 powered on VMs, you can select **specific VMs** option and define the rule **Power State > Equal to > On** . You can specify multiple such rules by clicking the **+** icon.

To generate a list of entities that are not associated with any category in a report, define the rule as **Categories > Equal to > <blank>** .

Important: Leaving the last field blank lists all the entities without a user-assigned category.

The screenshot shows a 'Rules' configuration interface. At the top, it says 'Rules' and 'In case of multiple rules, a conjunction (AND operator) will be applied between them'. Below this, there is a rule configuration row with three input fields: 'Categories', 'Equal to', and an empty field. Each field has a dropdown arrow on its right. To the right of these fields is a blue button labeled 'OR' and a red minus sign. Below the first field, there is a blue plus sign icon for adding more rules.

Figure 103: Rules Filter for Uncategorized Entities

- i. To configure advanced setting, click the **Show Advanced Settings** link. Select the **Group Data** checkbox to select the entities based on which the data in the view should be grouped and the views should be repeated.

6. Click Add.

The configured view gets added to the report.

Adding a Data Table View to a Report (Nutanix)

The data table view provides you with the point in time data of a particular entity. To include a data table view to the report, perform the following procedure.

Procedure

1. Log in to Prism Central.
2. Select the **Intelligent Operations** application from **Application Switcher** function and navigate to **Reports** from the **Navigation Bar**.

The Reports page displays the **Configuration** tab by default. The **Configuration** tab provides information about custom and predefined system report configurations and saved report configurations about the Nutanix and vCenter entities.
3. Click **+ Report Configurations**.
The **Report Preview** window opens.
4. Select the **Data Table** view from the **Custom** view pane.

5. Configure the Data Table view.

Add Data Table

?

×

TITLE

data-table-report

DESCRIPTION

This is a test data table report

Entities

Select the entities that need to be reported in the view.

ENTITY TYPE

Nutanix Entities

VM

☐ All vms

☒ Specific vms

Rules

In case of multiple rules, a conjunction (AND operator) will be applied between them

+

Columns

FOCUS

Custom Columns

General

Column Name	Aggregation
Name	-
Power State	-
Project	-
IP Addresses	-
Owner	-

Sorting

SORT BASED ON

Name

SORT ORDER

Ascending

Show Advanced Settings

Cancel

Add

Figure 104: Data Table View

category

- a. Enter the title of the data table view in the **Title** field.
- b. Enter the description of the view in the **Description** field.
- c. Select **Nutanix Entities** and the entity type (**Cluster**, **Container**, **Disk**, **Host**, or **VM**) for which you want to configure the view.
- d. Select the entities that you want add.
 - » **All Entities**: Select this option if you want to include all the entities that you have selected in the view.
 - » **Specific Entities**: Select this option to define rules on a particular entity that you have selected from the **Entity Type** dropdown menu. For example, if you select entity type as VM and desire to limit the

view for the 50 powered on VMs, you can select **specific VMs** option and define the rule **Power State** > **Equal to** > **On** . You can specify multiple such rules by clicking the **+** icon.

To generate a list of entities that are not associated with any category in a report, define the rule as **Categories** > **Equal to** > **<blank>** .

Important: Leaving the last field blank lists all the entities without a user-assigned category.




Figure 105: Rules Filter for Uncategorized Entities

- e. For an entity that you have selected in the **Entity Type** dropdown menu, you can select a list of metrics and each of these metrics are displayed as columns in the table.

The metric columns are further divided in to **General** or **Performance** depending on the entity that you have selected. Depending on your selection, pre-defined columns are displayed that will be included as part of the view. However, you can create your own custom column.

- f. To create your own custom column, click the **Custom Columns** link and click the **+** icon to add desired columns.

Depending on your selection the focus area is updated.

Note:

- If any metric has time-series data, you can define the aggregation parameter also. If you do not define the aggregation, by default Average is configured for the metric.
- For tables with more than 10 columns, you can export the table only in CSV format.

- g. Define the sorting by selecting the entity and order on which sorting should be performed by selecting the entity from the **Sort Based On** dropdown menu and **Sort Order** dropdown menu.

Note: You can define sorting only on selected entities from the entity list.

- h. To configure advanced setting, click the **Show Advanced Settings** link.
- i. Select the **Limit number of rows in the table** checkbox to limit the number of rows that you want to include in the view and enter the value in the text box that is provided.
- j. Select the **Group Data** checkbox to select the entities based on which the data in the view should be grouped and the views should be repeated.

- » **All Entities:** Select this option if you want to perform the grouping for all the entities.
- » **Specific Entities:** Select this option if you want to perform the grouping for selected entities depending on your selection from the **Entity Type** dropdown menu. For example, if you select entity type as VM

and desire to limit the view for the 50 powered on VMs, you can select **specific VMs** option and define the rule **Power State > Equal to > On** . You can specify multiple such rules by clicking the **+** icon.

For example, if you want to perform grouping for all the VMs at the cluster level and you have three clusters registered, you will get three tables and each table will have VMs for that particular cluster. You can define multiple such views. The entire set of views as a group will be repeated for each of the registered cluster.

If you select *Clusters*, *VMs*, or *Hosts* as entities along with the specific entities option, and then select *Categories* in the first dropdown menu for defining the rules, you can select *Equal to* or *Not Equal to* in the second dropdown menu, and then enter a combination of category name and category value in the format *category name:category value* in the last text box to filter the relevant data. For example, you must enter *AppType:Hadoop* to filter the VMs that are assigned *AppType* category and *Hadoop* as its value (see the image below). You can see the category name and category value correlation in the Categories view. You can leave the last text box blank to filter all the entities that do not have a user-defined category assigned to them.

Add Data Table ? X

TITLE

DESCRIPTION

Entities
 Select the entities that need to be reported in the view.

ENTITY TYPE
 Nutanix Entities : VM :
☐ All vms
☒ Specific vms

Rules
 In case of multiple rules, a conjunction (AND operator) will be applied between them

Categories : Equal to : AppType:Hadoop **OR** **-**
☒

Columns

FOCUS Custom Columns

Column Name	Aggregation
Name	

Cancel Add

Figure 106: Using Category with Data Table Rules

- Click **Add**.
 The configured view gets added to the report.

Adding a Data Table View to a Report (vCenter)

The data table view provides you with the point in time data of a particular entity. To include a data table view to the report, perform the following procedure.

Procedure

1. Log in to Prism Central.
2. Select the **Intelligent Operations** application from **Application Switcher** function and navigate to **Reports** from the **Navigation Bar**.

The Reports page displays the **Configuration** tab by default. The **Configuration** tab provides information about custom and predefined system report configurations and saved report configurations about the Nutanix and vCenter entities.

3. Click **+ Report Configurations**.
The **Report Preview** window opens.
4. Select the **Data Table** view from the **Custom View** pane.

5. Configure the Data Table view.

Add Data Table

TITLE
vcenter-data-table-report

DESCRIPTION
This is a test data table report.

Entities
Select the entities that need to be reported in the view.

ENTITY TYPE
vCenter Entities External vCenter VM

☒ All external vcenter vms
☐ Specific external vcenter vms

Columns
[Custom Columns](#)

FOCUS
General

Column Name	Aggregation
Name	-
Power State	-
IP Addresses	-

Sorting

SORT BASED ON
Name

SORT ORDER
Ascending

[Show Advanced Settings](#)

Cancel Add

Figure 107: Data Table View

- Enter the title of the data table view in the **Title** field.
- Enter the description of the view in the **Description** field.
- Select **vCenter Entities** and the entity type (**External vCenter Cluster**, **External vCenter Host**, or **External vCenter VM**) from the **Entity Type** dropdown menu for which you want to configure the view.
- Select the entities that you want add.
 - » **All Entities:** Select this option if you want to include all the entities that you have selected in the view.

- » **Specific Entities:** Select this option to define rules on a particular entity that you have selected from the **Entity Type** dropdown menu. For example, if you select entity type as VM and desire to limit the view for the 50 powered on VMs, you can select **Specific external vcenter vms** option and define the rule **Power State > Equal to > On** . You can specify multiple such rules by clicking the **+** icon.
- e. For an entity that you have selected in the **Entity Type** dropdown menu, you can select a list of metrics and each of these metrics are displayed as columns in the table.

The metric columns are further divided in to **General** or **Performance** depending on the entity that you have selected. Depending on your selection, pre-defined columns are displayed that will be included as part of the view. However, you can create your own custom column.
- f. To create your own custom column, click the **Custom Columns** link and click the **+** icon to add desired columns.

Depending on your selection the focus area is updated.

Note:

- If any metric has time-series data, you can define the aggregation parameter also. If you do not define the aggregation, by default Average is configured for the metric.
- For tables with more than 10 columns, you can export the table only in CSV format.

- g. Define the sorting by selecting the entity and order on which sorting should be performed by selecting the entity from the **Sort Based On** dropdown menu and **Sort Order** dropdown menu.

Note: You can define sorting only on selected entities from the entity list.

- h. To configure advanced setting, click the **Show Advanced Settings** link.
- i. Select the **Limit number of rows in the table** checkbox to limit the number of rows that you want to include in the view and enter the value in the text box that is provided.
- j. Select the **Group Data** checkbox to select the entities based on which the data in the view should be grouped and the views should be repeated.
 - » **All Entities:** Select this option if you want to perform the grouping for all the entities.
 - » **Specific Entities:** Select this option if you want to perform the grouping for selected entities depending on your selection from the **Entity Type** dropdown menu. For example, if you select entity type as External vCenter VM and desire to limit the view for the 50 powered on VMs, you can select **Specific**

external vcenter cms option and define the rule **Power State > Equal to > On** . You can specify multiple such rules by clicking the **+** icon.

To generate a list of entities that are not associated with any category in a report, define the rule as **Categories > Equal to > <blank>** .

Important: Leaving the last field blank lists all the entities without a user-assigned category.

The screenshot shows a 'Rules' section with a greyed-out instruction: 'In case of multiple rules, a conjunction (AND operator) will be applied between them'. Below this is a rule configuration row. It starts with a blue '+' icon. The first field is 'Categories' followed by a dropdown arrow. The second field is 'Equal to' followed by a dropdown arrow. The third field is an empty text box. To the right of the text box is a blue button labeled 'OR' and a red '-' icon.

Figure 108: Rules Filter for Uncategorized Entities

For example, if you want to perform grouping for all the VMs at the cluster level and you have three clusters registered, you will get three tables and each table will have VMs for that particular cluster. You can define multiple such views. The entire set of views as a group will be repeated for each of the registered cluster.

6. Click **Add.**

The configured view gets added to the report.

Adding Configuration Summary View

You can add this view to include the static cluster configuration information, for example, cluster IP address, hardware model, license information, etc.

Procedure

1. Log in to Prism Central.
2. Select the **Intelligent Operations** application from **Application Switcher** function and navigate to **Reports** from the **Navigation Bar**.

The Reports page displays the **Configuration** tab by default. The **Configuration** tab provides information about custom and predefined system report configurations and saved report configurations about the Nutanix and vCenter entities.

3. Click **+ Report Configurations**.
The **Report Preview** window opens.
4. Select the **Configuration Summary** view from the **Custom View** pane.

5. Configure the Configuration Summary view.

Add Configuration Summary ? X

In case multiple entities are selected, every entity will have one view in the report. Metrics for multiple entities cannot be combined in a single metric summary view. X

CONFIGURATION INFORMATION

Cluster IP

TITLE

Cluster IP

Entities

Select the entities that need to be reported in the view.

☐ All clusters

☒ Specific clusters

Rules

In case of multiple rules, a conjunction (AND operator) will be applied between them

-

+

Cancel Add

Figure 109: Configuration Summary View

- a. Select the configuration information that you want to include from the **Configuration Information** dropdown menu.

You can include the configuration information for the **Cluster IP**, **Prism Central IP**, **PC Version**, **Cluster Version**, **Cluster Hypervisor**, **Number of Blocks**, **Hardware Model**, **Cluster License**, or **Prism Central License**.

Specific information is included depending on the cluster information that you have added. For example, if you have selected **Cluster IP**, the IP addresses of all the clusters that are registered is included in the view.

- b. Enter the title of the report in the **Title** field.
 - c. Select the entities that you want to add.
 - » **All Clusters:** Select this option if you want to include this view for all the clusters that are registered.
 - » **Specific Clusters:** Select this option to define rules on a particular cluster. For example, if you want to include this view for AHV, you can select **Hypervisors** option and define the rule **Equal to** > **AHV** . You can specify multiple such rules by clicking the **+** icon.
- To generate a list of entities that are not associated with any category in a report, define the rule as **Categories** > **Equal to** > **<blank>** .

Important: Leaving the last field blank lists all the entities without a user-assigned category.

Rules

In case of multiple rules, a conjunction (AND operator) will be applied between them

Categories Equal to

OR

+

Figure 110: Rules Filter for Uncategorized Entities

6. Click **Add**.
The configured view gets added to the report.

Adding Metric Summary View (Nutanix)

You can add the metric summary view to get information on a metric for a defined aggregation value.

Procedure

1. Log in to Prism Central.
2. Select the **Intelligent Operations** application from **Application Switcher** function and navigate to **Reports** from the **Navigation Bar**.

The Reports page displays the **Configuration** tab by default. The **Configuration** tab provides information about custom and predefined system report configurations and saved report configurations about the Nutanix and vCenter entities.
3. Click **+ Report Configurations**.
The **Report Preview** window opens.
4. Select the **Metric Summary** view from the **Custom View** pane.

5. Configure the Metric Summary view.

Add Metric Summary

?

×

In case multiple entities are selected, every entity will have one view in the report. Metrics for multiple entities cannot be combined in a single metric summary view.

×

ENTITY TYPE

Nutanix Entities

VM

METRIC

Storage Controller IOPS

TITLE

VM Storage Controller IOPS

AGGREGATION

Average

DISPLAY MODE

Textual

Entities

Select the entities that need to be reported in the view.

☐ All vms

☒ Specific vms

Rules

In case of multiple rules, a conjunction (AND operator) will be applied between them

+

Cancel

Add

Figure 111: Metric Summary View

- a. Select **Nutanix Entities** and entity type (**Cluster**, **Container**, **Disk**, **Host**, or **VM**) for which you want to define the metric summary.
- b. For an entity that you have selected in the **Entity Type** dropdown menu, you can select a metric for which you want to add the information in the metric summary view.
- c. Enter the title of the view in the **Title** field.
- d. Select the aggregation from the **Aggregation** dropdown menu.
For the metrics that has time-series data, you can define the aggregation. For example, if you want to include maximum CPU usage for a metric, you can use this field and select the **CPU Usage** as the metric select aggregation as **Max**. Depending on the report schedule that you have defined, maximum value of the CPU usage is captured and included as part of the view.
You can define the aggregation value of **Sum**, **Max**, **Min**, **Average**, **Last**, or **Count**.
- e. Select the mode of display from the **Display Mode** dropdown menu.
You can select either **Textual** or **Graphical (Gauge Chart)**. The graphical chart is only displayed for consumption related metrics, for example, CPU usage.
- f. Select the entities that you want to add.
 - » **All Entities**: Select this option if you want to include all the entities that you have selected in the view.
 - » **Specific Entities**: Select this option to define rules on a particular entity that you have selected from the **Entity Type** dropdown menu. For example, if you select entity type as VM and desire to limit the view for the 50 powered on VMs, you can select **Specific VMs** option and define the rule **Power State > Equal to > On** . You can specify multiple such rules by clicking the **+** icon.

To generate a list of entities that are not associated with any category in a report, define the rule as **Categories > Equal to > <blank>** .

Important: Leaving the last field blank lists all the entities without a user-assigned category.

Rules

In case of multiple rules, a conjunction (AND operator) will be applied between them

Categories	:	Equal to	:		OR	+

+

Figure 112: Rules Filter for Uncategorized Entities

6. Click **Add**.
The configured view gets added to the report.

Adding Metric Summary View (vCenter)

You can add the metric summary view to get information on a metric for a defined aggregation value.

Procedure

1. Log in to Prism Central.

2. Select the **Intelligent Operations** application from **Application Switcher** function and navigate to **Reports** from the **Navigation Bar**.

The Reports page displays the **Configuration** tab by default. The **Configuration** tab provides information about custom and predefined system report configurations and saved report configurations about the Nutanix and vCenter entities.

3. Click **+ Report Configurations**.
The **Report Preview** window opens.
4. Select the **Metric Summary** view from the **Custom View** pane.

5. Configure the Metric Summary view.

Add Metric Summary

In case multiple entities are selected, every entity will have one view in the report. Metrics for multiple entities cannot be combined in a single metric summary view.

ENTITY TYPE

vCenter Entities External vCenter VM

METRIC

Write rate

TITLE

External vCenter VM Write rate

AGGREGATION

Average

DISPLAY MODE

Textual

Entities

Select the entities that need to be reported in the view.

☒ All external vcenter vms

☐ Specific external vcenter vms

Cancel Add

Figure 113: Metric Summary View

- Select **vCenter Entities** and the entity type (**External vCenter Cluster**, **External vCenter Host**, or **External vCenter VM**) from the **Entity Type** dropdown menu for which you want to configure the view.
- For an entity that you have selected in the **Entity Type** dropdown menu, you can select a metric for which you want to add the information in the metric summary view.

c. Enter the title of the view in the **Title** field.

d. Select the aggregation from the **Aggregation** dropdown menu.

For the metrics that has time-series data, you can define the aggregation. For example, if you want to include maximum CPU usage for a metric, you can use this field and select the **CPU Usage** as the metric select aggregation as **Max**. Depending on the report schedule that you have defined, maximum value of the CPU usage is captured and included as part of the view.

You can define the aggregation value of **Sum, Max, Min, Average, Last, or Count**.

e. Select the mode of display from the **Display Mode** dropdown menu.

You can select either **Textual** or **Graphical (Gauge Chart)**. The graphical chart is only displayed for consumption related metrics, for example, CPU usage.

f. Select the entities that you want to add.

» **All Entities:** Select this option if you want to include all the entities that you have selected in the view.

» **Specific Entities:** Select this option to define rules on a particular entity that you have selected from the **Entity Type** dropdown menu. For example, if you select entity type as External vCenter VM and desire to limit the view for the 50 powered on VMs, you can select **Specific external vcenter vms** option and define the rule **Power State > Equal to > On**. You can specify multiple such rules by clicking the **+** icon.

To generate a list of entities that are not associated with any category in a report, define the rule as **Categories > Equal to > <blank>**.

Important: Leaving the last field blank lists all the entities without a user-assigned category.

Rules

In case of multiple rules, a conjunction (AND operator) will be applied between them

Categories ▾ Equal to ▾ OR -

+

Figure 114: Rules Filter for Uncategorized Entities

6. Click **Add**.

The configured view gets added to the report.

Adding Entity Count View (Nutanix)

You can add the entity count view to get the number of entities for all the clusters or a particular cluster.

Procedure

1. Log in to Prism Central.

2. Select the **Intelligent Operations** application from **Application Switcher** function and navigate to **Reports** from the **Navigation Bar**.

The Reports page displays the **Configuration** tab by default. The **Configuration** tab provides information about custom and predefined system report configurations and saved report configurations about the Nutanix and vCenter entities.

3. Click **+ Report Configurations**.
The **Report Preview** window opens.
4. Select the **Entity Count** view from the **Custom View** pane.
5. Configure the Entity Count view.

Add Entity Count?||X

ENTITY TYPE

Nutanix Entities

:

VM

:

TITLE

VM count

Entities

Select the entities that need to be reported in the view.

☐ All vms

☒ Specific vms

Rules

In case of multiple rules, a conjunction (AND operator) will be applied between them

:

:

-

+

Cancel

Add

Figure 115: Entity Count View

- a. Select **Nutanix Entities** and the entity type (**Cluster**, **Container**, **Disk**, **Host**, or **VM**) for which you want to get the count information.
- b. Enter the title of the view in the **Title** field.
- c. Select the entities that you want to add.
 - » **All Entities**: Select this option if you want to include all the entities that you have selected in the view.
 - » **Specific Entities**: Select this option to define rules on a particular entity that you have selected from the **Entity Type** dropdown menu. For example, if you select entity type as VM and desire to limit the view for

the 50 powered on VMs, you can select **Specific VMs** option and define the rule **Power State > Equal to > On** . You can specify multiple such rules by clicking the **+** icon.

