

ESPECIFICAÇÕES TÉCNICAS – CARACTERÍSTICAS DOS ITENS

1. Licenciamento de software de	Referência
1.1. Ser fornecida licença de software para gerenciamento centralizado avançado de 1 (um) ou mais cluster, facilitando a tarefa de administração diária dos clusters localizados localmente ou distantes geograficamente. A	https://portal.nutanix.com/page/documents/details/?targetId=Prism-Central-Guide-Prism-v5_15:mul-pc-overview-c.html#nconcept_vjc_5y1_nfb
1.1.1. Saúde dos Sistema clusters;	https://portal.nutanix.com/page/documents/details/?targetId=Prism-Central-Guide-Prism-v5_15:mul-pc-overview-c.html#nconcept_vjc_5y1_nfb
1.1.2. Máquinas Virtuais;	https://portal.nutanix.com/page/documents/details/?targetId=Prism-Central-Guide-Prism-v5_15:mul-pc-overview-c.html#nconcept_vjc_5y1_nfb
1.1.3. Armazenamento;	https://portal.nutanix.com/page/documents/details/?targetId=Prism-Central-Guide-Prism-v5_15:mul-pc-overview-c.html#nconcept_vjc_5y1_nfb
1.1.4. Situação do Hardware;	https://portal.nutanix.com/page/documents/details/?targetId=Prism-Central-Guide-Prism-v5_15:mul-pc-overview-c.html#nconcept_vjc_5y1_nfb
1.1.5. Painel de Análise de Performance;	https://portal.nutanix.com/page/documents/details/?targetId=Prism-Central-Guide-Prism-v5_15:mul-pc-overview-c.html#nconcept_vjc_5y1_nfb
1.1.6. Painel de Alertas e Eventos.	https://portal.nutanix.com/page/documents/details/?targetId=Prism-Central-Guide-Prism-v5_15:mul-pc-overview-c.html#nconcept_vjc_5y1_nfb
1.2. O licenciamento será por entidade gerenciada;	https://portal.nutanix.com/page/documents/details/?targetId=Prism-Central-Guide-Prism-v5_15:mul-pc-overview-c.html#nconcept_vjc_5y1_nfb

1.2.1. Ter uma ferramenta de planejamento de capacidade disponível, de forma a permitir a análise dos recursos e indicar máquinas virtuais subdimensionadas, superdimensionadas e inativas, para que seja possível identificação e remediação/otimização através da própria interface de gerenciamento. A ferramenta de planejamento de capacidade deve permitir
1.2.2. Oferecer funcionalidade de planejamento de capacidade para crescimento baseado na carga de trabalho
1.2.3. Permitir a customização dos painéis de informação;
1.2.4. Possuir funcionalidade de busca (<i>search</i>) que suporte busca contextualizada;
1.2.5. Possuir um algoritmo que determina o desvio padrão de recursos e situações (anomalias) e que são desta forma gatilhos
1.2.6. Possibilitar a criação e customização de relatório de consumo sobre os recursos da infraestrutura;
1.2.7. Oferecer a configuração de agendamento e envio dos relatórios via email.
2. Licenciamento de software de
2.1.1. Ser fornecida licença de rede de software para servidor de hiperconvergência
2.1.2. Ser baseada em, ou compatível com Open Virtual Switch (OVS);

https://portal.nutanix.com/page/documents/details/?targetId=Prism-Central-Guide-Prism-v5_15:mul-resource-planning-pc-c.html#nconcept_bcx_ldr_mx
https://portal.nutanix.com/page/documents/details/?targetId=Prism-Central-Guide-Prism-v5_15:mul-resource-planning-pc-c.html#nconcept_bcx_ldr_mx
https://portal.nutanix.com/page/documents/details/?targetId=Prism-Central-Guide-Prism-v5_15:mul-dashboard-widgets-pc-r.html#nref_gdf_wzw_c5
https://portal.nutanix.com/page/documents/details/?targetId=Prism-Central-Guide-Prism-v5_15:mul-search-options-pc-r.html
https://portal.nutanix.com/page/documents/details/?targetId=Prism-Central-Guide-Prism-v5_15:mul-behavioral-learning-pc-c.html#nconcept_icv_zgz_1bb
https://portal.nutanix.com/page/documents/details/?targetId=Prism-Central-Guide-Prism-v5_15:ssp-report-management-ssp-pc-c.html#nconcept_yym_c5f_f1b
https://portal.nutanix.com/page/documents/details/?targetId=Prism-Central-Guide-Prism-v5_15:ssp-report-management-ssp-pc-c.html#nconcept_yym_c5f_f1b
https://portal.nutanix.com/page/documents/solutions/details/?targetId=TN-2094-Flow:TN-2094-Flow
https://portal.nutanix.com/page/documents/solutions/details/?targetId=TN-2094-Flow:TN-2094-Flow

2.1.3.	Suportar microssegmentação para prover controle granular e governança de todo o tráfego de entrada e saída de uma máquina

2.1.4. Permitir a associação de políticas de rede a VMs e aplicativos em vez de segmentos de rede específicos (por exemplo

2.1.5. Prover visualização de todo tráfego e relacionamentos com a descoberta

2.1.6. Prover uma estrutura de segurança orientada por políticas que inspeciona o

2.1.6.1. As políticas de segurança inspecionam o tráfego originado e terminado dentro de um datacenter, ajudando a eliminar

2.1.6.2. Utilizar uma abordagem centrada na carga de trabalho em vez de uma abordagem centrada na rede, permitindo examinar o tráfego de, e para as VMS,

2.1.7. Prover uma abordagem agnóstica a estrutura de rede, centrada na carga de trabalho, permitindo que a equipe de virtualização implemente essas políticas de

2.1.8. As políticas de segurança deverão ser aplicadas às categorias (um agrupamento lógico de VMs) e não às próprias VMs, não importando quantas VMs são inicializadas em uma determinada categoria. O tráfego

2.1.9. A interface de gerenciamento deve oferecer uma abordagem baseada em visualização para configurar políticas e

2.1.9.1.	Política de Segurança de Aplicação: quando for necessário proteger um aplicativo especificando origens e destinos de

<https://portal.nutanix.com/page/documents/solutions/details/?targetId=TN-2094-Flow:TN-2094-Flow>

<https://portal.nutanix.com/page/documents/solutions/details/?targetId=TN-2094-Flow>: TN-2094-Flow

<https://portal.nutanix.com/page/documents/solutions/details/?targetId=TN-2094-Flow>: TN-2094-Flow

<https://portal.nutanix.com/page/documents/solutions/details/?targetId=TN-2094-Flow>: TN-2094-Flow

<https://portal.nutanix.com/page/documents/solutions/details/?targetId=TN-2094-Flow>:TN-2094-Flow

<https://portal.nutanix.com/page/documents/solutions/details/?targetId=TN-2094-Flow>:TN-2094-Flow

<https://portal.nutanix.com/page/documents/solutions/details/?targetId=TN-2094-Flow>: TN-2094-Flow

<https://portal.nutanix.com/page/documents/solutions/details/?targetId=TN-2094-Flow>: TN-2094-Flow

<https://portal.nutanix.com/page/documents/solutions/details/?targetId=TN-2094-Flow>:TN-2094-Flow

<https://portal.nutanix.com/page/documents/solutions/details/?targetId=TN-2094-Flow>:TN-2094-Flow

2.1.9.2.	Política de Isolamento do Ambiente: quando for necessário bloquear todo o tráfego, independentemente da direção, entre dois grupos de VMS

2.1.9.3. Política de Quarentena: quando for necessário isolar uma VM comprometida ou infectada e, opcionalmente, desejar

2.1.10. Garantir que seja apenas permitido o tráfego entre camadas de aplicativos ou outros limites lógicos, garantindo a proteção contra ameaças

2.1.11.	Permitir a atualização automática durante todo o ciclo de vida da VM, eliminando a carga do gerenciamento de

2.1.12.	Permitir categorizar as Máquinas Virtuais de forma a permitir a criação políticas de segurança com no mínimo

2.1.12.1. Isolar o tráfego de dados entre Máquinas Virtuais de Diferentes categorias;

2.1.12.2. Isolar o tráfego de dados de Máquinas Virtuais específicas para modo de quarentena, tanto forense quanto restrita, de forma a prover uma rápida reação ao time de infraestrutura em caso de Máquinas Virtuais

<https://portal.nutanix.com/page/documents/solutions/details/?targetId=TN-2094-Flow>: TN-2094-Flow

<https://portal.nutanix.com/page/documents/solutions/details/?targetId=TN-2094-Flow>:TN-2094-Flow

<https://portal.nutanix.com/page/documents/solutions/details/?targetId=TN-2094-Flow>: TN-2094-Flow

<https://portal.nutanix.com/page/documents/solutions/details/?targetId=TN-2094-Flow:TN-2094-Flow>

<https://portal.nutanix.com/page/documents/solutions/details/?targetId=TN-2094-Flow:TN-2094-Flow>

<https://portal.nutanix.com/page/documents/solutions/details/?targetId=TN-2094-Flow:TN-2094-Flow>

<https://portal.nutanix.com/page/documents/solutions/details/?targetId=TN-2094-Flow:TN-2094-Flow>

2.1.12.3. Mapear o tráfego de entrada, entre as camadas e de saída de aplicações, permitindo ao administrador determinar quais servidores tem acesso de entrada na aplicação, o tipo de protocolo e o número da porta que o fluxo de dados pode ocorrer, permitir ou restringir também o fluxo de dados entre as camadas, máquinas virtuais, pertencentes à aplicação, através da
2.1.13. Possuir integração com software de terceiros que permita o redirecionamento do tráfego das VMs para ferramentas terceiras, como por exemplo, mas não limitado a softwares d
3. Licença de software para estrutura de
3.1.1. Possibilizar a entrega de serviço de arquivos compartilhados dos tipos SMB e
3.1.2. Permitir integração com o <i>Nutanix Operating System</i> para criar serviços
3.1.3. Suportar funcionalidades de balanceamento de carga, replicação nativa e
3.1.3.1. Diretório de arquivos do usuário, onde apenas o dono do perfil terá acesso
3.1.3.2. Diretórios departamentais que poderão ser acessados por qualquer usuário com as devidas permissões (<i>Access Based</i>
3.1.4. Suportar tierização, compressão, deduplicação e erasure-coding;
3.1.5. Suportar o protocolo - ICAP (<i>Internet Content Adaptation Protocol</i>);

https://portal.nutanix.com/page/documents/solutions/details/?targetId=TN-2094-Flow:TN-2094-Flow
https://portal.nutanix.com/page/documents/solutions/details/?targetId=TN-2094-Flow:TN-2094-Flow
https://portal.nutanix.com/page/documents/details/?targetId=Files-v3_6:Files-v3_6
https://portal.nutanix.com/page/documents/details/?targetId=Files-v3_6:Files-v3_6
https://portal.nutanix.com/page/documents/details/?targetId=Files-v3_6:file-server-load-balance-c.html
https://portal.nutanix.com/page/documents/details/?targetId=Files-v3_6:file-server-create-share-export-c.html#nconcept_wyw_2zp_scb
https://portal.nutanix.com/page/documents/details/?targetId=Files-v3_6:file-server-enable-abe-wc-t.html
https://portal.nutanix.com/page/documents/details/?targetId=Files-v3_6:file-share-compression-c.html
https://portal.nutanix.com/page/documents/details/?targetId=Files-v3_6:file-server-anti-virus-c.html#nconcept_on4_xnh_n1b

4.1.5. Possibilidade de versionamento, gerenciamento de ciclo de vida.
5. Expansão de suporte e garantia -
5.1.1. A garantia e suporte técnico deverão ser estendidas por mais 3 (três) anos, pelo fabricante, na modalidade 24 (horas) x 7 (dias por semana) x 365 (dias por anos) e troca de peças no próximo dia útil;
5.1.2. A renovação deverá abranger o <i>hardware</i> e <i>software</i> existentes em cada servidor;
5.1.3. Os chamados técnicos serão categorizados em níveis de severidade, conforme a seguir:
5.1.3.1. Nível 1: Equipamento total ou parcialmente inoperante, com degradação de desempenho ou ocorrência de mau
5.1.3.2. Nível 2: Equipamento operante com ocorrência de alarmes; consultas gerais sobre administração, configuração,
5.1.4. Os chamados técnicos deverão ser atendidos nos seguintes prazos:

https://portal.nutanix.com/page/documents/details/?targetId=Objects-v21:Objects-v21
Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications
Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications
Conforme Proposta Comercial. https://www.nutanix.com/support-services/product-support/product-support-programs
Conforme Proposta Comercial. https://www.nutanix.com/support-services/product-support/product-support-programs
Conforme Proposta Comercial. https://www.nutanix.com/support-services/product-support/product-support-programs
Conforme Proposta Comercial. https://www.nutanix.com/support-services/product-support/product-support-programs

5.1.4.1. Nível 1: Até 2h (duas horas) para o início do atendimento;
5.1.4.2. Nível 2: Até 24h (vinte e quatro horas) para o início do atendimento.
5.1.5. O atendimento de um chamado técnico será considerado solucionado conforme os seguintes critérios:
5.1.5.1. Nível 1: O equipamento estiver operante, sem degradação de desempenho ou ocorrência de mau funcionamento;
5.1.5.2. Nível 2: O equipamento estiver operante sem ocorrência de alarmes e as consultas realizadas tenham sido
5.1.6. Cada servidor possui as seguintes características:
5.1.6.1. Dois processadores Intel(R) Xeon(R) v3 ou v4 com 10 (dez) cores cada;
5.1.6.2. Memória RAM de 256GB (duzentos e cinquenta e seis);

Conforme Proposta Comercial. https://www.nutanix.com/support-services/product-support/product-support-programs
Conforme Proposta Comercial. https://www.nutanix.com/support-services/product-support/product-support-programs
Conforme Proposta Comercial. https://www.nutanix.com/support-services/product-support/product-support-programs
Conforme Proposta Comercial. https://www.nutanix.com/support-services/product-support/product-support-programs
Conforme Proposta Comercial. https://www.nutanix.com/support-services/product-support/product-support-programs
Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications
Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications
Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications

5.1.6.3. Até 1 (um) disco SSD de 960GB padrão SATA de 6.0 Gb/s e hot-swap;
5.1.1.1. Até 2 (dois) discos HDD de 8TB cada padrão SATA de 6.0 Gb/s e hot-swap;
5.1.1.2. 2 (duas) interfaces de rede de 10Gbps SFP+;
5.1.1.3. Licenciamento Pro ou Ultimate.
6. Expansão de suporte e garantia -

Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications
Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications
Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications
Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications , https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=Licensing-Guide:lic-licensing-prism-element-c.html

6.1.5.1.	Nível 1: O equipamento estiver operante, sem degradação de desempenho ou ocorrência de mau funcionamento;
6.1.5.2.	Nível 2: O equipamento estiver operante sem ocorrência de alarmes e as consultas realizadas tenham sido
6.1.6.	Cada servidor possui as seguintes características:
6.1.6.1.	Dois processadores Intel(R) Xeon(R) v3 ou v4 com 12 (doze) cores cada;
6.1.6.2.	Memória RAM de 512GB (duzentos e cinquenta e seis);
6.1.6.3.	Até 2 (dois) disco SSD de 1200GB padrão SATA de 6.0 Gb/s e hot-swap;
6.1.6.4.	Até 4 (dois) discos HDD de 8TB cada padrão SATA de 6.0 Gb/s e hot-swap;

Conforme Proposta Comercial. https://www.nutanix.com/support-services/product-support/product-support-programs
Conforme Proposta Comercial. https://www.nutanix.com/support-services/product-support/product-support-programs
Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications
Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications
Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications
Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications
Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications

6.1.6.5.	2 (duas) interfaces de rede de 10Gbps SFP+;
6.1.6.6.	Licenciamento Pro ou Ultimate.
7. Expansão de suporte e garantia -	
7.1.1.	A garantia e suporte técnico deverão ser estendidas por mais 3 (três) anos, pelo fabricante, na modalidade 24 (horas) x 7 (dias por semana) x 365 (dias por anos) e troca de peças no próximo dia útil;
7.1.2.	A renovação deverá abranger o <i>hardware</i> e <i>software</i> existentes em cada servidor;
7.1.3.	Os chamados técnicos serão categorizados em níveis de severidade, conforme a seguir:
7.1.3.1.	Nível 1: Equipamento total ou parcialmente inoperante, com degradação de desempenho ou ocorrência de mau

Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications
Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications , https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=Licensing-Guide:lic-licensing-prism-element-c.html
Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications
Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications
Conforme Proposta Comercial. https://www.nutanix.com/support-services/product-support/product-support-programs
Conforme Proposta Comercial. https://www.nutanix.com/support-services/product-support/product-support-programs

7.1.3.2. Nível 2: Equipamento operante com ocorrência de alarmes; consultas gerais sobre administração, configuração,
7.1.4. Os chamados técnicos deverão ser atendidos nos seguintes prazos:
7.1.4.1. Nível 1: Até 2h (duas horas) para o início do atendimento;
7.1.4.2. Nível 2: Até 24h (vinte e quatro horas) para o início do atendimento.
7.1.5. O atendimento de um chamado técnico será considerado solucionado Conforme os seguintes critérios:
7.1.5.1. Nível 1: O equipamento estiver operante, sem degradação de desempenho ou ocorrência de mau funcionamento;
7.1.5.2. Nível 2: O equipamento estiver operante sem ocorrência de alarmes e as consultas realizadas tenham sido
7.1.6. Cada servidor possui as seguintes características:
7.1.6.1. Dois processadores Intel(R) Xeon(R) v3 ou v4 com 8 (oito) cores cada;

Conforme Proposta Comercial. https://www.nutanix.com/support-services/product-support/product-support-programs
Conforme Proposta Comercial. https://www.nutanix.com/support-services/product-support/product-support-programs
Conforme Proposta Comercial. https://www.nutanix.com/support-services/product-support/product-support-programs
Conforme Proposta Comercial. https://www.nutanix.com/support-services/product-support/product-support-programs
Conforme Proposta Comercial. https://www.nutanix.com/support-services/product-support/product-support-programs
Conforme Proposta Comercial. https://www.nutanix.com/support-services/product-support/product-support-programs
Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications
Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications

7.1.6.2. Memória RAM de 64GB (sessenta e quatro);	Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications
7.1.6.3. Até 1 (um) disco SSD de 800GB padrão SATA de 6.0 Gb/s e hot-swap;	Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications
7.1.6.4. Até 5 (cinco) discos HDD de 8TB (oito) cada padrão SATA de 6.0 Gb/s e hot-swap;	Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications
7.1.6.5. 2 (duas) interfaces de rede de 10Gbps SFP+;	Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications
7.1.6.6. Licenciamento Pro ou Ultimate.	Conforme Proposta Comercial. https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G4-Multinode:System%20Specifications , https://portal.nutanix.com/page/documents/details/?targetId=System-Specs-G5-Multinode:System%20Specifications e https://portal.nutanix.com/page/documents/details/?targetId=Licensing-Guide:lic-licensing-prism-element-c.html

Nutanix Files

Nutanix Tech Note

Version 5.2 • October 2019 • TN-2041

Copyright

Copyright 2019 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 is a trademark of Nutanix, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

Contents

- 1. Executive Summary.....5
- 2. Audience and Purpose.....6
- 3. Nutanix Enterprise Cloud Overview..... 7
 - 3.1. Nutanix Acropolis Architecture.....8
- 4. Nutanix Files Architecture..... 9
 - 4.1. Nutanix Files: File Server Virtual Machine..... 9
 - 4.2. Exports and Shares..... 13
 - 4.3. Load Balancing and Scaling.....20
 - 4.4. High Availability.....22
 - 4.5. Active Directory and SMB Operations.....26
 - 4.6. Active Directory, LDAP, and NFS Operations.....32
 - 4.7. Multiprotocol..... 35
 - 4.8. Quotas.....39
 - 4.9. Selective File Blocking.....41
 - 4.10. File Analytics.....42
 - 4.11. Hypervisor-Specific Support.....47
- 5. Backup and Disaster Recovery.....48
 - 5.1. Self-Service Restore..... 48
 - 5.2. Protection Domains and Consistency Groups.....49
 - 5.3. Cluster Migration, Failure, and Restoration.....49
 - 5.4. Cloning..... 50
- 6. Third-Party Integration..... 51
 - 6.1. Antivirus..... 51
 - 6.2. File Operations Monitoring.....54
 - 6.3. Intelligent Backup.....55



7. Conclusion.....59

Appendix.....60

 About Nutanix.....60

List of Figures.....61

List of Tables.....64



1. Executive Summary

Nutanix Files is a software-defined, scale-out file storage solution that provides a repository for unstructured data, such as home directories, user profiles, departmental shares, application logs, backups, and archives. Flexible and responsive to workload requirements, Files is a fully integrated, core component of the Nutanix Enterprise Cloud.

You can deploy Nutanix Files on an existing cluster or a standalone cluster. Unlike standalone NAS appliances, Files consolidates VM and file storage, eliminating the need to create an infrastructure silo. Administrators can manage Files with Nutanix Prism, just like VM services, unifying and simplifying management. Integration with Active Directory enables support for quotas and access-based enumeration, as well as self-service restores with the Windows Previous Versions feature. Nutanix Files also supports file server cloning, which lets you back up Files off-site and run antivirus scans and machine learning without affecting production.

Nutanix Files can run on a dedicated cluster or be colocated on a cluster running user VMs. Nutanix supports Files with both ESXi and AHV. Files includes native high availability and uses the Acropolis Distributed Storage Fabric (DSF) for intracluster data resiliency and intercluster asynchronous disaster recovery. The DSF also provides data efficiency techniques such as erasure coding (EC-X) and compression.

2. Audience and Purpose

This tech note is part of the Nutanix Solutions Library. We wrote it for architects and systems engineers who want to use Nutanix Files as a simple way to deliver user and group file management. This document describes how to implement and operate Files in your datacenter.

We cover the following subject areas:

- Overview of the Nutanix architecture with Files.
- Load balancing of standard and distributed shares (SMB) and exports (NFS).
- High availability.
- Backup and recovery.
- Quotas and permission management.
- Antivirus.

Table 1: Document Version History

Version Number	Published	Notes
1.0	December 2016	Original publication.
1.1	May 2017	Updated Backup and Disaster Recovery section.
1.2	September 2017	Updated for version 2.2 features.
2.0	February 2018	Updated for version 3.0.
2.1	April 2018	Solution overview update.
2.2	August 2018	SMB share and NFS export updates.
3.0	October 2018	Updated for version 3.1 and updated product naming.
4.0	January 2019	Updated for version 3.2 and AOS 5.9 and later.
5.0	April 2019	Updated for version 3.5 and AOS 5.10 and later.
5.1	August 2019	Updated for version 3.5.1.
5.2	October 2019	Updated for version 3.6.

3. Nutanix Enterprise Cloud Overview

Nutanix delivers a web-scale, hyperconverged infrastructure solution purpose-built for virtualization and cloud environments. This solution brings the scale, [resilience](#), and economic benefits of web-scale architecture to the enterprise through the Nutanix Enterprise Cloud Platform, which combines three product families—Nutanix Acropolis, Nutanix Prism, and Nutanix Calm.

Attributes of this Enterprise Cloud OS include:

- Optimized for storage and compute resources.
- Machine learning to plan for and adapt to changing conditions automatically.
- Self-healing to tolerate and adjust to component failures.
- API-based automation and rich analytics.
- Simplified one-click upgrade.
- Native file services for user and application data.
- Native backup and disaster recovery solutions.
- Powerful and feature-rich virtualization.
- Flexible software-defined networking for visualization, automation, and security.
- Cloud automation and life cycle management.

Nutanix Acropolis provides data services and can be broken down into three foundational components: the Distributed Storage Fabric (DSF), the App Mobility Fabric (AMF), and AHV. Prism furnishes one-click infrastructure management for virtual environments running on Acropolis. Acropolis is hypervisor agnostic, supporting two third-party hypervisors—ESXi and Hyper-V—in addition to the native Nutanix hypervisor, AHV.




Figure 1: Nutanix Enterprise Cloud

3.1. Nutanix Acropolis Architecture

Acropolis does not rely on traditional SAN or NAS storage or expensive storage network interconnects. It combines highly dense storage and server compute (CPU and RAM) into a single platform building block. Each building block delivers a unified, scale-out, shared-nothing architecture with no single points of failure.

The Nutanix solution requires no SAN constructs, such as LUNs, RAID groups, or expensive storage switches. All storage management is VM-centric, and I/O is optimized at the VM virtual disk level. The software solution runs on nodes from a variety of manufacturers that are either all-flash for optimal performance, or a hybrid combination of SSD and HDD that provides a combination of performance and additional capacity. The DSF automatically tiers data across the cluster to different classes of storage devices using intelligent data placement algorithms. For best performance, algorithms make sure the most frequently used data is available in memory or in flash on the node local to the VM.

To learn more about the Nutanix Enterprise Cloud, please visit [the Nutanix Bible](#) and [Nutanix.com](#).

4. Nutanix Files Architecture

Nutanix Files is a scale-out approach that provides Server Message Block (SMB) and Network File System (NFS) file services to clients. Nutanix Files server instances are composed of a set of VMs (called FSVMs). Files requires at least three FSVMs running on three nodes to satisfy a quorum for high availability.




Figure 2: Nutanix Files Server Instances Run as VMs for Isolation from the DSF

4.1. Nutanix Files: File Server Virtual Machine

The File Server VM (FSVM) is based on CentOS and incorporates all the security and hardening that goes into the Nutanix Controller VM (CVM). All the FSVMs have the same configuration, starting with four vCPUs and 12 GiB of RAM. You can add more vCPUs, RAM, and FSVMs to the cluster. For each file server the number of FSVMs must be less than or equal to the number of nodes in the Nutanix cluster; however, you can create multiple file server deployments if needed. Each Nutanix Files cluster can support up to 16 FSVMs. Nutanix Files 3.1 introduced single FSVM deployments intended for one- and two-node Nutanix clusters. You can also have single FSVM deployments for larger clusters with 5.10.1 or later.




Figure 3: Data Path Architecture of Nutanix Files

Nutanix Files can support SMB and NFS from the same FSVM, and with Files 3.5 you can enable simultaneous SMB and NFS access to the same share and data, commonly referred to as multiprotocol support. Both SMB and NFS share a common library, allowing a modular approach. InsightsDB is a NoSQL database that maintains the statistics and alerts for Files. Zookeeper is a centralized service that maintains configuration information, such as domain, share or export, and IP information. The Minerva NVM Service talks to the local CVM and sends heartbeats to share health information and to help with failover.

Each FSVM stores file server data on multiple file systems that store share-specific data. The individual file system provides the snapshot capability used to provide Windows Previous Versions (WPV) support to clients. By using separate file systems for each share or export, Nutanix Files can scale to support billions of files in one cluster.




Figure 4: FSVM Internal Communication on One Node

The above diagram shows one FSVM running on a node, but you can put multiple FSVMs on a node for multitenancy.