To generate a list of entities that are not associated with any category in a report, define the rule as **Categories > Equal to > <blank>** .

Important: Leaving the last field blank lists all the entities without a user-assigned category.

The screenshot shows a 'Rules' configuration interface. At the top, it says 'Rules' and 'In case of multiple rules, a conjunction (AND operator) will be applied between them'. Below this, there is a rule configuration row. It starts with a dropdown menu showing 'Categories', followed by a colon separator, another dropdown menu showing 'Equal to', another colon separator, and an empty text input field. To the right of the input field is a blue button labeled 'OR' and a red minus icon. Below the rule configuration row is a blue plus icon.

Figure 116: Rules Filter for Uncategorized Entities

6. Click Add.

The configured view gets added to the report.

Adding Entity Count View (vCenter)

You can add the entity count view to get the number of entities for all the clusters or a particular cluster.

Procedure

1. Log in to Prism Central.
2. Select the **Intelligent Operations** application from **Application Switcher** function and navigate to **Reports** from the **Navigation Bar**.

The Reports page displays the **Configuration** tab by default. The **Configuration** tab provides information about custom and predefined system report configurations and saved report configurations about the Nutanix and vCenter entities.

3. Click **+ Report Configurations**.
The **Report Preview** window opens.
4. Select the **Entity Count** view from the **Custom View** pane.

5. Configure the Entity Count view.

Add Entity Count

ENTITY TYPE

vCenter Entities External vCenter Datastore

TITLE

External vCenter Datastore count

Entities

Select the entities that need to be reported in the view.

☒ All external vcenter datastores

☐ Specific external vcenter datastores

Cancel Add

Figure 117: Entity Count View

- Select **vCenter Entities** and entity type (**External vCenter Cluster**, **External vCenter Host**, or **External vCenter VM**) from the **Entity Type** dropdown menu for which you want to configure the view.
- Enter the title of the view in the **Title** field.
- Select the entities that you want to add.
 - All Entities:** Select this option if you want to include all the entities that you have selected in the view.
 - Specific Entities:** Select this option to define rules on a particular entity that you have selected from the **Entity Type** dropdown menu. For example, if you select entity type as External vCenter VM and desire to limit the view for the 50 powered on VMs, you can select **Specific external vcenter vms** option and

define the rule **Power State > Equal to > On** . You can specify multiple such rules by clicking the **+** icon.

To generate a list of entities that are not associated with any category in a report, define the rule as **Categories > Equal to > <blank>** .

Important: Leaving the last field blank lists all the entities without a user-assigned category.



Figure 118: Rules Filter for Uncategorized Entities

6. Click Add.

The configured view gets added to the report.

Adding Title and Description View

You can add the title and description view to add any miscellaneous information, for example, disclaimer etc.

Procedure

1. Log in to Prism Central.
2. Select the **Intelligent Operations** application from **Application Switcher** function and navigate to **Reports** from the **Navigation Bar**.

The Reports page displays the **Configuration** tab by default. The **Configuration** tab provides information about custom and predefined system report configurations and saved report configurations about the Nutanix and vCenter entities.

3. Click + Report Configurations.

The **Report Preview** window opens.

4. Click New Report.

5. In the window that opens, select the Title and Description view from the Custom View pane.

6. Configure the Title and Description view.

- a. Enter the title of the view in the **Title** field.
- b. Enter the description in the **Description** field.

7. Click Add.

The configured view gets added to the report.

Adding a Group View

About this task

Grouping is always performed in an hierarchical manner. For example, if you have created a group for the cluster entity, you can add all the views to this group. However, if you have created the group for the VM

entity, you cannot add a view to this group that you have created at the cluster level. You cannot group Nutanix and vCenter entities in a single view. For example, a Nutanix group for VM entity cannot be added to a vCenter group for VM entity.

Note: This view takes the full page width.

Perform the following steps to add a group view.

Procedure

1. Log in to Prism Central.
2. Select the **Intelligent Operations** application from **Application Switcher** function and navigate to **Reports** from the **Navigation Bar**.

The Reports page displays the **Configuration** tab by default. The **Configuration** tab provides information about custom and predefined system report configurations and saved report configurations about the Nutanix and vCenter entities.

3. Click **+ Report Configurations**.
The **Report Preview** window opens.
4. Select the **Group** view from the **Custom** view pane.

5. Configure the Group view.

Add Group

?

X

TITLE

Group1

DESCRIPTION

☒ Repeat views in this group

Select the entities based on which the data in the view will be grouped and the views will be repeated.

ENTITY TYPE

Nutanix Entities

Storage Container

☒ All storage containers

☐ Specific storage containers

Cancel

Add

Figure 119: Group View (Nutanix)

Add Group

?

X

TITLE

DESCRIPTION

☒ Repeat views in this group

Select the entities based on which the data in the view will be grouped and the views will be repeated.

ENTITY TYPE

vCenter Entities

:

External vCenter VM

:

☒ All external vcenter vms

☐ Specific external vcenter vms

Cancel

Add

Figure 120: Group View (vCenter)

- Enter the title of the data table view in the **Title** field.
- Enter the description of the view in the **Description** field.
- Select the **Repeat views in this group** checkbox to select the entities based on which the data in the view should be grouped and the views should be repeated.
- Select the entity type from the **Entity Type** dropdown menu.
 - » **All Entities:** Select this option if you want to perform the grouping for all the entities.
 - » **Specific Entities:** Select this option if you want to perform the grouping for selected entities depending on your selection from the **Entity Type** dropdown menu. For example, for Nutanix entities, if you select entity type as VM and desire to limit the view for the 50 powered on VMs, you can select **specific VMs** option and define the rule **Power State > Equal to > On** . You can specify multiple such rules by clicking the **+** icon. Similarly, for vCenter entities, if you select entity type as External vCenter VM and desire to limit the view for the 50 powered on VMs, you can select **Specific external vcenter vms** option and define the rule **Power State > Equal to > On** . You can specify multiple such rules by clicking the **+** icon.

6. Click **Add**.

The configured view gets added to the report. Depending on the entity type that you have configured, you can create grouping of the views. For example, if you have created the group view for the entity type VM, you can add all the views to this except the view with entity type cluster.

What to do next

You can only add group view by following this procedure. If you want to add any views to this group, see [Adding Views into a Group View](#) on page 257

Adding Views into a Group View

You can add a view to a group in multiple ways. The addition of views depends on whether you have created the group view or not.

About this task

Perform the following steps to add views to a already created group view.

Procedure

1. If you have created the group view, perform the following procedure. For more information on creating the group view, see [Adding a Group View](#) on page 253

- a. Log in to Prism Central.

- b. Select the **Intelligent Operations** application from **Application Switcher** function, and navigate to **Reports** from the **Navigation Bar**.

The Reports page displays the **Configuration** tab by default. The **Configuration** tab provides information about all the report configurations created about the Nutanix and vCenter entities.

- c. Select the checkbox associated with the report in which you want to add a view to the group view.

- d. From the **Actions** dropdown menu, select **Edit**.

- e. In the window that opens, select the checkbox of the view that you want to add to the group view from the **Report Preview** section, and click the **Actions** dropdown menu in the extreme right of the view.

The group to which you can add this view is displayed in the dropdown menu. You can perform the group only in an hierarchical manner. For example, if you have created group for the cluster entity, you can add all the

views to this group. However, if you have created the group for the VM entity, you cannot add a view to this group that you have created at the cluster level.

- f. Select the group that you want to move this view.

For example, the bar chart view with title Cluster CPU Usage can be moved to the group named **Cluster Group**.




Figure 121: Adding Views to Group view

The view is moved in the group view.

2. If you have not created the group view, perform the following procedure.
 - a. Access the Edit report page by following the Steps a through d in Step 1.
 - b. In the window that opens, select the checkbox of the view that you want to add to the group view from the **Report Preview** section, and click **Create Group** button at the top.
 - c. Create the group by following the procedure [Adding a Group View](#) on page 253.
The view is added to the newly created group.

Depending on the grouping, the data in the view is grouped and the views are repeated. For example, if you want to perform grouping for all the VMs at the cluster level and you have three cluster registered, you will get three tables and each table will have VMs for that particular cluster. You can define multiple such views. The entire set of views as a group is repeated for the each cluster.

Adding a Predefined View

In addition to adding and configuring views according to your requirements, you can also add some predefined views. You cannot modify these views.

About this task

As part of the **Predefined** views, you can add the following views.

Note:

- You can only include the **License Details Table** view as part of the group view.
- You can only view DR specific widgets when a DR report is edited.

Procedure

1. Log in to Prism Central.
2. Select the **Intelligent Operations** application from **Application Switcher** function and navigate to **Reports** from the **Navigation Bar**.

The Reports page displays the **Configuration** tab by default. The **Configuration** tab provides information about custom and predefined system report configurations and saved report configurations about the Nutanix and vCenter entities.

3. Click **+ Report Configurations**.
The **Report Preview** window opens.

4. Select the view that you want to add from the **Predefined Views** pane.

You can add the both detailed and summary licensing information and block summary information.

- **License Details Table:** Provides detailed information about the Prism Central licenses and includes information about license ID, model that are using the licenses, type of licenses, displays whether the licenses are used or not, and displays when the licenses are going to expire.
- **Prism Central License Summary:** Provides summary information of the Prism Central licenses and includes information about the type of licenses, number of used licenses, and number of licenses that are available.
- **Alerts Histogram:** Provides distribution of alerts over a specific time interval as configured during the execution of the report. X-axis plots the time period of the report as configured in the **Time Period of Report** field. Y-axis plots the number of alerts generated in the specific time period.
- **Blocks Summary:** Provides summarized information about the different block models, number of nodes, and number of VMs running in the Prism Central.
- **CPU Runway:** Provides visualized information about the historical CPU usage, predicted CPU usage, and predicted runway of the cluster.
- **Memory Runway:** Provides visualized information about the historical memory usage, predicted memory usage, and predicted runway of the cluster.
- **Storage Runway:** Provides visualized information about the historical storage usage, predicted storage usage, and predicted runway of the cluster.
- **Inactive VMs List:** Provides detailed information about the inactive VMs on the cluster; including VM name, efficiency status, reasons for being marked as inactive.
- **Constrained VMs List:** Provides detailed information about the constrained VMs on the cluster (that require more resources to function properly); including the VM name, efficiency status and the reasons for being marked as constrained.
- **Overprovisioned VMs List:** Provides detailed information about the over-provisioned VMs on the cluster (that have more resources than they deserve); including the VM name, efficiency status, reasons for being marked as over-provisioned.
- **Bully VMs List:** Provides detailed information about the VMs whose activities are affecting other VMs that is not directly related to their sizes; including VM name, efficiency status and the reasons for being marked as bully.
- **Potential CPU Reclaim:** Provides summarized information about the potential CPU resource reclaim on the cluster from the inactive VMs and the over-provisioned VMs.
- **Potential Memory Reclaim:** Provides summarized information about the potential memory resource reclaim on the cluster from the inactive VMs and the over-provisioned VMs.
- **Potential Storage Reclaim:** Provides summarized information about the potential storage resource reclaim on the cluster from the inactive VMs and the over-provisioned VMs.
- **vCenter Blocks Summary:** Provides summarized information about the different block models, number of nodes, and number of VMs running in vCenter.
- **vCenter CPU Runway:** Provides summarized information about the potential CPU resource reclaim on the cluster from inactive VMs and over-provisioned VMs.
- **vCenter Memory Runway:** Provides summarized information about the potential memory resource reclaim on the cluster from inactive VMs and over-provisioned VMs.

- **vCenter Storage Runway:** Provides summarized information about the potential storage resource reclaim on the cluster from inactive VMs and over-provisioned VMs.
- **vCenter Inactive VMs List:** Provides detailed information about the inactive VMs on the vCenter cluster; including VM name, efficiency status, reasons for being marked as inactive.
- **vCenter Constrained VMs List:** Provides detailed information about the constrained VMs on the vCenter cluster (that require more resources to function properly); including the VM name, efficiency status and the reasons for being marked as constrained.
- **vCenter Overprovisioned VMs List:** Provides detailed information about the over-provisioned vCenter registered VMs on the cluster (that have more resources than they deserve); including the VM name, efficiency status, reasons for being marked as over-provisioned.
- **vCenter Bully VMs List:** Provides summarized information about the potential storage resource reclaim on the cluster from inactive VMs and over-provisioned VMs.
- **vCenter Potential CPU Reclaim:** Provides summarized information about the potential CPU resource reclaim on the vCenter cluster from inactive VMs and over-provisioned VMs.
- **vCenter Potential Memory Reclaim:** Provides summarized information about the potential memory resource reclaim on the vCenter cluster from inactive VMs and over-provisioned VMs.
- **vCenter Potential Storage Reclaim:** Provides summarized information about the potential storage resource reclaim on the vCenter cluster from inactive VMs and over-provisioned VMs.
- **Configuration Alerts:** Provides summarized information about the potential storage resource reclaim on the cluster from inactive VMs and over-provisioned VMs.
- **DR Alerts Summary Table:** Provides summarized information about the potential storage resource reclaim on the cluster from inactive VMs and over-provisioned VMs.
- **Recovery Plan Events:** Provides summarized information about the potential storage resource reclaim on the cluster from inactive VMs and over-provisioned VMs.
- **Primary:** Provides summarized information about the potential storage resource reclaim on the cluster from inactive VMs and over-provisioned VMs.
- **Recovery:** Provides summarized information about the potential storage resource reclaim on the cluster from inactive VMs and over-provisioned VMs.
- **Entity Protection Summary:** Provides summarized information about the potential storage resource reclaim on the cluster from inactive VMs and over-provisioned VMs.
- **List of Entities with RPO Not Met:** Provides summarized information about the potential storage resource reclaim on the cluster from inactive VMs and over-provisioned VMs.
- **List of Entities Meeting RPO:** Provides summarized information about the potential storage resource reclaim on the cluster from inactive VMs and over-provisioned VMs.
- **Recovery Plan Events Details:** Provides summarized information about the potential storage resource reclaim on the cluster from inactive VMs and over-provisioned VMs.

Adding a Saved View

This task describes how to add saved views while creating a new report or to an existing report.

About this task

The views you saved from the Charts, Playbooks, Alerts, and Events page are available as saved views to add to a report. The saved views cannot be modified.

Procedure

1. Log in to Prism Central.
2. Select the **Intelligent Operations** application from **Application Switcher** function and navigate to **Reports** from the **Navigation Bar**.

The Reports page displays the **Configuration** tab by default. The **Configuration** tab provides information about custom and predefined system report configurations and saved report configurations about the Nutanix and vCenter entities.

3. Click **+ Report Configurations**.
The **Report Preview** window opens.
4. Select **Saved** tab from the **ADD VIEWS** pane.
5. Select the saved views that you want to add to the report.
The saved view gets added to the report.

Configuring Report Settings

You can configure the report settings by configuring its appearance, scheduling, email settings, retention policy, format, and time zone.

About this task

You can configure the report settings for individual reports or for all the reports that you have generated. Depending on where you configure the settings, the reports settings are applied.

Note: If you apply the report settings both the global level (for all the reports) or for report level (when you are creating a new report), the setting applied at the report level takes precedence.

Perform the following procedure to configure the report settings.

Procedure

1. Log in to Prism Central.
2. Select the **Intelligent Operations** application from **Application Switcher** function and navigate to **Reports** from the **Navigation Bar**.

The Reports page displays the **Configuration** tab by default. The **Configuration** tab provides information about custom and predefined system report configurations and saved report configurations about the Nutanix and vCenter entities.

3. Go to the report settings page.
 - » To configure the report settings at the global level, click **Report Settings** in the report **Configurations** tab.
 - » To configure the report settings at the report level, in the **New Report** wizard, click **Report Settings**.

Report Settings ? X

Appearance Email Settings Report Retention

Cover Page Settings

LOGO ?
Please select file... Upload

BACKGROUND COLOR

☒ Include Copyright Information

COPYRIGHT TEXT

Content Page Settings

HEADER COLOR

Reset to Default Cancel Save

Report Name
Date and Time of Generation
Owner Name

Figure 122: Global Report Settings Configuration

Report Settings ? X

< **Appearance** Add Schedule Email Settings Report Retention Report >

Cover Page Settings

LOGO ?
 Please select file... Upload

BACKGROUND COLOR
 [Dropdown]

☐ Include Copyright Information

Content Page Settings

HEADER COLOR
 [Dropdown]

Reset to Default Cancel Save

Figure 123: Report Level Settings Configuration

4. Configure the appearance of the report.

The **Appearance** tab of the report is divided into two parts; cover page settings and content page settings. Configuration of logo and the background color is included in the cover page settings. Configuration of color of the header is included in the content page settings. For colors, you can specify any of the 16 color names supported for HTML: aqua, black, blue, fuchsia, gray (or grey), green, lime, maroon, navy, olive, purple, red, silver, teal, white, and yellow.

Note: The copyright information will not be displayed in the footer section, therefore the ability to customize copyright information is not allowed.

- a. To upload the logo for the report, click the **Upload** button.

Note: The logo size must not be more than 1 MB.

- b. To change the background color of the report, select the background color from the **Background Color** field.
- c. To change the header color of the report, select the header color from the **Header Color** field.

5. Configure the Report schedule.
 - a. Click the **Add Schedule** tab.
 - b. Select when you want to run the report from the **Report Runs** dropdown menu.

You can run the report on daily, weekly, monthly, or yearly basis.
 - c. Select the time when you want to repeat the report.

For example, if you select **Yearly** from the **Report Runs** dropdown menu, you can repeat the report by selecting the month when you want to run the report from the **Every** dropdown menu, date of the month from the **On** dropdown menu, time of the day when you want to run the report from the **Generated Time** dropdown menu.
 - d. Select the duration for the data that you want to display in the report from the **Report Time Period** dropdown menu.

You can display the data for last 24 hours, last week, or last month. You also have an option to select a custom data range.
 - e. If you select a **Custom Date Range** from the **Report Time Period**, you need to select the date range from the **From** and **To** option.

If you select a date range from 10 a.m. (Monday) to 10 a.m. (Tuesday), the data for this time is displayed in the report. For example, in a line chart, the X-axis will display the time range from 10 a.m. (Monday) to 10 a.m. (Tuesday) and the Y-axis will display the variation in data.
 - f. Select the user impersonation to run the report from the **Run Report on Behalf of** option.
 - If you are the creator of the report, select **Self** button.
 - If you want to assign another user as an executor while creating the report, select **Other** button and choose the impersonated user from the dropdown list.
6. Configure the Email settings.
 - a. Click the **Email Settings** tab.
 - b. Enter the subject of the Email in the **Enter Content to Prepend to the Email Subject** text box.

You can view the preview of the subject that you are adding in the Preview pane and is prepended to the subject *subject_email* Prism Report *report_name*.
 - c. Enter the body of the Email in the **Enter Content to Append to the Email Body** text box.

You can view the preview of the body that you are adding in the Preview pane and is appended to the body Please find attached the generated report *email_body*.
 - d. Enter the recipients (comma-separated) of the report in the **Email Recipients** text box.

The recipients that you have configured will get the report with the subject and body of the Email that you have entered.

7. Configure the retention policy of the report.

You can define the retention policy by entering the number of instances of the report that you want to retain or the duration of time that you want to retain the report.

a. Click the **Report Retention** tab.

b. To configure the retention policy in instances form, select the **Number of Instances** option and enter the desired number of instances.

Note: You cannot retain more than 25 instances of a report. If you have 25 instances of a report and you generate a new instance of the same report, the first instance that was generated gets deleted.

c. To configure the retention policy in duration form, select the **Time Duration** option and enter the desired number and duration in terms of **Days, Weeks, or Months** from the dropdown menu.

For the time-based retention, you can retain a report for maximum of 3 months.

Select the **Do not retain reports** checkbox to delete all the configured retention policies.

Note: If you do not define any retention policy, by default 10 instances of a report are retained.

Perform the following step to choose the report format.

Note: You can choose the report format at the report level only.

8. Click the **Report Format** tab, and select the **PDF** checkbox or **CSV** checkbox or both to set a report format.

A PDF generated report displays a graphical representation of infrastructure resources. A CSV generated report contains information separated by commas. You can use this CSV file to load CSV data into other systems.

Note: You cannot set the report format at the global level. You can only set the report format at the report level while creating a new report or editing an existing report.

9. Click the **Time Zone** tab, and select your time zone from the dropdown list.

10. Click **Save**.

The report configuration is saved and next time you generate the report, the report generation follows the configured setting. If you apply the report settings both the global level (for all the reports) or for report level (when you are creating a new report), the setting applied at the report level takes precedence.

11. (Optional) Click **Reset to Default** to discard all the changes that you have made in the Report Settings page.

Role Based Access Control

Role-Based Access Control (RBAC) restricts Reports Management access to authorized users. For any role you are assigned, privileges are given to you to modify certain configurations in Reports Management. The following table shows the roles and certain privileges given to that role.

Table 58: Roles and Permissions

Access/ Permission	Consumer	Developer	Operator	Prism Admin	Prism Viewer	Project Admin	Self- Service Admin	Super Admin
Create Common Report Config				#		#	#	#
Create Report Config				#		#	#	#
Create Report Instance				#		#	#	#
Delete Common Report Config				#		#	#	#
Delete Report Config				#		#	#	#
Delete Report Instance				#		#	#	#
Notify Report Instance	#	#	#	#	#	#	#	#
Update Common Report Config				#		#	#	#
Update Report Config				#		#	#	#
View Common Report Config	#	#	#	#	#	#	#	#
View Report Config	#	#	#	#		#	#	#
View Report Instance	#	#	#	#	#	#	#	#

SETTINGS & CONFIGURATIONS

You can use **Intelligent Operation settings** to configure or update the following:

- Capacity Configurations: Update the capacity configurations. For more information, see [Updating Capacity Configurations](#) on page 118.
- Credentials: Create and manage your user access credentials to all the integrated non-Nutanix systems. These credentials are leveraged by all the Intelligent Operations features that require authentication into those non-Nutanix systems. For more information, see [Storing Non-Nutanix Credentials in Repository](#) on page 268.
- Operations Policies: Update the system-defined operations policies that by default run on all the entities in your environment. For more information, see [Operations Policy Management](#) on page 269.

Storing Non-Nutanix Credentials in Repository

About this task

To store a non-Nutanix system user access credential, perform the following steps:

Procedure

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher. The system displays the **Intelligent Operations** dashboard.
3. Navigate to **Operations Settings** from the **Navigation Bar**. The **Settings** page opens.
4. Click **Credentials**. The **Credentials** window appears.



Figure 124: Settings and Configurations page

5. Click **New Credential**. The system displays the **Configure Credential** window.

6. In the **Configure Credential** window, do the following in the indicated fields:
 - a. **External System Type:** Select the external system type.
Prism Central shows only vCenter external system type.
 - b. **External System IP:** Type the IP address of the external system.
 - c. **Username:** Type the user name of the external system.
 - d. **Password:** Type the password of the external system.
7. Click **Save**.
The non-Nutanix system user access credential is saved in the Prism Central repository.

Operations Policy Management

Prism Central provides the system services to analyze cluster integrity and performance. It uses the following machine-learning capabilities to analyze resource usage and identify the potential problems:

- *Anomaly Detection* - This capability records the performance or resource usage when it is outside an expected range.
- *VM Right-sizing or VM Inefficiency Detection* - This capability is used to identify the VMs that exhibit inefficient profiles.

For more information about machine-learning capabilities in Prism Central, see [Behavioral Learning Tools](#) on page 119.

Prism Central allows you to define the following operation policies to exclude these machine-learning capabilities for cases where you do not want the usage surge to be counted as an anomaly or to exclude selected VMs from the inefficiency analysis:

- **Anomaly Detection Exclusion** - This policy allows you to specify VMs to ignore when looking for anomalies. For example, you might have some VMs with expected peak load on certain days, and you do not want the usage surge of these VMs to be counted as an anomaly. In this case, you can either define a category for those VMs and enable anomaly detection exclusion for that category or enable the exclusion for each VM individually. To enable anomaly detection exclusion for each VM individually, select the target VM in the VMs **List** tab , and choose **Disable Anomaly Detection** from **Actions** drop down menu. For more information, see [Managing a VM through Prism Central \(AHV\)](#) in *Prism Central Infrastructure Guide*.
- **Inefficiency Measurement Exclusion** - This policy allows you to exclude selected VMs from the inefficiency analysis using the VM category. To enable inefficient measurement exclusion for each VM individually, select the target VM in the VMs **List** tab , and choose **Disable Efficiency Measurement Detection** from **Actions** drop down menu. For more information, see [Managing a VM through Prism Central \(AHV\)](#) in *Prism Central Infrastructure Guide*.

Operations Policies Summary View

The operations policies summary view allows you to access all the system-defined operation policies across all registered clusters.

To access the operations policies summary view, perform the following steps:

1. Log in to Prism Central.

2. Select **Intelligent Operations** from application switcher. The system displays the **Intelligent Operations** dashboard. Navigate to **Operations Settings** from the **Navigation Bar**.

The system displays the **List** tab by default with consolidated view of operations policies across all registered clusters.




Figure 125: Operations Policies List Tab

Table 59: Operations Policies Page - Field Description

Field	Description	Values
Name	Displays the operations policy name. Click the name to display the policy details for Anomaly Detection Exclusion or Inefficiency Measurement Exclusion operations policies.	(Anomaly Detection Exclusion Inefficiency Measurement Exclusion)
Purpose	Describes the purpose of the operation policy.	(description, string)
No. of excluded categories	Displays the number of categories that are excluded for anomaly detection or inefficient measurements checks.	(integer)
Policy Type	Describes the type of operations policy.	(Anomaly Detection Exclusion Inefficiency Measurement Exclusion)
Last Updated	Displays the date and time when the operations policy is last updated.	(date and time)

You can perform the following actions for the operations policies in the **Operations Policies** summary view:

- Export the table that contains the list of operations policies and their information to a file in a CSV format.
- View operations policies based on pre-defined criteria or create a custom view.
- Filter the operations policies list based on operations policy name using **Filters** pane.
- Update the operations policies to include VM category. For more information, see [Managing Operations Policies](#) on page 271.

Managing Operations Policies

This section describes how to customize the system-defined operations policies to apply them to the target VM categories.

Before you begin

Ensure that the VM categories are created, and the target VMs are added to the VM categories.

For information about how to create a category, see [Creating a Category](#) in *Prism Central Infrastructure Guide*.

For information about how to assign a category to the VMs, see [Assigning a Category](#) in *Prism Central Infrastructure Guide*.

About this task

System categories are assigned to the Anomaly Detection Exclusion and Inefficiency Measurement Exclusion policies by default.

Procedure

To modify (customize) the operations policies to include the selected VM categories, perform the following steps:

1. Log in to Prism Central.
2. Select **Intelligent Operations** from application switcher. The system displays the **Intelligent Operations** dashboard. Navigate to **Operations Settings** from the **Navigation Bar**.
The system displays the **List** tab by default with consolidated view of operations policies across all registered clusters.
3. Select the target operations policy (**Anomaly Detection Exclusion** or **Inefficiency Detection**) checkbox, and click **Update**.
The system displays the **Anomaly Detection Exclusion** or **Inefficiency Detection** window based on your selection.
4. Do one of the following:
 - To modify an Anomaly Detection Exclusion Policy, see [Modifying an Anomaly Detection Exclusion Policy](#) on page 271.
 - To modify an Inefficiency Detection Policy, see [Modifying an Inefficiency Detection Policy](#) on page 273.

Modifying an Anomaly Detection Exclusion Policy

About this task

For an Anomaly Detection Exclusion Policy, you can modify the list of categories that are excluded from anomaly detection.

Procedure

- 1. Click the **+ Add** link under **Categories** field to associate the VM category with the selected policy.

Anomaly Detection Exclusion

Name

Anomaly Detection Exclusion

Purpose

Excludes entities from anomaly detection

Categories

Click to select the VM category and add it to the policy

+ Add

Name	Value	
Testjit_VM_cat1	Test	
AnalyticsExclusions	AnomalyDetection	System ⓘ Click to delete the category
VM-Cat-Site-A	Site-A-VMs	


Cancel

Save

Figure 126: Operations Policies - Anomaly Detection Exclusion

2. Type the name of the VM category in the **Add** window. The system displays the list of matching categories based on the typed entry.

Use the Add icon and Remove icon to add and remove the required categories.



The screenshot shows a modal window titled "Add" with a close button (X) in the top right corner. Below the title bar, the text "Select categories" is displayed. A list of categories is shown, with "VM-Cat-Site-A:Site-A-VMs" selected and highlighted in blue. To the right of this category is a red minus icon. Below the list is a search bar with the placeholder text "Search for a category" and a blue plus icon to its right. At the bottom right of the window is a large blue "Save" button.

Figure 127: Add Window

3. Click **Save** to add the VM category.
The system returns to the selected operations policy window.
4. Click **Save**.
The policy is now enabled for entities in the selected categories.

Modifying an Inefficiency Detection Policy

About this task

For an Inefficiency Detection Policy, you can modify the list of categories that are excluded from inefficiency calculation. Additionally, you can update the lookback period for all the applicable VMs. The time frame for data used for Inefficiency Detection is called lookback period. By default, the lookback period is set to 30 days for all inefficiency calculations, you can change it to 60 days or 90 days.

Note: Updating lookback period for vCenter-monitored VMs is not supported.

Procedure

1. [Optional] In **Category** under **Exclude Entities from Inefficiency Detection** section, select the categories to be excluded from the inefficiency calculation.

Inefficiency Detection

Name
Inefficiency Detection

Purpose
Update lookback period & exclude entities

Exclude Entities from Inefficiency Detection
Select all the categories which needs to be excluded from the inefficiency calculation.

Category
AnalyticsExclusions.EfficiencyMeasurement

Update Lookback Period
By default the lookback period is set to 30 days for all inefficiency calculations. You can change the lookback period for all applicable VMs by updating the value below.

Lookback Period
30 days
30 days ✓
60 days
90 days

Figure 128: Operations Policies - Inefficiency Detection

2. [Optional] In **Lookback Period** dropdown, update the lookback period for all the applicable VMs.
3. Click **Save**.
The policy now excludes all the selected entities. If you have updated the lookback period, the updated lookback period is used for inefficient calculation for all VMs.

COPYRIGHT

Copyright 2025 Nutanix, Inc.

Nutanix, Inc.

1740 Technology Drive, Suite 150

San Jose, CA 95110

All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. Nutanix and the Nutanix logo are registered trademarks of Nutanix, Inc. in the United States and/or other jurisdictions. All other brand and product names mentioned herein are for identification purposes only and may be trademarks of their respective holders.

NX-3155-G9 System Specifications

Platform ANY

NX-3155-G9

September 16, 2025

Contents

- System Specifications..... 3**
 - Node Naming (NX-3155-G9).....3
 - NX-3155-G9 System Specifications..... 4
 - NX-3155-G9 GPU Specifications..... 10
- Component Specifications..... 12**
 - Controls and LEDs for Single-node Platforms..... 12
 - Network Card LED Description..... 14
 - Power Supply Unit Redundancy and Node Configuration.....15
- Memory Configurations..... 17**
 - Supported Memory Configurations..... 17
- Nutanix Hardware Naming Convention..... 20**
- Copyright..... 22**

SYSTEM SPECIFICATIONS

Node Naming (NX-3155-G9)

Nutanix assigns a name to each node in a block, which varies by product type.

The NX-3155-G9 block contains a single node named Node A. It supports hard disk drives (HDD) and non-volatile memory express (NVMe) drives.

Only drive slots 4, 5, 7, 8, 10, and 11 can be populated. All other drive slots must remain empty for thermal reasons.

The front panel contains the control panel, service tag, and the drives.



Figure 1: NX-3155-G9 Front Panel

Table 1: Supported Drive Configurations

Hybrid NVMe + HDD	Two NVMe drives in slots 4 and 5; four HDDs in slots 7, 8, 10, and 11; all other drive slots empty
All-NVMe	Two, four, or six NVMe drive in slots 4, 5, 7, 8, 10, and 11; all other drive slots empty

NX-3155-G9 supports partial population of drives for all-NVMe configurations, with a minimum of two drives. When adding drives in the field, populate the drives two by two, from left to right, in the active drive slots. The NX-3155-G9 does not support odd-numbered drive configurations. Hybrid NVMe + HDD configurations do not support partial population of drives.

Hybrid NVMe + HDD configuration is offered only as a factory installation. You cannot convert an all-NVMe node to a hybrid NVMe + HDD node in the field.




Figure 2: NX-3155-G9 Back Panel

The back panel maps the placement of these components:

- PSU 1 and PSU 2
- 1 x 100/1000 (IPMI)
- 2 x USB 3.0
- 15 Pin VGA
- UID LED
- AIO ports
 - Port 1 (Shared IPMI)
 - Port 2
- Slot 1 (NIC 3 or GPU 2)
- Slot 3 (NIC 2 or GPU 3)
- Slot 4 (NIC 1 or GPU 4)
- Slot 5 (GPU 1)
- Slot 6 (M.2 air shroud)
- Slot 7 (HBA or NIC 3)
- Slot 8 (M.2 air shroud, M.2 riser card, and M.2 Drives)

NX-3155-G9 System Specifications

Table 2: System Characteristics

Boot Device	Boot Drive
	• 2 x 512GB M.2 Boot Device

Chassis	Chassis • 2U1N LFF Chassis
CPU	Processor • Dual Intel Xeon® 5th Gen. (Emerald Rapids) • 2 x Intel Xeon® Silver 4510 [12 cores / 2.40 GHz] • 2 x Intel Xeon® Silver 4514Y [16 cores / 2.00 GHz] • 2 x Intel Xeon® Silver 4516Y+ [24 cores / 2.20 GHz] • 2 x Intel Xeon® Gold 5515+ [8 cores / 3.20 GHz], Excluded from ENERGY STAR certification. • 2 x Intel Xeon® Gold 5520+ [28 cores / 2.20 GHz] • 2 x Intel Xeon® Gold 6526Y [16 cores / 2.80 GHz] • 2 x Intel Xeon® Gold 6542Y [24 cores / 2.90 GHz] • 2 x Intel Xeon® Gold 6548Y+ [32 cores / 2.50 GHz] • Dual Intel Xeon® 4th Gen. (Sapphire Rapids) • 2 x Intel Xeon® Silver 4410T [10 cores / 2.70 GHz] • 2 x Intel Xeon® Silver 4410Y [12 cores / 2.00 GHz] • 2 x Intel Xeon® Silver 4416+ [20 cores / 2.00 GHz] • 2 x Intel Xeon® Gold 5415+ [8 cores / 2.90 GHz] • 2 x Intel Xeon® Gold 5420+ [28 cores / 2.00 GHz] • 2 x Intel Xeon® Gold 6416H [18 cores / 2.20 GHz] • 2 x Intel Xeon® Gold 6426Y [16 cores / 2.50 GHz] • 2 x Intel Xeon® Gold 6442Y [24 cores / 2.60 GHz] • 2 x Intel Xeon® Gold 6448Y [32 cores / 2.10 GHz]
GPU	Graphics Card • 0, 1 or 2 x NVIDIA A16 64GB • 0, 1 or 2 x NVIDIA A40 48GB • 0, 1 or 2 x NVIDIA H100 80GB (Qualified but not sold by Nutanix) • 0, 1 or 2 x NVIDIA H100 NVL 94GB • 0, 1 or 2 x NVIDIA L40 48GB (Qualified but not sold by Nutanix) • 0, 1 or 2 x NVIDIA L40S 48GB • 0, 1, 2, 3 or 4 x NVIDIA A2 16GB (ESXi Only) • 0, 1, 2, 3 or 4 x NVIDIA L4 24GB

Note:

- The 128 GB DDR5 5600 DIMMs require BIOS version 31.001 or later, and must not be mixed with DDR5 4800 DIMMs within the same node.
- 2304 GB capacity is available via 96GB DIMMs which are only supported with Emerald Rapids Processors
- The 96 GB memory configurations are only available with Emerald Rapids processors
- The minimum memory configuration required for Energy STAR 4.0 certification is 512 GB. If memory is less than 512 GB, the platform is not Energy Star certified, even if the CPU is certified.

32GB DIMM

- 4 x 32 GB = 128 GB
- 8 x 32 GB = 256 GB
- 12 x 32 GB = 384 GB
- 16 x 32 GB = 512 GB
- 24 x 32 GB = 768 GB
- 32 x 32 GB = 1024 GB

64GB DIMM

- 4 x 64 GB = 256 GB
- 8 x 64 GB = 512 GB
- 12 x 64 GB = 768 GB
- 16 x 64 GB = 1024 GB
- 24 x 64 GB = 1536 GB
- 32 x 64 GB = 2048 GB

96GB DIMM

- 16 x 96 GB = 1536 GB
- 24 x 96 GB = 2304 GB
- 32 x 96 GB = 3072 GB

128GB DIMM

- 4 x 128 GB = 512 GB
- 8 x 128 GB = 1024 GB
- 12 x 128 GB = 1536 GB
- 16 x 128 GB = 2048 GB
- 24 x 128 GB = 3072 GB
- 32 x 128 GB = 4096 GB

Network	Ports on Serverboard • 1x 1GbE Dedicated IPMI AIOM • 1 x 2P 10GBase-T + 2P SFP+ (Port 1 is shared IPMI) • 1 x 2P 10GBase-T (Port 1 is shared IPMI) Add on NICs in PCIe slots • 0, 1, 2 or 3 x 25GbE 2P NIC • 0, 1, 2 or 3 x 10GBaseT 2P NIC • 0, 1, 2 or 3 x 10GBaseT 4P NIC • 0, 1, 2 or 3 x 10GbE 4P NIC
Network Cables	Network Cable • 1M-SFP28 • 3M-SFP28 • 3M-SFP+ • 5M-SFP28 • 5M-SFP+
Power Cable	Power Cable • C13/14 4ft Power Cable
Power Supply	Power Supply • 2 x PWS, 2000W
Server	Server • 1 x NX-3155-G9 Server
Storage	Note: • AOS 7.3 (or higher) is required for NVMe SED support. • NVMe SED is supported with only Emerald Rapids processors.
All NVMe	2, 4 or 6 x NVMe • 1.92TB • 3.84TB • 7.68TB • 15.36TB

NVMe+HDD	2 x NVMe • 3.84TB • 7.68TB • 15.36TB 4 x HDD • 8TB • 12TB • 18TB
All NVMe SED	2, 4 or 6 x NVMe • 3.84TB • 7.68TB • 15.36TB
NVMe SED+HDD SED	2 x NVMe • 3.84TB • 7.68TB • 15.36TB 4 x HDD • 8TB • 12TB
Storage Controller	Storage Controller • 1 x Broadcom 3816 Host Bus Controller (HBA), required for SATA/SAS based drive support
TPM	TPM • 0 or 1 x Unprovisioned Trusted Platform Module
Transceiver	Transceiver • SR SFP+ Transceiver
VGA	1 x VGA connector per node (15-pin female)
Chassis fans	4 x 80 mm heavy-duty fans with PWM fan speed controls.