Networking

The FSVM has two network interfaces: the storage interface and the client interface. The FSVM service that talks to the CVMs uses the storage interface, which also provides access to Nutanix Volumes iSCSI vDisks in volume groups. The storage interface helps manage deployment, failover, and maintenance, and enables control over one-click upgrades. Integration with the CVM lets the FSVM determine if a storage fault has occurred and if so, whether you must take action. The FSVM service sends a heartbeat to its local CVM service each second indicating its state.



Tip: We recommend placing the FSVM storage interface on the same network VLAN as the Nutanix CVM management interface (eth0 of the CVM). Using the same VLAN helps ensure maximum performance for the iSCSI sessions.

The client interface allows clients to connect to SMB shares and NFS exports hosted on an FSVM. A client can connect to any FSVM client network interface to access their file data. If a different FSVM provides the data, the client connection automatically redirects to the correct

FSVM interface. If an FSVM fails, the client network address for the failed FSVM moves to another to preserve data access.

Storage

As shown in the following figure, each FSVM uses three separate vDisks: a 12 GiB boot disk that contains the boot image, a 45 GiB disk (/home/nutanix) that contains logs and software state, and a 45 GiB disk for Cassandra data.




Figure 5: FSVM vDisks and Volume Groups

A Nutanix Files cluster is a single namespace that includes a collection of file systems used to support the SMB and NFS shares. These file systems use Nutanix volume groups as storage. Volume groups enable both high availability and scale-out for the file systems. A volume group is a collection of logically related vDisks (or volumes) attached to the guest via iSCSI. When an FSVM is down for maintenance or if a fault occurs, one of the surviving FSVMs takes over volume group ownership and continues servicing requests.

With the 2.x release of Nutanix Files, volume groups contain six virtual disks: two disks used for metadata and four disks used for data. The four data disks represent 40 TB of available storage for a given share using that volume group.

With the 3.x release of Nutanix Files, volume groups contain 16 virtual disks. Files 3.x versions prior to version 3.2 use 2 disks for metadata and 4 disks for data, as in the 2.x release, with the remaining disks reserved for expansion in later versions. From the 3.2 release of Files onward, all disks in the volume group are available for use by the assigned shares, representing a maximum standard share size of 140 TB.



Figure 6: FSVM Volume Group with 16 vDisks

The initial file system pool created on the volume group uses six of the drives as previously noted and supports 40 TB of data.



Figure 7: Initial 40 TB File System Pool

Once this file system pool reaches 80 percent space utilization, Nutanix Files automatically expands it to 80 TB by incorporating another four data disks from the volume group.



Figure 8: 80 TB File System Pool Following First Expansion

This expansion process can occur two more times, increasing the file system capacity to 120 TB with the second expansion and 140 TB with the final expansion. Each expansion is triggered at 80 percent space utilization of the current pool.



Figure 9: 120 TB File System Pool After Second Expansion




Figure 10: 140 TB File System Pool After Final Expansion

4.2. Exports and Shares

There are two types of SMB shares and NFS exports.

SMB shares:

1. Distributed (previously called home).
2. Standard (previously called general).

NFS exports:

1. Distributed (previously called sharded).
2. Standard (previously called nonsharded).

A standard share is an SMB share or NFS export hosted by a single FSVM. A distributed share spreads the workload by distributing the hosting of top-level directories across FSVMs, which also simplifies administration.

Distributed shares and exports don't allow files in the root.

Distributed Shares

Distributed shares and exports distribute data by dividing the top-level directories across all the FSVMs that make up the file server. Nutanix Files maintains the directory mapping for each responsible FSVM using InsightsDB. FSVMs use DFS referrals for SMB and junctions for NFSv4 to make sure that clients can connect to the right top-level directories.




Figure 11: Distributed Directory Shares

Distributed share directories work well for home shares and exports because Nutanix Files automatically spreads the workload over multiple FSVMs per user (see the preceding figure, Distributed Directory Shares). If a user creates a share called “\\FileServer1\Users” that contains top-level directories \\Bob, \\Becky, and \\Kevin, \\Bob may be on FSVM-1, \\Becky on FSVM-2, \\Kevin on FSVM-3, and so on. The FSVMs use a string hashing algorithm based on the directory names to distribute the top-level directories.

This distribution can accommodate a very large number of users or directories in a single share or export. The scaling limits of more traditional designs can force administrators to create multiple shares or exports in which, for example, one set of users whose last names begin with A through M run off one controller, and users whose names begin with N through Z run off another controller. This design limitation leads to management overhead headaches and unnecessary Active Directory complexity. For these reasons, Nutanix Files expects to have one

SMB distributed share for the entire cluster. If you need to have more than one home directory share, you can create additional shares as needed.

The top-level directories act as a reparse point—essentially a shortcut. Consequently, administrators must create directories at the root of the share for optimal load balancing. Because it appears as a shortcut, we don't allow user files in the root level of the share; we recommend setting permissions at the share or export root before deploying user folders. This step allows newly created top-level directories to inherit permissions, so you don't have to adjust them after the fact using the Nutanix Files Microsoft Management Console (MMC) plug-in.

Standard shares and exports do not distribute top-level directories. A single file server always owns the files and subfolders for standard shares and exports. The following diagram illustrates two standard shares (for example, accounting and IT) on the same file server.




Figure 12: Two Standard Shares on the Same File Server

Unlike distributed shares and exports, standard shares and exports can store files in the root of the directory.

Managing NFS Distributed Share Directories

Distributed share directories with NFSv4 introduce some unique behaviors. To balance performance across FSVMs, each top-level directory you create becomes an automatically generated export. Nutanix Files mounts the export on demand when the directory is accessed. Since these are exports, rather than standard directories, it takes a few steps to remove a top-

level directory. Below we provide an example to demonstrate how distributed top-level directories behave and then walk through the process of deleting a top-level directory.

If you create a distributed share and mount it as **/projects**, you can see that mount on Linux using the **df** command.

```
# df /projects

Filesystem 1K-blocks Used Available Use% Mounted on
1.1.1.10:/projects 1073741824 839455744 234286080 79% /projects
```

Next, create some project directories and access them.

```
mkdir /projects/project1
mkdir /projects/project2
ls /projects/project1
ls /projects/project2
```

Accessing the directory, using **ls** in this case, mounts the automatically created top-level directory export. If you run **df** again you see two additional mount points.

```
# df | grep project

1.1.1.10:/projects 1073741824 839455744 234286080 79% /projects
1.1.1.11:/projects/project1 1073741824 839455744 234286080 79% /projects/project1
1.1.1.12:/projects/project2 1073741824 839455744 234286080 79% /projects/project2
```

These additional mount points allow a different FSVM to serve each export.

This behavior introduces additional steps when deleting a top-level directory. You can delete the project2 directory from any NFS client with the export mounted if you log on as a user with the appropriate permissions. There are three steps to deleting a distributed share export.

- Delete the contents of the share:

```
rm -rf /projects/project2/*
```

- Unmount the project2 share:

```
umount /projects/project2
```

- Delete the project2 directory:

```
rmdir /projects/project2
```

At this point you've deleted the top-level directory and it becomes inaccessible to other clients mounting the export. Processes that access the export after you delete it receive a Stale File Handle error.

Distributed shares with NFSv3 clients behave differently than those with NFSv4 clients. NFSv3 connections do not automatically mount the top-level directories of distributed shares. Because

the top-level directories are not submounted, you can rename and delete them without the additional steps required for NFSv4.

Nested Shares


Nutanix Files 3.2 introduced support for nested shares. A nested share allows you to create a folder within an existing standard or distributed share and turn that folder into a directly accessible share. Both SMB and NFS protocols support nested shares. Nested shares inherit some attributes from the parent share, but you can modify other attributes.

Nested share inherited attributes:

- Protocol type
- Maximum size
- Self-service restore
- Quota policy

Nested share modifiable attributes:

- SMB access-based enumeration (ABE)
- NFS authentication and access
- NFS advanced settings



Create a share/export ? X

Basics Settings Summary

Name

NestedShare

Description (Optional)

File Server

ADPAFS

Share Path (Optional) ?

/home1/tld1/newnestedshare

Cancel Next

Figure 13: Create Share Menu with Nested Share Path

vDisk and Volume Group Allocation

Distributed shares and exports begin with 5 volume groups for each FSVM (for example, 15 for a three-node cluster). Nutanix Files distributes the volume groups to different FSVMs in the file server cluster. When you have a large number of users, multiple volume groups improve load balancing across the FSVMs.




Figure 14: Distributed Share and Top-Level Directories

A single FSVM and volume group serves each standard share or export.




Figure 15: Three Standard Shares

Once you reach the limit of 10 volume groups per FSVM, new shares and exports use existing volume groups of the same type. For example, if you have deployed a distributed share (15 volume groups) and 15 standard shares (creating 15 additional volume groups) on a three-node physical cluster, each FSVM hosts 10 volume groups: 5 for the distributed share and 5 for the standard shares. In this situation, because each FSVM is serving the maximum of 10 volume groups, the next share created uses an existing volume group.

Every file server maps one-to-one to a container. This mapping allows you to manage compression and erasure coding individually for each file server deployed. Inline compression and erasure coding are turned on by default to save capacity. Deduplication is not recommended.



Tip: Use distributed shares for most of your use cases to ensure that user connections and data ownership are distributed across the FSVMs in the cluster.

From Files 3.6 onward, you can control compression on a share-by-share basis. When you create a share, compression is enabled by default.

Directory Layout

Nutanix Files can store millions of files within a single share and billions of files across a multinode cluster with multiple shares. To achieve good response times for high file and directory count environments, be sure to give some thought to directory design. Placing millions of files or directories in a single directory leads to slow file enumeration, which may impact some applications.

For performance-sensitive applications, Nutanix recommends limiting directory width—the number of files or directories in the root of a folder—to 100,000 objects. Increasing FSVM memory to cache metadata can help improve performance for environments with high object counts, where directory enumeration is common.

Nutanix also recommends limiting the number of top-level directories in distributed shares. The recommended number of top-level directories depends on the memory assigned to the FSVMs in the cluster. Your directory count should not exceed 3,000 times the FSVM memory of one node. For example, for a three-node cluster with 12 GB of memory assigned to each FSVM, the maximum top-level directory count is $3,000 \times 12 = 36,000$.


4.3. Load Balancing and Scaling

Load balancing occurs on two levels. First, a client can connect to any one of the FSVMs and users can add FSVMs as needed. Second, on the storage side, Nutanix Files can redistribute volume groups to different FSVMs for better load balancing across nodes. The following situations prompt load balancing:

1. When removing an FSVM from the cluster, Files automatically load balances all its volume groups across the remaining FSVMs.
2. During normal operation, the distribution of top-level directories becomes poorly balanced because of changing client usage patterns or suboptimal initial placement.
3. When increased user demand necessitates adding a new FSVM, its volume groups are initially empty and may require rebalancing.

Nutanix Files addresses the second and third situations by maintaining usage statistics and patterns to detect per-FSVM load (in terms of CPU and memory utilization) and per-volume

group load (in terms of user connections and latency of operations). Nutanix Files uses these statistics to make a load balancing recommendation, but the administrator must accept the recommendation before Files carries out the action. Nutanix calls this feature one-click performance optimization.



Available Actions

- ☒ **Scale Up (Recommended)**
Scale up can help performance by adding more vCPUs and Memory to the existing File Server VMs.
- ☐ **Rebalance**
Rebalance can eliminate performance bottlenecks by distributing the workload among the existing File Server VMs.
- ☐ **Scale Out**
Scale out can help performance by adding additional File Server VMs to the existing File Server.

Cancel Continue

Figure 16: Performance Optimization Recommendation

Load balancing through volume group redistribution may not always improve performance. For example, if clients target a low-level share directory that cannot be further distributed among FSVMs, performance does not improve. In such cases, Nutanix Files supports scaling up the FSVMs by adding vCPUs and memory. Scaling up is seamless to end users.

In the current implementation, there is a brief outage during volume group migration and FSVM scale-out. The file share or export requires a client reconnect after migration and scaling out. When you move a volume group, the time to live (TTL) slowly decreases to limit the disconnection from the client. Today most clients try to reconnect for 50–60 seconds, limiting the overall impact.



Note: With Files 3.6, we recommend a scale-up operation for performance optimization if SMB connection limits reach 95 percent utilization over a two-hour time window. Scale-out and rebalance recommendations are disabled with Files 3.5 and later versions. Nutanix plans to enable scale-out and rebalance recommendations in a later release.

Load balancing also occurs for each virtual disk in a volume group. Nutanix Volumes requires the administrator to configure an iSCSI data service IP. The data service IP is a highly available virtual IP address used as an iSCSI discovery portal. Each virtual disk in a volume group represents its own iSCSI target that any controller VM (CVM) in the cluster can own. Nutanix Volumes uses iSCSI redirection to place and automatically load balance these sessions as

needed. The load balancing feature for Nutanix Volumes is called Acropolis Dynamic Scheduling (ADS).




Figure 17: Load Balancing Volume Groups with Nutanix Volumes

For more detailed information, refer to the [Nutanix Volumes best practices guide](#).

4.4. High Availability

Nutanix designed Files to recover from a range of service disruptions, including when a local CVM or FSVM restarts or fails.

CVM Failure or Upgrade

If a CVM goes offline because of failure or planned maintenance, any active sessions against that CVM are disconnected, triggering the iSCSI client to log on again. The new logon occurs through the external data services IP, which redirects the session to a healthy CVM. The following figure shows this general process.




Figure 18: Nutanix Volumes Load Balancing for File Server Volume Groups

When the failed CVM returns to operation, the iSCSI session fails back. In the case of a failback, the FSVM is logged off and redirected to the appropriate CVM.

Node Failure

When a physical node fails completely, Nutanix Files uses leadership elections and the local Minerva CVM service to recover. The FSVM sends heartbeats to its local Minerva CVM service once per second, indicating its state and that it's alive. The Minerva CVM service keeps track of this information and can take action during a failover.

When an FSVM goes down, the Minerva CVM service unlocks the files from the downed FSVM and releases the external address from eth1. The downed FSVM's resources then appear on a running FSVM. The internal Zookeeper instances store this information so that they can send it to other FSVMs if necessary.




Figure 19: Each FSVM Controls Its Own Volume Groups

When an FSVM is unavailable, the remaining FSVMs volunteer for ownership of the shares and exports that were associated with the failed FSVM. The FSVM that takes ownership of the volume group informs the CVM that the volume group reservation has changed. If the FSVM that attempts to take control of the volume group is already the leader for a different volume group it has volunteered for, it relinquishes leadership for the new volume group immediately. This arrangement ensures distribution of volume groups, even if multiple FSVMs fail.




Figure 20: FSVM-1 Failure

The Nutanix Files Zookeeper instance tracks the original FSVM's ownership using the storage IP address (eth0), which does not float from node to node. Because FSVM-1's client IP address from eth1 is now on FSVM-2, client connections persist. The volume group and its shares and exports are reregistered and locked to FSVM-2 until FSVM-1 can recover and a grace period has elapsed.

When FSVM-1 comes back up and finds its shares and exports locked, it assumes that an HA event has occurred. After the grace period expires, FSVM-1 regains control of the volume group through the Minerva CVM service.

To summarize, the process Nutanix Files goes through to reinstate control is:

1. Stop SMB and NFS services.
2. Disconnect the volume group.
3. Release the IP address and share and export locks.
4. Register the volume group with FSVM-1.
5. Present new shares and exports to FSVM-1 with eth1.

A node failure drops the client connection, but the client can recover based on common timeout values.

4.5. Active Directory and SMB Operations

Nutanix Files SMB works with Active Directory. To deploy a cluster, you must have domain privileges; these privileges are necessary to create the machine account and the DNS entries for DFS referrals. The file server does not store these credentials.

Files 3.6 and later versions have simplified the permission requirements to add domains when creating a file server. The delegated domain user permissions required are as follows:

- Create computer objects
- Read servicePrincipalName
- Write servicePrincipalName

Nutanix Files can also create the required DNS entries automatically during file server creation. Automated DNS entry creation requires Microsoft Windows DNS and a user account with DNS admin permissions.

The maximum number of client connections depends primarily on the amount of RAM in the FSVM. See the following table for Nutanix configuration recommendations as of the Files 3.6 release.

Table 2: Supported Active Client Connections

RAM per FSVM	Supported Client Connections per FSVM	Supported Client Connections for Four-FSVM File Server Cluster
12 GB	500	2,000
16 GB	1,000	4,000
24 GB	1,500	6,000
32 GB	2,000	8,000
40 GB	2,750	11,000
60 GB	3,250	13,000
96 GB	4,000	16,000

You can continue deploying additional FSVMs if you have free nodes; you can also deploy multiple file servers.

Nutanix Files uses DFS referrals to direct clients to the FSVM that owns the targeted share or top-level directory. The following diagram shows what happens behind the scenes when a client sends a file access request.




Figure 21: DNS Request for SMB

1. When the user Dodi accesses their files, they click on a shortcut that triggers a DNS request. The DNS request is first sent for the file server name.
2. Using DNS round robin, DNS replies with an FSVM IP address. In this example, the IP address for FSVM-1 returned first.
3. The client sends a create or open request to FSVM-1.
4. The \\homes\dodi folder doesn't exist on this file server, so a STATUS_PATH_NOT_COVERED is returned.
5. The client then requests a DFS referral for the folder.
6. FSVM-1 looks up the correct mapping in the file server's Zookeeper and refers the client to FSVM-2.
7. A DNS request goes out to resolve FSVM-2.
8. The DNS request returns the IP address of FSVM-2.
9. The client gets access to the correct folder.

Managing Shares

You can manage Nutanix Files shares much the same way that you manage traditional file servers. For standard shares and top-level directories in distributed shares, you can assign permissions with native tools such as Windows Explorer.

The distributed share has some special requirements because of the way we use DFS referrals. DFS referrals have a single namespace even though the data contained in the share may be spread out over many FSVMs. Typically, to delete a user's home directory, you select the top-level directory and delete the entire directory subtree. However, when using a distributed export or share you cannot simply delete the top-level directory, because it is a separate share created internally by Nutanix Files. Since you cannot directly delete a top-level directory, removing it requires additional steps.

There are three options for renaming or deleting home share folders:

- Find out which FSVM is hosting the folder and delete the folder directly.
- For SMB shares, you can use Nutanix PowerShell scripts (found on the [support portal](#)) to handle mass operations.
- Use the Nutanix Files MMC snap-in to manage top-level directories. Any file server administrator can perform the MMC operations; there is no need to assign privileges manually. Establish a connection to the Files namespace, and perform the following SMB share management tasks:
 - # Create, delete, rename, and change permissions for top-level directories.
 - # Change permissions for shares.



Tip: Use the Nutanix Files MMC when modifying permissions at the root of a distributed share. The Files MMC ensures that permissions are propagated to the top-level directories based on inheritance.

You can manage open file locks and user sessions with the Windows-native Shared Folders MMC. Simply launch the shared folders MMC and point it to the name of the Nutanix Files instance under the **Another computer** option.

Console1 - [Console Root\Shared Folders (\FILESVR1)\Sessions]

File Action View Favorites Window Help

User	Comp...	Type	# Open Files	Connected Time	Idle Time	Guest
TME\fsctclient1u33	fsctclie...	Windows	0	00:00:10	00:00:00	No
TME\fsctclient1u77	fsctclie...	Windows	0	00:00:10	00:00:00	No
TME\fsctclient1u27	fsctclie...	Windows	0	00:00:11	00:00:00	No
TME\fsctclient1u47	fsctclie...	Windows	0	00:00:11	00:00:00	No
TME\fsctclient1u87	fsctclie...	Windows	0	00:00:11	00:00:00	No
TME\fsctclient1u89	fsctclie...	Windows	0	00:00:11	00:00:00	No
TME\fsctclient1u91	fsctclie...	Windows	0	00:00:11	00:00:00	No
TME\fsctclient1u91	fsctclie...	Windows	0	00:00:11	00:00:00	No
TME\fsctclient1u24	fsctclie...	Windows	0	00:00:12	00:00:00	No
TME\fsctclient1u46	fsctclie...	Windows	0	00:00:12	00:00:00	No
TME\fsctclient1u68	fsctclie...	Windows	0	00:00:12	00:00:00	No
TME\fsctclient1u88	fsctclie...	Windows	0	00:00:12	00:00:00	No
TME\fsctclient1u81	fsctclie...	Windows	0	00:00:13	00:00:00	No
TME\fsctclient1u17	fsctclie...	Windows	0	00:00:14	00:00:00	No
TME\fsctclient1u78	fsctclie...	Windows	0	00:00:14	00:00:00	No
TME\fsctclient1u85	fsctclie...	Windows	0	00:00:14	00:00:00	No
TME\fsctclient1u21	fsctclie...	Windows	0	00:00:17	00:00:00	No
TME\fsctclient1u5	fsctclie...	Windows	0	00:00:18	00:00:00	No
TME\fsctclient1u92	fsctclie...	Windows	0	00:00:18	00:00:00	No
TME\fsctclient1u26	fsctclie...	Windows	0	00:00:19	00:00:00	No

Figure 22: Shared Folders MMC

Access-Based Enumeration

Access-based enumeration (ABE) is a Windows feature (SMB protocol) that filters the list of available files and folders on the file server to include only those the requesting user can access. This filter helps both users and administrators, saving time for the user and worry for the administrator concerned about users accessing files not meant for them.

ABE does not control security permissions, and running ABE has associated overhead. Every time a user requests a browse operation, ABE must filter out objects for which the user doesn't have permission. Even if the user has permission to access all contents of the share, ABE still runs, causing additional CPU cycles and increased latency.



Tip: Home directories, where a distributed share is used, are a great example of where you shouldn't enable ABE. Most users get their home mapping when they log on and always have access to their own contents, therefore we don't recommend enabling ABE for home directory use cases.

SMB Signing

SMB signing is a feature of the SMB protocol that enables a digital signature against each network packet. Digital signing helps ensure the packet's origin and authenticity and prevent tampering, such as "man in the middle" attacks against data in-flight.

Nutanix Files supports SMB signing and honors the signing request as configured from the SMB client. Nutanix Files doesn't require a configuration setting; just enable SMB signing from the client and Nutanix Files negotiates the appropriate setting.

The SMB protocol uses AES-CMAC to compute signatures for SMB signing. This computation process can impact performance against the SMB share where you enabled signing. With Nutanix Files 3.2 and later versions, the signature computation uses the Intel processor AES-NI instruction set. The Intel processor hardware acceleration helps reduce the overhead associated with SMB signing.



Note: SMB signing lowers the maximum performance possible for SMB shares.

SMB Encryption

Nutanix Files 3.6 introduces support for in-flight encryption of SMB data between clients and shares. SMB encryption is disabled by default with Nutanix Files. You can enable encryption as needed on a share-by-share basis.

The screenshot shows a web-based configuration window titled "Create a share/export" with a question mark and a close button (X) in the top right. Below the title bar are three tabs: "Basics", "Settings" (which is selected and underlined), and "Summary". Under the "Settings" tab, there are five checkboxes, each followed by a help icon (question mark):

- ☐ Use "Distributed" share/export type instead of "Standard"
- ☐ Enable Self Service Restore
- ☐ Enable Access Based Enumeration (ABE)
- ☐ Blocked File Types
- ☒ Encrypt SMB3 Messages

Below the "Blocked File Types" checkbox, there is a small text note: "You can also block file types on the file server(all shares) from file server update". At the bottom of the dialog, there are three buttons: "Back", "Cancel", and "Next" (which is highlighted in blue).

Figure 23: Enable SMB Encryption

Clients must support encryption in order to access shares with encryption enabled. Nutanix Files offloads encryption processing using the Intel processor AES-NI instruction set to limit performance impact.



Note: SMB encryption lowers the maximum performance possible for SMB shares.

Durable SMB File Handles

Durable handles let SMB clients survive a temporary client-side connection loss after opening a file, allowing transparent client reconnection within a timeout. Starting with the 3.6 release, Nutanix Files supports durable handles for SMB clients with no configuration required. SMB 2.x and SMB 3.0 clients can use durable handles to reconnect transparently if there's a client-side network interruption. The system does not maintain durable handles through file server events such as FSVM failures or upgrades.



Note: Durable SMB file handles are not equivalent to SMB 3.0 transparent failover (continuously available shares). Nutanix Files does not currently support SMB 3.0 transparent failover.

DFS Namespace

DFS Namespaces (DFS-Ns) are commonly used to logically organize shared folders on different servers or in different geographic locations into a single namespace. These shared folders appear to the user as a unified hierarchical directory they can navigate using any Windows client. The server names and their locations are completely hidden from the user, enabling large scale-out architectures.

The most common architecture is an Active Directory–integrated DFS-N, where the namespace is hosted on two or more domain controllers (namespace servers) and the file data is stored on member servers. Nutanix Files supports the use of DFS-N where Nutanix Files is used as a member server and Nutanix Files shares are used as folder targets in the namespace. When using Files 3.5.1 or later versions, you can use either distributed shares or standard shares.

DFS-N is also commonly used in active-active replication scenarios. DFS-N allows multiple file servers hosting the same data to support a common folder. DFS-N helps provide site affinity, based on Active Directory, for users to connect to the file server acting as a folder target. You need a replication engine to maintain the data when active-active scenarios are present. Nutanix supports the use of Peer Software as the replication engine for active-active scenarios. To learn more about the Peer Software and Nutanix integration, see [Active-Active File Services for the Nutanix Enterprise Cloud](#).

4.6. Active Directory, LDAP, and NFS Operations

NFSv4

Nutanix Files supports NFSv4, which is more stateful than NFSv3, and includes advanced features such as strongly mandated security and DFS-like referrals. Moreover, most recent distributions of Linux, Solaris, and AIX use NFSv4 as the default client protocol.

Nutanix Files supports Active Directory, LDAP, and unmanaged access to NFSv4 exports. To make the transition from NFSv3 easier, Files does not require administrators to configure Active Directory or LDAP. You can use AUTH_SYS or AUTH_NONE authentication. AUTH_SYS authenticates at the client, just like NFSv3.

There are three different levels of Kerberos authentication when you enable Active Directory or LDAP support. Each of the following options uses Kerberos version 5:

1. krb5, DES symmetric key encryption, and an MD5 one-way hash for Nutanix Files credentials.
2. krb5i, in addition to krb5, uses MD5-based MAC on every request and response.
3. krb5p, on top of krb5 and krb5i, makes the connection between client and server private by applying DES encryption.