Table 3: Block, power and electrical

Block	• Height : 88.1 mm • Depth : 799 mm • Width : 445.6 mm Maximum configuration (without rails and accessory kits) • Weight : 25.5 kg
Package	Package maximum weight with rails and accessory kits • Weight : 36.36 kg
Shock	5G, half-sine, one shock on each side • Operating : 10 ms 20G, square wave, one shock on each side • Non-Operating : 10 m
Thermal Dissipation	• Typical : 4665 BTU/hr • Maximum : 6663 BTU/hr
Vibration (Random)	5 to 500 Hz., approx. 15 min./axis • Operating : 0.21 Grms 5 to 200 Hz., approx. 30 min./axis • Non-Operating : 0.98 Grms
Vibration (Sinusoidal)	5 to 200 Hz., approx. 15 min./axis • Operating : 0.25 G 5 to 200 Hz., approx. 15 min./axis • Non-Operating : 0.5 G
Power consumption (calculated)	Note: • For the power consumption calculations, Max NIC and disk configurations are considered. Max Config • Typical: 1368 VA • Maximum: 1954 VA

Operating environment	Operating temperature : 10-25C
	Non-Operating temperature : -20-70C
	Operating relative humidity : 20-85%
	Non-operating relative humidity : 5-95%

Certifications

- BIS
- BSMI
- CE
- CSA
- CSAus
- EAC
- Energy Star
- FCC
- ICES
- KCC
- RCM
- Reach
- RoHS
- S-MARK
- SABS
- SII
- UKCA
- UKRSEPRO
- UL
- VCCI-A
- WEEE
- cUL

NX-3155-G9 GPU Specifications

The following table provides minimum firmware and software versions for GPUs supported by NX-3155-G9:

Table 4: Minimum Firmware and Software Versions When Using a GPU Card

Firmware or Software	NVIDIA A2	NVIDIA A16	NVIDIA L40S/L4	NVIDIA H100/H100 NVL
BIOS	EH10.201	EH10.201	EH10.201	EH10.201
BMC	01.00.07	01.00.07	01.00.07	01.00.07
Hypervisor	ESXi only	ESXi and AHV	ESXi and AHV	ESXi and AHV
Foundation	5.4.2 with Foundation-platforms-2.13.3	5.4.2 with Foundation-platforms-2.13.3	5.4.2 with Foundation-platforms-2.13.3	5.4.2 with Foundation-platforms-2.13.3
AOS	6.5.3.5	6.5.3.5	6.5.3.5	6.5.3.5
NCC	4.6.5.2	4.6.5.2	4.6.5.2	4.6.5.2
NVIDIA vGPU driver	GRID 15.2	GRID 15.2	GRID 16.1	GRID 15.2

COMPONENT SPECIFICATIONS

Controls and LEDs for Single-node Platforms



Figure 3: Front Panel Controls and Indicators

Table 5: Front Panel Controls and Indicators

LED or Button	Function
Power button	Power button applies or removes primary power from the power supply to the system (System standby power maintains)
Power LED	Solid green: Indicates power is supplied to the system power supply units. This LED is illuminated when the system is operating normally.
UID/Reset button	Turns information LED <i>i</i> On/Off (Single press) Reboots the system (Press and hold for 6 seconds)
Drive LED	Not used
NIC LED (AIOM Port 1)	Indicates network activity on AIOM 10Gbase-T port 1
NIC LED (AIOM Port 2)	Indicates network activity on AIOM 10Gbase-T port 2
Multiple function <i>i</i> LED	Solid blue: UID activated locally Blinking blue: UID activated remotely Solid red: Overheat condition Blinking red at 0.25Hz: Power supply failure Blinking red at 1Hz: Fan failure

Note: NIC1 and NIC2 LED in the table above are specific to the AIOM only. If it is a dual port AIOM, then NIC1 LED is controlled by one port and NIC2 LED is controlled by the other port.

Table 6: Drive LEDs

Top LED (Activity)	Blinking blue or green: I/O activity Off: Idle
Bottom LED (Status)	Red solid: Failed drive On five seconds after boot: Power on

Back Panel LED Indication




Figure 4: Back Panel LEDs

Table 7: Back Panel Controls and Indicators

Name	LED status	Function
IPMI, Link LED (Left)	Solid green	100 Mbps
	Solid amber	1 Gbps
IPMI, Activity LED (Right)	Blinking amber	Active
AIO, Link LED (Left)	Off	No link
	Solid green	Linked at 10 Gbps
	Solid amber	Linked at 1 Gbps
AIO, Activity LED (Right)	Off	No activity
	Blinking green	Link up (traffic flowing)
Locator LED (UID)	Blinking blue	Node identified

Table 8: Power Supply LED Indicators

LED status	Power supply condition
Off	No AC power to all power supplies

LED status	Power supply condition
Solid amber	Power supply critical events that cause a shutdown: Failure, Over Current Protection, Over Voltage Protection, Fan Fail, Over Temperature Protection, Under Voltage Protection.
Blinking amber (1 Hz)	Power supply warning events. Power supply continues to operate. High temperature, over voltage, under voltage and other conditions.
Blinking green (1 Hz) minute	When AC is present only: 12VSB on (PS off) or PS in sleep state
Solid green	Output on and OK
Solid amber	AC cord unplugged
Blinking green (2 Hz)	Power supply firmware updating mode

For add-on NIC LED indications, see [Network Card LED Description](#) on page 14.

Network Card LED Description

Different NIC manufacturers use different LED colors and blink states. Not all NICs are supported for every Nutanix platform. See the system specifications for your platform to verify which NICs are supported.

Table 9: SuperMicro NICs

NIC	Link (LNK) LED	Activity (ACT) LED
Dual/Quad Port 25GbE	Green: 25 GbE Amber: Less than 25 GbE Off: No link	Blinking: Activity Off: No activity
Dual Port 100GbE	Green: 100 GbE Amber: Less than 100 GbE	Blinking: Activity Off: No activity

Table 10: Intel NICs

NIC	Link (LNK) LED	Activity (ACT) LED
Quad port 10G BaseT	Green: 10 Gbps Yellow: 5/2.5/1Gbps Off: 100 Mbps	Blinking green: Transmitting or receiving data Off: No link
Quad port 10G SFP+	Green: 10 Gbps Yellow: 1 Gbps	Blinking: Activity Off: No activity
Dual Port 25G	Green: 10 Gbps Amber: 10 Gbps	Green: SFP LAN port active

Table 11: Broadcom NICs

NIC	Link (LNK) LED	Activity (ACT) LED
Dual Port 25G	Green: Linked at 25 Gb/s Yellow: Linked at 10 Gb/s or 1 Gb/s Off: No link	Blinking green: Traffic flowing activity Off: No activity
Dual Port 10G	Green: Linked at 10 Gb/s Amber: Linked at 1 Gb/s Off: No link	Blinking green: Link up (traffic flowing) Off: No activity

Table 12: Mellanox NICs

NIC	Bi-color LED (Yellow/Green)	Single color LED (Green)	Indicates
Dual port CX-6 25G	1 Hz blinking yellow	Off	Beacon command for locating the adapter card
Dual port CX-6 100G	4 Hz blinking yellow	On	Error with the link
	Blinking green	Blinking	Physical activity
	Solid green	On	Link up

Power Supply Unit Redundancy and Node Configuration

Note: Carefully plan your AC power source needs, especially in cases where the cluster consists of mixed models. Nutanix recommends that you use a 208 V - 240 V AC power source to ensure Power Supply Unit (PSU) redundancy.

Table 13: PSU Redundancy and Node Configuration

Nutanix model	Number of nodes	Redundancy at 110 V	Redundancy at 208-240 V
NX-1065-G9	1, 2, or 3	No	Yes
NX-3035-G9	1 or 2	No	Yes
NX-3060-G9	1, 2, 3, or 4	No	Yes
NX-1150S-G9	1	Yes	Yes
NX-1175S-G9	1	Yes	Yes
NX-3155-G9	1	No	Yes
NX-8150-G9	1	No	Yes

Nutanix model	Number of nodes	Redundancy at 110 V	Redundancy at 208-240 V
NX-8155-G9	1	No	Yes
NX-8155A-G9	1	No	Yes
NX-8170-G9	1	No	Yes
NX-9151-G9	1	Not supported	Yes (2+1 PSU redundancy) Two PSUs must remain functional at all times.

Caution:

For all G9 platforms except NX-1175S-G9 and NX-1150S-G9: When the input power source is 110 V, a single PSU failure will cause all nodes in a block to power off.

For NX-1175S-G9 and NX-1150S-G9: The block can tolerate a single PSU failure when connected to a 110 V input power source.

MEMORY CONFIGURATIONS

Supported Memory Configurations

DIMM installation information for all Nutanix G9 Hyper platforms.

DIMM Restrictions

DIMM capacity

Each G9 node must contain only DIMMs of the same capacity. For example, you cannot mix 32 GB DIMMs and 64 GB DIMMs in the same node.

DIMM speed

G9 platforms that use Intel Sapphire Rapids processors support 4800 MT/s DIMMs. The speed of the CPU / DIMM interface is 4000 MT/s or 4800 MT/s based on the CPU SKU used.

G9 platforms that use Intel Emerald Rapids processors ship with 5600 MT/s DIMMs. The speed of the CPU / DIMM interface depends on the CPU class and on whether you have installed one DIMM per memory channel (1DPC) or two DIMMs per memory channel (2DPC).

- Platinum-8xxx: max 5600 MT/s at 1DPC; max 4400 MT/s at 2DPC.
- Gold-6xxx: max 5200 MT/s at 1DPC; max 4400 MT/s at 2DPC.
- Gold-5xxx: max 4800 MT/s at 1DPC; max 4400 MT/s at 2DPC.
- Silver-4xxx: max 4400 MT/s at both 1DPC and 2DPC.
- Bronze-3xxx (single socket boards only): max 4400 MT/s at both 1DPC and 2DPC.

128 GB 5600 MT/s DIMMs require BIOS Ex30.001 or later. 128 GB 5600 MT/s DIMMs cannot mix in the same node with 4800 MT/s DIMMs.

If you install a 5600 MT/s DIMM in a G9 platform that uses Intel Sapphire Rapids processors, it runs at a max of 4800 MT/s.

DIMM manufacturer

Mixing of DIMMs from different manufacturers in a node is not allowed. You must use DIMMs from the same manufacturer in a node.

Memory Installation Order for Single-Node Hyper G9 Platforms

A *memory channel* is a group of DIMM slots.

For G9 single-node hyper platforms, each CPU is associated with eight memory channels. Each memory channel contains two DIMM slots, one blue slot and one black slot, for a total of 32 DIMM slots.




Figure 5: DIMM Slots for a Single-Node Hyper G9 Motherboard

Table 14: DIMM Installation Order for Single-Node Hyper G9 Platforms

Number of DIMMs	Slots to Use	Supported Capacities
4	CPU1: A1, G1 (blue slots) CPU2: A1, G1 (blue slots)	32 GB, 64 GB, 128 GB

Number of DIMMs	Slots to Use	Supported Capacities
8	CPU1: A1, C1, E1, G1 (blue slots) CPU2: A1, C1, E1, G1 (blue slots)	32 GB, 64 GB, 128 GB
12	CPU1: A1, C1, D1, E1, F1, G1 (blue slots) CPU2: A1, C1, D1, E1, F1, G1 (blue slots)	32 GB, 64 GB, 128 GB
16	CPU1: A1, B1, C1, D1, E1, F1, G1, H1 (blue slots) CPU2: A1, B1, C1, D1, E1, F1, G1, H1 (blue slots)	32 GB, 64 GB, 128 GB
24	CPU1: A1, B1, C1, D1, E1, F1, G1, H1 (blue slots) CPU1: A2, C2, E2, G2 (black slots) OR B2, D2, F2, H2 (black slots) CPU2: A1, B1, C1, D1, E1, F1, G1, H1 (blue slots) CPU2: A2, C2, E2, G2 (black slots) OR B2, D2, F2, H2 (black slots)	32 GB, 64 GB, 128 GB
32	Fill all slots.	32 GB, 64 GB, 128 GB

NUTANIX HARDWARE NAMING CONVENTION

Every Nutanix block has a unique name based on the standard Nutanix naming convention.

The Nutanix hardware model name uses the format *prefix-body-suffix*.

For all Nutanix platforms, the prefix is NX to indicate that the platform is sold directly by Nutanix and support calls are handled by Nutanix.

body uses the format *ABCD* | *Y*. The following table describes the body values.




Figure 6: Nutanix Hardware Naming Convention

Table 15: Body

Body	Description
A	Indicates the product series and is one of the following values. 1 – Entry-level and ROBO 3 – Balanced compute and storage 8 – High performance 9 – Accelerated system

Body	Description
<i>B</i>	Indicates the number of nodes. - For single-node platforms, <i>B</i> is always 1. - For multi-node platforms, <i>B</i> can be 1, 2, 3, or 4. Note: For multi-node platforms, the documentation always uses a generic zero for <i>B</i>.
<i>C</i>	Indicates the chassis form factor and is one of the following values. 1 – 1U1N (one rack unit high with one node) 3 – 2U2N (two rack units high with two nodes) 5 – 2U1N (two rack units high with one node) 6 – 2U4N (two rack units high with four nodes) 7 – 1U1N (one rack unit high with one node)
<i>D</i>	Indicates the drive form factor and is one of the following values. 0 – 2.5 inch drives 1 – E1.S drives 3 – E3.S drives 5 – 3.5 inch drives
<i>Y</i>	Indicates platform types, and takes one of the following values: - S – Single socket - G – GPU (Not used in G9 since GPU is available on multiple models) - A – AMD

Table 16: Suffix

Suffix	Description
G5	The platform uses the Intel Broadwell CPU.
G6	The platform uses the Intel Skylake CPU.
G7	The platform uses the Intel Cascade Lake CPU.
G8 or N-G8	The platform uses the Intel Ice Lake CPU.
G9	The platform uses the Intel Sapphire Rapids, Emerald Rapids or AMD Genoa CPU.

COPYRIGHT

Copyright 2025 Nutanix, Inc.

Nutanix, Inc.

1740 Technology Drive, Suite 150

San Jose, CA 95110

All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. Nutanix and the Nutanix logo are registered trademarks of Nutanix, Inc. in the United States and/or other jurisdictions. All other brand and product names mentioned herein are for identification purposes only and may be trademarks of their respective holders.

AOS - AOS Storage

[PDF generated October 28 2025. For all recent updates please see the Nutanix Bible releases notes located at https://nutanixbible.com/release_notes.html. Disclaimer: Downloaded PDFs may not always contain the latest information.]

Nutanix AOS storage is the scale-out storage technology that appears to the hypervisor like any centralized storage array, however all of the I/Os are handled locally to provide the highest performance. More detail on how these nodes form a distributed system can be found in the next section.

Data Structure

Nutanix AOS storage is composed of the following high-level struct:

Storage Pool

- Key Role: Group of physical devices
- Description: A storage pool is a group of physical storage devices including PCIe SSD, SSD, and HDD devices for the cluster. The storage pool can span multiple Nutanix nodes and is expanded as the cluster scales. In most configurations, only a single storage pool is leveraged.

Container

- Key Role: Group of VMs/files
- Description: A container is a logical segmentation of the Storage Pool and contains a group of VM or files (vDisks). Some configuration options (e.g., RF) are configured at the container level, however are applied at the individual VM/file level. Containers typically have a 1 to 1 mapping with a datastore (in the case of NFS/SMB).


vDisk

- Key Role: vDisk
- Description: A vDisk is any file over 512KB on AOS including .vmdks and VM hard disks. vDisks are logically composed of vBlocks which make up the 'block map.'

Maximum AOS vDisk Size

No artificial limits are imposed on the vdisk size on the AOS/stargate side. As of 4.6, the vdisk size is stored as a 64 bit signed integer that stores the size in bytes. This means the theoretical maximum vDisk size can be $2^{63}-1$ or 9E18 (9 Exabytes). Any limits below this value would be due to limitations on the client side, such as the maximum vmdk size on ESXi.

The following figure shows how these map between AOS and the hypervisor:



vBlock

- Key Role: 1MB chunk of vDisk address space
- Description: A vBlock is a 1MB chunk of virtual address space composing a vDisk. For example, a vDisk of 100MB will have 100 x 1MB vBlocks, vBlock 0 would be for 0-1MB, vBlock 1 would be from 1-2MB, and so forth. These vBlocks map to extents which are stored as files on disk as extent groups.


Extent

- Key Role: Logically contiguous data
- Description: An extent is a 1MB piece of logically contiguous data which consists of n number of contiguous blocks (varies depending on guest OS block size). Extents are written/read/modified on a sub-extent basis (aka slice) for granularity and efficiency. An extent's slice may be trimmed when moving into the cache depending on the amount of data being read/ cached.


Extent Group

- Key Role: Physically contiguous stored data
- Description: An extent group is a 1MB or 4MB piece of physically contiguous stored data. This data is stored as a file on the storage device owned by the CVM. Extents are dynamically distributed among extent groups to provide data striping across nodes/disks to improve performance. NOTE: as of 4.0, extent groups can now be either 1MB or 4MB depending on deduplication.

The following figure shows how these structs relate between the various file systems:



Here is another graphical representation of how these units are related:



Graphical Filesystem Breakdown


I/O Path and Cache

For a visual explanation, you can watch the following video: [LINK](#)

The typical hyperconverged storage I/O path can be characterized into the following core layers:

1. Guest OS (UVM) to virtual disk(s)
 - This remains unchanged with Nutanix. Depending on the hypervisor the guest OS will use a device driver to talk to a virtual disk device. Depending on the hypervisor this could be virtio-scsi (AHV), pv-scsi (ESXi), etc. The virtual disks will also vary based upon the hypervisor (e.g. vmdk, vhd, etc.)
2. Hypervisor to AOS (via CVM)
 - Communication between the hypervisor and Nutanix occurs via standard storage protocols (e.g. iSCSI, NFS, SMBv3) over the local interface of the CVM and hypervisor. At this point all communication has been local to the host (there are scenarios where I/O will be remote (e.g. local CVM down, etc.).
3. Nutanix I/O path
 - This is all transparent to the hypervisor and UVMs and it native to the Nutanix platform.

The following image shows a high-level overview of these layers:




High-level I/O Path - Traditional

Communication and I/O


Within the CVM the Stargate process is responsible for handling all storage I/O requests and interaction with other CVMs / physical devices. Storage device controllers are passed through directly to the CVM so all storage I/O bypasses the hypervisor.

The following image shows a high-level overview of the traditional I/O path:




High-level I/O Path

Nutanix BlockStore (shipped in AOS 5.18) is an AOS capability which creates an extensible filesystem and block management layer all handled in user space. This eliminates the filesystem from the devices and removes the invoking of any filesystem kernel driver. The introduction of newer storage media (e.g. NVMe), devices now come with user space libraries to handle device I/O directly (e.g. SPDK) eliminating the need to make any system calls (context switches). With the combination of BlockStore + SPDK all Stargate device interaction has moved into user space eliminating any context switching or kernel driver invocation.



Stargate - Device I/O Path

The following image shows a high-level overview of the updated I/O path with BlockStore + SPDK:




High-level I/O Path - BlockStore

To perform data replication the CVMs communicate over the network. With the default stack this will invoke kernel level drivers to do so.

However, with RDMA these NICs are passed through to the CVM, bypassing the hypervisor and reducing interrupts. Also, within the CVM all network traffic using RDMA only uses a kernel level driver for the control path, then all actual data I/O is done in user-space without any context switches.

The following image shows a high-level overview of the I/O path with RDMA:



High-level I/O Path - RDMA


To summarize, the following enhancements optimize with the following:

1. PCI passthrough bypasses the hypervisor for device I/O
2. SPDK + Blockstore eliminates kernel storage driver interactions and moves them to user-space
3. RDMA bypasses the hypervisor and all data transfer is done in CVM user-space

Stargate I/O Logic

Within the CVM the Stargate process is responsible for handling all I/O coming from user VMs (UVMs) and persistence (RF, etc.). When a write request comes to Stargate, there is a write characterizer which will determine if the write gets persisted to the OpLog for bursty random writes, or to Extent Store for sustained random and sequential writes. Read requests are satisfied from Oplog or Extent Store depending on where the data is residing when it is requested.

The Nutanix I/O path is composed of the following high-level components:



AOS I/O Path

*As of AOS 5.10, the Autonomous Extent Store (AES) can be used to handle sustained random workloads when requisite conditions are met.

As:

1. In all-flash node configurations (All NVMe SSD, All SATA/SAS SSD, NVMe+SATA/SAS SSD) the Extent Store will only consist of SSD devices and no tier ILM will occur as only a single flash tier exists.
2. In cases where Optane is used (e.g. Intel Optane + NVMe/SATA SSD) the highest performance media will be Tier 0 and the lower performance media will be Tier 1.
3. For hybrid scenarios with flash and HDD, the flash would be Tier 0 with HDD being Tier 1.
4. OpLog is always on SSDs in Hybrid Clusters and on both Optane and NVMe SSDs in Optane+NVMe clusters.
5. Data is moved between tiers by Intelligent Lifecycle Management (ILM) based on access patterns.

OpLog

- Key Role: Persistent write buffer
- Description: The OpLog is similar to a filesystem journal and is built as a staging area to handle bursts of random writes, coalesce them, and then sequentially drain the data to the extent store. Upon a write, the OpLog is synchronously replicated to another n number of CVM's OpLog before the write is acknowledged for data availability purposes. All CVM OpLogs partake in the replication and are dynamically chosen based upon load. The OpLog is stored on the SSD tier on the CVM to provide extremely fast write I/O performance, especially for random I/O workloads. All SSD devices participate and handle a portion of OpLog storage. For sequential workloads, the OpLog is bypassed and the writes go directly to the extent store. If data is currently sitting in the OpLog and has not been drained, all read requests will be directly fulfilled from the OpLog until they have been drained, where they would then be served by the extent store/unified cache. For containers where fingerprinting (aka Dedupe) has been enabled, all write I/Os will be fingerprinted using a hashing scheme allowing them to be deduplicated based upon fingerprint in the unified cache.

Dynamic OpLog Sizing

The OpLog is a shared resource and, however allocation is done on a per-vDisk basis to ensure each vDisk has an equal opportunity to leverage. Prior to AOS 5.19, Oplog size was capped to 6GB per vdisk. Starting with AOS 5.19, OpLog for individual vdisks can keep growing beyond 6GB if required based on IO patterns until a cap is reached for OpLog index memory used per node. This design decision allows flexibility such that if there are VMs with fewer vdisks that are more active from an I/O perspective, they can keep growing their OpLog as required at the expense of other vdisks that are not as active.

Extent Store

- Key Role: Persistent data storage
- Description: The Extent Store is the persistent bulk storage of AOS and spans all device tiers (Optane SSD, PCIe SSD, SATA SSD, HDD) and is extensible to facilitate additional devices/tiers. Data entering the extent store is either being A) drained from the OpLog or B) is sequential/sustained in nature and has bypassed the OpLog directly. Nutanix ILM will determine tier

placement dynamically based upon I/O patterns, number of accesses of data and weight given to individual tiers and will move data between tiers.

Sequential Write Characterization

Write IO is deemed as sequential when there is more than 1.5MB of outstanding write IO to a vDisk (as of 4.6). IOs meeting this will bypass the OpLog and go directly to the Extent Store since they are already large chunks of aligned data and won't benefit from coalescing.

This is controlled by the following Gflag: `vdisk_distributed_oplog_skip_min_outstanding_write_bytes`.

All other IOs, including those which can be large (e.g. >64K) will still be handled by the OpLog.


Autonomous Extent Store (AES)

- Key Role: Persistent data storage
- Description: The Autonomous Extent Store (AES) is a new method for writing / storing data in the Extent Store introduced in AOS 5.10. It leverages a mix of primarily local + global metadata (more detail in the 'Scalable Metadata' section following) allowing for much more efficient sustained performance due to metadata locality. For sustained random write workloads, these will bypass the OpLog and be written directly to the Extent Store using AES. For bursty random workloads these will take the typical OpLog I/O path then drain to the Extent Store using AES where possible. As of AOS 5.20 (LTS), AES is enabled by default for new containers on All Flash Clusters and as of AOS 6.1 (STS) if requirements are met, AES is enabled on new containers created on Hybrid(SSD+HDD) Clusters.

Unified Cache

- Key Role: Dynamic read cache
- Description: The Unified Cache is a read cache which is used for data, metadata and deduplication and stored in the CVM's memory. Upon a read request of data not in the cache (or based upon a particular fingerprint), the data will be read from the extent store and will also be placed into the single-touch pool of the Unified Cache which completely sits in memory, where it will use LRU (least recently used) until it is evicted from the cache. Any subsequent read request will "move" (no data is actually moved, just cache metadata) the data into the multi-touch pool. Any read request for data in the multi-touch pool will cause the data to go to the peak of the multi-touch pool where it will be given a new LRU counter. Cache size can be calculated using the following formula: $((\text{CVM Memory} - 12 \text{ GB}) * 0.45)$. For example a 32GB CVM would have the following cache size: $((32 - 12) * 0.45) == 9 \text{ GB}$.

The following figure shows a high-level overview of the Unified Cache:



AOS Unified Cache

Cache Granularity and Logic


Data is brought into the cache at a 4K granularity and all caching is done real-time (e.g. no delay or batch process data to pull data into the cache).

Each CVM has its own local cache that it manages for the vDisk(s) it is hosting (e.g. VM(s) running on the same node). When a vDisk is cloned (e.g. new clones, snapshots, etc.) each new vDisk has its own block map and the original vDisk is marked as immutable. This allows us to ensure that each CVM can have its own cached copy of the base vDisk with cache coherency.


In the event of an overwrite, that will be re-directed to a new extent in the VM's own block map. This ensures that there will not be any cache corruption.

Single vDisk Sharding

AOS was designed and architected to deliver performance for applications at scale. Inside Stargate, I/O is processed by threads for every vdisk created by something called vdisk controller. Every vdisk gets its own vdisk controller inside Stargate responsible for I/O for that vdisk. The expectation was that workloads and applications would have multiple vdisks each having its own vdisk controller thread capable of driving high performance the system is capable of delivering.



This architecture worked well except in cases of traditional applications and workloads that had VMs with single large vdisk. These VMs were not able to leverage the capabilities of AOS to its fullest. As of AOS 6.1, enhancement was done to vdisk controller such that requests to a single vdisk are now distributed to multiple vdisk controllers by creating shards of the controller each having its own thread. I/O distribution to multiple controllers is done by a primary controller so for external interaction this still looks like a single vdisk. This results in effectively sharding the single vdisk making it multi-threaded. This enhancement along with other technologies talked above like Blockstore, AES allows AOS to deliver consistent high performance at scale even for traditional applications that use a single vdisk.



Scalable Metadata

For a visual explanation, you can watch the following YouTube video: [Tech TopX by Nutanix University: Scalable Metadata](#)

Metadata is at the core of any intelligent system and is even more critical for any filesystem or storage array. For those unsure about the term 'metadata'; essentially metadata is 'data about data'. In terms of AOS, there are a few key principles that are critical for its success:

- Must be right 100% of the time (known as "strictly consistent")
- Must be ACID compliant
- Must have unlimited scalability


- Must not have any bottlenecks at any scale (must be linearly scalable)

As of AOS 5.10 metadata is broken into two areas: global vs. local metadata (prior all metadata was global). The motivation for this is to optimize for "metadata locality" and limit the network traffic on the system for metadata lookups.

The basis for this change is that not all data needs to be global. For example, every CVM doesn't need to know which physical disk a particular extent sits on, they just need to know which node holds that data, and only that node needs to know which disk has the data.

By doing this we can limit the amount of metadata stored by the system (eliminate metadata RF for local only data), and optimize for "metadata locality."

The following image shows the differentiation between global vs. local metadata:



Global vs. Local Metadata

Local Metadata

- Description:
 - Local metadata store per CVM containing information only needed by the local node. This is leveraged by the Autonomous Extent Store (AES) introduced in 5.10.
- Storage Mechanism:
 - AES DB (based on Rocksdb)
- Types of data stored:
 - Physical extent / extent group placement (e.g. egroup to disk mappings), etc.

Global Metadata


- Description:
 - Metadata that is globally available to any CVM and shared across CVMs in the cluster. All metadata prior to 5.10.
- Storage Mechanism:
 - Medusa Store (based on Cassandra)
- Types of data stored:
 - vDisk block maps, extent to node mappings, time series stats, configurations, etc.

The section below covers how global metadata is managed:

As mentioned in the architecture section above, AOS utilizes a "ring-like" structure as a key-value store which stores essential global metadata as well as other platform data (e.g., stats, etc.). In order to ensure global metadata availability and redundancy a replication factor (RF) is utilized among an odd amount of nodes (e.g., 3, 5, etc.). Upon a global metadata write or update, the row

is written to a node in the ring that owns that key and then replicated to n number of peers (where n is dependent on cluster size). A majority of nodes must agree before anything is committed, which is enforced using the Paxos algorithm. This ensures strict consistency for all data and global metadata stored as part of the platform.


The following figure shows an example of a global metadata insert/update for a 4 node cluster:



Cassandra Ring Structure

Performance at scale is also another important struct for AOS global metadata. Contrary to traditional dual-controller or “leader/worker” models, each Nutanix node is responsible for a subset of the overall platform’s metadata. This eliminates the traditional bottlenecks by allowing global metadata to be served and manipulated by all nodes in the cluster. A consistent hashing scheme is utilized for key partitioning to minimize the redistribution of keys during cluster size modifications (also known as “add/remove node”). When the cluster scales (e.g., from 4 to 8 nodes), the nodes are inserted throughout the ring between nodes for “block awareness” and reliability.

The following figure shows an example of the global metadata “ring” and how it scales:



Cassandra Scale Out

Data Protection

For a visual explanation, you can watch the following video: [LINK](#)


The Nutanix platform currently uses a resiliency factor, also known as a replication factor (RF), and checksum to ensure data redundancy and availability in the case of a node or disk failure or corruption. As explained above, the OpLog acts as a staging area to absorb incoming writes onto a low-latency SSD tier. Upon being written to the local OpLog, the data is synchronously replicated to another one or two Nutanix CVM’s OpLog (dependent on RF) before being acknowledged (Ack) as a successful write to the host. This ensures that the data exists in at least two or three independent locations and is fault tolerant. NOTE: For RF3, a minimum of 5 nodes is required since metadata will be RF5.

OpLog peers are chosen for every episode (1GB of vDisk data) and all nodes actively participate. Multiple factors play into which peers are chosen (e.g. response time, business, capacity utilization, etc). This eliminates any fragmentation and ensures every CVM/OpLog can be used concurrently.

Data RF is configured via Prism and is done at the container level. All nodes participate in OpLog replication to eliminate any "hot nodes", ensuring linear performance at scale. While the data is being written, a checksum is computed and stored as part of its metadata. Data is then asynchronously drained to the extent store where the RF is implicitly maintained. In the case of a node or disk failure, the data is then re-replicated among all nodes in the cluster to maintain the RF. Any time the data is read, the checksum is computed to ensure the data is valid. In the event where the checksum and data don't match, the replica of the data will be read and will replace the non-valid copy.

Data is also consistently monitored to ensure integrity even when active I/O isn't occurring. Stargate's scrubber operation will consistently scan through extent groups and perform checksum validation when disks aren't heavily utilized. This protects against things like bit rot or corrupted sectors.

The following figure shows an example of what this logically looks like:



AOS Data Protection

Availability Domains

For a visual explanation, you can watch the following video: [LINK](#)

Availability Domains (aka node/block/rack awareness) is a key struct for distributed systems to abide by for determining component and data placement. Nutanix refers to a "block" as the chassis which contains either one, two, or four server "nodes" and a "rack" as a physical unit containing one or more "block". NOTE: A minimum of 3 blocks must be utilized for block awareness to be activated, otherwise node awareness will be used.


Nutanix currently supports the following levels or awareness:

- Disk (always)
- Node (always)
- Block (as of AOS 4.5)
- Rack (as of AOS 5.9)

It is recommended to utilize uniformly populated blocks / racks to ensure the awareness is enabled and no imbalance is possible. Common scenarios and the awareness level utilized can be found at the bottom of this section. The 3-block requirement is due to ensure quorum. For example, a 3450 would be a block which holds 4 nodes. The reason for distributing roles or data across blocks is to ensure if a block fails or needs maintenance the system can continue to run without interruption. NOTE: Within a block, the redundant PSU and fans are the only shared components.

NOTE: Rack awareness requires the administrator to define "racks" in which the blocks are placed.

The following shows how this is configured in Prism:



Rack Configuration


Awareness can be broken into a few key focus areas:

- Data (The VM data)
- Metadata (Cassandra)
- Configuration Data (Zookeeper)

Data


With AOS, data replicas will be written to other [blocks/racks] in the cluster to ensure that in the case of a [block/rack] failure or planned downtime, the data remains available. This is true for both RF2 and RF3 scenarios, as well as in the case of a [block/rack] failure. An easy comparison would be “node awareness”, where a replica would need to be replicated to another node which will provide protection in the case of a node failure. Block and rack awareness further enhances this by providing data availability assurances in the case of [block/rack] outages.

The following figure shows how the replica placement would work in a 3-block deployment:



Block/Rack Aware Replica Placement

In the case of a [block/rack] failure, [block/rack] awareness will be maintained (if possible) and the data will be replicated to other [blocks/racks] within the cluster:



Block/Rack Failure Replica Placement

Rack/Block Awareness vs. Metro clustering

A common question is can you span a cluster across two locations (rooms, buildings, etc.) and use block / rack awareness to provide resiliency around a location failure.

While theoretically possible this is not the recommended approach. Let's first think about what we're trying to achieve with this:

1. Low RPO
2. Low RTO (HA event instead of a DR event)

If we take the first case where we're trying to achieve an RPO ~0, it is preferred to leverage synchronous or near-synchronous replication. This will provide the same RPOs with less risk.

To minimize the RTO one can leverage a metro-cluster on top of synchronous replication and handle any failures as HA events instead of doing DR recoveries.

In summary it is preferred to leverage synchronous replication / metro clustering for the following reasons:

- The same end result can be achieved with sync rep / metro clustering, avoiding any risks and keeping isolated fault domains
- If network connectivity goes down between the two locations in a non-supported "stretched" deployment, one side will go down as quorum must be maintained (e.g. majority side will stay up). In the metro cluster scenario, both sides can continue to operate independently.
- Availability domain placement of data is best effort in skewed scenarios
- Additional Latency / reduced network bandwidth between both sites can impact performance in the "stretched" deployment

Awareness Conditions and Tolerance

Below we breakdown some common scenarios and the level of tolerance:

Desired Awareness Type	FT Level	EC Enabled?	Min. Units	Simultaneous failure tolerance
Node	1	No	3 Nodes	1 Node
Node	1	Yes	4 Nodes	1 Node
Node	2	No	5 Nodes	2 Node
Node	2	Yes	6 Nodes	2 Nodes
Block	1	No	3 Blocks	1 Block
Block	1	Yes	4 Blocks	1 Block
Block	2	No	5 Blocks	2 Blocks
Block	2	Yes	6 Blocks	2 Blocks
Rack	1	No	3 Racks	1 Rack

Rack	1	Yes	4 Racks	1 Rack
Rack	2	No	5 Racks	2 Racks
Rack	2	Yes	6 Racks	2 Racks

As of AOS base software version 4.5 and later block awareness is best effort and doesn't have strict requirements for enabling. This was done to ensure clusters with skewed storage resources (e.g. storage heavy nodes) don't disable the feature. With that stated, it is however still a best practice to have uniform blocks to minimize any storage skew.

Prior to 4.5 the following conditions must be met for block awareness:

- If SSD or HDD tier variance between blocks is > max variance: **NODE** awareness
- If SSD and HDD tier variance between blocks is < max variance: **BLOCK + NODE** awareness


Max tier variance is calculated as: $100 / (RF+1)$

- E.g., 33% for RF2 or 25% for RF3

Metadata

As mentioned in the Scalable Metadata section above, Nutanix leverages a heavily modified Cassandra platform to store metadata and other essential information. Cassandra leverages a ring-like structure and replicates to n number of peers within the ring to ensure data consistency and availability.


The following figure shows an example of the Cassandra's ring for a 12-node cluster:



12 Node Cassandra Ring

Cassandra peer replication iterates through nodes in a clockwise manner throughout the ring. With [block/rack] awareness, the peers are distributed among the [blocks/racks] to ensure no two peers are on the same [block/rack].


The following figure shows an example node layout translating the ring above into the [block/rack] based layout:



Cassandra Node Block/Rack Aware Placement

With this [block/rack]-aware nature, in the event of a [block/rack] failure there will still be at least two copies of the data (with Metadata RF3 – In larger clusters RF5 can be leveraged).

The following figure shows an example of all of the nodes replication topology to form the ring (yes – it's a little busy):



Full Cassandra Node Block/Rack Aware Placement

Metadata Awareness Conditions

Below we breakdown some common scenarios and what level of awareness will be utilized:


- FT1 (Data RF2 / Metadata RF3) will be block aware if:
 - ≥ 3 blocks
 - Let X be the number of nodes in the block with max nodes. Then, the remaining blocks should have at least 2X nodes.
 - Example: 4 blocks with 2,3,4,2 nodes per block respectively.
 - The max node block has 4 nodes which means the other 3 blocks should have 2×4 (8) nodes. In this case it **WOULD NOT** be block aware as the remaining blocks only have 7 nodes.
 - Example: 4 blocks with 3,3,4,3 nodes per block respectively.
 - The max node block has 4 nodes which means the other 3 blocks should have $2 \times 4 = 8$ nodes. In this case it **WOULD** be block aware as the remaining blocks have 9 nodes which is above our minimum.
- FT2 (Data RF3 / Metadata RF5) will be block aware if:
 - ≥ 5 blocks

- Let X be the number of nodes in the block with max nodes. Then, the remaining blocks should have at least 4X nodes.
- Example: 6 blocks with 2,3,4,2,3,3 nodes per block respectively.
 - The max node block has 4 nodes which means the other 3 blocks should have $4 \times 4 = 16$ nodes. In this case it **WOULD NOT** be block aware as the remaining blocks only have 13 nodes.
- Example: 6 blocks with 2,4,4,4,4,4 nodes per block respectively.
 - The max node block has 4 nodes which means the other 3 blocks should have $4 \times 4 = 16$ nodes. In this case it **WOULD** be block aware as the remaining blocks have 18 nodes which is above our minimum.

Configuration Data


Nutanix leverages Zookeeper to store essential configuration data for the cluster. This role is also distributed in a [block/rack]-aware manner to ensure availability in the case of a [block/rack] failure.

The following figure shows an example layout showing 3 Zookeeper nodes distributed in a [block/rack]-aware manner:



Zookeeper Block/Rack Aware Placement

In the event of a [block/rack] outage, meaning one of the Zookeeper nodes will be gone, the Zookeeper role would be transferred to another node in the cluster as shown below:



Zookeeper Placement Block/Rack Failure

When the [block/rack] comes back online, the Zookeeper role would be transferred back to maintain [block/rack] awareness.

NOTE: Prior to 4.5, this migration was not automatic and must be done manually.

Data Path Resiliency

Reliability and resiliency are key, if not the most important concepts within AOS or any primary storage platform.


Contrary to traditional architectures which are built around the idea that hardware will be reliable, Nutanix takes a different approach: it expects hardware will eventually fail. By doing so, the system is designed to handle these failures in an elegant and non-disruptive manner.

NOTE: That doesn't mean the hardware quality isn't there, just a concept shift. The Nutanix hardware and QA teams undergo an exhaustive qualification and vetting process.

As mentioned in the prior sections metadata and data are protected using a RF which is based upon the cluster FT level. As of 5.0 supported FT levels are FT1 and FT2 which correspond to metadata RF3 and data RF2, or metadata RF5 and data RF3 respectively.

To learn more about how metadata is sharded refer to the prior 'Scalable Metadata' section. To learn more about how data is protected refer to the prior 'Data protection' section.

In a normal state, cluster data layout will look similar to the following:



Data Path Resiliency - Normal State

As you can see the VM/vDisk data has 2 or 3 copies on disk which are distributed among the nodes and associated storage devices.