The screenshot shows the 'Create File Server' wizard with the 'User Management' tab selected. The wizard has five tabs: Basics, Client Network, Storage Network, User Management, and Summary. The 'User Management' tab contains the following content:

- Instruction: "Select the protocols you plan to use and configure user management options. You can also skip this step and complete it later (from 'user management' action)"
- Option: ☐ Use SMB Protocol
- Note: "Active Directory is necessary for using SMB protocol"
- Option: ☒ Use NFS Protocol
- Note: "You can use AD or LDAP for NFS user management and authentication"
- Section: "USER MANAGEMENT AND AUTHENTICATION"
- Dropdown menu: "Active Directory" (with an upward arrow icon)
- Option: ☐ Enable RFC-2307
- Form fields: "REALM NAME" (containing "343.com") and "User account location to join" (empty)
- Form fields: "USERNAME" (empty) and "PASSWORD" (empty)
- Note: "Username should have admin privileges in the domain to join."
- Option: ☐ Show Active Directory Advanced Options
- Option: ☐ Show Advanced Options
- Buttons: "Back", "Cancel", and "Next" (highlighted in blue)

Figure 24: Nutanix Files User Management with NFS

When you deploy Nutanix Files for NFS, you select AD, LDAP, or leave it unmanaged. The following diagram shows what happens behind the scenes when a client sends a file access request using NFS.




Figure 25: DNS Request for NFSv4

1. When the server wants to access files, the client first sends a DNS request for the file server name.
2. Using DNS round robin, a DNS reply returns with an FSVM address. In this example, the IP address for FSVM-1 returned first.
3. The client sends a create/open request to FSVM-1.
4. The \\Logs mount doesn't exist on this file server, so it returns NFS4ERR_MOVED.
5. The client then requests a GETATTR(FS_LOCATIONS).
6. FSVM-1 looks up the correct mapping in the file server's Zookeeper and refers the client to FSVM-2.
7. A DNS request goes out to resolve FSVM-2.
8. The DNS request returns the IP address of FSVM-2.
9. The client gets access to the correct mount point.

NFSv3

With the 3.5 release, Nutanix Files supports NFSv3. NFSv3 is enabled by default for each file server but you can manually disable it. All NFS exports have NFSv3 enabled or disabled based on this file server setting. Files 3.5 supports LDAP and unmanaged exports with NFSv3 but does not support Active Directory and Kerberos for NFSv3.



Note: You must mount exports with the TCP protocol if clients use NFSv3.

Nutanix Files with NFSv3 includes support for both distributed and standard exports. Unlike NFSv4, NFSv3 does not support DFS-like referrals. Clients are not redirected to the FSVMs that host a given distributed top-level directory or standard export; they connect to the first FSVM resolved using DNS. An internal remote procedure call (RPC) manages any file access requests to exports and top-level directories not hosted by the client-connected FSVM. This RPC sends or receives data between the client-connected FSVM and the FSVM that owns the required exports and volume groups.




Figure 26: DNS Request for NFS v3

4.7. Multiprotocol

From the Nutanix Files 3.5.1 release onward, you can create file shares that are accessible from both SMB and NFS clients, referred to as multiprotocol-enabled shares. Multiprotocol shares allow simultaneous read access to the underlying file data from either protocol. Write access can also occur from either protocol, but not simultaneously to the same file. Authentication support for multiprotocol shares includes all protocols and authentication options available with Nutanix Files (such as Active Directory for SMB and LDAP, AUTH_SYS, AUTH_NONE, and Active Directory support for NFS).




Figure 27: Multiprotocol Shares

Multiprotocol introduces the concept of a native protocol (either SMB or NFS) for a given file share. You specify the native protocol and non-native protocol when you create the share.

Create a share/export ? X

Basics Settings Summary

Description (Optional)

Multi-protocol Native SMB

File Server

filesvr1

Share Path (Optional) ?

/<sharename>/<dir>

Max Size (Optional)

GiB

Select Protocol

☐ NFS

☒ SMB

☒ Also enable non-native NFS Access

Cancel Next

Figure 28: Enable Non-Native Protocol Access

You manage all access control for a multiprotocol-enabled share using the native protocol. For SMB, the native protocol is Windows access control lists (ACLs), and for NFS, it's Unix mode bits. When non-native protocol access occurs, Nutanix Files maps user access to the permission applied with the native protocol. A mapping must exist between the user accounts using the non-native protocol and the user accounts with permission applied through the native protocol.

If you use Active Directory for both SMB and NFS with Kerberos, you don't need any explicit mappings because both have the same users and groups. You can use Identity Management for Unix (RFC 2307) to map NFS users to your Active Directory users with attributes. For LDAP or unmanaged accounts, or for Active Directory without Kerberos, you must create a mapping between the users and groups in Active Directory. You can use Prism or the NCLI to manage Nutanix Files user mappings.

Figure 29: Multiprotocol User Mapping

You can configure a default index to map all non-native users and groups to a specific native user or group. You can also use a rule-based mapping for AD and LDAP users—specifically, a template where the SMB name matches the NFS username. Additionally, you can configure explicit mapping, which overrides rule-based mapping. Explicit mapping consists of two mapping subcategories: one-to-one mapping lists and wildcard mapping. You can use the one-to-one mapping list to manually enter or upload a CSV file that maps users or groups across protocols. Use wildcards for many-to-one mapping. You can also deny share access for a specific user or group. For more details on how to create user mappings, see the [user mapping section of the Nutanix Files guide](#). The following table shows the user mapping requirements depending on the chosen directory service and authentication type.

Table 3: User Mapping Requirements

SMB Directory Service	NFS Directory Service	Export Authentication Type	Supported	User Mapping Required
AD	AD	Kerberos5*	Yes	No
AD	AD	System	Yes	Yes (Name to ID mapping)
AD	AD	None	Yes (Primary NFS only)	Yes (Name to ID mapping)
AD	AD + RFC2307	Kerberos5*	Yes	No

SMB Directory Service	NFS Directory Service	Export Authentication Type	Supported	User Mapping Required
AD	AD + RFC2307	System	Yes	Yes (Name to ID mapping)
AD	AD + RFC2307	None	Yes (Primary NFS only)	Yes (Name to ID mapping)
AD	LDAP	System	Yes	Yes (Name to Name mapping)
AD	LDAP	None	Yes (Primary NFS only)	Yes (Name to Name mapping)
AD	Unmanaged	System	Yes	Yes (Name to ID mapping)
AD	Unmanaged	None	Yes (Primary NFS only)	Yes (Name to ID mapping)

4.8. Quotas

Administrators can configure the default, user, and group quota space for any share. The default level is the quota limit for every user unless the administrator specifies otherwise. A user-level quota policy sets a specific amount of storage for a single user. For example, if an administrator allocates 1 GB, the user cannot take more than 1 GB. A group-level quota policy extends a user policy to include all users for an entire Active Directory group, where each user can use the assigned quota value. For example, if the administrator sets a group's quota to 10 GB, then each member of that group can use 10 GB. You can assign quotas for SMB, NFS, or multiprotocol-enabled shares. In the case of multiprotocol shares, administrators apply quotas using the native protocol. User or group policies are enforced for the non-native protocol based on the user mapping.

Quota Policy

?

✕

Add a quota policy

Add a quota policy for a specific user or each user in a group.

☐ User

☒ Group (Quota will be applied to each member of group individually)

USER OR GROUP

domain users

QUOTA

20

GiB

ENFORCEMENT TYPE

☐ Hard Limit

User will be put in read only mode and will receive email notification daily until remediation.

☒ Soft Limit

User will receive daily email notification but will not be put in read only mode.

Alert Emails

Quota alert emails are sent at 90% and 100% of quota limit.

Cancel

Save

Table 4: Order of Precedence for Quotas

Order of Precedence	Policy
1	User policy
2	Group policy (the group policy with the highest quota wins)
3	Default user policy

Enforcement

The administrator can also configure enforcement types for each quota-level category. Enforcement types determine if a user or group can continue to use the share after they've consumed their quota. A hard enforcement type prevents the user from writing on the share when they reach their quota limit. If a user or application is writing to the share after it has reached the quota, the operation fails. A soft enforcement type allows a user to write even if they exceed the quota limit. Under either enforcement type, users over their quota receive an email notification every 24 hours until the issue is resolved.

4.9. Selective File Blocking

With Nutanix Files 3.6 you can define a list of file types, based on file extension, to block from being stored on SMB, NFS, and multiprotocol enabled shares. You can define a list of files to block at the server level and at the share level.





Figure 31: Blocked File Types for the File Server

A comma-separated list of file extensions defines the blocked file types. File types blocked at the server level apply to all shares. When you define a list of file types at the share level, the share-level setting overrides the server-level setting.



Create a share/export ? x

Basics Settings Summary

☐ Use "Distributed" share/export type instead of "Standard" ?

☐ Enable Self Service Restore ?

☐ Enable Access Based Enumeration (ABE) ?

☒ Blocked File Types ?

File Extensions Blocked Here Will Override The File Server Settings

.mp4, .mp3, .mpg, .mpeg, .m4v, .avi, .mov

Provide a comma separated list of extensions like .db, .txt to be blocked

☐ Encrypt SMB3 Messages ?

Back Cancel Next

Figure 32: Blocked File Types for a Share

When you attempt to create a file with or rename a file to a blocked extension, an access denied message appears. You can read, edit, or delete existing files, created prior to a blocked file type policy. You can also use File Analytics to discover if any unwanted file types exist on a file server before applying the policy.

4.10. File Analytics

Nutanix Files offers a rich auditing API to which applications can subscribe and receive real-time notifications of file-related events, including file creation, deletion, read, write, and permissions changes. A common use of these APIs is forwarding these events to a syslog server for retention and audit trails. While logging audit trails is useful, simplifying insights into this data is a must. We developed Nutanix File Analytics to consume this native auditing API while delivering additional insights to the underlying data and user activity.

You can deploy File Analytics using Prism: you simply download the Analytics bundle and perform a one-click deployment. Once deployed, launch the File Analytics page and enable analytics against the desired file server instances running on your cluster. Enabling analytics requires a user account with administrative permissions to do an initial scan of the file server.

Intelligent Insights

File Analytics scans and then continuously captures all file activity for registered file server instances. This logging helps to form a repository of information so administrators can review what operations have occurred against specific data and by specific users. Logged events are analyzed to provide an initial dashboard of information:

- Capacity trend, which shows what's being consumed and how it's changed over time.
- Data age, which is the calculation of the last time a file was modified and the percentage of data at varying age ranges.
- Anomaly alerts, which show all file operations that exceed a given anomaly threshold, like the deletion of many files.
- Permission denials, which are the number of permission denial events for specific users over the selected time range.
- File distribution by size, which shows the number of files within a given size range.
- File distribution by type, which details storage consumption by file type, such as log files, image files, and video files.
- The top five active users based on total operations over the selected time period.
- The top five accessed files based on total operations over the selected time period.
- File operations, which detail the most frequent types of operations, like file create, read, or write, over a selected time period, including the trend over that time.




Figure 33: File Analytics Dashboard

The dashboard provides a quick and easy health check so you better understand capacity trends, data age, and file activity that could indicate malicious activity.

Audit Trails

Audit trails allow you to search for a specific file or a specific end user to find all file or user activity for a given timeframe. You can search for a given file or user based on wildcards. The frequency and types of operations, by users and against files, including the time they occurred, are displayed over your specified time period. You can further filter the audit trail based on operation type, such as open, read, write, delete, and other events.




Figure 34: File Analytics Audit Trails

Being able to search user and file activity helps you monitor access activity against sensitive data. You can also download the queried activity to a JSON or CSV formatted files for further reporting activities.

Anomaly Detection

Define Anomaly Rules

×

Setup your anomaly policies and anomaly email recipients for fileserver "filesvr1" here.

Anomaly Email Recipients

Add one or more comma separated email addresses if you want to receive email alerts.

administrator@tme.local

+ Configure new anomaly

EVENTS	MINIMUM OPERATION %	MINIMUM OPERATION COUNT	USER	TYPE	INTERVAL	ACTIONS
Create	▼ 1	100	All Users	▼ Hourly	▼ 1	✎ · ✕
Delete	▼ 1	100	All Users	▼ Hourly	▼ 1	✎ · ✕

Cancel

Save

Figure 35: File Analytics Anomaly Detection

You can specify different events, such as permission changes, permission denials, file deletes, and file creates; define an operation percentage and operation count as a part of the alert; specify the interval for when these thresholds apply to the operations in question; and specify email recipients for the anomaly events.



Tip: Minimum operation percentage is based on the number of files. For example, if there are 1,000 files and the minimum operation percentage is 5, it means the defined event (such as create or delete) impacted 50 files within the specified interval.

As anomaly events occur, the email addresses defined in the anomaly rule receive detailed information. You can also monitor each anomaly alert and track it over time with the usage anomalies dashboard. The dashboard shows you the users who caused the anomalies and the folders the anomalies impacted and details the types of anomalies and trends.




Figure 36: File Analytics Usage Anomalies Dashboard

Strange access patterns that trigger anomalies can indicate malicious user, virus, or ransomware activity. Seeing this activity and being alerted in real time can help you stop these operations and better understand user access patterns across your organization.

4.11. Hypervisor-Specific Support

Nutanix supports ESXi and AHV for Files. For ESXi support, you need vCenter credentials to deploy Nutanix Files and to create DRS rules to make sure the FSVs are on different nodes. You must register vCenter with the Nutanix cluster instances where you deployed Files. The deployment process generates the DRS rules for AHV automatically.

Data-at-Rest Encryption

Nutanix Files supports data-at-rest encryption (DARE) with self-encrypting drives (SEDs) or software-defined DARE using AOS. Official support for AOS software encryption with Nutanix Files begins with AOS 5.10.2 in combination with Nutanix Files 3.5.

For AHV, you enable software encryption at the cluster level. For ESXi, you enable software encryption at either the cluster level or the storage container level. To enable storage container-level software encryption with Nutanix Files you must use NCLI:

```
<ncli> storage-container edit enable-software-encryption=true name=<Files_Container_Name>
force=true
```

5. Backup and Disaster Recovery

Following modern data protection methodologies, Nutanix provides administrators and users quick restore access using Self-Service Restore (SSR) and site recovery with Nutanix-based snapshots.

5.1. Self-Service Restore

Administrators can enable SSR at any time for SMB or NFS shares. Windows Previous Version in each folder exposes SSR for SMB shares. SSR for NFS shares is exposed as a hidden snapshot directory in each folder.

Protection configuration: filesvr1

Self Service Restore

Snapshots will be created based on the configured schedule for shares/exports with self service restore enabled. End users can access these snapshots through their native interface.

Enabled for

Self service restore is enabled for all shares/exports.✓

1 hour RPO • Total :38 Snapshots [+ Add schedule](#)

TYPE	FREQUENCY	SNAPSHOTS	ACTIONS
Hourly	Every 1 hour	24	✎ ✕
Daily	Every 1 day	7	✎ ✕
Weekly	Every week on Sun	4	✎ ✕
Monthly	Every month on 1	3	✎ ✕

Note: Self Service Restore can be disabled for specific shares/exports if required

Figure 37: Protection Using SSR

SSR allows you to view snapshots of shares while the share is in use. The share snapshots are read-only, point-in-time copies (that is, copies created at a specified time).

You can view and restore removed or overwritten files, which allows you to choose a share snapshot from the same file at different times during the file's history. Administrators can configure a snapshot schedule at the file server level that applies to all shares in the file server.

Nutanix Files supports 24-hour, daily, weekly, and monthly snapshots on a fixed schedule. The default snapshot policy includes:

- Hourly, 24 per day.
- Once daily for seven days.
- Four weekly.
- Three monthly.



Note: You can configure a maximum of 50 total snapshots in the SSR schedule.

Schedule frequency can vary to suit your requirements, including shorter intervals for same-day protection against accidental deletions. You can enable SSR during or after share creation. After share creation, the administrator can change the current settings using the share update workflow feature. SSR supports both standard and distributed SMB shares.

5.2. Protection Domains and Consistency Groups

Nutanix provides integrated, automated disaster recovery between Nutanix clusters. A Nutanix Files cluster can be protected with Prism and uses the same asynchronous replication with protection domains and consistency groups as any other Nutanix cluster. A protection domain is a defined group of entities (VMs and volume groups) that you back up locally on a cluster and that may replicate to one or more remote sites. A consistency group is a subset of the entities in a protection domain. Consistency groups are configured to snapshot a group of VMs or volume groups in a crash-consistent manner.



Note: Nutanix Files supports both asynchronous (RPO of one hour or longer) and NearSync (down to one minute RPO) schedules. NearSync support requires a minimum of AOS 5.11.1 and Files 3.6.

When you create a file server, Prism automatically sets up a corresponding protection domain, which it annotates with the Nutanix Files cluster name. Prism also creates multiple consistency groups in a protection domain, including a group that includes all FSVMs.

Once you've protected Nutanix Files, all future operations on it (such as adding or removing FSVMs or adding or deleting volume groups) automatically update the corresponding consistency group in the protection domain.

5.3. Cluster Migration, Failure, and Restoration

In the event of a Nutanix Files cluster failure, restore Files in a remote cluster by initiating the **Activate** workflow, which restores from the last good snapshot. If you're moving your file services because you need to shut the cluster down, as with a planned outage, the **Migrate** workflow

shuts down all the FSVMs and takes a final snapshot for replication. You run the Migrate workflow from the Nutanix cluster that owns the active protection domain. You initiate both the Activate and Migrate workflows from the **Data Protection** menu in Prism.

After you run either the Activate or Migrate workflow, you also need to activate the file server instance, which you can do from the **File Server** menu in Prism. Activating the file server may require you to configure network VLANs on the replica site before Nutanix Files becomes operational again. These VLANs can be in subnets that are the same as or different from the source site. The recovery process is like creating the file server, but in this context you can change networks and IP addresses if necessary.

5.4. Cloning

Because Nutanix Files cloning does not affect the original Files cluster, it offers improved support for a variety of use cases, including:

- Backups at the primary and secondary sites.
- DR at the secondary site.
- File server recovery from a specific point in time.
- File server creation at the primary or remote site for testing or development.
- File server clone copies.




Figure 38: Nutanix Files Cloning Use Cases

Files uses Nutanix native snapshots to clone entire file servers. The clone is a thin copy that consumes minimal storage space. File server clones reside on the same container as the original and maintain the original security permissions. During the clone process you can specify new IP addresses and give the cloned file server a new name.

6. Third-Party Integration

6.1. Antivirus

To protect users from malware and viruses, you need to address both the client and the file server. Nutanix currently supports third-party vendors that use Internet Content Adaptation Protocol (ICAP) servers. ICAP, which is supported by a wide range of security vendors and products, is a standard protocol that allows file and web servers to be integrated with security products. Nutanix chose this method to give customers wide latitude in selecting the antivirus solution that works best for their specific environment.

Following is the workflow for an ICAP-supported antivirus solution:

1. An SMB client submits a request to open or close a file.
2. The file server determines if the file needs to be scanned, based on the metadata and virus scan policies. If a scan is needed, the file server sends the file to the ICAP server and issues a scan request.
3. The ICAP server scans the file and reports the scan results back to the file server.
4. The file server takes an action based on the scan results:
 - a. If the file is infected, the file server quarantines it and returns an “access denied” message to the SMB client.
 - b. If the file is clean, it returns the file handle to the SMB client.




Figure 39: ICAP Workflow

The ICAP service runs on each Nutanix Files file server and can interact with more than one ICAP server in parallel to support horizontal scale-out of the antivirus server. The scale-out nature of Files and one-click optimization greatly mitigates any antivirus scanning performance overhead. If the scanning affects Nutanix Files file server VM performance, one-click optimization recommends increasing the virtual CPU resources or scaling out the file server VMs. This feature also allows both the ICAP server and Files to scale out, ensuring fast responses from the customer’s antivirus vendor.



Tip: We recommend configuring two or more ICAP servers for production.

IP ADDRESS OR HOSTNAME	PORT	DESCRIPTION	CONNECTION STATUS	ACTIONS
10.7.69.224	1344		✓ OK	
10.7.69.225	1344		✓ OK	
10.5.236.201	1344		✓ OK	
10.7.69.226	1344		✓ OK	

Figure 40: Configure Multiple ICAP Servers

Nutanix Files sets scanning defaults across the entire file server. You can enable scan on write and scan on read. Scan on write begins when the file is closed, and scan on read occurs when the file is opened. You can also exclude certain file types and files over a certain size. Share scan policies can override any defaults set for the file server.

Antivirus Setup SennaFS ? X

ICAP Servers · **Scan Settings**

Scan Settings (For all shares)
Settings can be overridden for individual share if required, through the share level antivirus settings.

On access scan type ?

- ☒ SCAN ON WRITE
- ☒ SCAN ON READ

EXCLUDE FILE TYPES
Comma separated extension like .db, .txt to be excluded from scanning

EXCLUDE FILES LARGER THAN ?
No file size limit MiB

☒ SHOW ADVANCED OPTIONS

SCAN TIMEOUT ? MAXIMUM 240 SECONDS
60 seconds

☒ BLOCK ACCESS TO FILES IF ANTIVIRUS SCAN CANNOT BE COMPLETED (RECOMMENDED) ?

Back Cancel **Save**

Figure 41: Default Scan Settings

For each ICAP server, Nutanix Files will spin up no more than 10 parallel connections per FSVM and randomly distribute the file scanning among all the ICAP servers. With heavier workloads, which may involve many scan requests and use all connections, the scan servers can be given more processing power to scan more files. As soon as the current scan finishes, the next file is picked up from the queue, which keeps the number of active connections at 10.

ICAP Servers

Reports

Quarantined Files

Unquarantined Files ?

Actions ▾

1 of 17 files selected.

1-10 of 17

< > ⚙

search in table

Q

Rescan

Unquarantine

Delete

	FILE PATH	THREAT DESCRIPTION	ICAP SERVER	SCAN TIME
	/vdi2.V2/Documents/virus - Copy (4).txt	EICAR-Test-File;	10.7.69.226	08/03/17, 2:03:00 PM
	/virus1.txt	EICAR-Test-File;	10.7.69.224	08/02/17, 1:52:34 PM
☐ General_1	/yifeng_virus2	EICAR-Test-File;	10.5.236.201	08/07/17, 3:35:25 PM
☐ roaming	/vdi3.V2/Music/roaming_virus - Copy (3).txt	EICAR-Test-File;	10.7.69.225	08/03/17, 2:14:00 PM
☐ General_1	/yifeng_virus3.txt	EICAR-Test-File;	10.5.236.201	08/07/17, 3:36:37 PM
☐ General_38	/virus9.txt	EICAR-Test-File;	10.7.69.226	08/09/17, 10:14:56 AM
☐ roaming	/vdi2.V2/Documents/virus - Copy (2).txt	EICAR-Test-File;	10.7.69.226	08/03/17, 2:03:00 PM
☐ General_38	/vsample04.txt	EICAR-Test-File;	10.7.69.225	08/07/17, 5:52:28 PM
☐ roaming	/vdi3.V2/Documents/virus.txt	EICAR-Test-File;	10.7.69.226	08/04/17, 7:02:19 PM
☐ roaming	/vdi3.V2/Music/roaming_virus - Copy (4).txt	EICAR-Test-File;	10.7.69.226	08/03/17, 2:14:00 PM

Figure 42: Quarantined Files

Once Nutanix Files quarantines a file, the administrator can rescan, unquarantine, or delete the file. You can search quarantined files if you need to restore a file quickly.

If your antivirus vendor doesn't support ICAP, you can scan the shares by installing an antivirus agent on a Windows machine, then mounting all the shares from the file server. This approach allows you to schedule scans for periods of low usage. At the desktop or client level, you can set your antivirus solution to scan on write or scan only when files are modified. You can configure high-security environments to scan inline for both reads and writes.

See the Nutanix Files release notes for the latest list of qualified ICAP server vendors. The current list of supported vendors includes Bitdefender, McAfee, Symantec, Kaspersky, and Sophos.

6.2. File Operations Monitoring

File operations monitoring is a native Nutanix Files API used to forward all SMB and NFS operations executed by clients to a registered target repository.

File operations monitoring can be broken down into two major areas for third-party vendors:

1. File activity
2. Audit

Partner software can register with the Nutanix Files instance to capture file activity events. Partner software can also make REST calls to Files to create a policy defining which file notification types you want to receive and the shares you want to receive them from. The partner

software uses a web client and communicates with the Nutanix Files file server using HTTPS requests. The HTTPS communication relies on SSL authentication. The HTTP server runs with either its unique self-signed SSL certificate or a Transport Layer Security (TLS) connection with the partner, then exchanges the keys between the two parties and sends messages to the partner server over this secure channel.

Third-party software can use different protocol communication methods such as syslog, Google, or Kafka protobuf. One common way to use the Nutanix file operations monitoring API is to forward events to a syslog server for any future auditing needs. Several vendors have integrated with this operations monitoring API:

- Peer Software

Peer Software has integrated with Nutanix, including the file operations monitoring API, to support active-active deployments with real-time replication and file locking between heterogeneous environments. To learn more about the Peer Software and Nutanix integration, please see the document [Active-Active File Services for the Nutanix Enterprise Cloud](#).

- Netwrix

Netwrix has integrated with the Nutanix file operations monitoring API. Netwrix provides comprehensive visibility into changes and data access across Nutanix Files instances. Netwrix offers a free add-on for Nutanix Files. See www.netwrix.com for more details.

6.3. Intelligent Backup

Traditional NAS backup approaches like Network Data Management Protocol (NDMP) have core limitations. Problems like complexity, performance (including the need for periodic full backups), and scale led Nutanix to develop a more modern backup solution: a changed file tracking (CFT) API that third-party backup vendors use. This API speeds backup times because it does not perform a metadata scan across your file server, which could contain millions of files and directories. The API enables scale as backups can occur across multiple FSVMs in parallel. The API also reassures customers aren't locked into any one solution as backups taken using the Nutanix Files API can be restored to other vendors.

CFT uses internal snapshots to track differences between backup windows. The backup software vendor requests the creation of these snapshots, and Files returns a list of the changed files in the form of a share. The backup software proceeds to mount the share onto the backup server to be read for backup.




Figure 43: CFT Backup Process

Backup software can specify multiple shares and their respective snapshot information. Nutanix Files returns the list of URLs that map to the number of client streams that can start in parallel. Because shares are distributed evenly across the FSVMs based on load, the backup software can take advantage of all the FSVMs to drive throughput.

HYCU Backup and Recovery is integrated with the Nutanix Files CFT API. See www.hycu.com for more details.

Commvault version 11 service pack 15 is also integrated with the Nutanix Files CFT API as of the 3.5 release. Please see [the Commvault blog post on the topic](#) for more details.

There are two ways to back up Files shares with software that doesn't support CFT. One option is to run the backup application on a Windows machine and map the UNC path of the share as a drive that needs to be backed up.

There are also vendors (such as Commvault, Rubrik, and Veritas) that provide support for backing up file shares without mounting to a guest VM. These applications can read directly from

the UNC path. Note for this scenario that Nutanix does not qualify mounting SMB shares from Linux machines or NFS shares from Windows machines.



Tip: Do not attempt to back up the FSVMs. FSVMs are stateless and data is stored in Nutanix volume groups. Back up Nutanix Files from the share level. You can protect FSVMs using Nutanix protection domains.

Because the system spreads different standard shares across the cluster, try to back up multiple shares at the same time with multiple subclients. The distributed share allows you to configure multiple data readers to drive throughput.