Importance of Data Distribution

By ensuring metadata and data is distributed across all nodes and all disk devices we can ensure the highest possible performance during normal data ingest and re-protection.

As data is ingested into the system its primary and replica copies will be distributed across the local and all other remote nodes. By doing so we can eliminate any potential hot spots (e.g. a node or disk performing slowly) and ensure a consistent write performance.

In the event of a disk or node failure where data must be re-protected, the full power of the cluster can be used for the rebuild. In this event the scan of metadata (to find out the data on the failed device(s) and where the replicas exist) will be distributed evenly across all CVMs. Once the data replicas have been found all healthy CVMs, disk devices (SSD+HDD), and host network uplinks can be used concurrently to rebuild the data.

For example, in a 4 node cluster where a disk fails each CVM will handle 25% of the metadata scan and data rebuild. In a 10 node cluster, each CVM will handle 10% of the metadata scan and data rebuild. In a 50 node cluster, each CVM will handle 2% of the metadata scan and data rebuild.

Key point: With Nutanix and by ensuring uniform distribution of data we can ensure consistent write performance and far superior re-protection times. This also applies to any cluster wide activity (e.g. erasure coding, compression, deduplication, etc.)

Comparing this to other solutions where HA pairs are used or a single disk holds a full copy of the data, they will face frontend performance issues if the mirrored node/disk is under duress (facing heavy IO or resource constraints).

Also, in the event of a failure where data must be re-protected, they will be limited by a single controller, a single node's disk resources and a single node's network uplinks. When terabytes of data must be re-replicated this will be severely constrained by the local node's disk and network bandwidth, increasing the time the system is in a potential data loss state if another failure occurs.

Potential levels of failure

Being a distributed system, AOS is built to handle component, service, and CVM failures, which can be characterized on a few levels:

- Disk Failure
- CVM "Failure"
- Node Failure

When does a rebuild begin?

When there is an unplanned failure (in some cases we will proactively take things offline if they aren't working correctly) we begin the rebuild process **immediately**.

Unlike some other vendors which wait 60 minutes to start rebuilding and only maintain a single copy during that period (very risky and can lead to data loss if there's any sort of failure), we are not willing to take that risk at the sacrifice of potentially higher storage utilization.

We can do this because of a) the granularity of our metadata b) choose peers for write RF dynamically (while there is a failure, all new data (e.g. new writes / overwrites) maintain their configured redundancy) and c) we can handle things coming back online during a rebuild and re-admit the data once it has been validated. In this scenario data may be "over-replicated" in which a Curator scan will kick off and remove the over-replicated copies.

Disk Failure

A disk failure can be characterized as just that, a disk which has either been removed, encounters a failure, or one that is not responding or has I/O errors. When Stargate sees I/O errors or the device fails to respond within a certain threshold it will mark the disk offline. Once that has occurred Hades will run S.M.A.R.T. and check the status of the device. If the tests pass the disk will be marked online, if they fail it will remain offline. If Stargate marks a disk offline multiple times (currently 3 times in an hour), Hades will stop marking the disk online even if S.M.A.R.T. tests pass.

VM impact:


- HA event: **No**
- Failed I/Os: **No**
- Latency: **No impact**

In the event of a disk failure, a Curator scan (MapReduce Framework) will occur immediately. It will scan the metadata (Cassandra) to find the data previously hosted on the failed disk and the nodes / disks hosting the replicas.

Once it has found that data that needs to be "re-replicated", it will distribute the replication tasks to the nodes throughout the cluster.

During this process a Drive Self Test (DST) is started for the bad disk and SMART logs are monitored for errors.

The following figure shows an example disk failure and re-protection:



Data Path Resiliency - Disk Failure

An important thing to highlight here is given how Nutanix distributes data and replicas across all nodes / CVMs / disks; all nodes / CVMs / disks will participate in the re-replication.

This substantially reduces the time required for re-protection, as the power of the full cluster can be utilized; the larger the cluster, the faster the re-protection.


Node Failure

VM Impact:

- HA event:**Yes**
- Failed I/Os:**No**
- Latency:**No impact**

In the event of a node failure, a VM HA event will occur restarting the VMs on other nodes throughout the virtualization cluster. Once restarted, the VMs will continue to perform I/Os as usual which will be handled by their local CVMs.

Similar to the case of a disk failure above, a Curator scan will find the data previously hosted on the node and its respective replicas. Once the replicas are found all nodes will participate in the re-protection.



Data Path Resiliency - Node Failure

In the event where the node remains down for a prolonged period of time (30 minutes as of 4.6), the down CVM will be removed from the metadata ring. It will be joined back into the ring after it has been up and stable for a duration of time.

Pro tip

Data resiliency state will be shown in Prism on the dashboard page.

You can also check data resiliency state via the cli:

Node Status

```
ncli cluster get-domain-fault-tolerance-status type=node
```

Block Status

```
ncli cluster get-domain-fault-tolerance-status type=rackable_unit
```

These should always be up to date, however to refresh the data you can kick off a Curator partial scan.

CVM "Failure"

A CVM "failure" can be characterized as a CVM power action causing the CVM to be temporarily unavailable. The system is designed to transparently handle these gracefully. In the event of a failure, I/Os will be re-directed to other CVMs within the cluster. The mechanism for this will vary by hypervisor.

The rolling upgrade process actually leverages this capability as it will upgrade one CVM at a time, iterating through the cluster.

VM impact:

- HA event:**No**
- Failed I/Os:**No**
- Latency:**Potentially higher given I/Os over the network**

In the event of a CVM "failure" the I/O which was previously being served from the down CVM, will be forwarded to other CVMs throughout the cluster. ESXi and Hyper-V handle this via a process called CVM Autopathing, which leverages HA.py (like "happy"), where it will modify the routes to forward traffic going to the internal address (192.168.5.2) to the external IP of other CVMs throughout the cluster. This enables the datastore to remain intact, just the CVM responsible for serving the I/Os is remote.

Once the local CVM comes back up and is stable, the route would be removed and the local CVM would take over all new I/Os.

In the case of AHV, iSCSI multi-pathing is leveraged where the primary path is the local CVM and the two other paths would be remote. In the event where the primary path fails, one of the other paths will become active.

Similar to Autopathing with ESXi and Hyper-V, when the local CVM comes back online, it'll take over as the primary path.

Resilient Capacity

Some refresher on terms used:

- Failure/Availability Domain(FD): Logical grouping of entities across which replicas are placed. Node is default domain on Nutanix clusters for node sizes > 3. Refer to Availability Domain section.
- Replication Factor(RF): Number of copies of data to be maintained on cluster for data availability.
- Fault Tolerance(FT): Number of failure(s) of entities at configured FD that can be handled by cluster, ensuring data on failed entity is rebuilt completely.

Resilient capacity is storage capacity in a cluster that can be consumed at the lowest availability/failure domain while maintaining the cluster's ability to self-heal and recover to desired replication factor (RF) after FT failure(s) at the configured availability/failure domain. So in simple terms, Resilient Capacity = Total Cluster Capacity - Capacity needed to rebuild from FT failure(s).

Resilient Capacity Examples

Homogeneous Cluster Capacity

- Configured Availability Domain: Node
- Lowest Availability Domain: Node
- FT=1
- RF=2
- Node Capacities: 10TB,10TB,10TB,10TB
- Resilient Capacity = $(40-10)*0.95 = 28.5\text{TB}$. 95% is threshold at which Stargate goes read-only mode not taking any more user writes


Non-Homogeneous Cluster Capacity

- Configured Availability Domain: Block(1 node/block)
- Lowest Availability Domain: Node
- FT=1
- RF=2
- Node Capacities: 10TB,10TB,40TB,40TB
- Resilient Capacity = $40*0.95 = 38\text{TB}$. 95% is threshold at which Stargate goes read-only mode not taking any more user writes


The resilient capacity in this case is 40TB and not 60TB because after losing the 40TB block, the cluster has node availability domain. At this level to maintain 2 data copies, the capacity available is 40TB which makes resilient capacity in this case to be 40TB overall.

It is recommended to keep clusters uniform and homogenous from capacity and failure domain perspective.


Starting with AOS 5.18, the resilient capacity is displayed in Prism Element storage summary widget with a gray line. Thresholds can be set to warn end users when cluster usage is reaching resilient capacity. By default that is set to 75%.



Prism can also show detailed storage utilizations on a per node basis which helps administrators understand resiliency on a per node basis. This is useful in clusters which have a skewed storage distribution.



When cluster usage is greater than resilient capacity for that cluster, the cluster might not be able to tolerate and recover from failures anymore. Cluster can possibly still recover and tolerate failure at a lower failure domain as resilient capacity is for configured failure domain. For example, a cluster with a node failure domain may still be able to self-heal and recover from a disk failure but cannot self-heal and recover from a node failure.



It is highly recommended to not exceed the resilient capacity of a cluster in any circumstances to ensure proper functioning of cluster and maintain its ability to self-heal and recover from failures.

Capacity Optimization

The Nutanix platform incorporates a wide range of storage optimization technologies that work in concert to make efficient use of available capacity for any workload. These technologies are intelligent and adaptive to workload characteristics, eliminating the need for manual configuration and fine-tuning.

The following optimizations are leveraged:

- Erasure Coding (EC-X)
- Compression
- Deduplication

More detail on how each of these features can be found in the following sections.

The table describes which optimizations are applicable to workloads at a high-level:

Data Transform	Best suited Application(s)	Comments
Erasure Coding (EC-X)	Most, Ideal for Nutanix Files/Objects	Provides higher availability with reduced overheads than traditional RF. No impact to normal write or read I/O performance. Does have some read overhead in the case of a disk / node / block failure where data must be decoded. Not suitable for workloads that are write or overwrite intensive, such as VDI
Inline Compression	All	No impact to random I/O, helps increase storage tier utilization. Benefits large or sequential I/O performance by reducing data to replicate and read from disk.
Offline Compression	None	Given inline compression will compress only large or sequential writes inline and do random or small I/Os post-process, that should be used instead.
Dedupe	Full Clones in VDI, Persistent Desktops, P2V/V2V, Hyper-V (ODX)	Greater overall efficiency for data which wasn't cloned or created using efficient AOS clones

Erasure Coding

The Nutanix platform leverages a replication factor (RF) for data protection and availability. This method provides the highest degree of availability because it does not require reading from more than one storage location or data re-computation on failure. However, this does come at the cost of storage resources as full copies are required.

To provide a balance between availability while reducing the amount of storage required, AOS provides the ability to encode data using erasure codes (EC).

Similar to the concept of RAID (levels 4, 5, 6, etc.) where parity is calculated, EC encodes a strip of data blocks on different nodes and calculates parity. In the event of a host and/or disk failure, the parity can be leveraged to calculate any missing data blocks (decoding). In the case of AOS, the data block is an extent group. Based upon the read nature of the data (read cold vs. read hot), the system will determine placement of the blocks in the strip.

For data that is read cold, we will prefer to distribute the data blocks from the same vDisk across nodes to form the strip (same-vDisk strip). This simplifies garbage collection (GC) as the full strip can be removed in the event the vDisk is deleted. For read hot data we will prefer to keep the vDisk data blocks local to the node and compose the strip with data from different vDisks (cross-vDisk strip). This minimizes remote reads as the local vDisk's data blocks can be local and other VMs/vDisks can compose the other data blocks in the strip. In the event a read cold strip becomes hot, AOS will try to recompute the strip and localize the data blocks.

The number of data and parity blocks in a strip is configurable based upon the desired failures to tolerate. The configuration is commonly referred to as the number of <number of data blocks>/<number of parity blocks>.

For example, "RF2 like" availability (N+1) could consist of 3 or 4 data blocks and 1 parity block in a strip (3/1 or 4/1). "RF3 like" availability (N+2) could consist of 3 or 4 data blocks and 2 parity blocks in a strip (3/2 or 4/2). The default strip sizes are 4/1 for RF2 like availability and 4/2 for RF3 like availability. These can be overridden using nCLI if desired.

```
ncli container [create/edit] ... erasure-code=<N>/<K>
```

where N is number of data blocks and K is number of parity blocks

EC + Block Awareness

As of 5.8, EC can place data and parity blocks in a block aware manner (prior to 5.8 this was done at a node level).

Pre-existing EC containers will not immediately change to block aware placement after being upgraded to 5.8. If there are enough blocks (strip size (k+n) + 1) available in the cluster these previously node aware strips will move to block aware. New EC containers will build block aware EC strips.

The expected overhead can be calculated as <# parity blocks> / <# data blocks>. For example, a 4/1 strip has a 25% overhead or 1.25X compared to the 2X of RF2. A 4/2 strip has a 50% overhead or 1.5X compared to the 3X of RF3.

The following table characterizes the encoded strip sizes and example overheads:


	FT1 (RF2 equiv.)		FT2 (RF3 equiv.)	
Cluster Size (nodes)	EC Strip Size (data/parity blocks)	EC Overhead (vs. 2X of RF2)	EC Strip Size (data/parity)	EC Overhead (vs. 3X of RF3)
4	2/1	1.5X	N/A	N/A
5	3/1	1.33X	N/A	N/A
6	4/1	1.25X	2/2	2X
7	4/1	1.25X	3/2	1.6X
8+	4/1	1.25X	4/2	1.5X

Pro tip

It is always recommended to have a cluster size which has at least 1 more node (or block for block aware data / parity placement) than the combined strip size (data + parity) to allow for rebuilding of the strips in the event of a node or block failure. This eliminates any computation overhead on reads once the strips have been rebuilt (automated via Curator). For example, a 4/1 strip should have at least 6 nodes in the cluster for a node aware EC strip or 6 blocks for a block aware EC strip. The previous table follows this best practice.

The encoding is done post-process and leverages the Curator MapReduce framework for task distribution. Since this is a post-process framework, the traditional write I/O path is unaffected.

A normal environment using RF would look like the following:




Typical AOS RF Data Layout

In this scenario, we have RF2 data whose primary copies are local and replicas are distributed to other nodes throughout the cluster.

When a Curator full scan runs, it will find eligible extent groups which are available to become encoded. Eligible extent groups must be "write-cold" meaning they haven't been written to for awhile. This is controlled with the following Curator Gflag: `curatoreraserasurecodethresholdseconds`. After the eligible candidates are found, the encoding tasks will be distributed and throttled via Chronos.


The following figure shows an example 3/1 strip:



AOS Encoded Strip - Pre-savings

Once the data has been successfully encoded (strips and parity calculation), the replica extent groups are then removed.

The following figure shows the environment after EC has run with the storage savings:



AOS Encoded Strip - Post-savings

Starting with AOS 5.18, inline erasure coding can be enabled on containers using nCLI. Inline erasure coding encodes and creates coding strips inline by erasure coding the data without waiting for it to become write-cold. In inline erasure coding, the erasure coded strips are created on the same vDisk data by default. With AOS 6.6 there is also an option to create strips from cross vdisks preserving data locality if that is desired. Due to this, it is only recommended for workloads that do not require data locality or perform numerous overwrites. Objects deployment have inline EC-X turned on by default.

Pro tip

Erasure Coding pairs perfectly with inline compression which will add to the storage savings.

Compression

For a visual explanation, you can watch the following video: [LINK](#)

The Nutanix Capacity Optimization Engine (COE) is responsible for performing data transformations to increase data efficiency on disk. Currently compression is one of the key features of the COE to perform data optimization. AOS provides both inline and offline flavors of compression to best suit the customer's needs and type of data. As of 5.1, offline compression is enabled by default.

Inline compression will compress sequential streams of data or large I/O sizes (>64K) when written to the Extent Store (SSD + HDD). This includes data draining from OpLog as well as sequential data skipping it.

OpLog Compression

As of 5.0, the OpLog will now compress all incoming writes >4K that show good compression (Gflag: `vdisk_distributed_oplog_enable_compression`). This will allow for a more efficient utilization of the OpLog capacity and help drive sustained performance.

When drained from OpLog to the Extent Store the data will be decompressed, aligned and then re-compressed at a 32K aligned unit size (as of 5.1).

This feature is on by default and no user configuration is necessary.

Offline compression will initially write the data as normal (in an un-compressed state) and then leverage the Curator framework to compress the data cluster wide. When inline compression is enabled but the I/Os are random in nature, the data will be written un-compressed in the OpLog, coalesced, and then compressed in memory before being written to the Extent Store.


Nutanix leverages LZ4 and LZ4HC for data compression with AOS 5.0 and beyond. Prior to AOS 5.0 the Google Snappy compression library is leveraged which provides good compression ratios with minimal computational overhead and extremely fast compression / decompression rates.

Normal data will be compressed using LZ4 which provides a very good blend between compression and performance. For cold data, LZ4HC will be leveraged to provide an improved compression ratio.

Cold data is characterized into two main categories:

- Regular data: No R/W access for 3 days (Gflag: `curatormediumcompressmutabledatadelaysecs`)
- Immutable data (snapshots): No R/W access for 1 day (Gflag: `curatormediumcompressimmutabledatadelaysecs`)

The following figure shows an example of how inline compression interacts with the AOS write I/O path:



Inline Compression I/O Path

Pro tip


Almost always use inline compression (compression delay = 0) as it will only compress larger / sequential writes and not impact random write performance.

This will also increase the usable size of the SSD tier increasing effective performance and allowing more data to sit in the SSD tier. Also, for larger or sequential data that is written and compressed inline, the replication for RF will be shipping the compressed data, further increasing performance since it is sending less data across the wire.

Inline compression also pairs perfectly with erasure coding.

For offline compression, all new write I/O is written in an un-compressed state and follows the normal AOS I/O path. After the compression delay (configurable) is met, the data is eligible to become compressed. Compression can occur anywhere in the Extent Store. Offline compression uses the Curator MapReduce framework and all nodes will perform compression tasks. Compression tasks will be throttled by Chronos.

The following figure shows an example of how offline compression interacts with the AOS write I/O path:



Offline Compression I/O Path

For read I/O, the data is first decompressed in memory and then the I/O is served.

You can view the current compression rates via Prism on the Storage > Dashboard page.

Elastic Dedupe Engine

For a visual explanation, you can watch the following video: [LINK](#)

The Elastic Dedupe Engine is a software-based feature of AOS which allows for data deduplication in the capacity (Extent Store) tiers. Streams of data are fingerprinted during ingest at a 16KB granularity (Controlled by: `stargatededupfingerprint`) within a 1MB extent. Prior to AOS 5.11, AOS used only SHA-1 hash to fingerprint and identify candidates for dedupe. Since AOS 5.11, AOS now uses logical checksums to select candidates for dedupe. This fingerprint is only done on data ingest and is then stored persistently as part of the written block's metadata.

Contrary to traditional approaches which utilize background scans requiring the data to be re-read, Nutanix performs the fingerprint inline on ingest. For duplicate data that can be deduplicated in the capacity tier, the data does not need to be scanned or re-read, essentially duplicate copies can be removed.

To make the metadata overhead more efficient, fingerprint refcounts are monitored to track dedupability. Fingerprints with low refcounts will be discarded to minimize the metadata overhead. To minimize fragmentation full extents will be preferred for deduplication.

Pro tip


Use deduplication on your base images (you can manually fingerprint them using `vdisk_manipulator`) to take advantage of the unified cache.

Fingerprinting is done during data ingest of data with an I/O size of 64K or greater (initial I/O or when draining from OpLog). AOS then looks at hashes/fingerprints of each 16KB chunk within a 1MB extent and if it finds duplicates for more than 40% of chunks, it dedupes the entire extent. That resulted in many dedupe extents with reference count of 1 (no other duplicates) within the 1MB extent from the remaining 60% of extent, that ended up using metadata.