Figure 44: Adding More Readers to Drive Throughput on a Home Share

As an example of backup using Commvault, we tested 400 users spread out on three FSVMs, placing the data on a home share. We found that adding more readers for the backup job could increase performance. The bottleneck was the media agent, which was the backup destination

for the files. The media agent was virtualized and configured with only eight vCPUs. Adding more vCPUs to the media agent achieved a faster backup time.




Figure 45: Backup Job Details

As the preceding figure shows, we achieved almost 600 GB average throughput per hour with a fairly small setup.

7. Conclusion

Nutanix Files delivers on-demand performance and automated provisioning to provide highly scalable file management. It reduces the administration and configuration time needed to deploy and maintain your environment, providing a public cloud experience within your private cloud.

The Acropolis Distributed Storage Fabric (DSF) and Prism make hosted file services highly resilient and easy to use—for example, you can configure the native data efficiency features of the DSF, such as compression and erasure coding (EC-X), for each individual file server. Prism lets you administer network and resource management, Active Directory, fault tolerance, and Nutanix Files share and export management all in one place, vastly improving operational efficiency.

You can deploy Files on a new or existing Nutanix environment, which means you have the flexibility to use your existing investment to take advantage of Files services.

Nutanix Files streamlines design, implementation, and maintenance, which provides an unrivaled user experience, an essential element of success for end-user computing projects that store data in shared environments.

Appendix

About Nutanix

Nutanix makes infrastructure invisible, elevating IT to focus on the applications and services that power their business. The Nutanix Enterprise Cloud OS leverages web-scale engineering and consumer-grade design to natively converge compute, virtualization, and storage into a resilient, software-defined solution with rich machine intelligence. The result is predictable performance, cloud-like infrastructure consumption, robust security, and seamless application mobility for a broad range of enterprise applications. Learn more at www.nutanix.com or follow us on Twitter [@nutanix](https://twitter.com/nutanix).

List of Figures

Figure 1: Nutanix Enterprise Cloud..... 7

Figure 2: Nutanix Files Server Instances Run as VMs for Isolation from the DSF.....9

Figure 3: Data Path Architecture of Nutanix Files..... 10

Figure 4: FSVM Internal Communication on One Node..... 11

Figure 5: FSVM vDisks and Volume Groups..... 12

Figure 6: FSVM Volume Group with 16 vDisks..... 12

Figure 7: Initial 40 TB File System Pool..... 13

Figure 8: 80 TB File System Pool Following First Expansion..... 13

Figure 9: 120 TB File System Pool After Second Expansion..... 13

Figure 10: 140 TB File System Pool After Final Expansion..... 13

Figure 11: Distributed Directory Shares..... 14

Figure 12: Two Standard Shares on the Same File Server..... 15

Figure 13: Create Share Menu with Nested Share Path..... 18

Figure 14: Distributed Share and Top-Level Directories..... 19

Figure 15: Three Standard Shares..... 19

Figure 16: Performance Optimization Recommendation..... 21

Figure 17: Load Balancing Volume Groups with Nutanix Volumes..... 22

Figure 18: Nutanix Volumes Load Balancing for File Server Volume Groups..... 23

Figure 19: Each FSVM Controls Its Own Volume Groups..... 24

Figure 20: FSVM-1 Failure..... 25

Figure 21: DNS Request for SMB..... 27

Figure 22: Shared Folders MMC.....	29
Figure 23: Enable SMB Encryption.....	30
Figure 24: Nutanix Files User Management with NFS.....	33
Figure 25: DNS Request for NFSv4.....	34
Figure 26: DNS Request for NFS v3.....	35
Figure 27: Multiprotocol Shares.....	36
Figure 28: Enable Non-Native Protocol Access.....	37
Figure 29: Multiprotocol User Mapping.....	38
Figure 30: Setting a Share Quota.....	40
Figure 31: Blocked File Types for the File Server.....	41
Figure 32: Blocked File Types for a Share.....	42
Figure 33: File Analytics Dashboard.....	44
Figure 34: File Analytics Audit Trails.....	45
Figure 35: File Analytics Anomaly Detection.....	46
Figure 36: File Analytics Usage Anomalies Dashboard.....	47
Figure 37: Protection Using SSR.....	48
Figure 38: Nutanix Files Cloning Use Cases.....	50
Figure 39: ICAP Workflow.....	51
Figure 40: Configure Multiple ICAP Servers.....	52
Figure 41: Default Scan Settings.....	53
Figure 42: Quarantined Files.....	54
Figure 43: CFT Backup Process.....	56
Figure 44: Adding More Readers to Drive Throughput on a Home Share.....	57

Figure 45: Backup Job Details.....58



List of Tables

Table 1: Document Version History..... 6

Table 2: Supported Active Client Connections..... 26

Table 3: User Mapping Requirements..... 38

Table 4: Order of Precedence for Quotas..... 41



Nutanix Flow

Nutanix Tech Note

Version 1.1 • May 2018 • TN-2094

Copyright

Copyright 2018 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 is a trademark of Nutanix, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

Contents

1. Executive Summary.....

4

2. Introduction.....

5

2.1. Audience.....

5

2.2. Purpose.....

5

3. Nutanix Enterprise Cloud Overview.....

6

3.1. Nutanix Acropolis Architecture.....

7

4. Nutanix Flow.....

8

4.1. Architecture.....

8

4.2. Enabling Flow Microsegmentation and Visualization.....

9

4.3. Categories.....

11

4.4. Security Policies.....

14

4.5. Using Quarantine Policies.....

16

4.6. Using Isolation Policies.....

19

4.7. Using Application Policies.....

23

5. Conclusion.....

36

Appendix.....

37

References.....

37

About Nutanix.....

37

List of Figures.....

38

List of Tables.....

40

1. Executive Summary

Flow is a software-defined networking product tightly integrated into Nutanix AHV and Prism. Flow provides rich visualization, automation, and security for VMs running on AHV. Microsegmentation is a component of Flow networking that simplifies policy management. Using multiple Prism Central categories (logical groups), you can create a powerful distributed firewall that gives administrators an application-centric policy management tool for securing VM traffic.

Categories enable flexible, attribute-based grouping of VMs to define security policies. By using logical categories, you no longer need to base policy definitions on network addresses or manually update policies to handle network changes. Your security policy can automatically apply to a VM independent of its network configuration. With Flow, administrators can visualize traffic between groups of VMs to create plain-language policies based on application behavior, rather than defining allowed traffic in the language of IP addresses and subnets.

With Flow, Nutanix extends Prism Central's consumer-grade simplicity and ease of use beyond virtualization and storage infrastructure management to the world of network security.

2. Introduction

2.1. Audience

This tech note is part of the Nutanix Solutions Library. It is intended for architects and administrators who are responsible for VM networking and security. Readers of this document should already be familiar with Nutanix AHV and Prism Central.

2.2. Purpose

In this document, we cover the following topics:

- Nutanix Flow.
- Flow microsegmentation.
- Flow visualization.
- Creating and managing categories.
- Creating and managing security policies.

Table 1: Document Version History

Version Number	Published	Notes
1.0	April 2018	Original publication.
1.1	May 2018	Added Service Chain KB article.

3. Nutanix Enterprise Cloud Overview

Nutanix delivers a web-scale, hyperconverged infrastructure solution purpose-built for virtualization and cloud environments. This solution brings the scale, [resilience](#), and economic benefits of web-scale architecture to the enterprise through the Nutanix Enterprise Cloud Platform, which combines three product families—Nutanix Acropolis, Nutanix Prism, and Nutanix Calm.

Attributes of this Enterprise Cloud OS include:

- Optimized for storage and compute resources.
- Machine learning to plan for and adapt to changing conditions automatically.
- Self-healing to tolerate and adjust to component failures.
- API-based automation and rich analytics.
- Simplified one-click upgrade.
- Native file services for user and application data.
- Native backup and disaster recovery solutions.
- Powerful and feature-rich virtualization.
- Flexible software-defined networking for visualization, automation, and security.
- Cloud automation and life cycle management.

Nutanix Acropolis provides data services and can be broken down into three foundational components: the Distributed Storage Fabric (DSF), the App Mobility Fabric (AMF), and AHV. Prism furnishes one-click infrastructure management for virtual environments running on Acropolis. Acropolis is hypervisor agnostic, supporting three third-party hypervisors—ESXi, Hyper-V, and XenServer—in addition to the native Nutanix hypervisor, AHV.




Figure 1: Nutanix Enterprise Cloud

3.1. Nutanix Acropolis Architecture

Acropolis does not rely on traditional SAN or NAS storage or expensive storage network interconnects. It combines highly dense storage and server compute (CPU and RAM) into a single platform building block. Each building block delivers a unified, scale-out, shared-nothing architecture with no single points of failure.

The Nutanix solution requires no SAN constructs, such as LUNs, RAID groups, or expensive storage switches. All storage management is VM-centric, and I/O is optimized at the VM virtual disk level. The software solution runs on nodes from a variety of manufacturers that are either all-flash for optimal performance, or a hybrid combination of SSD and HDD that provides a combination of performance and additional capacity. The DSF automatically tiers data across the cluster to different classes of storage devices using intelligent data placement algorithms. For best performance, algorithms make sure the most frequently used data is available in memory or in flash on the node local to the VM.

To learn more about the Nutanix Enterprise Cloud, please visit [the Nutanix Bible](#) and [Nutanix.com](#).

4. Nutanix Flow

Nutanix Flow delivers advanced networking and security services, providing visibility into the virtual network, application-centric protection from network threats, and automation of common networking operations.

Fully integrated into Nutanix AHV virtualization and the Nutanix Enterprise Cloud OS, Flow allows organizations to deploy software-defined virtual networking without the complexity of installing and managing additional products that have separate management and independent software maintenance requirements.

Flow provides application-centric policies that enable complete visibility and traffic control. This policy model allows administrators to implement fine-grained rules regarding traffic sources and destinations, or microsegmentation. These same policies make it possible to visualize traffic flowing within and between application VMs. This granular level of control is an important part of a defense-in-depth strategy against modern datacenter threats.

Nutanix Flow protects against new threats that are designed to spread laterally from one system to another within 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. Flow applies security rules between all applications and VMs within the datacenter, thus adding internal protection behind your perimeter firewall.

4.1. Architecture

The Nutanix management plane, Prism Central, provides policy administration for Nutanix Flow. Prism Central and the registered Nutanix clusters combine to form the control plane for Flow. The data plane is the Nutanix AHV host Open vSwitch, which processes VM network traffic.




Figure 2: Flow Microsegmentation Architecture

4.2. Enabling Flow Microsegmentation and Visualization

Nutanix Flow is disabled by default. To enable Flow in Prism Central, click the question mark, select **New**, and then click **Microsegmentation**. Refer to the requirements section in the [Policies](#)

[Management documentation](#) before enabling Flow. Flow is a licensed feature, but Nutanix offers customers a free 60-day trial period.




Figure 3: Enable Microsegmentation

4.3. Categories

Categories allow you to group entities such as VMs together logically. When creating a category, you can assign it multiple potential values. The administrator selects from this list of potential values when adding the category to an entity. For example, the Environment key could have several values, such as Production, Development, Staging, or Test. When assigned to a VM, each category consists of a text-based key-value pair—for example, Environment: Production, with a key of Environment and a value of Production.

Categories are extremely flexible, and you can create new high-level key-value pairs to group VMs based on application requirements. Categories are the building blocks of security policies, so think carefully about what groups you may need to segment network traffic along existing logical boundaries. For example, production may be isolated from development, or a perimeter network may have special access from the Internet.

There are a number of system categories built into Prism Central that have a special meaning when used in Nutanix Flow. Start with the existing categories and add your own if needed. The following sections describe several important system categories that Flow requires for building policies.

AppType

The category AppType defines a group of VMs that belong to the same application. You can use AppType: Exchange, for example, to define all of the entities that belong to the Exchange application. Nutanix recommends creating new values inside the AppType category to match the applications running in your AHV environment.





Figure 4: AppType Category List

When creating Application Security Policies, Nutanix Flow uses AppType categories to manage traffic to and from an application, so be sure to define your applications with the appropriate AppType values. Select **AppType**, then select **Update** to add your own custom application names.




custom-app Unused - +

Cancel Save

Figure 5: AppType Custom Application Name

AppTier

The category AppTier defines a set of entities that serve the same function within an application. For example, AppTier: Exchange_Mailbox and AppTier: Exchange_Edge_Transport are two tiers of VMs within the Exchange application, each serving a different function.



AppTier Update Delete ? X

Policies	Values	Entities
There are no associated policies for this category.	App	VMs: 2
	DB	VMs: 2
	Default	No associated entities
	Exchange_Edge_Transport	No associated entities
	Exchange_Mailbox	No associated entities

Figure 6: AppTier Categories

When creating application security policies, Nutanix Flow uses AppTier categories to secure communication to, from, and between application tiers. Create custom tiers for your application by selecting **AppTier**, then **Update**.

Creating Custom Categories

Nutanix recommends updating the existing system categories where possible and expanding these with values that suit your needs. For example, expand the AppType and AppTier categories by adding your own applications and tiers as values. You can create entirely new categories as needed, but Nutanix recommends keeping things as simple as possible while still achieving the desired result. Creating the smallest logical number of categories helps ensure that policy creation remains quick and easy for the administrator.

The creation of branch office policies offers an example of the tradeoff between the number of categories and the number of policies. Consider the categories and policies required based on the two following approaches: one category per site type, and one category per site.




Figure 7: Approach 1: One Category per Site Type




Figure 8: Approach 2: One Category per Site

With the first approach, it's easy to create a security policy as long as all branches have the same relationship to the HQ application. The second approach gives more control per site, but we must define the relationship between each site and HQ. If we have a large number of sites, this definition process could be time consuming.

Quarantine

By default, Quarantine categories isolate a VM from all other network devices. In addition to the standard Quarantine category—Default, which denies all traffic—a Forensic category allows quarantined VMs access to and from a policy-defined set of destinations and sources to assist in analyzing the quarantined VM. The Quarantine: Default and Quarantine: Forensic categories

cannot be modified, and you cannot manually apply them to a VM. Nutanix Flow uses these categories when you select the quarantine action for a VM.

Applying Categories to VMs

To apply categories to VMs, navigate to **Explore** in Prism Central, then click **VMs** in the **Entities** menu. Select the VMs you want, then click **Manage Categories** in the **Actions** drop-down menu.

You can apply more than one category to each VM, so it's possible to apply both of the approaches described in the Creating Custom Categories section above. With VM-level categories, we could allow connectivity between branches based on one attribute (such as Site: Branch-001), and connectivity from branches to HQ based on another attribute (such as SiteType: Branch).




Figure 9: Site and SiteType Categories on a VM

With the Site category, you can define complex relationships between sites. For simpler policy creation, you can use the SiteType category to define relationships between sites and HQ. This flexibility demonstrates the power of categories and shows that thinking about the relationships in an application can help you create more useful categories.

4.4. Security Policies

Security policies managed in Prism Central use categories to define the network traffic allowed or denied between VMs. Different security policy types describe the relationships between categories. In AOS 5.6, we offer the following policy types: quarantine, isolation, and application.

The quarantine policy has two available tiers: Strict and Forensic. Strict quarantine blocks all traffic to and from VMs. Forensic quarantine allows specific predefined sources and destinations for quarantined VMs. Use Strict quarantine to completely block all inbound and outbound traffic

to a VM. Use Forensic quarantine to allow a defined set of administrators to troubleshoot and analyze a VM.

Isolation policies define specific categories of VMs that should be disallowed from communicating with each other. For example, you could create an isolation policy to ensure that the development environment can never communicate with the production environment.

Finally, application security policies allow you to define a multitiered application and control traffic to and from as well as within the application tiers. For example, you could create an application policy to define a mailbox and an edge tier inside the Exchange application. Furthermore, you can then set this policy to allow sources such as Marketing and Sales access to the mailbox tier, with both Exchange tiers allowed access to AD and Windows Update servers.

Nutanix Flow evaluates security policies in the order shown in the figure below to build traffic rules. If traffic encounters a matching quarantine or isolation rule, Flow blocks it, and further policy processing stops. If no quarantine or isolation rule blocks the traffic, then the application rule evaluates it. If the traffic source or destination does not match any rules, Flow allows the traffic.




Figure 10: Policy Evaluation Order

Combining Policies

Combine multiple policies and policy types to build a complete security solution for a set of applications. First, use quarantine policies to temporarily isolate a specific VM from all other

VMs. Next, use isolation policies to define which groups of VMs must never be allowed to talk with other groups. Finally, use application policies to control traffic allowed to, from, and within applications. Application policies can allow traffic between two or more applications as well, but remember to ensure that the policy of each application allows this traffic.

Policy Mode

The two selectable modes for security policies are Apply and Monitor. A policy in Monitor mode allows traffic, even if the policy does not specifically define that traffic as allowed. In contrast, Apply mode only allows specifically defined traffic.

Monitor mode is the default state for a newly created isolation or application policy. Use Monitor mode in combination with flow visualization to track flows on a newly created policy, ensuring that it contains the expected traffic.

Once you have confirmed policy accuracy, you can move policies from Monitor to Apply mode. Use flow visualization in Apply mode to see blocked traffic.

4.5. Using Quarantine Policies

The option to place a VM into a quarantine category is in the VM action menu.




Figure 11: Quarantine VM Action

You can place VMs into either Strict or Forensic quarantine.

Quarantine VMs ? X

1 VM will be quarantined.

Quarantine Method

☒ **Strict**
Isolate VMs from everything on the network.

☐ **Forensic**
Isolate VMs from everything except for forensic tools specified in the quarantine policy.

Cancel Quarantine

Figure 12: Quarantine Methods

Select **Strict** to place the VM into the default quarantine category, which restricts all traffic to and from the VM. Select **Forensic** to place the VM into the forensic category, which allows a defined set of sources and destinations.

The forensic quarantine policy controls which forensic tools have access to VMs in the forensic category. Modify the quarantine policy to add your own tools as needed. In this example, the security team is allowed access to the forensic quarantined VMs on specific ports. You can add inbound sources and even outbound sources to the forensic category; however, you cannot add any sources or destinations to the default strict category.




Figure 13: Quarantine Policy Definition

Quarantine Flow Visualization

Use the quarantine policy view to visualize the traffic to and from quarantined VMs. Traffic blocked by the quarantine policy displays a red blocked symbol.




Figure 14: Quarantine Flow Visualization

4.6. Using Isolation Policies

Isolation policies prevent two defined groups of entities (VMs) from communicating with each other. This kind of restriction is useful in instances with a limited number of groups, where these groups require absolute isolation from other defined groups. Isolation policies still allow traffic to undefined groups and traffic within a group. To restrict traffic to undefined groups and to control traffic within the group, use application policies instead of isolation policies.

Environment isolation is a great example of how to use isolation policies, where VMs in category **Environment: Production** should not communicate with VMs in **Environment: Dev**.

Create isolation policies by selecting **Explore**, then **Security Policies**. Click **Create Security Policy**, then select **Isolate Environments**.

Entities

Disks12

GPUs0

POLICIES

Security Policies10

☐

Create Security Policy

Name contains

Secure an Application


Isolate Environments

1 Total Security Policies

☐	NAME	PURPOSE

Figure 15: Isolation Policy Menu

Give the isolation policy a name and select the two categories that should be separated from each other. You can use any two categories to create an isolation policy.



The screenshot shows a web form for creating an isolation policy. It has several sections: 'NAME' with a text input containing 'prod-dev-isolation'; 'PURPOSE' with a text input containing 'Isolate Production VMs from Development VMs.'; 'ISOLATE THIS CATEGORY' with a dropdown menu showing 'Environment:Production'; 'FROM THIS CATEGORY' with a dropdown menu showing 'Environment:Dev'; and a checkbox labeled 'Apply the isolation only within a subset of the data center' which is currently unchecked. At the bottom right are three buttons: 'Cancel', 'Apply Now', and 'Save and Monitor'.

NAME

prod-dev-isolation

PURPOSE

Isolate Production VMs from Development VMs.

ISOLATE THIS CATEGORY

Environment:Production

FROM THIS CATEGORY

Environment:Dev

☐ Apply the isolation only within a subset of the data center

Cancel Apply Now Save and Monitor

Figure 16: Isolation Policy Creation

If you have more than two groups that require separation (for example, **Production**, **Dev**, **Staging**, and **Testing**), then create an isolation policy for each unique pair of groups. For a large list of groups, this list of isolation policies can grow long, so you may find application policies to be more effective than isolation policies if you have a large number of groups to isolate.

For example, to separate four groups, you would need the following six policies.




Figure 17: Isolation Policies for Four Categories

Adding one more group that required isolation, such as **Environment: Backup**, would require four new isolation policies—one to isolate each of the previously existing groups from the new group.

Isolation Flow Visualization

Use the flow visualization feature in an isolation policy to view connection attempts between isolated VMs.




Figure 18: Isolation Flow Visualization

Isolate Within a Subset

In addition to simple isolation between two groups, Nutanix Flow also offers the option for the isolation policy to function within only a subset of VMs in the matching pair. For example, our previous isolation policy separated all Environment: Production from all Environment: Development. If we only wanted to separate Production from Development when the VMs also had the category **businessUnit : payroll** applied, we would add **businessUnit : payroll** as a subset to the isolation policy.

☒ Apply the isolation only within a subset of the data center

businessUnit:payroll

Figure 19: Isolate Within a Subset

4.7. Using Application Policies

Application policies are the most flexible type of security policy, defining both inbound traffic sources and outbound destinations for a single- or multitiered application. An application policy can also control traffic to, from, and within individual application tiers. An application policy must be configured with a single AppType category that acts as the target group, because the mapping between application policies and AppType categories is 1:1. For example, you can use the category AppType: Exchange as the AppType in only one application security policy.

Use application policies whenever traffic must be allowed on specific ports and protocols between sources and destinations. Application policies can also isolate one group of AppType VMs from all other AppType VMs, without needing to use isolation policies. The key difference between these two policy types is that application policies allow configurable sources and destinations between apps, while isolation policies enforce strict separation with no exceptions.

The diagram below shows the allowed inbound sources on the left and outbound destinations on the right, with a single-tiered application in the center. The central application is labeled Your App in the diagram, but in some contexts Nutanix also calls it the target group.




Figure 20: Application Policy Source, Application, and Destination

To create application policies in Prism Central, navigate to **Explore**, select **Security Policies**, click **Create Security Policy**, and then select **Secure an Application**.




Figure 21: Create Application Policy

Enter a name for the policy and a helpful text description. The application drop-down menu shows a list of all available AppType categories.



Note: You can only use an AppType category to create an application policy.

NAME

app_exchange

PURPOSE

Application Policy for Exchange

Secure this app

AppType: Exchange



Figure 22: Application Policy Creation

Sources and Destinations

After selecting the application, choose whether to allow or whitelist the sources and destinations associated with it. With the **Allow All** option, all sources or destinations can communicate with the application. To control which sources and destinations are allowed, use the **Whitelist Only** option to allow explicit categories or IP subnets and block any unspecified traffic.




Figure 23: Allow All vs. Whitelist

Add sources on the left side of the application policy (refer to the Application Policy Source, Application, and Destination figure above) by category and IP subnet. Select **Category** from the drop-down menu and use the text box to search for and select the desired category.

Inbound

Whitelist Only ▾

Deny all incoming traffic

No inbound rules have been configured. All inbound traffic will be blocked.


Add source by: Category ▾

Search for a category

Add

Figure 24: Add Source by Category

Select **Subnet/IP** and enter the IPv4 network address in CIDR notation (IP/mask length). For example, use a mask length of 24 to allow an entire Class C subnet, or a mask length of 32 to allow only a single host IP address. Using IP subnets is helpful when the source or destination you're adding is external to the Nutanix cluster.




Add source by: Subnet/IP ▾

aaa.bbb.ccc.ddd/mm

Add

Figure 25: Add Source by Subnet/IP

After adding a source, click the plus sign to select the traffic's destination, allowing traffic to a specific tier within the application.



Create App Security Policy

← Back

✓ Define Policy 2 Secure Application 3 Review

Inbound Whitelist Only ▾

Deny all incoming traffic

No inbound rules have been configured. All inbound traffic will be blocked.

Site Branch-001

+ Add Source

AppType Exchange

+ AppType Exchange

Set rules on App Tiers, instead

Figure 26: Connecting a Source to AppType

Add allowed protocols and ports such as TCP, UDP, and ICMP to the flow.

Create Inbound Rule

Site Branch-001

AppType Exchange

Specify the protocol details for this rule

+ Add Port/Protocol

PROTOCOL	PORTS	ACTIONS
Any ✓ TCP UDP ICMP	25	✕

Cancel

Save

Figure 27: Allowed Flow Specification

Adding categories and IP subnets as destinations in an application policy works in a similar way, but you would select the plus sign on the source tier of the application after adding a destination.

Application Tiers

Add tiers to an application that requires granular traffic control for the different VMs within the application. In the previous example, we combined the Exchange AppType category with two AppTier categories: Exchange_Mailbox and Exchange_Edge_Transport. To add these AppTier categories in the application policy, select **Set rules on App Tiers, instead** to allow traffic to and from individual tiers of the application.




Figure 28: Add Application Tiers

Now you can control sources and destinations at the individual tier level.




Figure 29: Source to AppTier

To control traffic between application tiers, click **Set Rules within App** at the top of the screen and select the tiers from which traffic should be allowed. Use the plus signs that appear on the other tiers to define the allowed traffic.

Set Rules to & from App Set Rules within App

AppType Exchange

AppTier Exchange_Mailbox

Can VMs in this tier talk to each other ?

☒ Yes ☐ No

AppTier Exchange_Edge_Transport

VMs in this tier can talk to each other

Figure 30: Set Rules Within App

To control traffic between VMs within a single application tier, select whether traffic is allowed within the same tier. For Exchange, it is desirable to allow traffic between servers in the same tier, but an administrator may want to disallow traffic between web VMs in the same front-end web application tier for additional security.

Service Chain Insertion

Each defined flow in an application policy can be directed through a service chain when a chain exists. Service chains define a set of network function VMs (NFVs) for advanced traffic processing. For example, the service chain can direct network traffic on a specific port to a VM for antivirus scanning, deep packet inspection, or packet capture. You can combine multiple NFVs in a chain to apply multiple functions to guest VM traffic.




Figure 31: Network Function VM Chain

Follow the documentation for service chain creation in [KB 5486](#). In Prism Central, create the allowed inbound or outbound flow as usual, then select the desired service chain from the drop-down menu.

Create Inbound Rule

Site Branch-001

AppTier Exchange_Edge...

Specify the protocol details for this rule

+ Add Port/Protocol

PROTOCOL	PORTS	ACTIONS
TCP	25	×

☒ Redirect through a service chain

Sec_Chain

Figure 32: Redirect Through a Service Chain

Flow Monitoring

Application policies provide flow visualization based on collected traffic statistics to monitor blocked and allowed traffic flows. Nutanix AHV hosts collect traffic information and send it to Prism Central, which builds a flow visualization view inside the application policy.




Figure 33: Flows Detected in Monitor Mode

When an application policy is in Monitor mode, traffic to the application that isn't allowed in the policy is shown in yellow. Hovering over the traffic flow shows a history of connection attempts along with details about port and protocol. Hovering over the traffic inbound source or outbound destination allows the administrator to add a specific flow to the policy. This flexibility helps create an accurate policy, ensure that no traffic is missed, and see what activity is taking place.

When an application policy is in Apply mode, denied traffic displays with a red X, denoting connection attempts that the policy blocked.




Figure 34: Flows Blocked in Apply Mode

Application Policies Filter

During application policy creation, you can also select a specific subset of categories where this application rule applies. Using the AppType: Exchange example and selecting a filter of Environment: Production, the application policy would only apply to AppType: Exchange VMs that also shared the category Environment: Production.

To avoid ambiguity, when the AppType policy uses the Environment filter, all sources and destinations must have an Environment category explicitly defined.

Secure this app ?

AppType: Exchange

☒ Filter the app type by category (e.g. environment, location, etc.)

Selecting Environment:Production means that all app security policies on AppType:Exchange must include a specific Environment category.

Environment:Production

Figure 35: Application Policies Within a Subset

5. Conclusion

Nutanix Flow is a software-defined networking solution for AHV that provides visualization, automation, and security. As part of Flow, microsegmentation uses Prism Central categories and security policies to protect VMs. By using categories to create flexible groups of VMs, Flow simplifies security policy definition. With security policies, administrators can protect or isolate applications or environments and quarantine infected or rogue VMs.

Through policy- and application-centric traffic monitoring, flow visualization answers the vital question of what VM traffic is actually sent and received in the virtual system.

The ease of creating and applying categories and policies opens the door for automation. Security is no longer tied directly to VMs or IP addresses, and administrators can respond quickly to datacenter changes.

Because Nutanix Flow is built into Prism Central and works natively in AHV, there are no additional components to install or manage. Network security is now just one click away.

For feedback or questions, please contact us using the [Nutanix NEXT Community forums](#).

Appendix

References

1. [Creating Service Chains with REST KB 5486](#)
2. [Nutanix AHV Administration Guide: Integration with Network Functions](#)
3. [Nutanix Prism Central Guide: Category Management](#)
4. [Nutanix Prism Central Guide: Policies Management](#)

About Nutanix

Nutanix makes infrastructure invisible, elevating IT to focus on the applications and services that power their business. The Nutanix Enterprise Cloud OS leverages web-scale engineering and consumer-grade design to natively converge compute, virtualization, and storage into a resilient, software-defined solution with rich machine intelligence. The result is predictable performance, cloud-like infrastructure consumption, robust security, and seamless application mobility for a broad range of enterprise applications. Learn more at www.nutanix.com or follow us on Twitter [@nutanix](https://twitter.com/nutanix).

List of Figures

Figure 1: Nutanix Enterprise Cloud.....	6
Figure 2: Flow Microsegmentation Architecture.....	9
Figure 3: Enable Microsegmentation.....	10
Figure 4: AppType Category List.....	11
Figure 5: AppType Custom Application Name.....	12
Figure 6: AppTier Categories.....	12
Figure 7: Approach 1: One Category per Site Type.....	13
Figure 8: Approach 2: One Category per Site.....	13
Figure 9: Site and SiteType Categories on a VM.....	14
Figure 10: Policy Evaluation Order.....	15
Figure 11: Quarantine VM Action.....	17
Figure 12: Quarantine Methods.....	18
Figure 13: Quarantine Policy Definition.....	19
Figure 14: Quarantine Flow Visualization.....	19
Figure 15: Isolation Policy Menu.....	20
Figure 16: Isolation Policy Creation.....	21
Figure 17: Isolation Policies for Four Categories.....	22
Figure 18: Isolation Flow Visualization.....	23
Figure 19: Isolate Within a Subset.....	23
Figure 20: Application Policy Source, Application, and Destination.....	24
Figure 21: Create Application Policy.....	24

Figure 22: Application Policy Creation..... 25

Figure 23: Allow All vs. Whitelist.....26

Figure 24: Add Source by Category.....27

Figure 25: Add Source by Subnet/IP..... 28

Figure 26: Connecting a Source to AppType.....28

Figure 27: Allowed Flow Specification.....29

Figure 28: Add Application Tiers..... 30

Figure 29: Source to AppTier..... 30

Figure 30: Set Rules Within App.....31

Figure 31: Network Function VM Chain..... 32

Figure 32: Redirect Through a Service Chain.....33

Figure 33: Flows Detected in Monitor Mode..... 34

Figure 34: Flows Blocked in Apply Mode.....34

Figure 35: Application Policies Within a Subset.....35

List of Tables

Table 1: Document Version History..... 5





Prism 5.15

Prism Web Console Guide

May 15, 2020

Contents

Help Organization.....	viii

1. Nutanix Platform Overview.....	9

Guest VM Data Management.....	9
MapReduce Tiering.....	10
Live Migration.....	10
High Availability.....	10
Virtualization Management VM High Availability.....	11
Data Path Redundancy.....	11
Cluster Components.....	12
Zeus.....	12
Zookeeper.....	12
Medusa.....	13
Cassandra.....	13
Stargate.....	13
Curator.....	13
Prism.....	13
Availability Domains.....	14
Node Failure.....	14
Drive Failures.....	15
Network Link Failure.....	17
Rack Fault Tolerance.....	17
Block Fault Tolerance.....	22
Configuring Node Fault Tolerance.....	27
Redundancy Factor 3.....	28
Degraded Node.....	30
Maximum System Values.....	35
Three Node Cluster Considerations.....	36

2. Cluster Management.....	37

Web Console Overview.....	37
Logging Into the Web Console.....	38
Logging Out of the Web Console.....	46
Main Menu.....	46
Settings Menu.....	50
Home Dashboard.....	53
Understanding Displayed Statistics.....	56
Modifying Cluster Details.....	58
Modifying UI Settings.....	60
Checking Version.....	61
Licensing.....	61
Software and Firmware Upgrades.....	61
Life Cycle Management (LCM).....	62
Multi-Cluster Management.....	62
Register (Unregister) Cluster with Prism Central.....	62
Installing Prism Central (1-Click Internet).....	69
Single-Node Clusters.....	78



Single-Node Cluster Guidelines.....	78
Overriding Read-Only Mode.....	79
Two-Node Clusters.....	81
Two-Node Cluster Guidelines.....	82
Configuring a Witness (Two-Node Cluster).....	84
Witness VM Dashboard (two-node cluster).....	85
Failure and Recovery Scenarios.....	87
Increasing the Cluster Fault Tolerance Level.....	88
Increasing the Controller VM Memory Size.....	93
Restarting the Nodes in the Cluster (AHV and XenServer).....	96
3. Storage Management.....	97
Storage Components.....	97
Compression.....	99
Deduplication.....	100
Erasure Coding.....	101
Capacity Reservation Best Practices.....	103
Storage Dashboard.....	103
Storage Overview View.....	105
Storage Diagram View.....	107
Storage Table View.....	114
Creating a Storage Pool.....	125
Modifying a Storage Pool.....	126
Creating a Storage Container.....	126
Modifying a Storage Container.....	131
Volume Group Configuration.....	132
Concurrent Access from Multiple Clients.....	133
Creating a Volume Group.....	133
Modifying a Volume Group.....	135
Flash Mode for Virtual Machines and Volume Groups.....	135
Removing Flash Mode for Virtual Disks of a Volume Group.....	137
4. Network Management.....	138
Network Configuration for VM Interfaces.....	138
Configuring a Virtual Network for Guest VM Interfaces.....	138
Network Segmentation.....	144
Configuring Network Switch Information.....	144
Creating SNMP Profiles.....	148
Modifying Switch Information.....	150
Network Visualization.....	151
Prerequisites.....	152
Network Visualizer.....	152
Viewing the Network Visualizer.....	153
Customizing the Topology View.....	153
Viewing VM Information.....	154
Viewing Host Information.....	155
Viewing Switch Information.....	156
Uplink Configuration (AHV Only).....	156
5. Hardware Management.....	161
Hardware Dashboard.....	161
Hardware Overview View.....	163
Hardware Diagram View.....	165

Hardware Table View.....	173
Expanding a Cluster.....	185
Cluster Expansion Guidelines.....	190
Starting Host Boot Disk Repair (Proactive or Graceful Replacement).....	194
Modifying a Cluster.....	198
Compute-Only Node Configuration (AHV Only).....	201
Adding a Compute-Only Node to an AHV Cluster.....	203
Storage-only Node Configuration.....	204
Imaging and Adding Storage-only Nodes.....	206
6. Nutanix Volumes.....	207
Example Use Cases Supported by Volumes.....	208
If You are Already Using Volume Groups.....	208
Volumes Requirements and Supported Clients.....	208
Enabling Nutanix Volumes.....	210
Creating a Volume Group for Use with Volumes.....	211
Configuring Windows Clients.....	215
Configuring Linux Clients.....	222
Configuring AIX Clients.....	225
Bootting Over iSCSI.....	227
Converting Volume Groups and Updating Clients to Use Volumes.....	231
Modifying Windows Client Settings to Use Volumes.....	232
Modifying Linux Client Settings to Use Volumes.....	232
Modifying a Volume Group for Use with Volumes.....	233
Creating a Linux Guest Cluster.....	234
7. File Server Management.....	237
8. Data Protection.....	238
9. Health Monitoring.....	239
Health Dashboard.....	239
Configuring Health Checks.....	244
Configuring NCC Frequency.....	246
Running Checks by Using Web Console.....	247
Collecting Logs by Using Web Console.....	248
10. Virtual Machine Management.....	252
VM Dashboard.....	252
VM Overview View.....	254
VM Table View.....	255
VM Management.....	268
Creating a VM (AHV).....	268
Managing a VM (AHV).....	275
Creating a VM (ESXi).....	283
Managing a VM (ESXi).....	285
Configuring Images.....	289
Virtual Machine Customization.....	292
Customizing Linux Virtual Machines with Cloud-Init.....	293
Cloud-Init Limitations and Guidelines.....	293
Customization of Windows Virtual Machines with System Preparation.....	294



VM High Availability in Acropolis.....	295
Enabling High Availability for the Cluster.....	296
Nutanix Guest Tools.....	297
Nutanix Guest Tools Requirements and Limitations.....	298
Enabling NGT and Mounting the NGT Installer in a VM.....	299
NGT Installation.....	301
Enabling NGT and Mounting the NGT Installer Simultaneously on Multiple Cloned VMs.....	304
NGA and Controller VM Communication.....	305
CD-ROM Eject Functionality of NGT.....	306
Nutanix Guest Tools Usage in Disaster Recovery.....	306
Upgrading NGT.....	306
Reconfiguring NGT.....	307
Disabling, Uninstalling, and Removing Nutanix Guest Tools.....	307
Affinity Policies for AHV.....	308
Configuring VM-VM Anti-Affinity Policy.....	309
Connect to Citrix Cloud.....	310
Connecting to the Citrix Cloud.....	311
 11. Performance Monitoring.....	 317
Analysis Dashboard.....	317
Creating an Entity Chart.....	319
Creating a Metric Chart.....	320
Chart Metrics.....	321
Exporting Performance Data.....	326
 12. Alert and Event Monitoring.....	 327
Alerts Dashboard.....	327
Alerts Summary View.....	327
Events Summary View.....	333
Configuring Alert Emails.....	335
Configuring Alert Policies.....	338
Alerts/Health checks.....	340
Availability Zone.....	340
Cluster.....	340
Controller VM.....	426
DR.....	439
Entity Sync.....	482
Guest VM.....	482
Hardware.....	492
License.....	517
Node.....	517
Network.....	539
Nutanix Files.....	547
Other.....	563
Storage.....	574
System Indicator.....	590
VM.....	591
 13. Task Status.....	 592
View Task Status.....	593

14. System Management.....	594
Configuring a Filesystem Whitelist.....	594
Configuring Name Servers.....	595
Cluster Time Synchronization.....	596
Recommendations for Time Synchronization.....	596
Configuring NTP Servers.....	597
Configuring an SMTP Server.....	598
Configuring SNMP.....	599
Nutanix MIB.....	608
Configuring a Banner Page.....	615
VM Management through Prism Element (ESXi).....	616
Registering a vCenter Server.....	617
Managing vCenter Server Registration Changes.....	619
In-Place Hypervisor Conversion.....	620
Requirements and Limitations for In-Place Hypervisor Conversion.....	621
In-Place Hypervisor Conversion Process.....	624
Converting Cluster (ESXi to AHV).....	625
Converting Cluster (AHV to ESXi).....	626
Aborting Cluster Conversion.....	628
Internationalization (i18n).....	629
Localization (L10n).....	629
Changing the Language Settings.....	630
Hyper-V Setup.....	631
Joining the Cluster and Hosts to a Domain.....	631
Creating a Failover Cluster for Hyper-V.....	633
Enabling Kerberos for Hyper-V.....	634
15. Security and User Management.....	639
16. Support Services.....	640
Pulse Health Monitoring.....	640
Configuring Pulse.....	642
Pulse Access Requirements.....	644
Controlling Remote Connections.....	645
Configuring HTTP Proxy.....	646
Accessing the Nutanix Support Portal.....	648
Nutanix REST API.....	651
Accessing the REST API Explorer.....	652
17. Help Resources.....	656
Accessing Online Help.....	656
Accessing the Nutanix Next Community.....	660
Glossary.....	661



Copyright.....	667
License.....	667
Conventions.....	667





HELP ORGANIZATION

This documentation is organized as follows:


- [Nutanix Platform Overview](#) on page 9 describes the Nutanix architecture.
- [Cluster Management](#) on page 37 describes how to access and use the web console, how to apply a Nutanix cluster license, how to upgrade the cluster to a later AOS release, and how to upgrade other software components such as disk firmware.
- [Multi-Cluster Management](#) on page 62 describes how to install a VM that runs an application called Prism Central and how to use Prism Central to monitor and manage multiple clusters.
- [Storage Management](#) on page 97 describes how to monitor storage use in a Nutanix cluster and how to create storage pools and containers.
- [Hardware Management](#) on page 161 describes how to monitor hardware configurations in a Nutanix cluster and how to expand the cluster.
- [Health Monitoring](#) on page 239 describes how to monitor the health of VMs, hosts, and disks across a Nutanix cluster.
- [Virtual Machine Management](#) on page 252 describes how to monitor status of the VMs across a Nutanix cluster.
- [Performance Monitoring](#) on page 317 describes how to monitor and analyze performance in a Nutanix cluster.
- [Alert and Event Monitoring](#) on page 327 describes how to monitor activity in a Nutanix cluster and how to configure alert policies and notification.
- [System Management](#) on page 594 describes how to configure various system settings such as for SNMP, NTP, and SMTP.
- [Security and User Management](#) on page 639 describes how to configure various security settings including authentication method, SSL certificates, and SSH keys. It also describes how to add, edit, and delete user accounts.
- [Support Services](#) on page 640 describes how to enable (or disable) Nutanix technical support access to your cluster, how to access the Nutanix support portal, and how to access the Nutanix REST API explorer.

NUTANIX PLATFORM OVERVIEW


The Nutanix Enterprise Cloud Platform is a converged, scale-out compute and storage system that is purpose-built to host and store virtual machines. All nodes in a Nutanix cluster converge to deliver a unified pool of tiered storage and present resources to VMs for seamless access. A global data system architecture integrates each new node into the cluster, allowing you to scale the solution to meet the needs of your infrastructure.

The foundational unit for the cluster is a Nutanix node. Each node in the cluster runs a standard hypervisor and contains processors, memory, and local storage (SSDs and hard disks).

A Nutanix Controller VM runs on each node, enabling the pooling of local storage from all nodes in the cluster.



Guest VM Data Management



Hosts read and write data in shared Nutanix datastores as if they were connected to a SAN. From the perspective of a hypervisor host, the only difference is the improved performance that results from data not traveling across a network. VM data is stored locally, and replicated on other nodes for protection against hardware failure.

When a guest VM submits a write request through the hypervisor, that request is sent to the Controller VM on the host. To provide a rapid response to the guest VM, this data is first stored on the metadata drive, within a subset of storage called the oplog. This cache is rapidly

distributed across the 10 GbE network to other metadata drives in the cluster. Oplog data is periodically transferred to persistent storage within the cluster. Data is written locally for performance and replicated on multiple nodes for high availability.


When the guest VM sends a read request through the hypervisor, the Controller VM reads from the local copy first, if present. If the host does not contain a local copy, then the Controller VM reads across the network from a host that does contain a copy. As remote data is accessed, the remote data is migrated to storage devices on the current host, so that future read requests can be local.

MapReduce Tiering

The Nutanix cluster dynamically manages data based on how frequently it is accessed. When possible, new data is saved on the SSD tier. Frequently-accessed, or "hot" data is kept on this tier, while "cold" data is migrated to the HDD tier. Data that is accessed frequently is again moved back to the SSD tier.

This automated data migration also applies to read requests across the network. If a guest VM repeatedly accesses a block of data on a remote host, the local controller VM migrates that data to the SSD tier of the local host. This migration not only reduces network latency, but also ensures that frequently-accessed data is stored on the fastest storage tier.


Live Migration



Live migration of VMs, whether it is initiated manually or through an automatic process like vSphere DRS, is fully supported by the Nutanix Enterprise Cloud Computing Platform. All hosts within the cluster have visibility into shared Nutanix datastores through the Controller VMs. Guest VM data is written locally, and is also replicated on other nodes for high availability.

If a VM is migrated to another host, future read requests are sent to a local copy of the data, if it exists. Otherwise, the request is sent across the network to a host that does contain the requested data. As remote data is accessed, the remote data is migrated to storage devices on the current host, so that future read requests can be local.

High Availability



The built-in data redundancy in a Nutanix cluster supports high availability provided by the hypervisor. If a node fails, all HA-protected VMs can be automatically restarted on other nodes in the cluster. The hypervisor management system, such as vCenter, selects a new host for the VMs, which may or may not contain a copy of the VM data.

If the data is stored on a node other than the VM's new host, then read requests are sent across the network. As remote data is accessed, the remote data is migrated to storage devices on the current host, so that future read requests can be local. Write requests are sent to the local storage, and also replicated on a different host. During this interaction, the Nutanix software also creates new copies of pre-existing data, to protect against future node or disk failures.


Virtualization Management VM High Availability

In virtualization management VM high availability, when a node becomes unavailable, VMs that are running on that node are restarted on another node in the same cluster.

Typically, an entity failure is detected by its isolation from the network (the failure to respond to heartbeats). Virtualization management ensures that at most one instance of the VM is running at any point during a failover. This property prevents concurrent network and storage I/O that could lead to corruption.

Virtualization management VM high availability may implement admission control to ensure that in case of node failure, the rest of the cluster has enough resources to accommodate the VMs.

Data Path Redundancy




The Nutanix cluster automatically selects the optimal path between a hypervisor host and its guest VM data. The Controller VM has multiple redundant paths available, which makes the cluster more resilient to failures.

When available, the optimal path is through the local Controller VM to local storage devices. In some situations, the data is not available on local storage, such as when a guest VM was recently migrated to another host. In those cases, the Controller VM directs the read request across the network to storage on another host through the Controller VM of that host.

Data Path Redundancy also responds when a local Controller VM is unavailable. To maintain the storage path, the cluster automatically redirects the host to another Controller VM. When the local Controller VM comes back online, the data path is returned to this VM.

Cluster Components



The Nutanix cluster has a distributed architecture, which means that each node in the cluster shares in the management of cluster resources and responsibilities. Within each node, there are software components that perform specific tasks during cluster operation.

All components run on multiple nodes in the cluster, and depend on connectivity between their peers that also run the component. Most components also depend on other components for information.

Zeus

A key element of a distributed system is a method for all nodes to store and update the cluster's configuration. This configuration includes details about the physical components in the cluster, such as hosts and disks, and logical components, like storage containers. The state of these components, including their IP addresses, capacities, and data replication rules, are also stored in the cluster configuration.

Zeus is the Nutanix library that all other components use to access the cluster configuration, which is currently implemented using Apache Zookeeper.

Zookeeper

Zookeeper runs on either three or five nodes, depending on the redundancy factor that is applied to the cluster. Using multiple nodes prevents stale data from being returned to other components, while having an odd number provides a method for breaking ties if two nodes have different information.

Of these three nodes, one Zookeeper node is elected as the leader. The leader receives all requests for information and confers with the two follower nodes. If the leader stops responding, a new leader is elected automatically.

Zookeeper has no dependencies, meaning that it can start without any other cluster components running.

Medusa

Distributed systems that store data for other systems (for example, a hypervisor that hosts virtual machines) must have a way to keep track of where that data is. In the case of a Nutanix cluster, it is also important to track where the replicas of that data is stored.

Medusa is a Nutanix abstraction layer that sits in front of the database that holds this metadata. The database is distributed across all nodes in the cluster, using a modified form of Apache Cassandra.

Cassandra

Cassandra is a distributed, high-performance, scalable database that stores all metadata about the guest VM data stored in a Nutanix datastore. In the case of NFS datastores, Cassandra also holds small files saved in the datastore. When a file reaches 512K in size, the cluster creates a vDisk to hold the data.

Cassandra runs on all nodes of the cluster. These nodes communicate with each other once a second using the Gossip protocol, ensuring that the state of the database is current on all nodes.

Cassandra depends on Zeus to gather information about the cluster configuration.

Stargate

A distributed system that presents storage to other systems (such as a hypervisor) needs a unified component for receiving and processing data that it receives. The Nutanix cluster has a large software component called Stargate that manages this responsibility.

From the perspective of the hypervisor, Stargate is the main point of contact for the Nutanix cluster. All read and write requests are sent across vSwitchNutanix to the Stargate process running on that node.

Stargate depends on Medusa to gather metadata and Zeus to gather cluster configuration data.

Tip: If Stargate cannot reach Medusa, the log files include an HTTP timeout. Zeus communication issues can include a Zookeeper timeout.

Curator

In a distributed system, it is important to have a component that watches over the entire process. Otherwise, metadata that points to unused blocks of data could pile up, or data could become unbalanced, either across nodes, or across disk tiers.

In the Nutanix cluster, each node runs a Curator process that handles these responsibilities. A Curator master node periodically scans the metadata database and identifies cleanup and optimization tasks that Stargate or other components should perform. Analyzing the metadata is shared across other Curator nodes, using a MapReduce algorithm.

Curator depends on Zeus to learn which nodes are available, and Medusa to gather metadata. Based on that analysis, it sends commands to Stargate.

Prism

A distributed system is worthless if users can't access it. Prism provides a management gateway for administrators to configure and monitor the Nutanix cluster. This includes the nCLI and web console.



Prism runs on every node in the cluster, and like some other components, it elects a leader. All requests are forwarded from followers to the leader using Linux iptables. This allows administrators to access Prism using any Controller VM IP address. If the Prism leader fails, a new leader is elected.

Prism communicates with Zeus for cluster configuration data and Cassandra for statistics to present to the user. It also communicates with the ESXi hosts for VM status and related information.

Availability Domains

Hardware failures are a part of any datacenter lifecycle. The Nutanix architecture was designed with this inevitability in mind. A cluster can tolerate one or two failures (depending on the replication factor of the cluster or container) of a variety of hardware components while still running guest VMs and responding to commands through the management console. Many of these failures also trigger an alert through that same management console in order to give the administrator a chance to respond to the situation.


Nutanix provides the ability to tolerate rack failures for extended data availability, in addition to drive, node, block, and network link failure.

- [Rack Fault Tolerance](#) on page 17
- [Block Fault Tolerance](#) on page 22
- [Configuring Node Fault Tolerance](#) on page 27

Node Failure

A Nutanix node is comprised of a physical host and a Controller VM. Either component can fail without impacting the rest of the cluster.

Controller VM Failure



The Nutanix cluster monitors the status of Controller VMs in the cluster. If any Stargate process fails to respond two or more times in a 30-second period, another Controller VM redirects the storage path on the related host to another Stargate. Reads and writes occur over the 10 GbE network until the missing Stargate comes back online.

To prevent constant switching between Stargates, the data path is not restored until the original Stargate has been stable for 30 seconds.

What Will Users Notice?


During the switching process, the host with a failed Controller VM may report that the shared storage is unavailable. Guest VMs on this host appear to "hang" until the storage path is restored. Although the primary copy of the guest VM data is unavailable because it is stored on disks mapped to the failed Controller VM, the replicas of that data are still accessible.

As soon as the redirection takes place, VMs can resume reads and writes. Performance may decrease slightly, because the IO is traveling across the network, rather than across the internal network. Because all traffic goes across the 10 GbE network, most workloads do not diminish in a way that is perceivable to users.

What Happens If Another One Fails?

A second Controller VM failure has the same impact on the VMs on the other host, which means there will be two hosts sending IO requests across the network. More importantly, however, is the additional risk to guest VM data. With two Controller VMs unavailable, there are now two sets of physical disks that are inaccessible. In a cluster with a replication factor 2, there is now a chance that some VM data extents are missing completely, at least until one of the Controller VMs resumes operation.

Host Failure



The built-in data redundancy in a Nutanix cluster supports high availability provided by the hypervisor. If a node fails, all HA-protected VMs can be automatically restarted on other nodes in the cluster.

Curator and Stargate responds to two issues that arise from the host failure. First, when the guest VM begins reading across the network, Stargate begins migrating those extents to the new host. This improves performance for the guest VM. Second, Curator notices that there is a missing replica of those extents, and instruct Stargate to begin creating a second replica.

What Will Users Notice?

Users who are accessing HA-protected VMs will notice that their VM is unavailable while it is restarted on the new host. Without HA, the VM will need to be manually restarted.

What Happens If Another One Fails?

Depending on how fully loaded the cluster is, a second host failure could leave the remaining hosts with insufficient processing power to restart the VMs from the second host. However, even in a lightly-loaded cluster, the bigger concern is additional risk to guest VM data. With two sets of physical disks that are inaccessible, there would be a chance that some VM data extents are missing completely and the IO requests are not served.

Drive Failures

Drives in a Nutanix node store four primary types of data: persistent data (hot-tier and cold-tier), storage metadata, oplog, and Controller VM boot files. Cold-tier persistent data is stored in the hard-disk drives of the node. Storage metadata, oplog, hot-tier persistent data, and Controller VM boot files are kept in the SATA-SSD in drive bay one. SSDs in a dual SSD system are used for storage metadata, oplog, hot-tier persistent data according to the replication factor of system and in a RAID-1 configuration for CVM files. In all-flash nodes, data of all types is stored in the SATA-SSDs.

Note: On hardware platforms that contain PCIe-SSD drives, the SATA-SSD holds the only the Controller VM boot files. Storage metadata, oplog, and hot-tier persistent data reside on the PCIe-SSD.

Boot Drive Failure

Each Controller VM boots from a SATA-SSD. During cluster operation, this drive also holds component logs and related files.

A boot drive failure eventually causes the Controller VM to fail. The host does not access the boot drive directly, so other guest VMs can continue to run. Data Path Redundancy redirects the storage path to another Controller VM. For more information, see [Controller VM Failure](#) on page 14.

Note: The Controller VM restarts if a boot drive fails, or if you remove a boot drive without marking the drive for removal and the data has not successfully migrated.