With AOS 6.6, the algorithm was further enhanced such that within a 1MB extent, only chunks that have duplicates will be marked for deduplication instead of entire extent reducing the metadata required. With AOS 6.6, changes were also made with the way dedupe metadata was stored. Before AOS 6.6, dedupe metadata was stored in a top level vdisk block map which resulted in dedupe metadata to be copied when snapshots were taken. This resulted in a metadata bloat. With AOS 6.6, that metadata is now stored in extent group id map which is a level lower than vdisk block map. Now when snapshots are taken of the vdisk, it does not result in copying of dedupe metadata and prevents metadata bloat. Once the fingerprinting is done, a background process will remove the duplicate data using the AOS MapReduce framework (Curator). For data that is being read, the data will be pulled into the AOS Unified Cache which is a multi-tier/pool cache. Any subsequent requests for data having the

same fingerprint will be pulled directly from the cache. To learn more about the Unified Cache and pool structure, please refer to the [Unified Cache](#) sub-section in the I/O path overview.

The following figure shows an example of how the Elastic Dedupe Engine interacts with the AOS I/O path:




EDE I/O Path

You can view the current deduplication rates via Prism on the Storage > Dashboard page.

Storage Tiering and Prioritization

The Disk Balancing section below talks about how storage capacity was pooled among all nodes in a Nutanix cluster and that ILM would be used to keep hot data local. A similar concept applies to disk tiering, in which the cluster's SSD and HDD tiers are cluster-wide and AOS ILM is responsible for triggering data movement events. A local node's SSD tier is always the highest priority tier for all I/O generated by VMs running on that node, however all of the cluster's SSD resources are made available to all nodes within the cluster. The SSD tier will always offer the highest performance and is a very important thing to manage for hybrid arrays.


The tier prioritization can be classified at a high-level by the following:



AOS Tier Prioritization


Specific types of resources (e.g. SSD, HDD, etc.) are pooled together and form a cluster wide storage tier. This means that any node within the cluster can leverage the full tier capacity, regardless if it is local or not.

The following figure shows a high level example of what this pooled tiering looks like:



AOS Cluster-wide Tiering


A common question is what happens when a local node's SSD becomes full? As mentioned in the Disk Balancing section, a key concept is trying to keep uniform utilization of devices within disk tiers. In the case where a local node's SSD utilization is high, disk balancing will kick in to move the coldest data on the local SSDs to the other SSDs throughout the cluster. This will free up space on the local SSD to allow the local node to write to SSD locally instead of going over the network. A key point to mention is that all CVMs and SSDs are used for this remote I/O to eliminate any potential bottlenecks and remediate some of the hit by performing I/O over the network.



AOS Cluster-wide Tier Balancing

The other case is when the overall tier utilization breaches a specific threshold [`curatortierusageilmthresholdpercent` (Default=75)] where AOS ILM will kick in and as part of a Curator job will down-migrate data from the SSD tier to the HDD tier. This will bring utilization within the threshold mentioned above or free up space by the following amount [`curatortierfreeupperpercentby_ilm` (Default=15)], whichever is greater. The data for down-migration is chosen using last access time. In the case where the SSD tier utilization is 95%, 20% of the data in the SSD tier will be moved to the HDD tier (95% → 75%).

However, if the utilization was 80%, only 15% of the data would be moved to the HDD tier using the minimum tier free up amount.



AOS Tier ILM

AOS ILM will constantly monitor the I/O patterns and (down/up) migrate data as necessary as well as bring the hottest data local regardless of tier. The logic for up-migration (or horizontal) follows the same as that defined for egroup locality: "3 touches for

random or 10 touches for sequential within a 10 minute window where multiple reads every 10 second sampling count as a single touch".


Disk Balancing

For a visual explanation, you can watch the following video: [LINK](#)

AOS is designed to be a very dynamic platform which can react to various workloads as well as allow heterogeneous node types: compute heavy (3050, etc.) and storage heavy (60X0, etc.) to be mixed in a single cluster. Ensuring uniform distribution of data is an important item when mixing nodes with larger storage capacities. AOS has a native feature, called disk balancing, which is used to ensure uniform distribution of data throughout the cluster. Disk balancing works on a node's utilization of its local storage capacity and is integrated with AOS ILM. Its goal is to keep utilization uniform among nodes once the utilization has breached a certain threshold.

NOTE: Disk balancing jobs are handled by Curator which has different priority queues for primary I/O (UVM I/O) and background I/O (e.g. disk balancing). This is done to ensure disk balancing or any other background activity doesn't impact front-end latency / performance. In this cases the job's tasks will be given to Chronos who will throttle / control the execution of the tasks. Also, movement is done within the same tier for disk balancing. For example, if data is skewed in the HDD tier, it can be moved amongst nodes in the same tier.


The following figure shows an example of a mixed cluster (3050 + 6050) in an "unbalanced" state:



Disk Balancing - Unbalanced State

Disk balancing leverages the AOS Curator framework and is run as a scheduled process as well as when a threshold has been breached (e.g., local node capacity utilization > n %). In the case where the data is not balanced, Curator will determine which data needs to be moved and will distribute the tasks to nodes in the cluster. In the case where the node types are homogeneous (e.g., 3050), utilization should be fairly uniform. However, if there are certain VMs running on a node which are writing much more data than others, this can result in a skew in the per node capacity utilization. In this case, disk balancing would run and move the coldest data on that node to other nodes in the cluster. In the case where the node types are heterogeneous (e.g., 3050 + 6020/50/70), or where a node may be used in a "storage only" mode (not running any VMs), there will likely be a requirement to move data.


The following figure shows an example the mixed cluster after disk balancing has been run in a "balanced" state:



Disk Balancing - Balanced State

In some scenarios, customers might run some nodes in a "storage-only" state where only the CVM will run on the node whose primary purpose is bulk storage capacity. In this case, the full node's memory can be added to the CVM to provide a much larger read cache.

The following figure shows an example of how a storage only node would look in a mixed cluster with disk balancing moving data to it from the active VM nodes:



Disk Balancing - Storage Only Node


Snapshots and Clones

For a visual explanation, you can watch the following video: [LINK](#)

AOS provides native support for offloaded snapshots and clones which can be leveraged via VAAI, ODX, ncli, REST, Prism, etc. Both the snapshots and clones leverage the redirect-on-write algorithm which is the most effective and efficient. As explained in the Data Structure section above, a virtual machine consists of files (vmdk/vhdx) which are vDisks on the Nutanix platform.


A vDisk is composed of extents which are logically contiguous chunks of data, which are stored within extent groups which are physically contiguous data stored as files on the storage devices. When a snapshot or clone is taken, the base vDisk is marked immutable and another vDisk is created as read/write. At this point, both vDisks have the same block map, which is a metadata mapping of the vDisk to its corresponding extents. Contrary to traditional approaches which require traversal of the snapshot chain (which can add read latency), each vDisk has its own block map. This eliminates any of the overhead normally seen by large snapshot chain depths and allows you to take continuous snapshots without any performance impact.

The following figure shows an example of how this works when a snapshot is taken (source: NTAP):



Example Snapshot Block Map


The same method applies when a snapshot or clone of a previously snapped or cloned vDisk is performed:



Multi-snap Block Map and New Write

The same methods are used for both snapshots and/or clones of a VM or vDisk(s). When a VM or vDisk is cloned, the current block map is locked and the clones are created. These updates are metadata only, so no I/O actually takes place. The same method applies for clones of clones; essentially the previously cloned VM acts as the "Base vDisk" and upon cloning, that block map is locked and two "clones" are created: one for the VM being cloned and another for the new clone. There is no imposed limit on the maximum number of clones.

They both inherit the prior block map and any new writes/updates would take place on their individual block maps.




Multi-Clone Block Maps

As mentioned previously, each VM/vDisk has its own individual block map. So in the above example, all of the clones from the base VM would now own their block map and any write/update would occur there.

In the event of an overwrite the data will go to a new extent / extent group. For example, if data exists at offset **o1** in extent **e1** that was being overwritten, Stargate would create a new extent **e2** and track that the new data was written in extent **e2** at offset **o2**. The Vblock map tracks this down to the byte level.

The following figure shows an example of what this looks like:




Clone Block Maps - New Write

Any subsequent clones or snapshots of a VM/vDisk would cause the original block map to be locked and would create a new one for R/W access.

Networking and I/O

The Nutanix platform does not leverage any backplane for inter-node communication and only relies on a standard 10GbE network. All storage I/O for VMs running on a Nutanix node is handled by the hypervisor on a dedicated private network. The I/O request will be handled by the hypervisor, which will then forward the request to the private IP on the local CVM. The CVM will then perform the remote replication with other Nutanix nodes using its external IP over the public 10GbE network. For all read requests, these will be served completely locally in most cases and never touch the 10GbE network. This means that the only traffic touching the public 10GbE network will be AOS remote replication traffic and VM network I/O. There will, however, be cases where the CVM will forward requests to other CVMs in the cluster in the case of a CVM being down or data being remote. Also, cluster-wide tasks, such as disk balancing, will temporarily generate I/O on the 10GbE network.

The following figure shows an example of how the VM's I/O path interacts with the private and public 10GbE network:



AOS Networking

Data Locality


Being a converged (compute+storage) platform, I/O and data locality are critical to cluster and VM performance with Nutanix. As explained above in the I/O path, all read/write IOs are served by the local Controller VM (CVM) which is on each hypervisor adjacent to normal VMs. A VM's data is served locally from the CVM and sits on local disks under the CVM's control. When a VM

is moved from one hypervisor node to another (or during a HA event), the newly migrated VM's data will be served by the now local CVM. When reading old data (stored on the now remote node/CVM), the I/O will be forwarded by the local CVM to the remote CVM. All write I/Os will occur locally right away. AOS will detect the I/Os are occurring from a different node and will migrate the data locally in the background, allowing for all read I/Os to now be served locally. The data will only be migrated on a read as to not flood the network.

Data locality occurs in two main flavors:

- Cache Locality
 - Pulling remote data into the local Stargate's Unified Cache. This is done at a 4K granularity.
 - For instances where there are no local replicas, the requests will be forward to the Stargate(s) containing the replicas which will return the data and the local Stargate will store this locally then return the I/O. All subsequent requests for that data will be returned from the cache.
- Extent Group (egroup) Locality
 - Migrating the vDisk extent group(s) (egroups) to be stored in the local Stargate's Extent Store.
 - If a replica egroup is already local, no movement is necessary.
 - In this scenario the actual replica egroup will be re-localized after certain I/O thresholds are met. We don't automatically re-localize / migrate egroups to ensure we're leveraging the network efficiently.
 - For AES enabled egroups the same horizontal migration occurs for cases where replicas aren't local and the patterns are met.

The following figure shows an example of how data will “follow” the VM as it moves between hypervisor nodes:



Data Locality

Thresholds for Data Migration

Cache locality occurs in real time and will be determined based upon vDisk ownership. When a vDisk / VM moves from one node to another the "ownership" of those vDisk(s) will transfer to the now local CVM. Once the ownership has transferred the data can be cached locally in the Unified Cache. In the interim the cache will be wherever the ownership is held (the now remote host). The previously hosting Stargate will relinquish the vDisk token when it sees remote I/Os for 300+ seconds at which it will then be taken by the local Stargate. Cache coherence is enforced as ownership is required to cache the vDisk data.

Egroup locality is a sampled operation and an extent group will be migrated when the following occurs: "3 touches for random or 10 touches for sequential within a 10 minute window where multiple reads every 10 second sampling count as a single touch".


Shadow Clones

The Distributed Storage Fabric has a feature called 'Shadow Clones', which allows for distributed caching of particular vDisks or VM data which is in a 'multi-reader' scenario. A great example of this is during a VDI deployment many 'linked clones' will be forwarding read requests to a central master or 'Base VM'. In the case of VMware View, this is called the replica disk and is read by all linked clones, and in XenDesktop, this is called the MCS Master VM. This will also work in any scenario which may be a multi-reader scenario (e.g., deployment servers, repositories, etc.). Data or I/O locality is critical for the highest possible VM performance and a key struct of AOS.

With Shadow Clones, AOS will monitor vDisk access trends similar to what it does for data locality. However, in the case there are requests occurring from more than two remote CVMs (as well as the local CVM), and all of the requests are read I/O, the vDisk will be marked as immutable. Once the disk has been marked as immutable, the vDisk can then be cached locally by each CVM

making read requests to it (aka Shadow Clones of the base vDisk). This will allow VMs on each node to read the Base VM's vDisk locally. In the case of VDI, this means the replica disk can be cached by each node and all read requests for the base will be served locally. NOTE: The data will only be migrated on a read as to not flood the network and allow for efficient cache utilization. In the case where the Base VM is modified, the Shadow Clones will be dropped and the process will start over. Shadow clones are enabled by default (as of 4.0.2) and can be enabled/disabled using the following NCLI command: `ncli cluster edit-params enable-shadow-clones=`.


The following figure shows an example of how Shadow Clones work and allow for distributed caching:



Shadow Clones

Storage Layers and Monitoring

The Nutanix platform monitors storage at multiple layers throughout the stack, ranging from the VM/Guest OS all the way down to the physical disk devices. Knowing the various tiers and how these relate is important whenever monitoring the solution and allows you to get full visibility of how the ops relate. The following figure shows the various layers of where operations are monitored and the relative granularity which are explained below:



Storage Layers

Virtual Machine Layer

- Key Role: Metrics reported by the hypervisor for the VM
- Description: Virtual Machine or guest level metrics are pulled directly from the hypervisor and represent the performance the VM is seeing and is indicative of the I/O performance the application is seeing.
- When to use: When troubleshooting or looking for VM level detail

Hypervisor Layer

- Key Role: Metrics reported by the Hypervisor(s)
- Description: Hypervisor level metrics are pulled directly from the hypervisor and represent the most accurate metrics the hypervisor(s) are seeing. This data can be viewed for one of more hypervisor node(s) or the aggregate cluster. This layer will provide the most accurate data in terms of what performance the platform is seeing and should be leveraged in most cases. In certain scenarios the hypervisor may combine or split operations coming from VMs which can show the difference in metrics reported by the VM and hypervisor. These numbers will also include cache hits served by the Nutanix CVMs.
- When to use: Most common cases as this will provide the most detailed and valuable metrics.

Controller Layer

- Key Role: Metrics reported by the Nutanix Controller(s)
- Description: Controller level metrics are pulled directly from the Nutanix Controller VMs (e.g., Stargate 2009 page) and represent what the Nutanix front-end is seeing from NFS/SMB/iSCSI or any back-end operations (e.g., ILM, disk balancing, etc.). This data can be viewed for one of more Controller VM(s) or the aggregate cluster. The metrics seen by the Controller Layer should normally match those seen by the hypervisor layer, however will include any backend operations (e.g., ILM, disk balancing). These numbers will also include cache hits served by memory. In certain cases, metrics like (IOPS), might not match as the NFS / SMB / iSCSI client might split a large IO into multiple smaller IOPS. However, metrics like bandwidth should match.
- When to use: Similar to the hypervisor layer, can be used to show how much backend operation is taking place.

Disk Layer

- Key Role: Metrics reported by the Disk Device(s)
- Description: Disk level metrics are pulled directly from the physical disk devices (via the CVM) and represent what the back-end is seeing. This includes data hitting the OpLog or Extent Store where an I/O is performed on the disk. This data can be viewed for one of more disk(s), the disk(s) for a particular node, or the aggregate disks in the cluster. In common cases, it is expected that the disk ops should match the number of incoming writes as well as reads not served from the memory portion of the cache. Any reads being served by the memory portion of the cache will not be counted here as the op is not hitting the disk device.
- When to use: When looking to see how many ops are served from cache or hitting the disks.

Metric and Stat Retention

Metrics and time series data is stored locally for 90 days in Prism Element. For Prism Central and Insights, data can be stored indefinitely (assuming capacity is available).

NX-8155-G9 System Specifications

Platform ANY

NX-8155-G9

August 6, 2025

Contents

- System Specifications..... 3**
 - Node Naming (NX-8155-G9).....3
 - NX-8155-G9 System Specifications..... 6
- Component Specifications..... 14**
 - Controls and LEDs for Single-node Platforms..... 14
 - Network Card LED Description..... 16
 - Power Supply Unit Redundancy and Node Configuration.....17
- Memory Configurations..... 19**
 - Supported Memory Configurations..... 19
- Nutanix Hardware Naming Convention..... 22**
- Copyright..... 24**

SYSTEM SPECIFICATIONS

Node Naming (NX-8155-G9)

Nutanix assigns a name to each node in a block, which varies by product type. The NX-8155-G9 block contains a single node named Node A.

The NX-8155-G9 block contains 12 drive slots, and it is supported as a hyper-converged infrastructure (HCI) node, Nutanix Unified Storage (NUS) node, or a compute-only node with no storage drives. When used as a HCI or NUS node, the NX-8155-G9 supports All-NVMe or hybrid NVMe + HDD configurations.



Figure 1: NX-8155-G9 Front Panel

Partial Population

The NX-8155-G9 supports partial population of drives according to the following restrictions:

- Partially populated configurations support only even numbers of drives.
- Populate the drive slots in order from bottom to top and left to right.

When partially populating a node, Nutanix recommends that you put an empty drive carrier in each unpopulated drive slot for proper airflow and to prevent dust contamination.

Hybrid NVMe + HDD Drive Configuration

Note: Hybrid NVMe + HDD configuration is available as a factory installation only. You cannot convert a hybrid NVMe + HDD node to an all-NVMe node in the field.

HDD 1 [SLOT3]	HDD 4 [SLOT6]	HDD 7 [SLOT9]	HDD 10 [SLOT12]
NVMe2 [SLOT2]	HDD 3 [SLOT5]	HDD 6 [SLOT8]	HDD 9 [SLOT11]
NVMe 1 [SLOT1]	HDD 2 [SLOT4]	HDD 5 [SLOT7]	HDD 8 [SLOT10]

Figure 2: Hybrid NVMe + HDD Configuration 1

- Two NVMe drives in slots 1 and 2; four HDDs in slots 3 through 6; all other slots empty
- Two NVMe drives in slots 1 and 2; six HDDs in slots 3 through 8; all other slots empty

- Two NVMe drives in slots 1 and 2; eight HDDs in slots 3 through 10; all other slots empty
- Two NVMe drives in slots 1 and 2; 10 HDDs in slots 3 through 12

NVMe3 [SLOT3]	HDD 2 [SLOT6]	HDD 5 [SLOT9]	HDD 8 [SLOT12]
NVMe2 [SLOT2]	HDD 1 [SLOT5]	HDD 4 [SLOT8]	HDD 7 [SLOT11]
NVMe 1 [SLOT1]	NVMe 4 [SLOT4]	HDD 3 [SLOT7]	HDD 6 [SLOT10]

Figure 3: Hybrid NVMe + HDD Configuration 2

- Four SSDs in slots 1 through 4; eight HDDs in slots 5 through 12

All-NVMe Configuration

NVMe 3 [SLOT3]	NVMe 6 [SLOT6]	NVMe 9 [SLOT9]	NVMe 12 [SLOT12]
NVMe 2 [SLOT2]	NVMe 5 [SLOT5]	NVMe 8 [SLOT8]	NVMe 11 [SLOT11]
NVMe 1 [SLOT1]	NVMe 4 [SLOT4]	NVMe 7 [SLOT7]	NVMe 10 [SLOT10]

Figure 4: All-NVMe Configuration

Note: A minimum of four NVMe drives are mandatory for all-NVMe configuration.