Metadata Drive Failure

The metadata "drive" (metadata spans multiple drives) serves many purposes. It holds the oplog for each host, which provides a fast response to VMs that send write requests. It is used as a persistent data tier. It is also used by the Cassandra database to provide fast read and write responses for cluster metadata. To protect against potential failure, the metadata is replicated to metadata drives in other hosts in the cluster.

Metadata is sharded, and Cassandra uses up to four SSDs to store metadata. If any metadata drive fails on a host, the Controller VM restarts. Once the Cassandra process restarts, the missing metadata is retrieved from the other Controller VMs and sharded across the remaining metadata drives. When the faulty drive recovers or is replaced, metadata is stored on that drive again.

If the Cassandra process is down on a node for more than 30 minutes, the surviving Cassandra nodes detach the node from the Cassandra database so that the unavailable metadata can be replicated on other nodes. The process of healing the database takes about 30-40 minutes. If the Cassandra process restarts and remains up for 30 minutes, the procedure to detach the node is canceled if the healing procedure is still running. If the process resumes and is stable after the healing procedure is complete, the node can be manually added to the database using the nCLI command:

```
ncli> host enable-metadata-store id=host_id
```

What Will Users Notice?

Performance may decrease slightly for user VMs on the host due to the reboot and the fact that some I/O is traveling across the network. However, most workloads should not diminish in a way that is perceivable to users (other than during the reboot).

What Happens If Another One Fails?

Multiple drive failures in a single selected domain (node, block, or rack) is tolerated.

Data Drive Failure

Each node contributes data drives to the cluster storage pool. Cold-tier data is stored in HDDs, while hot-tier data is stored in SSDs for faster performance. Because the HDDs have moving parts, and outnumber any other hardware component, this is the most likely component to experience a failure. Data is replicated across the cluster, so a single hard-disk drive failure does not result in data loss. In all-flash nodes, data of all types is stored in SSDs.



The cluster software receives a hardware alert from the host that a data drive (HDD or SSD) has failed, and immediately begin working to reduce the impact of a second failure. Curator instructs Stargate to create a second replica of any guest VM data that was stored on the drive.

Note: The Controller VM restarts if a data drive fails, or if you remove a data drive without marking the drive for removal and the data has not successfully migrated.

What Will Users Notice?

For a brief period of time, guest VMs with files on the failed data drive will need to read across the network. Curator will eventually instruct Stargate to migrate the relevant data to another drive on the current host.

What Happens If Another One Fails?

In a cluster with a replication factor 2, losing two drives on different different domains (node, block, or rack) means that some VM data extents could lose both replicas. Although a single drive failure does not have the same impact as a host failure, it is important to replace the failed drive as soon as possible.

Network Link Failure

The physical network adapters on each host are grouped together on the external network . Failure of a network link is tolerated with no impact to users if both 10 GbE ports are connected to the customer's network. If the cluster must fail over to a 1 GbE link, then write performance decreases.

What Will Users Notice?

Each Nutanix node is configured at the factory to use one 10 GbE port as the primary pathway for vSwitch0. The other 10 GbE port is configured in standby mode. In this configuration, guest VM performance does not decrease. If a 10 GbE port is not configured as the failover path, then traffic will fail over to a 1 GbE port. This failover reduces the throughput of storage traffic, and decreases the write performance for guest VMs on the host with the failed link. Other hosts may experience a slight decrease as well, but only on writes to extents that are stored on the host with the link failure.

What Happens if Another One Fails?

If both 10 GbE links are down, then the host will fail over to a 1 GbE port, if it is configured as a standby path. This failover reduces the throughput of storage traffic, and decreases the write performance for guest VMs on the host with the failed link. Other hosts may experience a slight decrease as well, but only on writes to extents that are stored on the host with the link failure.

Rack Fault Tolerance

Rack fault tolerance is the ability to provide rack level availability domain. With rack fault tolerance, redundant copies of data are made and placed on the nodes that are not in the same rack.

Rack failure can occur in the following situations:

- All power supplies fail within a rack
- Top-of-rack (TOR) switch fails
- Network partition; where one of the racks becomes inaccessible from other racks



When rack fault tolerance is enabled, the cluster has rack awareness and the guest VMs can continue to run with failure of one rack (RF2) or two racks (RF3). The redundant copies of guest VM data and metadata exist on other racks when one rack fails.

Note:

- If sufficient number of working racks are not present after failure, the data is rebuilt in a non-rack-aware manner to get node fault-tolerance level to one. The rack-fault-tolerance stays 0.
- Rack Fault Tolerance is supported for AHV and ESXi only.

Configuring Rack Fault Tolerance

To configure rack fault tolerance, do the following.

About this task

To enable rack fault tolerance, you must specify the mapping of the blocks to the racks based on the actual placement of the blocks in the datacenter.

Before you begin


- You must have information on the actual physical mapping of racks and blocks in the datacenter.
- Minimum cluster requirements:
 - Replication factor 2 - 3 racks (4 with Erasure Coding)
 - Replication factor 3 - 5 racks (6, with Erasure Coding), 1 node in each rack

Procedure

1. Go to **Settings > Setup > Rack Configuration..**

Note: Directory List (AD and OpenLDAP) users can view and configure Rack Awareness only if a service account is configured for the directory service. To configure a service account for the directory service, see "" () for details.

2. Select **Rack** for **Choose the fault tolerance domain**.



The screenshot shows a configuration window with two tabs at the top: '1 Fault Tolerance Level' (active) and '2 Block Assignment'. Below the tabs, the text 'Choose the fault tolerance domain' is displayed. Underneath, there are three radio button options: 'Node', 'Block', and 'Rack'. The 'Rack' option is selected, indicated by a blue dot inside its radio button.

Figure 1: Select Fault Tolerance Level

Click **Add New Rack** and click **Next**.

3. Click **Add Blocks** to assign the discovered blocks to the rack. Click the plus icon to add the block to the rack based on the actual placement of the block in the datacenter.




Figure 2: Add blocks to rack

Click **Add New Rack** and click **Next**.

Repeat steps 3 to 5 for adding more rack configurations. Verify the configuration and click **Save**.

Note:

- Once rack fault tolerance is enabled; create, edit, delete operations are not allowed. To make any new modifications, you must re-configure the fault tolerance level as **Node**, see [Configuring Node Fault Tolerance](#) on page 27.
- Enabling rack fault tolerance may trigger cassandra ring change operations and trigger curator scans to redistribute replicas for rack fault tolerance. The time

taken for these depend on existing data on cluster, workload on the cluster, and configurations.

The **Data Resiliency** widget on the Prism dashboard shows the current state of the fault tolerance.




Figure 3: Data Resiliency widget

Data Resiliency Levels for Rack Fault Tolerance

Table 1: Data Resiliency Levels

The following table shows the level of data resiliency (simultaneous failure) provided for the following combinations of replication factor, minimum number of nodes, minimum number of blocks, and minimum number of racks.

Fault Domain	Replication Factor	Minimum Number of Nodes	Minimum Number of Blocks	Minimum Number of Racks	Data Resiliency
Node	2	3*	1	1	1 node or 1 disk failure
	3	5*	2	1	2 nodes or 2 disk failures
Block	2	3	3*	1	1 node or 1 block or 1 disk failure
	3	5	5*	1	2 nodes or 2 blocks or 2 disk failures
Rack	2	3	3	3*	1 node or 1 block or 1 rack or 1 disk

Fault Domain	Replication Factor	Minimum Number of Nodes	Minimum Number of Blocks	Minimum Number of Racks	Data Resiliency
	3	5	5	5*	2 nodes or 2 blocks or 2 racks or 2 disks

* Erasure Coding with Rack Awareness - Erasure coding is supported on a rack aware cluster. You can enable erasure coding on new containers in rack aware clusters provided certain minimums are met that are shown in the table above.

Block Fault Tolerance

Block fault tolerance lets a Nutanix cluster make redundant copies of data and metadata and place the copies on nodes in different blocks.

A block is a rack-mountable enclosure that contains one to four Nutanix nodes. All nodes in a block share power supplies, front control panels (ears), backplane, and fans.

Nutanix offers block fault tolerance as an opt-in procedure, as in [Configuring Block Fault Tolerance](#) on page 22, or a best-effort procedure, as in [Block Fault Tolerance in Best Effort mode](#) on page 23. The opt-in block fault tolerance feature offers guaranteed data resiliency when required conditions are met. For best-effort fault tolerance mode, data copies remain on the same block when there is insufficient space across all blocks.

With block fault tolerance enabled, guest VMs can continue to run after a block failure because redundant copies of guest VM data and metadata exist on other blocks.

Configuring Block Fault Tolerance

To configure block fault tolerance, do the following.

About this task

Procedure

1. Go to **Settings > Setup > Rack Configuration..**
2. Select **Block** for **Choose the fault tolerance domain**.




Figure 4: Select Fault Tolerance Level

3. Optionally, click **Add New Rack** to specify rack and block mapping.

Block Fault Tolerance in Best Effort mode

When certain conditions are met, Nutanix clusters become block fault tolerant. Block fault tolerance is applied automatically when:

- Every storage tier in the cluster contains at least one drive on each block.
- Every storage container in the cluster has replication factor of at least two.
- For replication factor 2, there are a minimum of three blocks in the cluster.
- There is enough free space in all the tiers, in at least replication factor number of blocks in the cluster. For example, if the replication factor of storage containers in the cluster is two, then at least two blocks require free space.
- A minimum of four blocks for RF2 or six blocks for RF3 is required to maintain block awareness if erasure coding is enabled on any storage container. (If the cluster has fewer blocks, block awareness is lost when erasure coding is enabled.)

Note: This is not applicable for single-node replication target clusters. For more information on how single-node replication target clusters handle failures, see Single-Node Replication Target Clusters in the guide titled *Data Protection and Recovery with Prism Element*.

Once block fault tolerance conditions are met, the cluster can tolerate a specific number of block failures:

- A replication factor two or replication factor three cluster with three or more blocks can tolerate a maximum failure of one block.
- A replication factor three cluster with five or more blocks can tolerate a maximum failure of two blocks.

Block fault tolerance is one part of a resiliency strategy. It does not remove other constraints such as the availability of disk space and CPU/memory resources in situations where a significant proportion of the infrastructure is unavailable.

Metadata Block Awareness Requirements for Block Fault Tolerance

Metadata block awareness is required for block fault tolerance. Metadata block awareness can be enabled for replication factor 2 and replication factor 3. To enable metadata block fault tolerance, your metadata must meet the following requirements.

Table 2: Minimum Cluster Requirements

Replication Factor	Minimum Number of Blocks and Nodes Per Block Required
Replication factor 2	• 3 blocks • 1 node per block
Replication factor 3	• 5 blocks • 1 node per block



Table 3: Additional Requirements when Adding Nodes to an Existing Block Aware Cluster

Replication Factor	Requirement	Example
Note: Be sure your cluster has met the previous minimum cluster requirements.		
Replication factor 2	There must be at least 3 blocks populated with a specific number of nodes to maintain block fault tolerance. To calculate the number of nodes required to maintain block fault tolerance when the cluster RF=2, you need twice the number of nodes as there are in the block with the most or maximum number of nodes.	If a block contains 4 nodes, you need 8 nodes distributed across the remaining (non-failing) blocks to maintain block fault tolerance for that cluster. $X = \text{number of nodes in the block with the most nodes}$. In this case, 4 nodes in a block. $2X = 8$ nodes in the remaining blocks.
Replication factor 3	There must be at least 5 blocks populated with a specific number of nodes to maintain block fault tolerance. To calculate the number of nodes required to maintain block fault tolerance when the cluster replication factor 3 you need four times the number of nodes as there are in the block with the most or maximum number of nodes.	If a block contains 4 nodes, you need 16 nodes distributed across the remaining (non-failing) blocks to maintain block fault tolerance for that cluster. $X = \text{number of nodes in the block with the most nodes}$. In this case, 4 nodes in a block. $4X = 16$ nodes in the remaining blocks

Block Fault Tolerant Data Placement

Stargate is responsible for placing data across blocks, and Curator makes data placement requests to Stargate to maintain block fault tolerance.

New and existing clusters can reach a block fault tolerant state. New clusters can be block fault tolerant immediately after being created if the configuration supports it. Existing clusters that were not previously block fault tolerant can be made tolerant by reconfiguring the cluster in a manner that supports block fault tolerance.

New data in a block fault tolerant cluster is placed to maintain block fault tolerance. Existing data that was not in a block fault tolerant state is moved and scanned by Curator to a block fault tolerant state.

Depending on the volume of data that needs to be relocated, it might take Curator several scans over a period of hours to distribute data across the blocks.


Block fault tolerant data placement is on a best effort basis but is not guaranteed. Conditions such as high disk usage between blocks may prevent the cluster from placing guest VM redundant copy data on other blocks.

Guest VM Data

Redundant copies of guest VM data are written to nodes in blocks other than the block that contains the node where the VM is running. The cluster keeps two copies of each write stored in the oplog.



Redundant copies of the guest VM data (designated by ■ ● ▲) are placed on different



blocks.

Figure 5: Block-aware placement of guest VM data

In the case of a block failure, the under-replicated guest VM data is copied to other blocks in the cluster, and one copy of the oplog contents is available.




Figure 6: Block-aware placement of guest VM data with block failure

Metadata

The Nutanix Medusa component uses Cassandra to store metadata. Cassandra uses a ring-like structure where data is copied to peers within the ring to ensure data consistency and availability. The cluster keeps at least three redundant copies of the metadata, at least half of which must be available to ensure consistency.

With block fault tolerance, the Cassandra peers are distributed among the blocks to ensure that no two peers are on the same block. In the event of a block failure, at least two copies of the metadata is present in the cluster .




Figure 7: Block-aware placement of metadata

Configuration Data

The Nutanix Zeus component uses Zookeeper to store essential configuration data for the cluster. The Zookeeper role is distributed across blocks to ensure availability in the case of a block failure.




Figure 8: Block-aware placement of configuration data

Data Resiliency Levels for Block Fault Tolerance

Table 4: Data Resiliency Levels

The following table shows the level of data resiliency (simultaneous failure) provided for the following combinations of replication factor, minimum number of nodes, and minimum number of blocks.

Replication Factor	Minimum Number of Nodes	Minimum Number of Blocks	Data Resiliency
2	3	1	1 node or 1 disk failure
2	3	3 (minimum 1 node each)	1 block or 1 node or 1 disk failure
3	5	2	2 nodes or 2 disk failures

Replication Factor	Minimum Number of Nodes	Minimum Number of Blocks	Data Resiliency
3	5	5 (minimum 1 node each)	2 blocks or 2 nodes or 2 disks
3	6	3 (minimum 2 nodes each)	1 block or 2 nodes or 2 disks
Metro Cluster (see See "Metro Availability" in the "Data Protection" chapter of the <i>Prism Web Console Guide</i> .)	3 nodes at each site	2	1 cluster failure

The state of block fault tolerance is available for viewing through the Prism Web Console and Nutanix CLI. Although administrators must set up the storage tiers or storage containers, they cannot determine where data is migrated. AOS determines where data is migrated.

Web Console Data Resiliency Status view on the **Home** screen

nCLI `nccli> cluster get-domain-fault-tolerance-status type="rackable_unit"`

Configuring Node Fault Tolerance

To configure node fault tolerance, do the following.

About this task

Procedure

1. Go to **Settings > Setup > Rack Configuration**.
2. Select **Node** for **Choose the fault tolerance domain**.




Figure 9: Select Fault Tolerance Level

3. Optionally, click **Add New Rack** to specify rack and block mapping.
It recommended to configure rack and block mapping if you want to enable rack awareness at a later stage.

Redundancy Factor 3

Redundancy factor 3 is a configurable option that allows a Nutanix cluster to withstand the failure of two nodes or drives in different blocks.

By default, Nutanix clusters have redundancy factor 2, which means they can tolerate the failure of a single node or drive. The larger the cluster, the more likely it is to experience multiple failures. Without redundancy factor 3, multiple failures cause cluster unavailability until the failures are repaired.

Redundancy factor 3 has the following requirements:

- A cluster must have at least five nodes, blocks, racks for redundancy factor 3 to be enabled.
- For guest VMs to tolerate the simultaneous failure of two nodes or drives in different blocks, the data must be stored on storage containers with replication factor 3.

The state of fault tolerance is available to view through the management interfaces.


Web Console Data Resiliency Status view on the **Home** screen

nCLI `nccli> cluster get-redundancy-state`

Guest VM Data

For storage containers with replication factor 3, the cluster stores three redundant copies of guest VM data and the oplog.

Redundant copies of the guest VM data (designated by ■ ● ▲) are stored on different



nodes in the cluster.

Figure 10: Replication factor 3 placement of guest VM data

In the case of two nodes failing, at least one copy of all guest VM data, including the oplog, is available. Under-replicated VM data is copied to other nodes.




Figure 11: Replication factor 3 placement of guest VM data with failure of 2 nodes

Metadata

At least half of the redundant copies of metadata must be available to ensure consistency. Without redundancy factor 3, the cluster keeps three copies of metadata. With redundancy factor 3, the cluster keeps five copies of metadata so that if two nodes fail, at least three copies are available.




Figure 12: Redundancy factor 3 metadata

Configuration Data

The Nutanix Zeus component uses Zookeeper to store essential configuration data for the cluster. Without redundancy factor 3, the cluster keeps three copies of configuration data. With redundancy factor 3, the cluster keeps five copies of configuration data so that if two nodes fail, at least three copies are available.




Figure 13: Redundancy factor 3 configuration data

Degraded Node

A degraded node contains a CVM that is not entirely in an unresponsive state, but cannot reliably perform normal cluster operations. A degraded (or partially available) node can adversely affect the performance of an entire cluster. Degraded node detection is introduced to mitigate the results of a degraded node. Continuing to run a degraded node can affect overall cluster and user VM performance. By enabling degraded node detection, ZooKeeper places the node into maintenance mode, which moves it from the leadership position and distributes tasks automatically.

The alert for a degraded node appears in the Alerts page of the web console.

The following events might contribute to node degradation.

- Network bandwidth reduction
- Network packet drops
- Soft lockups
- Partially bad disks
- Hardware issues (such as unreliable DIMM with ECC errors)

Note: In the event of a degraded node, contact Nutanix Support for help investigating the issue. A support case will be automatically generated if Pulse is enabled.

Overview

Degraded node detection consists of a peer health global database. These services run on each node of the cluster. Each node publishes scores or votes for the service running on other nodes. The peer scores evaluate metrics such as:

- Remote Procedure Call (RPC) latency
- RPC failures or timeouts
- Network latency

If one node consistently receives poor scores for approximately 10 minutes then the peers mark that node as a degraded node. Clustering algorithms are used to identify outlier scores.

Note:

- A node is not marked as *Degraded* if the current cluster fault tolerance is less than the desired value.
- Containers or data stores might be unavailable for 10 minutes until the node is marked as degraded.
- Upgrades and break fixes are not allowed until the degraded node is fixed.

Detection

The web console alerts the user when a node begins to receive poor peer scores. The node is marked as degraded after a few minutes of poor scores.

The degraded node detection software mitigates the effects of failures and reduces the overall impact to the cluster. To mitigate the impact, this software can perform one of the following actions. When prompted, select one of the following.

- Prevent components on the degraded node from acquiring leadership roles
- Place the degraded node CVM in maintenance mode and reboots the CVM to stop the services
- Shut down the host

Enabling Degraded Node Detection

Enable the degraded node detection settings in the web console.

About this task

Procedure

1. Log into the web console for your cluster.
2. Click the gear icon in the main menu and then select **Degraded Node Settings** in the **Settings** page.
The Degraded Node Settings dialog box appears.




Figure 14: Degraded Node Settings Window

3. Set the degraded node detection status, **Enable Degraded Node Detection**.
 - a. Be sure to check the box to enable degraded node detection.
 - b. Otherwise, be sure to clear the box to disable degraded node detection.
4. Click **Save** to confirm the settings.

The Degraded Node Settings window closes and the web console displays a pop-up confirmation.

Managing a Degraded Node

Manage a node if it is marked as degraded.

About this task

Once a node is detected as degraded, the leadership and critical services will not be hosted on that node.

Until a degraded node is unmarked from the degraded state, Cassandra services remain in the forwarding state, see [Cluster Components](#) on page 12. Once the degraded node is marked as fixed in Prism, Cassandra restarts its services on the node.

The alert for a degraded node appears in the Alerts page of the web console.



Procedure

1. Open the degraded node window.
 - a. Click the **Alerts** icon at the top of the page.
 - b. Click the alert for the degraded node.




Figure 15: Possible Degraded Node in Alerts

The degraded node window is displayed.

2. Set the status for the degraded node.

a. **Resolve:** Select this option after you have fixed the underlying issue.

Note: Clicking **Resolve** does not resolve the degraded node issue. It merely marks the alert status as **Resolved**, indicating you or Nutanix Support have resolved the underlying issue that triggered the alert.

b. **Acknowledge:** Select this option to leave the degraded node in its current state, keep the alert unresolved, and mark the alert as acknowledged.




Figure 16: Possible Degraded Node Window

3. You can also view the degraded node in the Hosts table.

a. Go to **Home > Hardware > Host**.

b. Select the degraded node which displays a critical alert.




Figure 17: Degraded Node Pop-up Alert

4. To run an NCC check for the degraded node, go to the command line interface and enter

```
nutanix@cvm$ ncc health_checks system_checks degraded_node_check
```

The web console does not remove the warning sign for the degraded node.

5.

Note: Do not unmark the degraded node if you have not confirmed that the issue causing the node or CVM degradation has been fixed. Unmarking a node can lead to production impacting issues in the cluster. Contact Nutanix Support for assistance.

To unmark the degraded node, in the **Summary** view, select the fixed node and click **Mark as Fixed**.

Maximum System Values

The figures listed here are the maximum tested and supported values for entities in a Nutanix cluster.

Entity	Supported Maximum
VMs or volume groups	Up to available hardware resources and hypervisor limitations (no known logical limit) • Async DR: 200 VMs or volume groups for each protection domain, or consistency group • Metro Availability: See the Maximum Limits for Metro Availability table in the "Data Protection Guidelines (Metro Availability)" section of the <i>Data Protection and Recovery with Prism Element</i> guide. • 1200 files for any single consistency group and 3600 files for all entities per protection domain
vDisks (including snapshots)	600,000
vDisk size (limited by hypervisor or guest OS)	ESXi 5.5 and above = 62 TB ESXi 5.1 and below = 2 TB minus 512 B Hyper-V vhdx format = 64TB HyperV vhd format = 2TB AHV = No known limit
vCPUs per VM	Constrained by the number of logical processors available to the VM.
Memory per VM	Constrained by physical RAM capacity.
Memory per host	AHV = 3 TB ESXi = Limit as stated by VMware. Hyper-V = Limit as stated by Microsoft.
RDM vDisks	256 per vZone or ESXi host
NFS datastores	Nutanix = 256 per host ESXi 5.x and later = 256 per host



Entity	Supported Maximum
NFS datastore size	Available disk capacity in the Nutanix cluster (no known logical limit)
Small files (<512 KB) on NFS datastores	2.5 GB per Nutanix node
Storage per node	120 TB
Storage pools	2 per Nutanix cluster
Storage pool size	Available disk capacity in the Nutanix cluster (no known logical limit)
Physical disks per storage pool	Number of disks in the Nutanix cluster (no known logical limit)
Storage Containers (created on Nutanix cluster)	256 per Nutanix cluster
Storage Container size	Available disk capacity in a storage pool (no known logical limit)
Replication factor	3
Nutanix management interface (Prism) simultaneous connections	20 clients per Controller VM
iSCSI	128 LUNs per node
	256 LUNS per Volume Group
	No limit for the number of volume groups if you do not exceed maximum number of LUNs per node

Nutanix clusters are also subject to the vSphere maximum values documented by VMware. For a list of the vSphere maximums, see *Configuration Maximums* for the version of vSphere you are running.

Nutanix supports up to 50 VMs for each storage container if you are using Microsoft VSS-based backups in Hyper-V. If the number of VMs in each storage container increases to more than 50, backup jobs start to get into NFS timeouts.

Three Node Cluster Considerations

A Nutanix cluster must have at least three nodes. Minimum configuration (three node) clusters provide the same protections as larger clusters, and a three node cluster can continue normally after a node failure. However, one condition applies to three node clusters only.

When a node failure occurs in a cluster containing four or more nodes, you can dynamically remove that node to bring the cluster back into full health. The newly configured cluster still has at least three nodes, so the cluster is fully protected. You can then replace the failed hardware for that node as needed and add the node back into the cluster as a new node. However, when the cluster has just three nodes, the failed node cannot be dynamically removed from the cluster. The cluster continues running without interruption on two healthy nodes and one failed node, but the failed node cannot be removed when there are only two healthy nodes. Therefore, the cluster is not fully protected until you fix the problem (such as replacing a bad boot disk) for the existing node.

CLUSTER MANAGEMENT

Managing a Nutanix cluster involves configuring and monitoring the entities within the cluster, including virtual machines, storage containers, and hardware components. You can manage a Nutanix cluster through a web-based management console or a command line interface (nCLI).

- The web console is a graphical user interface (GUI) that allows you to monitor cluster operations and perform a variety of configuration tasks (see [Web Console Overview](#) on page 37).
- Nutanix employs a license-based system to enable your entitled Nutanix features, and you can install or regenerate a license through the web console (see [Licensing](#) on page 61).
- You can upgrade a cluster when a new AOS release is available through the web console (see [Software and Firmware Upgrades](#) on page 61).
- If you have multiple clusters, you can manage them all through a single web interface (see [Multi-Cluster Management](#) on page 62).

Note: You can perform most administrative actions using either the web console or nCLI. However, some tasks are only supported in the nCLI either because a new feature has not yet been incorporated into the web console or the task is part of an advanced feature that most administrators do not need to use. See the Nutanix *Command Reference* for information about using the nCLI. See the *AHV Administration Guide* and hypervisor-specific guides for information about platform configuration and hypervisor-specific tasks that are not performed through the web console.

Web Console Overview

The web console, also known as Prism Element, provides a graphical user interface to manage most activities in a Nutanix cluster.

Display Features

The web console screens are divided into the following sections:

- Main menu bar. The main menu bar appears at the top of every screen in the web console. The cluster name appears on the far left of the main menu bar. To the right of the cluster name, you can select an entity from the pull-down list (Home, Health, VM, Storage, Hardware, DR, Analysis, or Alerts) to display information about that entity. You can also search for specific topics or select various tasks from the pull-down lists on the right side of the main menu bar (see [Main Menu](#) on page 46). In addition, the main menu bar includes status icons for quick access to health, alert, and event information.
- Entity views. There is a dashboard view for each entity. Some entities (VM, Storage, Hardware, and Data Protection) include additional views such as a diagram or table view that you can select from the dashboard of that entity.
- Screen menu bar. Some entity dashboards include another menu bar below the main menu that provides options specific to that screen. In the following example from the Storage

dashboard, three view tabs (**Overview**, **Diagram**, and **Table**) and three task buttons (**+ Storage Container**, **+ Volume Group**, and **+ Storage Pool**) appear on this menu bar.

- Usage and performance/health statistics. Most views include fields that provide usage and either performance or health (or both) statistics. The usage and performance/health statistics vary based on the entity that you are viewing. For example, virtual machine usage statistics are displayed in terms of CPU and memory, while disk usage statistics show disk capacity in TB. In most cases, performance statistics are displayed in IOPS, latency, and bandwidth.
- Alert and event messages. Several views include fields that list current event and alert messages. The listed messages are context specific, so for example only storage-related alerts and events appear on the Storage screen. Clicking on a message opens the alert or event view at that message (see [Alerts Dashboard](#) on page 327).



Figure 18: Overview Screen Sections

Logging Into the Web Console

About this task

To log into the web console, do the following. Knowledge base article [KB 1661](#) lists default cluster credentials.

Procedure

1. Open a web browser.

Note: The supported browsers are the current version and two major versions back of Firefox, Chrome, and Safari, plus Internet Explorer version 11 and Microsoft Edge.

2. Enter **`http://management_ip_addr`** in the address field and press **Enter**. Replace **`management_ip_addr`** with the cluster virtual IP address (if configured) or the IP address of any Nutanix Controller VM in the cluster.

Note: If you are logging into Prism Central, enter the Prism Central VM IP address.

The browser redirects to the encrypted port (9440) and may 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 (see “” ()).

3. If a welcome screen appears (see [Configuring a Banner Page](#) on page 615), read the message and then click the "Accept terms and conditions" bar at the bottom.
4. In the login screen, enter your Nutanix login credentials and press **Enter** or click the right arrow icon.




Figure 19: Login Screen

Note: If LDAP authentication is used, enter the user name in the *samAccountName@domain* format; the *domain\user* format is not supported. (Authentication does not use the user principle name [UPN]; *user@domain* is simply a concatenation of the user and domain names specified in "" () .)

Note: The login page includes background animation that is enabled by default. Click the **Freeze space time continuum!** link at the bottom right of the login screen to disable the animation (or the **Engage the warp drive!** link to enable the animation). To permanently disable (or enable) the animation, see [Modifying UI Settings](#) on page 60.

5. If you are logging in as an administrator (**admin** user name and password) for the first time, which requires that the default password (Nutanix/4u) be changed, enter a new password in the **password** and **re-type password** fields and then press **Enter** or click the right arrow icon.




Figure 20: Login Screen (first admin login)

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
- Must not be among the last 5 passwords
- Must not have more than 2 consecutive occurrences of a character
- Must not be longer than 199 characters

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

Note:

- You are prompted to change the password when logging in as the **admin** user for the first time after upgrading AOS. If the first login after upgrade is to the Controller VM through SSH (instead of Prism), you must log in using the default **admin** user password (Nutanix/4u) and then change the password when prompted.
- When you change the **admin** user password, update any applications and scripts using the **admin** user credentials for authentication. Nutanix recommends that

you create a user assigned with the admin role instead of using the **admin** user for authentication. The *Prism Web Console Guide* describes authentication and roles.

6. If a license agreement screen appears (typically on the first login or if the EULA changed since the last login), which indicates the current EULA has not been acknowledged yet, do the following:
 - a. Read the license agreement (on the left).
 - b. Enter appropriate information in the **Name**, **Company**, and **Job Title** fields (on the right).
 - c. Check the "I have read and agree to the terms ..." box.
 - d. Click the **Accept** button.



Figure 21: EULA Screen

7. If a "Pulse will be Enabled" screen appears (typically on the first login or after an upgrade), which refers to the Pulse feature that alerts Nutanix customer support regarding the health of the cluster (see [Configuring Pulse](#) on page 642), read the statement and then do one of the following:
 - » Click the **Continue** button to enable the Pulse feature (recommended).
 - » Click the **Disable Pulse (not recommended)** button to disable the Pulse feature.

CAUTION: If Pulse is not enabled, alerting Nutanix customer support when problems occur is disabled.

Pulse provides Nutanix customer support with analytic information that allows them to dynamically monitor your cluster and identify potential issues before they become problems

(see [Configuring Pulse](#) on page 642). Enabling Pulse is recommended unless providing cluster information to Nutanix customer support violates your security policy.




Figure 22: Pulse Screen

8. If a screen about enhanced cluster health monitoring appears (typically either after enabling Pulse in the previous step or after upgrading a cluster), read the statement and then do one of the following:

- » Click the **Yes (recommended)** button to enable enhanced cluster health monitoring.
- » Click the **Not Now** button to disable this feature.

The enhanced (on top of standard Pulse) cluster health monitoring provides Nutanix customer support with more detailed (but more transparent) information that allows them to better monitor the health of your cluster. Enhanced cluster health monitoring is

recommended unless providing information such as entity names violates your security policy.




Figure 23: Enhanced Monitoring Screen (upgrade example)





Figure 24: Enhanced Monitoring Screen Example (initial example)

9. [2-node clusters only] If a screen about registering with a Witness appears, see [Configuring a Witness \(Two-Node Cluster\)](#) on page 84.

Logging Out of the Web Console

Procedure

To log out from the web console, click the user icon



in the main menu and then select the **Sign Out** option from the pull-down list. You are logged out immediately after selecting the option (no prompt or message).

Main Menu

The main menu at the top of every screen provides access to all the features of the web console. This section describes each of the main menu options.




Figure 25: Main Menu with Expanded View List

Cluster Information

The cluster name appears on the far left of the main menu bar. Clicking the cluster name opens the Cluster Details window. This window displays the cluster ID number, cluster name, and cluster virtual IP address (if set). You can modify the name or virtual IP address at any time (see [Modifying Cluster Details](#) on page 58).

View Options

Selecting a view (entity name) from the pull-down list on the left displays information about that entity. Select from the following options:

- **Home.** Displays the main dashboard (see [Home Dashboard](#)).
- **Health.** Displays the health dashboard (see [Health Dashboard](#) on page 239).
- **VM.** Displays a virtual machine information dashboard (see [VM Dashboard](#)).
- **Storage.** Displays a storage information dashboard (see [Storage Dashboard](#) on page 103).
- **Network** (AHV only). Displays the network visualiser.
- **Hardware.** Displays a hardware information dashboard (see [Hardware Dashboard](#) on page 161).
- **File Server.** Displays a file server dashboard (see [Nutanix Files Guide](#)).
- **Data Protection.** Displays a data protection information dashboard (see [Data Protection and Recovery with Prism Element](#)).
- **Analysis.** Displays a screen to create and run performance monitors (see [Analysis Dashboard](#) on page 317).
- **Alerts.** Displays a screen of alert and events messages (see [Alerts Dashboard](#)).
- **Tasks.** Displays a screen of task messages (see [Task Status](#) on page 592).
- **LCM.** Displays the life cycle manager (LCM) dashboard (see [Life Cycle Management \(LCM\)](#) on page 62).
- **Settings.** Displays the Settings menu, as does clicking the gear icon on the right of the main menu (see [Settings Menu](#) on page 50).

Informational and Search Features

There are multiple ways to access information from the main menu:

- A health (heart) icon appears on the left of the main menu. It can be green (healthy), yellow (warning), or red (unhealthy) indicating the current health status. Clicking the icon displays the health details view (see [Health Dashboard](#) on page 239).
- An alerts (bell) icon appears on the left of the main menu when critical (red), warning (yellow), or informational (gray) alert messages have been generated and have not been marked as resolved. The number of unresolved alerts is displayed in the icon. Click the icon to display a drop-down list of the most recent unresolved alerts. Click a message or the right arrow link to open the alerts view (see [Alerts Dashboard](#)).
- A tasks (circle) icon appears to the right of the alerts icon. The number of active tasks (running or completed within the last 48 hours) is displayed in the icon. Click the icon to display a drop-down list of the active tasks, and then click the **View All Tasks** button to open the tasks view (see [Task Status](#) on page 592).
- A search (magnifying glass) icon appears on the right side of the main menu. Click this icon to display a search field. You can search for information about entities or actions by entering a string in this field. For example, you can enter an action name such as **add** that returns a list of add actions or an entity name such as **MyVM** that returns a list of links related to that entity.






Figure 26: Search Box Example

Help Menu

A question mark icon appears on the right side of the main menu. Clicking the question mark displays a list of help resource options that you can select. The following table describes each option in the pull-down list.

Table 5: User Menu List

Name	Description
Help with this page	Opens the online help at the page that describes this screen (see Accessing Online Help on page 656).
Health Tutorial	Opens the Health dashboard tutorial that takes you through a guided tour of the health analysis features (see Health Dashboard on page 239).
General Help	Opens the online help at the introduction page.
Support Portal	Opens a new browser tab (or window) at the Nutanix support portal login page (see Accessing the Nutanix Support Portal on page 648).

Name	Description
Nutanix Next Community	Opens a new browser tab (or window) at the Nutanix Next Community entry page (see Accessing the Nutanix Next Community on page 660). This is an online community site for customers and partners to exchange ideas, tips, and information about Nutanix technologies and related datacenter topics.

User Menu

A user icon appears on the far right side of the main menu with the current user login name. Clicking the user icon displays a list of options to update your user account, log out from the web console, and other miscellaneous tasks. The following table describes each option in the pull-down list.

Table 6: User Menu List

Name	Description
Update Profile	Opens the Update Profile window to update your user name and email address (see "" ()).
Change Password	Opens the Change Password window to update your password (see "" ()).
REST API Explorer	Opens a new browser tab (or window) at the Nutanix REST API Explorer web page (see Accessing the REST API Explorer on page 652).
Download nCLI	Downloads the Nutanix command line interface (nCLI) as a zip file to your local system. The download occurs immediately after clicking this option (no additional prompts). See the Nutanix <i>Command Reference</i> for information about installing the nCLI locally and for nCLI command descriptions.
Download Cmdlets Installer	Downloads the PowerShell installer for the Nutanix cmdlets. See the Nutanix <i>Command Reference</i> for information about installing the cmdlets locally and for cmdlet descriptions.
Download Prism Central	Opens a new browser tab (or window) at the Support Tools page of the Nutanix support portal from which you can download the files to install Prism Central. If a login page appears, enter your Nutanix support portal credentials to access the portal. See the <i>Prism Central Guide</i> for more information.
About Nutanix	Opens a window that displays Nutanix operating system (AOS) and other version information (see Checking Version on page 61).
Nothing To Do?	Opens a game that is strictly for entertainment. To quit the game, click the "X" at the upper right of the screen.
Sign Out	Logs out the user from the web console (see Logging Out of the Web Console on page 46).



Settings Menu

The web console includes a Settings page from which you can configure a variety of system services.

You can access the Settings page by doing either of the following:

- Click the gear icon on the right of the main menu (see [Main Menu](#) on page 46).
- Select **Settings** from the pull-down list on the left of the main menu.

The Settings page displays a menu of tasks (on the left) you can perform. Click the task to open the window or page for that option in the pane to the right. The following table describes each menu option.




Figure 27: Settings Page

Table 7: Settings Menu List

Name	Description
General	
Cluster Details	Opens the Cluster Details window to view or modify certain cluster parameters (see Modifying Cluster Details on page 58).
Configure CVM	Opens the Configure CVM window to increase the Controller VM memory size (see Increasing the Controller VM Memory Size on page 93).
Convert Cluster	Opens the Convert Cluster window to convert the cluster from ESXi to AHV and then from AHV to ESXi (see In-Place Hypervisor Conversion on page 620).

Name	Description
Expand Cluster	Opens the Expand Cluster window to add new nodes to the cluster (see Expanding a Cluster on page 185).
Image Configuration [AHV only]	Opens the Image Configuration window to import and manage image files that can be used to create VMs (see Configuring Images on page 289).
Licensing	Opens the Licensing window to install or update the cluster license that enables entitled Nutanix features (see Licensing on page 61).
Reboot [AHV and XenServer only]	Opens the Request Reboot window to gracefully restart the nodes in the cluster one after the other. You can select the nodes you want to restart (see Restarting the Nodes in the Cluster (AHV and XenServer)).
Remote Support	Opens the Remote Support Services window, which enables (or disables) Nutanix remote support access (see Controlling Remote Connections on page 645).
Upgrade Software	Opens the Upgrade Software window to upgrade the cluster to a newer AOS version, or update other upgradeable components (see Software and Firmware Upgrades on page 61).
vCenter Registration [ESXi only]	Opens the vCenter Registration window to register (or unregister) the cluster with the vCenter instance (see Registering a vCenter Server on page 617).
Setup	
Connect to Citrix Cloud [AHV and XenServer only]	Opens the Connect to Citrix Cloud window to connect to the Citrix Workspace Cloud (see Connect to Citrix Cloud on page 310).
Prism Central Registration	Opens the Prism Central Registration window to add the cluster into a central registration for multicluster connection and support (see Register (Unregister) Cluster with Prism Central on page 62).
Pulse	Opens the Pulse window to enable the sending of cluster information to Nutanix customer support for analysis (see Configuring Pulse on page 642).
Rack Configuration	Opens the Rack Configuration page to configure the fault tolerant domain for node, block, and rack awareness (see Rack Fault Tolerance on page 17).
Network	
HTTP Proxy	Opens the HTTP Proxy window to configure an HTTP proxy to which the Nutanix software can connect (see Configuring HTTP Proxy on page 646).
Name Servers	Opens the Name Servers window to configure name servers for the cluster (see Configuring Name Servers on page 595).
Network Configuration [AHV only]	Opens the Network Configuration window to configure network connections for the cluster (see Network Configuration for VM Interfaces on page 138).



Name	Description
Network Switch	Opens the Network Switch Configuration window to configure network switch information needed for collecting network traffic statistics (see Configuring Network Switch Information on page 144). This option does not appear when running a hypervisor that does not support this feature.
NTP Servers	Opens the NTP Servers window to specify which NTP servers to access (see Configuring NTP Servers on page 597).
SNMP	Opens the SNMP Configuration window to enable and configure SNMP for the cluster (see Configuring SNMP on page 599).
Security	
Cluster Lockdown	Opens the Cluster Lockdown window, which allows you to delete (or add) public authentication keys used for SSH access into the cluster (see Cluster Lockdown). Removing all public keys locks down the cluster from external access.
Data at Rest Encryption [SEDs only]	Opens the Data-at-Rest Encryption screen to configure key management for self encrypting drives (SEDs) and enable data encryption across the cluster (see Data at Rest Encryption). This menu option appears only when the data drives in the cluster are SEDs.
Filesystem Whitelists	Opens the Filesystem Whitelist window to specify whitelist addresses (see Configuring a Filesystem Whitelist on page 594).
SSL Certificate	Opens the SSL Certificates window to create a self-signed certificate (see SSL Certificates).
Users and Roles	
Authentication	Opens the Authentication Configuration window to configure authentication for the cluster (see Authentication Configuration).
Local User Management	Opens the Local User Management window. This window lists current users and allows you to add, update, and delete user accounts. See Local User Management .
Role Mapping	Opens the Role Mapping window to configure role mappings that apply in the user authentication process (see Role Mapping).
Email and Alerts	
Alert Email Configuration	Opens the Alert Email Configuration window, which enables (or disables) the e-mailing of alerts (see Configuring Alert Emails on page 335).
Alert Policies	Opens the Alert Policies window, which allows you to specify what events should generate an alert and how frequently the system should check for each event type (see Configuring Alert Policies on page 338).
SMTP Server	Opens the SMTP Server Settings window to configure an SMTP server (see Configuring an SMTP Server on page 598).
Data Resiliency	



Name	Description
Configure Witness [ESXi and AHV only]	Opens the Configure Witness window to add a witness VM for Metro Availability and two-node clusters (see Configuring a Witness (Two-Node Cluster) on page 84).
Degraded Node Settings	Opens the Degraded Node Settings window to enable or disable Degraded Node Detection (see Degraded Node on page 30).
Manage VM High Availability [AHV only]	Opens the Manage VM High Availability window to enable high availability for guest VMs in the cluster (see Enabling High Availability for the Cluster on page 296).
Redundancy State	Opens the Redundancy Factor Readiness window to configure the redundancy factor of the cluster (see Increasing the Cluster Fault Tolerance Level on page 88).
Appearance	
Language Settings	Opens the Languages window, which allows you to select the language of the web console (see Localization (L10n) on page 629).
UI Settings	Opens the UI Settings window to disable (or re-enable) the login screen (see Modifying UI Settings on page 60).
Welcome Banner	Opens the Edit Welcome Banner window to create a welcome banner message that appears before users login to the web console (see Configuring a Banner Page on page 615).

Home Dashboard

The Home dashboard is the opening screen that appears after logging into the web console. It provides a dynamic summary view of the cluster status. To view the Home dashboard at any time, select **Home** from the pull-down list on the far left of the main menu.

Menu Options

The Home dashboard does not include menu options other than those available from the main menu (see [Main Menu](#) on page 46).

Home Screen Details

The Home dashboard displays cluster-level performance and usage statistics on the left, health status information in the middle, and the most recent alert and event messages on the right. The following figure is a sample dashboard, and the following table describes each field in this screen. Several fields include a slide bar on the right to view additional information in that field. The displayed information is dynamically updated to remain current.






Figure 28: Home Dashboard

Note: See [Understanding Displayed Statistics](#) on page 56 for information about how the statistics are derived.

Table 8: Home Screen Fields

Name	Description
Hypervisor Summary	Displays the name and version number of the hypervisor.
Prism Central	Displays whether you have registered the cluster to a Prism Central instance or not. Click Register to register the cluster to a Prism Central instance. If you have already registered, you can click OK to launch the Prism Central instance in a new tab of your browser.
Storage Summary	Displays the total, used, and unused storage space in the cluster.
VM Summary	Displays the total number of VMs in the cluster broken down by on, off, and suspended states.
Hardware Summary	Displays the number of hosts and blocks in the cluster, plus one or more Nutanix block model numbers.
Cluster-wide Controller IOPS	Displays I/O operations per second (IOPS) in the cluster. The displayed time period is a rolling interval 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. (These display features also apply to the I/O bandwidth and I/O latency monitors.) For more in depth analysis, you can add this chart (and any other charts on the page) to the analysis page by clicking the blue link in the upper right of the chart (see Analysis Dashboard on page 317).
Cluster-wide Controller IO Bandwidth	Displays I/O bandwidth used per second in the cluster. The value is displayed in an appropriate metric (MBps, KBps, and so on) depending on traffic volume.
Cluster-wide Controller Latency	Displays the average I/O latency (in milliseconds) in the cluster.
Cluster CPU Usage	Displays the current CPU utilization percentage along with the total available capacity (in GHz).

Name	Description
Cluster Memory Usage	Displays the current memory utilization percentage along with the total available capacity (in GB).
Health	Displays health status for the cluster as a whole (good, warning, critical) and summary health status for the VMs, hosts, and disks in the cluster. Clicking on the VMs, hosts, or disks line displays detailed information about that object in the Health page (see Health Dashboard on page 239).
Data Resiliency Status	Displays information indicating whether the cluster is protected currently from potential data loss due to a component failure. Click anywhere in this field to display a dialog box. Data Resiliency Status window with more information (see the following "Data Resiliency Status" section). • Resiliency Status. Indicates whether the cluster can safely handle a node failure, that is whether a copy exists somewhere in the cluster of all data in any node. If the status is not OK, the Data Resiliency Status window includes a message about the problem. Note: The resiliency status for single-node backup cluster is at the disk level and not at the node level. For more information, see Single-Node Replication Target Clusters in the guide titled <i>Data Protection and Recovery with Prism Element</i>. • Rebuild Capacity Available. Indicates whether there is sufficient unused storage in the cluster to rebuild a data copy after a node is lost. If the status is not Yes, the Data Resiliency Status window includes a message about the problem. Note: This option does not appear for single-node replication target clusters.
Critical Alerts	Displays the most recent unresolved critical alert messages. Click a message to open the Alert screen at that message (see Alerts Dashboard).
Warning Alerts	Displays the most recent unresolved warning alert messages. Click a message to open the Alert screen at that message.
Info Alerts	Displays a summary of informational alerts.
Events	Displays a summary of events.

Data Resiliency Status

The Data Resiliency Status window displays more detailed cluster resiliency status information. (Click the **OK** button at the bottom or **X** icon at the top to close the window.) This window provides information about the number and type of failures the cluster can withstand safely. The following figure is an example. The **Failures Tolerable** column indicates the number of simultaneous failures of that component type that can be tolerated (0, 1, or 2). When no failure can be tolerated, a message is displayed to highlight the limitation, as in this example where there are not enough blocks in the cluster to support a block failure.

Note: When a node goes down, extent group (egroup) fault tolerance status remains unchanged as the node is assumed (initially) to be unavailable just temporarily. However, Stargate (see



[Cluster Components](#) on page 12) fault tolerance goes down by one until all data has been migrated off that node. The egroup fault tolerance status goes down only when a physical copy of the egroup replica is permanently bad.

Data Resiliency Status		
Fault Domain Type: Host		
Component	Failures Tolerable	Message
Static Configuration	1	
ZooKeeper	1	
Oplog	1	
Erasure Code Strip Size	1	
Metadata	1	
		OK

Figure 29: Data Resiliency Status Window

Fault Domain Type: Block		
Component	Failures Tolerable	Message
Static Configuration	0	Not enough rackable units (blocks) in the cluster

Understanding Displayed Statistics

Various statistics are displayed in the web console and Prism Central screens. There are three possible sources for a statistic:

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. When the hypervisor can provide usage statistics, those figures appear where appropriate. ESXi provides such statistics, but Hyper-V and AHV do not. Getting the statistics from ESXi means numbers displayed in the Prism display screens (both Prism Central and Prism Element) should match the corresponding ones in vCenter.

2. Controller (Stargate). When hypervisor statistics are unavailable or inappropriate, the Controller VM provides statistics from Stargate. 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, operations not counted by the controller.
3. Disk (Stargate). Stargate can provide statistics from both the controller and disk perspective. The difference is that the controller perspective includes read I/O operations from memory and disk I/O operations, while the disk perspective includes just the disk I/O operations.

Note: The distinction between hypervisor, controller, and disk statistics applies only to storage-related statistics such as IOPS, latency, and bandwidth.

Field labels in the Prism display screens help identify the information source:

- A field name that does not include either "Controller" or "Disk" indicates that statistic derives from the hypervisor (for example "IOPS").
- A field name that includes the word "Controller" indicates that statistic derives from the controller (for example "Controller IOPS").
- A field name that includes the word "Disk" indicates that statistic derives from the disk (for example "Disk IOPS").

The following table identifies the information source for various statistics. Overview, VM, and storage statistics come from either the hypervisor or controller. IFor VM statistics in a mixed ESXi/AHV cluster, the source depends on the hypervisor hosting that VM (hypervisor for ESXi-based VMs and controller for AHV-based VMs). Hardware statistics come from the disk. Metrics in the analysis page can come from any of the sources (hypervisor, controller, or disk) depending on the type of metric.

Table 9: Source for Displayed Statistics

Hypervisor Type	Overview, VM, and Storage	Hardware	Analysis
ESXi	hypervisor (controller for some storage stats)	disk	metric dependent
Hyper-V	controller	disk	metric dependent
AHV	controller	disk	metric dependent
Citrix Hypervisor	controller	disk	metric dependent
Mixed (ESXi + AHV)	hypervisor	disk	metric dependent
Prism Central	cluster dependent (hypervisor or controller)	disk	metric dependent



Modifying Cluster Details

About this task

You can add a cluster name or virtual facing IP address at any time. To create or modify either value, do the following:

Procedure

1. In the main menu, either click the cluster name at the far left or click the gear icon in the main menu and then select **Cluster Details** in the **Settings** page.
The Cluster Details window appears. It displays the cluster UUID (universally unique identifier), ID, and incarnation ID values, along with the cluster name, virtual IP address, and iSCSI data services IP address (if configured). The cluster ID remains the same for the life of the cluster, but the incarnation ID is reset (typically to the wall time) each time the cluster is re-initialized.




Figure 30: Cluster Details Window

2. Enter (or update) a name for the cluster in the **Cluster Name** field.
The default name is simply Unnamed. Providing a custom name is optional but recommended.
3. Enter (or update) an IP address for the cluster in the **Cluster Virtual IP Address** field.
A Controller VM runs on each node and has its own IP address, but this field sets a logical IP address that always connects to one of the active Controller VM in the cluster (assuming at least one is active), which removes the need to enter the address of a specific Controller VM. The virtual IP address is normally set during the initial cluster configuration (see the *Field Installation Guide*), but you can update the address at any time through this field. Providing an virtual IP address is optional but recommended.

CAUTION: All the features that use the cluster virtual IP address will be impacted if you change that address. See [Virtual IP Address Impact](#) on page 59 for more information.

4. Enter (or update) an IP address for use with Nutanix Volumes in the **iSCSI Data Services IP Address** field.

CAUTION: For certain features, changing the external data services IP address can result in unavailable storage or other issues. The features in question include Volumes, Calm, Leap,

Karbon, Objects, and Files. See [KB 8216](#) and [iSCSI Data Services IP Address Impact](#) on page 59 for more information.

5. When the field entries are correct, click the **Save** button to save the values and close the window.

Virtual IP Address Impact

Any Nutanix feature that uses the virtual IP address might be adversely affected if you change the cluster virtual IP address.

- You can no longer manage a Nutanix cluster (Hyper-V) through the System Center Virtual Machine Manager (SCVMM).
- Nutanix data protection service might lose connection if you configured the remote site using the virtual IP address of that cluster.
- All the VMs running the NGT (Nutanix guest tools) instance will be affected. See [Reconfiguring NGT](#) on page 307.
- External machines mounting shares from Nutanix might fail if the virtual IP address is used for HA functionality (as recommended).
- Some products such as Nutanix Objects use the virtual IP address for access. Check your product-specific documentation before changing the virtual IP address.

iSCSI Data Services IP Address Impact

To provide access to cluster storage, Nutanix Volumes utilizes an iSCSI data services IP address to clients for target discovery which simplifies external iSCSI configuration on clients. This iSCSI data services IP address acts as an iSCSI target discovery portal and initial connection point.

Nutanix does not recommend configuring iSCSI client sessions to connect directly to Controller VM IP addresses. The data services address is owned by one Controller VM at a time. If the owner Controller VM becomes unavailable, the address moves to another Controller VM, ensuring that it is always available.

This IP address is also used as a cluster-wide address by clients configured as part of Nutanix Files and other products. This IP address:

- Should be in the same subnet as the cluster Controller VM IP **eth0** network interface addresses
- Helps load balance storage requests
- Enables path optimization in the cluster, preventing bottlenecks
- Eliminates the need for configuring a multipathing service such as Microsoft multipath I/O (MPIO)

Nutanix recommends setting the iSCSI data services IP address once for each cluster, but you can change it if needed through the Cluster Details window in the Prism web console. See [Modifying Cluster Details](#) on page 58.

If you change the iSCSI data services IP address, you will need to reconfigure any clients to use the new IP address.