- Four NVMe drives in slots 1 through 4; all other slots empty
- Six NVMe drives in slots 1 through 6; all other slots empty
- Eight NVMe drives in slots 1 through 8; all other slots empty
- 10 NVMe drives in slots 1 through 10; all other slots empty
- 12 NVMe drives in slots 1 through 12 (fill all slots)



Figure 5: NX-8155-G9 Control Panel




Figure 6: NX-8155-G9 Back Panel

The back panel maps the placement of these components:

- PSU 1 and PSU 2
- NIC 1, NIC 3 and NIC 3 slots
- HBA slot
- M2 air shroud location
- 1 x 100/1000 (IPMI)
- 2 x USB 3.0
- 15 Pin VGA
- UID LED

- AIOM ports
 - Port 1 (Shared IPMI)
 - Port 2

NX-8155-G9 System Specifications

Table 1: System Characteristics

Boot Device	Boot Drive
	• 2 x 512GB M.2 Boot Device
Chassis	Chassis
	• 2U1N LFF Chassis

CPU**Processor****Dual Intel Xeon® 5th Gen. (Emerald Rapids)**

- 2 x Intel Xeon® Silver 4509Y [8 cores / 2.60 GHz / 125 W]
- 2 x Intel Xeon® Silver 4510 [12 cores / 2.40 GHz / 150 W]
- 2 x Intel Xeon® Silver 4514Y [16 cores / 2.00 GHz / 150 W]
- 2 x Intel Xeon® Silver 4516Y+ [24 cores / 2.20 GHz / 185 W]
- 2 x Intel Xeon® Gold 5515+ [8 cores / 3.20 GHz / 165 W]
- 2 x Intel Xeon® Gold 6526Y [16 cores / 2.80 GHz / 195 W]
- 2 x Intel Xeon® Gold 6544Y [16 cores / 3.60 GHz / 270 W]
- 2 x Intel Xeon® Gold 6542Y [24 cores / 2.90 GHz / 250 W]
- 2 x Intel Xeon® Gold 5520+ [28 cores / 2.20 GHz / 205 W]
- 2 x Intel Xeon® Gold 6548Y+ [32 cores / 2.50 GHz / 250 W]
- 2 x Intel Xeon® Gold 6554S [36 cores / 2.20 GHz / 270 W]

Dual Intel Xeon® 4th Gen. (Sapphire Rapids)

- 2 x Intel Xeon® Silver 4410T [10 cores / 2.70 GHz / 150 W], Excluded from ENERGY STAR certification.
 - 2 x Intel Xeon® Silver 4410Y [12 cores / 2.00 GHz / 150 W]
 - 2 x Intel Xeon® Silver 4416+ [20 cores / 2.00 GHz / 165 W]
 - 2 x Intel Xeon® Gold 5415+ [8 cores / 2.90 GHz / 150 W], Excluded from ENERGY STAR certification.
 - 2 x Intel Xeon® Gold 5416S [16 cores / 2.00 GHz / 150 W]
 - 2 x Intel Xeon® Gold 5418Y [24 cores / 2.00 GHz / 185 W]
 - 2 x Intel Xeon® Gold 5420+ [28 cores / 2.00 GHz / 205 W]
 - 2 x Intel Xeon® Gold 6416H [18 cores / 2.20 GHz / 165 W]
 - 2 x Intel Xeon® Gold 6426Y [16 cores / 2.50 GHz / 185 W]
 - 2 x Intel Xeon® Gold 6434 [8 cores / 3.70 GHz / 195 W]
 - 2 x Intel Xeon® Gold 6442Y [24 cores / 2.60 GHz / 225 W]
 - 2 x Intel Xeon® Gold 6444Y [16 cores / 3.60 GHz / 270 W]
 - 2 x Intel Xeon® Gold 6454S [32 cores / 2.20 GHz / 270 W]
-

Note:

- The minimum memory configuration required for Energy STAR 4.0 certification is 512 GB. If memory is less than 512 GB, the platform is not Energy Star certified, even if the CPU is certified.
- The 128 GB DDR5 5600 DIMMs require BIOS version 31.001 or later, and must not be mixed with DDR5 4800 DIMMs within the same node.
- 2304 GB capacity is available via 96 GB DIMMs which are only supported with Emerald Rapids processors.

32 GB DIMM

- 4 x 32 GB = 128 GB
- 8 x 32 GB = 256 GB
- 12 x 32 GB = 384 GB
- 16 x 32 GB = 512 GB
- 24 x 32 GB = 768 GB
- 32 x 32 GB = 1024 GB

64 GB DIMM

- 4 x 64 GB = 256 GB
- 8 x 64 GB = 512 GB
- 12 x 64 GB = 768 GB
- 16 x 64 GB = 1024 GB
- 24 x 64 GB = 1536 GB
- 32 x 64 GB = 2048 GB

96 GB DIMM

- 16 x 96 GB = 1536 GB
- 24 x 96 GB = 2304 GB
- 32 x 96 GB = 3072 GB

128 GB DIMM

- 4 x 128 GB = 512 GB
- 8 x 128 GB = 1024 GB
- 12 x 128 GB = 1536 GB
- 16 x 128 GB = 2048 GB
- 24 x 128 GB = 3072 GB
- 32 x 128 GB = 4096 GB

Network	Ports on Serverboard • 1x 1GbE Dedicated IPMI AIOM • 1 x 2P 10GBase-T (Port 1 is shared IPMI) • 1 x 2P 10GBase-T + 2P SFP+ (Port 1 is shared IPMI) Add on NICs in PCIe slots • 0, 1, 2 or 3 x 10GbE 4P NIC • 0, 1, 2 or 3 x 10GBaseT 2P NIC • 0, 1, 2 or 3 x 10GBaseT 4P NIC • 0, 1, 2 or 3 x 25GbE 2P NIC • 0, 1, 2 or 3 x 25GbE 4P NIC • 0, 1, 2 or 3 x 100GbE 2P NIC
Network Cables	Network Cable • 1M-QSFP28 • 1M-QSFP+ • 1M-SFP28 • 3M-QSFP28 • 3M-SFP28 • 3M-SFP+ • 5M-QSFP28 • 5M-QSFP+ • 5M-SFP28 • 5M-SFP+
Power Cable	Power Cable • C13/14 4ft Power Cable
Power Supply	Power Supply • 2 x PWS, 2000W
Server	Server • 1 x NX-8155-G9 Server

Storage

Note:

- NX-8155-G9 supports the following storage solutions: (1) HCI node or NUS node with up to 12 drives installed. (2) Compute Only node with zero drives installed.
- AOS 7.3 (or higher) is required for NVMe SED support.
- NVMe SED is supported with only Emerald Rapids processors.

All NVMe	4, 6, 8, 10 or 12 x NVMe

- 1.92 TB
- 3.84 TB
- 7.68 TB
- 15.36 TB

NVMe+HDD	2 x NVMe

- 1.92 TB
- 3.84 TB
- 7.68 TB
- 15.36 TB

4, 6, 8 or 10 x HDD

- 8 TB
- 12 TB
- 18 TB
- 24 TB

NVMe+HDD	4 x NVMe

- 1.92 TB
- 3.84 TB
- 7.68 TB
- 15.36 TB

8 x HDD

- 8 TB
 - 12 TB
 - 18 TB
 - 24 TB
-

All NVMe SED	4, 6, 8, 10 or 12 x NVMe • 3.84 TB • 7.68 TB • 15.36 TB
NVMe SED+HDD SED	2 x NVMe • 3.84 TB • 7.68 TB • 15.36 TB 4, 6, 8 or 10 x HDD • 8 TB • 12 TB
NVMe SED+HDD SED	4 x NVMe • 3.84 TB • 7.68 TB • 15.36 TB 8 x HDD • 8 TB • 12 TB
Storage Controller	Storage Controller • 1 x Broadcom 3816 Host Bus Controller (HBA), required for SATA/SAS based drive support
TPM	TPM • 0 or 1 x Unprovisioned Trusted Platform Module
Transceiver	Transceiver • SR SFP+ Transceiver
VGA	1 x VGA connector per node (15-pin female)
Chassis fans	4 x 80 mm heavy-duty fans with PWM fan speed controls.

Table 2: Block, power and electrical

Block	• Height : 88.1 mm • Depth : 799 mm • Width : 445.6 mm Maximum configuration (without rails and accessory kits) • Weight : 27.3 kg
Package	Package maximum weight with rails and accessory kits • Weight : 36.36 kg
Shock	10ms, half-sine, one shock on each side • Operating : 5 G 10ms, square wave, one shock on each side • Non-Operating : 20 G
Thermal Dissipation	• Typical : 3705 BTU/hr • Maximum : 4941 BTU/hr
Vibration (Random)	at 5 to 500 Hz, Approx. 15 min./axis • Operating : 0.1 Grms
Power consumption (calculated)	Note: • For the power consumption calculations, Max NIC and disk configurations are considered. Max Config • Maximum: 1448 VA • Typical: 1086 VA
Operating environment	Operating temperature : 10-30 C Non-Operating temperature : -20-70 C Operating relative humidity : 0-85% Non-operating relative humidity : 0-95%

Certifications

- BIS
- BSMI

- CE
- CSA
- CSAus
- EAC
- Energy Star
- FCC
- ICES
- KCC
- RCM
- Reach
- RoHS
- S-MARK
- SABS
- SII
- UKCA
- UKRSEPRO
- UL
- VCCI-A
- WEEE
- cUL

COMPONENT SPECIFICATIONS

Controls and LEDs for Single-node Platforms



Figure 7: Front Panel Controls and Indicators

Table 3: Front Panel Controls and Indicators

LED or Button	Function
Power button	Power button applies or removes primary power from the power supply to the system (System standby power maintains)
Power LED	Solid green: Indicates power is supplied to the system power supply units. This LED is illuminated when the system is operating normally.
UID/Reset button	Turns information LED <i>i</i> On/Off (Single press) Reboots the system (Press and hold for 6 seconds)
Drive LED	Not used
NIC LED (AIOM Port 1)	Indicates network activity on AIOM 10Gbase-T port 1
NIC LED (AIOM Port 2)	Indicates network activity on AIOM 10Gbase-T port 2
Multiple function <i>i</i> LED	Solid blue: UID activated locally Blinking blue: UID activated remotely Solid red: Overheat condition Blinking red at 0.25Hz: Power supply failure Blinking red at 1Hz: Fan failure

Note: NIC1 and NIC2 LED in the table above are specific to the AIOM only. If it is a dual port AIOM, then NIC1 LED is controlled by one port and NIC2 LED is controlled by the other port.

Table 4: Drive LEDs

Top LED (Activity)	Blinking blue or green: I/O activity Off: Idle
Bottom LED (Status)	Red solid: Failed drive On five seconds after boot: Power on

Back Panel LED Indication




Figure 8: Back Panel LEDs

Table 5: Back Panel Controls and Indicators

Name	LED status	Function
IPMI, Link LED (Left)	Solid green	100 Mbps
	Solid amber	1 Gbps
IPMI, Activity LED (Right)	Blinking amber	Active
AIO, Link LED (Left)	Off	No link
	Solid green	Linked at 10 Gbps
	Solid amber	Linked at 1 Gbps
AIO, Activity LED (Right)	Off	No activity
	Blinking green	Link up (traffic flowing)
Locator LED (UID)	Blinking blue	Node identified

Table 6: Power Supply LED Indicators

LED status	Power supply condition
Off	No AC power to all power supplies

LED status	Power supply condition
Solid amber	Power supply critical events that cause a shutdown: Failure, Over Current Protection, Over Voltage Protection, Fan Fail, Over Temperature Protection, Under Voltage Protection.
Blinking amber (1 Hz)	Power supply warning events. Power supply continues to operate. High temperature, over voltage, under voltage and other conditions.
Blinking green (1 Hz) minute	When AC is present only: 12VSB on (PS off) or PS in sleep state
Solid green	Output on and OK
Solid amber	AC cord unplugged
Blinking green (2 Hz)	Power supply firmware updating mode

For add-on NIC LED indications, see [Network Card LED Description](#) on page 16.

Network Card LED Description

Different NIC manufacturers use different LED colors and blink states. Not all NICs are supported for every Nutanix platform. See the system specifications for your platform to verify which NICs are supported.

Table 7: SuperMicro NICs

NIC	Link (LNK) LED	Activity (ACT) LED
Dual/Quad Port 25GbE	Green: 25 GbE Amber: Less than 25 GbE Off: No link	Blinking: Activity Off: No activity
Dual Port 100GbE	Green: 100 GbE Amber: Less than 100 GbE	Blinking: Activity Off: No activity

Table 8: Intel NICs

NIC	Link (LNK) LED	Activity (ACT) LED
Quad port 10G BaseT	Green: 10 Gbps Yellow: 5/2.5/1Gbps Off: 100 Mbps	Blinking green: Transmitting or receiving data Off: No link
Quad port 10G SFP+	Green: 10 Gbps Yellow: 1 Gbps	Blinking: Activity Off: No activity
Dual Port 25G	Green: 10 Gbps Amber: 10 Gbps	Green: SFP LAN port active

Table 9: Broadcom NICs

NIC	Link (LNK) LED	Activity (ACT) LED
Dual Port 25G	Green: Linked at 25 Gb/s Yellow: Linked at 10 Gb/s or 1 Gb/s Off: No link	Blinking green: Traffic flowing activity Off: No activity
Dual Port 10G	Green: Linked at 10 Gb/s Amber: Linked at 1 Gb/s Off: No link	Blinking green: Link up (traffic flowing) Off: No activity

Table 10: Mellanox NICs

NIC	Bi-color LED (Yellow/Green)	Single color LED (Green)	Indicates
Dual port CX-6 25G	1 Hz blinking yellow	Off	Beacon command for locating the adapter card
Dual port CX-6 100G	4 Hz blinking yellow	On	Error with the link
	Blinking green	Blinking	Physical activity
	Solid green	On	Link up

Power Supply Unit Redundancy and Node Configuration

Note: Carefully plan your AC power source needs, especially in cases where the cluster consists of mixed models. Nutanix recommends that you use a 208 V - 240 V AC power source to ensure Power Supply Unit (PSU) redundancy.

Table 11: PSU Redundancy and Node Configuration

Nutanix model	Number of nodes	Redundancy at 110 V	Redundancy at 208-240 V
NX-1065-G9	1, 2, or 3	No	Yes
NX-3035-G9	1 or 2	No	Yes
NX-3060-G9	1, 2, 3, or 4	No	Yes
NX-1150S-G9	1	Yes	Yes
NX-1175S-G9	1	Yes	Yes
NX-3155-G9	1	No	Yes
NX-8150-G9	1	No	Yes

Nutanix model	Number of nodes	Redundancy at 110 V	Redundancy at 208-240 V
NX-8155-G9	1	No	Yes
NX-8155A-G9	1	No	Yes
NX-8170-G9	1	No	Yes
NX-9151-G9	1	Not supported	Yes (2+1 PSU redundancy) Two PSUs must remain functional at all times.

Caution:

For all G9 platforms except NX-1175S-G9 and NX-1150S-G9: When the input power source is 110 V, a single PSU failure will cause all nodes in a block to power off.

For NX-1175S-G9 and NX-1150S-G9: The block can tolerate a single PSU failure when connected to a 110 V input power source.

MEMORY CONFIGURATIONS

Supported Memory Configurations

DIMM installation information for all Nutanix G9 Hyper platforms.

DIMM Restrictions

DIMM capacity

Each G9 node must contain only DIMMs of the same capacity. For example, you cannot mix 32 GB DIMMs and 64 GB DIMMs in the same node.

DIMM speed

G9 platforms that use Intel Sapphire Rapids processors support 4800 MT/s DIMMs. The speed of the CPU / DIMM interface is 4000 MT/s or 4800 MT/s based on the CPU SKU used.

G9 platforms that use Intel Emerald Rapids processors ship with 5600 MT/s DIMMs. The speed of the CPU / DIMM interface depends on the CPU class and on whether you have installed one DIMM per memory channel (1DPC) or two DIMMs per memory channel (2DPC).

- Platinum-8xxx: max 5600 MT/s at 1DPC; max 4400 MT/s at 2DPC.
- Gold-6xxx: max 5200 MT/s at 1DPC; max 4400 MT/s at 2DPC.
- Gold-5xxx: max 4800 MT/s at 1DPC; max 4400 MT/s at 2DPC.
- Silver-4xxx: max 4400 MT/s at both 1DPC and 2DPC.
- Bronze-3xxx (single socket boards only): max 4400 MT/s at both 1DPC and 2DPC.

128 GB 5600 MT/s DIMMs require BIOS Ex30.001 or later. 128 GB 5600 MT/s DIMMs cannot mix in the same node with 4800 MT/s DIMMs.

If you install a 5600 MT/s DIMM in a G9 platform that uses Intel Sapphire Rapids processors, it runs at a max of 4800 MT/s.

DIMM manufacturer

Mixing of DIMMs from different manufacturers in a node is not allowed. You must use DIMMs from the same manufacturer in a node.

Memory Installation Order for Single-Node Hyper G9 Platforms

A *memory channel* is a group of DIMM slots.

For G9 single-node hyper platforms, each CPU is associated with eight memory channels. Each memory channel contains two DIMM slots, one blue slot and one black slot, for a total of 32 DIMM slots.




Figure 9: DIMM Slots for a Single-Node Hyper G9 Motherboard

Table 12: DIMM Installation Order for Single-Node Hyper G9 Platforms

Number of DIMMs	Slots to Use	Supported Capacities
4	CPU1: A1, G1 (blue slots) CPU2: A1, G1 (blue slots)	32 GB, 64 GB, 128 GB

Number of DIMMs	Slots to Use	Supported Capacities
8	CPU1: A1, C1, E1, G1 (blue slots) CPU2: A1, C1, E1, G1 (blue slots)	32 GB, 64 GB, 128 GB
12	CPU1: A1, C1, D1, E1, F1, G1 (blue slots) CPU2: A1, C1, D1, E1, F1, G1 (blue slots)	32 GB, 64 GB, 128 GB
16	CPU1: A1, B1, C1, D1, E1, F1, G1, H1 (blue slots) CPU2: A1, B1, C1, D1, E1, F1, G1, H1 (blue slots)	32 GB, 64 GB, 128 GB
24	CPU1: A1, B1, C1, D1, E1, F1, G1, H1 (blue slots) CPU1: A2, C2, E2, G2 (black slots) OR B2, D2, F2, H2 (black slots) CPU2: A1, B1, C1, D1, E1, F1, G1, H1 (blue slots) CPU2: A2, C2, E2, G2 (black slots) OR B2, D2, F2, H2 (black slots)	32 GB, 64 GB, 128 GB
32	Fill all slots.	32 GB, 64 GB, 128 GB

NUTANIX HARDWARE NAMING CONVENTION

Every Nutanix block has a unique name based on the standard Nutanix naming convention.

The Nutanix hardware model name uses the format *prefix-body-suffix*.

For all Nutanix platforms, the prefix is NX to indicate that the platform is sold directly by Nutanix and support calls are handled by Nutanix.

body uses the format *ABCD* | *Y*. The following table describes the body values.




Figure 10: Nutanix Hardware Naming Convention

Table 13: Body

Body	Description
A	Indicates the product series and is one of the following values. 1 – Entry-level and ROBO 3 – Balanced compute and storage 8 – High performance 9 – Accelerated system

Body	Description
<i>B</i>	Indicates the number of nodes. - For single-node platforms, <i>B</i> is always 1. - For multi-node platforms, <i>B</i> can be 1, 2, 3, or 4. Note: For multi-node platforms, the documentation always uses a generic zero for <i>B</i>.
<i>C</i>	Indicates the chassis form factor and is one of the following values. 1 – 1U1N (one rack unit high with one node) 3 – 2U2N (two rack units high with two nodes) 5 – 2U1N (two rack units high with one node) 6 – 2U4N (two rack units high with four nodes) 7 – 1U1N (one rack unit high with one node)
<i>D</i>	Indicates the drive form factor and is one of the following values. 0 – 2.5 inch drives 1 – E1.S drives 3 – E3.S drives 5 – 3.5 inch drives
<i>Y</i>	Indicates platform types, and takes one of the following values: - S – Single socket - G – GPU (Not used in G9 since GPU is available on multiple models) - A – AMD

Table 14: Suffix

Suffix	Description
G5	The platform uses the Intel Broadwell CPU.
G6	The platform uses the Intel Skylake CPU.
G7	The platform uses the Intel Cascade Lake CPU.
G8 or N-G8	The platform uses the Intel Ice Lake CPU.
G9	The platform uses the Intel Sapphire Rapids, Emerald Rapids or AMD Genoa CPU.

COPYRIGHT

Copyright 2025 Nutanix, Inc.

Nutanix, Inc.

1740 Technology Drive, Suite 150

San Jose, CA 95110

All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. Nutanix and the Nutanix logo are registered trademarks of Nutanix, Inc. in the United States and/or other jurisdictions. All other brand and product names mentioned herein are for identification purposes only and may be trademarks of their respective holders.