- Log out of or disconnect from existing iSCSI or Nutanix Files sessions.
- Delete the old IP address if necessary.
- Re-discover the new target or reestablish Nutanix Files sessions.



CAUTION: For certain features, changing the external data services IP address can result in unavailable storage or other issues. The features in question include Volumes, Calm, Leap, Karbon, Objects, and Files. See [KB 8216](#) for more information.

Modifying UI Settings

About this task

By default, the login page includes background animation, and users are logged out automatically after being idle for 15 minutes. You can disable the background animation, change the session timeout for users, and configure an override to the session timeout.

Procedure

1. Click the gear icon in the main menu and then select **UI Settings** in the **Settings** page. The UI Settings dialog box appears.




Figure 31: UI Settings Window

2. To disable the login page background animation, under **Login Customization**, uncheck the **Enable animated background particles** box (or check it to enable).
3. To configure session timeout, do the following under **Security Settings**:
 - Select the session timeout for the current user from the **SESSION TIMEOUT FOR CURRENT USER** drop-down list.
 - Select the default session timeout for all users (except an administrator) from the **DEFAULT SESSION TIMEOUT FOR NON-ADMIN USERS** drop-down list.
 - Select the appropriate option from the **SESSION TIMEOUT OVERRIDE FOR NON-ADMIN USERS** drop-down list to override the session timeout.

Note: The timeout interval for an administrator cannot be set for longer than 1 hour.

4. Uncheck the **Disable 2048 game** box to disable the 2048 game.
5. Click **Save** to save your changes.

Checking Version

Procedure

1. To view the Nutanix version running in the cluster, click the user icon in the main menu and then select the **About Nutanix** option from the pull-down list.

An About Nutanix window appears that includes the AOS and Nutanix cluster check (NCC) version numbers. It also includes a link to Nutanix patent information.




Figure 32: About Nutanix Window

2. Click the **OK** button to close the window.

Licensing

Nutanix provides licenses you can apply to enable a variety of features. See the [Nutanix Licensing Guide](#) for information about how to license your cluster, apply add-on licenses, and manage your licenses in a variety of situations.

Software and Firmware Upgrades

You can upgrade the AOS, hypervisor, NCC, or Foundation version through the Prism web console **Upgrade Software** feature (also known as 1-click upgrade).

Before you attempt a software upgrade, review the "" () for Nutanix cluster and IT administrators. This guide describes how to upgrade your Prism Element or Prism Central cluster through the web console. It includes upgrade best practices and upgrade work flows as recommended by Nutanix, including the recommended order of installation.

Life Cycle Management (LCM)

The life cycle manager (LCM) tracks software and firmware versions of the various components in a cluster. It allows you to view information about the current inventory and update the versions as needed. To view the LCM dashboard, select **LCM** from the pull-down list on the left of the main menu. See the [Life Cycle Manager Guide](#) for more information.

Multi-Cluster Management

Nutanix provides an option to monitor and manage multiple clusters through a single web console. The multi-cluster view, known as Prism Central, is a centralized management tool that runs as a separate VM. From the Prism web console, you can either register the cluster with an existing Prism Central instance or deploy a Prism Central instance and register the cluster with it.

Prism Central provides the following features:

- Single sign on for all registered clusters
- Summary dashboard across clusters that can be customized as desired
- Summary views for major entity types with drill-down options to view detailed information about individual entities
- Multi-cluster analytics capabilities
- Multi-cluster alerts summary with drill-down options to view information about possible causes and corrective actions for each alert
- Ability to configure individual clusters through direct Prism Central actions (for selected administrative tasks) or through one-click access to the web console for a cluster

Note: See the *Prism Central Guide* for more information about Prism Central.

Register (Unregister) Cluster with Prism Central

Before you begin

- If you have never logged into Prism Central as the user **admin**, you need to log in and change the password before attempting to register a cluster with Prism Central (see [Logging Into Prism Central](#)).
- Do not enable client authentication in combination with ECDSA certificates on a registered cluster since it causes interference when communicating with Prism Central.
- Ports 9440 and 80 need to 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.
- A cluster can register with just one Prism Central instance at a time. To register with a different Prism Central instance, first unregister the cluster.

About this task

To register a cluster with Prism Central, do the following:

Procedure

1. Log in to the Prism web console on the target cluster as the user **admin**.



2. Run NCC as described in [Run NCC Checks](#).
3. Do one of the following:
 - » On the **Home** dashboard, click **Register or create new** from the **Prism Central** widget.
 - » Click the gear icon and select **Prism Central Registration** from the **Settings** menu.




Figure 33: Home Dashboard




Figure 34: Settings Menu

4. In the first screen of the Prism Central dialog box, click the **Connect** button.

This screen includes two options, **Deploy** and **Connect**. This procedure describes how to connect to an existing Prism Central instance. See "" () for instructions on how to deploy a new Prism Central instance.




Figure 35: Prism Central Dialog Box (deploy/connect screen)

5. A screen appears outlining what services are available through Prism Element and Prism Central. After reviewing the message, click the **Next** button.




Figure 36: Prism Central Dialog Box (services screen)

6. In the registration screen, do the following in the indicated fields:
 - a. **Prism Central IP:** Enter the IP address of the Prism Central VM.
 - b. **Username:** Enter **admin** as the Prism Central user name.
 - c. **Password:** Enter the Prism Central **admin** user password.


A screenshot of the Prism Central registration dialog box. The title bar is dark grey with the text "Prism Central" and a question mark icon. The main area is white and contains the text "Connect to an existing Prism Central instance" and "Please fill in the information needed to establish a connection". Below this are three input fields: "PRISM CENTRAL IP", "USERNAME", and "PASSWORD". At the bottom are three buttons: "Back", "Cancel", and "Connect".

Figure 37: Prism Central Dialog Box (registration screen)

7. When all the fields are correct, click the **Connect** button to save the values and close the window.

This registers the cluster on the specified Prism Central VM and allows the passing of information between the cluster and Prism Central.

Note: If the Prism Central IP address changes after registering, you must unregister and then re-register the cluster(s) with the new IP address. However, changing the Prism Central administrator password after registering any clusters does not require re-registration.

Unregistering from Prism Central

About this task

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

Note: Unregistering a cluster through the Prism GUI is no longer available. This option was removed to reduce the risk of accidentally unregistering a cluster because several features (including 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, not only will these features not be available but the configuration for them may also be erased. Therefore, only the following procedure is available to unregister a cluster. See [KB 4944](#) for additional details if you have enabled Self Service, Calm, or other special features in Prism Central.

Procedure

1. Log on to any Controller VM of the registered cluster through an SSH session.
2. Run the `cluster status` command and verify that all services are in a healthy state.

```
acropolis: exit
nutanix@MTNX-1:~$ ssh -A CVM:10.5.25.30:~$
nutanix@MTNX-1:~$ ssh -A CVM:10.5.25.30:~$ cluster status
2017-11-19 21:31:20 INFO zookeeper_session.py:110 cluster is attempting to connect to Zookeeper
2017-11-19 21:31:20 INFO cluster:2420 Executing action status on SVMs 10.5.25.30
The state of the cluster: start
Lockdown mode: Disabled

CVM: 10.5.25.30 Up, ZeusLeader
  Zeus UP [1966, 1995, 1996, 1997, 2043, 2061]
  Scavenger UP [2708, 2736, 2737, 2738]
  SSITerminator UP [2104, 2136, 2137, 2138]
  SecureFileSync UP [2108, 2162, 2163, 2164]
  Medusa UP [2278, 2312, 2313, 2318, 2421]
  DynamicRingChanger UP [2520, 2660, 2661, 2665]
  Pitheas UP [2520, 2586, 2587, 2615]
  Mantle UP [2524, 2590, 2591, 2620]
  Hera UP [2533, 2608, 2609, 2610]
  Stargate UP [2823, 2893, 2894, 3250, 3258]
  InsightsDB UP [2826, 2876, 2877, 3027]
  InsightsDataTransfer UP [2834, 2916, 2917, 3001, 3004, 3005, 3006]
  Ergon UP [2842, 2963, 2964, 2965]
  Cerebro UP [2901, 2984, 2985, 3448]
  Chronos UP [2920, 2999, 3000, 3074]
  Curator UP [2942, 3047, 3048, 3506]
  Prism UP [3024, 3094, 3095, 3201, 3707, 3776]
  CIM UP [3037, 3154, 3155, 3214]
  AlertManager UP [3046, 3283, 3286, 3422]
  Arithmos UP [3083, 3215, 3216, 3488]
  Catalog UP [3106, 3189, 3190, 3191]
  Acropolis UP [3173, 3294, 3295, 3296]
  Uhura UP [3192, 3389, 3390, 3396]
  Smp UP [3198, 3391, 3392, 3395]
  SysStatCollector UP [3208, 3445, 3446, 3447]
  Tunnel UP [3231, 3582, 3583]
  Janus UP [3304, 3485, 3486]
  NutanixGuestTools UP [3517, 3617, 3618, 3683]
  MinervaCVM UP [4029, 4064, 4065, 4066, 4213]
```

Figure 38: Cluster Status Command

3. 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 force=true
```


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. This step can take some time (though typically just a few seconds). To check if the unregistration has completed, run the following command.

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

4. Get the UUID for the cluster.

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

Find the "Cluster UUID" value in the displayed information.



```
Cluster Id      : 00055875-0cf5-0f0f-0000-0000000097fc::38908
Cluster UUID    : 00055875-0cf5-0f0f-0000-0000000097fc
Cluster Name    : johny5-1
Cluster Version : 5.5
External IP Address :
External Data Services... :
Node Count      : 1
Block Count     : 1
Support Verbosity Level : BASIC_COREDUMP
Lock Down Status : Disabled
Shadow Clones Status : Enabled
Password Remote Login ... : Enabled
Timezone        : America/Los_Angeles
On-Disk Dedup   : Disabled
Has Self Encrypting Disk : no
Common Criteria Mode : Enabled
Degraded Node Monitoring : Enabled
```

Figure 39: Cluster Info Output

5. Log on to the Prism Central VM through an SSH session (as the "nutanix" user) and do the following:

- 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 5. This script removes all remaining registration information about that cluster and completes the unregistration process for the Prism Central VM.

- b. Get the UUID for Prism Central.

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

Find the "Cluster UUID" value in the displayed information (see step 5), which in this case is the UUID for Prism Central.

6. Go back to the Controller VM 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 step 6c.

What to do next

After performing these steps you can re-register the cluster with a new or re-created Prism Central instance.

If the clean up process does not complete successfully, try the following:

- Check the logs to indicate if there are any input errors when calling the script. The logs for the unregistration cleanup script can be found under **~/data/logs/unregistration_cleanup.log**.

- If errors occur during script execution, run the **cluster status** command and check that the cluster services are up and running. Rerun the script and check if it succeeds.

Note: If Prism Self Service is enabled and the unregistered cluster is set as the default cluster, you must update the default project to select a new (registered) cluster. See the topic *Modifying A Project* in the Prism Central Guide.

Installing Prism Central (1-Click Internet)

Before you begin

The following requirements must be met before you can install Prism Central (this procedure), expand Prism Central, or add a Prism Central VM:

- The specified gateway must be reachable.
- No duplicate IP addresses can be used.
- The storage container used for deployment is mounted on all hypervisor hosts.
- When installing on an ESXi cluster:
 - vCenter and the ESXi cluster must be configured properly. See the *vSphere Administration Guide for Acropolis (using vSphere Web Client)* for details.
 - vCenter must be registered in Prism.
 - DRS must be enabled in vCenter.
 - vCenter is up and reachable during the deployment.

About this task

To create a new Prism Central instance from an Internet-connected cluster, do the following:

Procedure

1. Log into the Prism web console for a cluster as the user **admin**.
2. Run NCC as described in [Run NCC Checks](#).



3. Do one of the following:

- » On the **Home** dashboard, click **Register or create new** from the **Prism Central** widget.
- » Click the gear icon in the main menu and then select **Prism Central Registration** from the **Settings** menu.




Figure 40: Home Dashboard




Figure 41: Settings Menu Launched from Gear Icon

4. In the first screen of the Prism Central dialog box, click the **Deploy** button.

This screen includes two options, **Deploy** and **Connect**. This procedure describes how to deploy a new Prism Central instance. The topic *Register (Unregister) Cluster with Prism*

Central in the *Prism Central Guide* describes how to connect to an existing Prism Central instance.

Note: On an ESXi cluster, you must first register a vCenter Server before you deploy a new Prism Central instance.




Figure 42: Prism Central Dialog Box (deploy/connect screen)

5. In the installation image screen, do one of the following:

- » Click the **Install** button to install that Prism Central version.
- » Click the **Download** button to download that Prism Central version from the Nutanix support portal. After it downloads, the button changes to **Install**. You can then install that Prism Central version.

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




Figure 43: Prism Central Dialog Box (installation image screen)

6. In the deploy screen, do one of the following:

- » To deploy a 1-VM instance of Prism Central, click the **Deploy 1 VM PC** button.
- » To deploy a 3-VM instance of Prism Central, click the **Deploy 3 VM PC** button and then skip to step 8.

A Prism Central instance can consist of either a single VM or a set of three VMs. A 3-VM instance increases both the capacity and resiliency of Prism Central at the cost of

maintaining the additional VMs. In addition, each VM can be either "large" or "small" in size. Thus, you may choose from four configurations.

Table 10: Prism Central Instance Configurations

VM Size	1-VM Instance	3-VM Instance
large	8 vCPUs, 32 GB of memory, and 2500 GiB of storage	24 vCPUs, 96 GB of memory, 7500 GiB of storage (combined)
small	4 vCPUs, 16 GB of memory, and 500 GiB of storage.	12 vCPUs, 48 GB of memory, and 1500 GiB of storage (combined)




Figure 44: Prism Central Dialog Box (deploy screen)

7. [1-VM instance only] In the general configuration screen, do the following in the indicated fields:
- VM Name:** Enter a name for the Prism Central VM.
 - Select a Container:** Select a container for the Prism Central VM from the pull-down list.
 - VM Sizing:** Select (click the radio button for) the Prism Central VM size based on the number of guest VMs it must manage across all the registered clusters:
 - **Small** - manage up to 2500 VMs
 - **Large** - manage up to 12500 VMs
 - 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. See the "Network Management" section in the *Prism Web Console Guide* for more information.
 - IP Address:** Enter a static IP address for the Prism Central VM.
 - Subnet Mask:** Enter the subnet mask value.
 - Default Gateway:** Enter the IP address of the gateway.
 - DNS Address(es):** Enter the IP address for one or more DNS servers. Enter multiple addresses in a comma separated list.
 - When all the fields are correct, click the **Deploy** button.
This begins 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.



Prism Central

?

×

General Configuration

VM NAME

SELECT A CONTAINER

SelfServiceContainer

VM Sizing

Select the size of your Prism Central VM.

SIZE	VCPUS	MEMORY (GB)	STORAGE (GIB)
☒ SMALL - (UP TO 2,500 VMs)	4	16	500
☐ LARGE - (UP TO 12,500 VMs)	8	32	2500

Network Config

AHV NETWORK

+ Create Network

Back

Close

Deploy

Figure 45: Prism Central Dialog Box (1-VM general configuration screen)

8. [3-VM instance only] In the deploy Prism Central screen, do the following in the indicated fields:
- How many VMs will you support:** Select (click the radio button for) the Prism Central VMs size based on the number of guest VMs they must manage across all the registered clusters:
 - **Up to 5K VMs** - manage up to 5000 VMs
 - **Up to 25K VMs** - manage up to 25000 VMs
 - Network:** Select an existing network for this Prism Central instance from the list.
If the target network is not listed, click the **Create New Network** option to create a new network. See the "Network Management" section in the *Prism Web Console Guide* for more information.
 - Subnet Mask:** Enter the subnet mask value.
 - Gateway:** Enter the gateway IP address.
 - Virtual IP** (optional): Enter a virtual IP address for the Prism Central instance.
A virtual IP can be used as a single point of access for Prism Central.
 - Select a Container:** Select a container for the Prism Central VMs from the pull-down list.
 - <IP address>:** Specify IP addresses for the three new Prism Central VMs.

Note: If you entered a virtual IP, the IP addresses for the three VMs are populated automatically. You can keep those addresses by change them as desired.

- 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 46: Edit IP Field

- Repeat this step for the other two VMs.

The vCPU count and memory size in the new VMs are determined by whether you chose the small or large VM support size (5K or 25K radio button).

- DNS Address(es)** (optional): Enter the IP addresses to one or more DNS servers. Enter multiple addresses in a comma separate list.
- When all the fields are correct, click the **Deploy** button.

This begins 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.

What to do next

Register this cluster with Prism Central. The management features are not available until Prism Central registers the cluster in which it is located. The topic *Register (Unregister) Cluster with Prism Central* in the *Prism Central Guide* describes how to connect to an existing Prism Central instance.

Single-Node Clusters

A traditional Nutanix cluster requires a minimum of three nodes, but Nutanix also offers the option of a single-node cluster for ROBO implementations and other situations that require a lower cost option and can accept lowered resiliency protections. Single-node clusters are capable of running a limited number of user VMs. These are unlike single-node replication targets which are for replication and backup purposes. The following table outlines features and limitations of a single-node cluster.

Note:

- Single-node clusters are supported only on a limited set of hardware models.
- This is an overview video about a single-node cluster. (Link does not appear in PDF format.)

Figure 50: Single-Node Cluster Overview

[A video on this topic is available https://youtu.be/qPys36C87mU on YouTube.](https://youtu.be/qPys36C87mU)

Table 11: Single-Node Cluster Features

Feature	Description
Controller VM minimum requirements	6 vCPU and 20 GB memory
Replication factor (RF)	RF 2 across drives within the same node
Hypervisors supported	AHV and ESXi
Unsupported features	• cluster expansion • deduplication • erasure coding • Volumes • near sync DR and Metro Availability (asynchronous DR is supported) • capacity analytics (in Prism)

Single-Node Cluster Guidelines

A single-node cluster is configured like a regular (three-node or more) cluster in many ways, but note the following for a single-node cluster:

- Do not exceed a maximum of 5 user VMs.



- Do not exceed a maximum of 1000 IOPS.
- Do not create a Prism Central instance (VM) in the cluster. There is no built-in resiliency for Prism Central in a single-node cluster, which means that a problem with the node takes out Prism Central with limited options to recover.
- LCM is not supported because firmware updates usually require services to be stopped (in a single-node cluster there is no other node to take over the workload from the node being updated).
- Async DR is supported for 6 hour RPO only.
- If an SSD fails, the cluster goes into a read-only state. In the read-only state, there are no create, read, update, or delete operations and no cluster changes. If the cluster experiences an unresponsive SSD, data loss is a risk. The Prism Central console offers an override option that allows users to continue writing to single-node clusters after it has entered read-only mode (see [Overriding Read-Only Mode](#) on page 79). Otherwise, be sure there is more than 1 SSD available for the system to place metadata.
- All user VMs must be shut down before upgrading a single-node cluster. If any user VMs are still running, a warning box appears.




Figure 51: Upgrading Single-Node Cluster Warning

Overriding Read-Only Mode

About this task

Single-node clusters enter read-only mode when certain requirements are not met or a disk fails. An alert (a yellow exclamation mark) is generated in the web console that the single-node cluster has entered read-only mode.

Disk failure can be categorized into two categories:

- **SSD Failure:** In an event on an SSD failure or when an SSD is marked for removal, the single-node replication target cluster becomes a Read Only cluster and stops accepting any incoming replication traffic. An alert is displayed in the primary and single-node replication target clusters. Also, the resiliency status of the cluster is changed to the Critical state. The recovery process also slows down. Hence, it is recommended to get the cluster back into the Normal state by adding a new SSD to the cluster. After you remove the faulty disk and replace with the new SSD, this new SSD is automatically detected and the cluster is moved back into the normal state (read/write mode).
- **HDD Failure:** The HDD failure is handled in the same way as that of a normal cluster and none of the situations that occur when an SSD fails (as described above) occur in case of HDD failure.

Note: If there is insufficient rebuild capacity when an HDD is unresponsive, then the cluster can fill up and IO can fail. For example, if you had an NX-1175S hybrid with 2 HDDs and 2 SSDs that has an HDD failure, the other disks do not have enough available space to bring up the data to disk fault tolerance.



Procedure

1. To perform this action through Prism, do the following:
 - a. Click the yellow exclamation mark to re-enable write mode on the cluster.




Figure 52: Re-enable Write Mode Message

A dialog box appears to confirm that you want to override read-only mode.

- b. Click **Override** to confirm.




Figure 53: Confirm Override Mode

2. To perform this action through the nCLI, do the following:
 - a. Log into the cluster using your Nutanix credentials.
 - b. Enter the following command.

```
nutanix@cvm$ ncli cluster set-operation-mode operation-mode=override
```

Two-Node Clusters

A traditional Nutanix cluster requires a minimum of three nodes, but Nutanix also offers the option of a two-node cluster for ROBO implementations and other situations that require a lower cost yet high resiliency option. Unlike a one-node cluster (see [Single-Node Clusters](#) on page 78), a two-node cluster can still provide many of the resiliency features of a three-node cluster. This is possible by adding an external Witness VM in a separate failure domain to the

configuration (see [Configuring a Witness \(Two-Node Cluster\)](#) on page 84). Nevertheless, there are some restrictions when employing a two-node cluster. The following table outlines the features and limitations of a two-node cluster.

Note:

- Two-node clusters are supported only on a limited set of hardware models.
- This is an overview video about a two-node cluster. (Link does not appear in PDF format.)

Figure 54: Two-Node Cluster Overview

A video on this topic is available <https://youtu.be/a-UpHAVEF60> on YouTube.

Table 12: Two-Node Cluster Features

Feature	Description
Controller VM minimum requirements	6 vCPU and 20 GB memory
Replication factor (RF)	RF2 spanned over two nodes and RF4 for metadata on SSDs over two nodes. RF4 for metadata helps during a node failure scenario to quickly transition the healthy node to run in single-node mode with the metadata remaining disk fault tolerant. (Metadata in a two-node cluster is typically small, so the storage need for four copies is modest.)
Single node failure effects	50% resource loss. Plan for 40% maximum resource usage to avoid read-only state on the remaining node. Data is made RF2 in the background so that data is resilient.
Drive failure effects	one node + one SSD failure (on other node) = read-only mode.
Hypervisors supported	AHV and ESXi
Unsupported features	• cluster expansion • dedupe (compression is supported) • erasure coding • near sync DR and Metro Availability (asynchronous DR is supported) • network segmentation

Two-Node Cluster Guidelines

A two-node cluster is configured and upgraded like a regular (three-node or more) cluster in most ways, but note the following for a two-node cluster:



- Size your implementation for $N + 1$ so that in the event of a node loss (50% loss of resources) the remaining node will have sufficient resources to allow the cluster to continue functioning.
- There is a heartbeat check (ping) between the nodes every two seconds. If a successful ping does not occur within 10 seconds (5 consecutive failed tries), a failover is initiated (see [Failure and Recovery Scenarios](#) on page 87). When the cluster recovers, it must remain in healthy status for at least 15 minutes before it will failback.
- The upgrade process in a two-node cluster may take longer than the usual process because of the additional step of syncing data while transitioning between single and two node state. Nevertheless, the cluster remains operational during upgrade.
- Witness VM considerations
 - A Witness VM for two-node clusters requires a minimum of 2 vCPUs, 6 GBs of memory, and 25 GBs of storage.
 - The same Witness VM can be used for both Metro Availability and two-node clusters, and a single Witness VM can support up to 50 instances (any combination of two-node and Metro Availability clusters).
 - You can bring up a two-node cluster without a Witness VM being present initially, but it is recommended that the Witness VM be alive and running before starting the cluster.
 - The Witness VM may reside on any supported hypervisor, and it can run on either Nutanix or non-Nutanix hardware.

Note: Nutanix does not support the deployment of a Witness VM on the AWS and Azure cloud platforms.

- The Witness VM must reside in a separate failure domain, which means independent power and network connections from any cluster it is managing. This separate platform can be deployed either on premise (including as an option the Nutanix one-node replication target NX-1155) or off premise (centrally, typically where Prism Central is hosted).
- During a node failure, the transition of a healthy node to single-node mode can take 30-60 seconds. The client VMs can experience I/O timeouts during this period. Therefore, it is recommended that the SCSI timeout of the client VM disks should be at least 60 seconds.
- The minimum recovery point objective (RPO) for a two-node (or one-node) cluster is six hours.
- Network latency between a two-node cluster and the Witness VM should not exceed 500 ms. (RPC timeouts are triggered if the network latency is higher.) During a failure scenario, nodes keep trying to reach (ping) the Witness VM until successful. Nodes ping the Witness VM every 60 seconds, and each Witness request has a two-second timeout, so it can tolerate up to one second of link latency.
- Node removal is not supported, but node replacement is supported (where the node remains part of the cluster). When replacing a node, the cluster remains active but will transition to standalone mode for a brief period of time.
- All node maintenance work flows (software upgrades, life cycle manager procedures, node and disk break-fix procedures, boot drive break-fix procedure) require that the cluster be registered with a Witness VM.

Configuring a Witness (Two-Node Cluster)

About this task

A two-node cluster requires a Witness VM that is located in a separate failure domain either off premise or in a different physical platform on premise. You can configure a Witness either through the **Configure Witness** option in the main menu (any cluster) or when logging in to the Prism web console (two-node cluster only).

Note: If the Witness VM does not exist yet, you must first create it. See "" () for instructions on how to install or upgrade a Witness VM. Two separate features, two-node clusters and metro availability, use a Witness VM. The installation and upgrade instructions (and the configuration instructions in this section) apply regardless of whether you are using a Witness VM for a two-node cluster or metro availability.

Procedure

- To configure a witness when logging in to the Prism web console, do the following:
 - a. Get the credentials of the target witness VM (IP address, user name, and password).
 - b. Click the gear icon in the main menu and then select **Configure Witness** in the **Settings** page (see [Settings Menu](#) on page 50).
 - c. In the Configure Witness window, enter the Witness VM IP address, user name, and password in the indicated fields, and then click the **Save** button.




Figure 55: Configure Witness Window

- To configure a Witness when logging in to the Prism web console, do the following:
 - a. Get the credentials of the target Witness VM (IP address, user name, and password).
 - b. Log in to the Prism web console for the cluster as described in [Logging Into the Web Console](#) on page 38.
 - c. When the Configure Witness screen appears, enter the witness VM IP address, user name, and password in the indicated fields, and then click the **Apply & Continue** button.

After clicking the **Apply & Continue** button, the cluster is registered with the witness and the Prism web console home dashboard is displayed (see [Home Dashboard](#) on page 53). If you check the **Skip for now (not recommended)** box, you still go to the home dashboard, but the cluster is not registered with a witness. This leaves the cluster vulnerable, so you should register it as soon as possible.

Note: The Configure Witness screen appears only for clusters that require a witness (such as a two-node cluster) and only when no witness is configured.




Figure 56: Configure Witness Screen

Witness VM Dashboard (two-node cluster)

The Witness provides a dashboard of status information about all the registered two-node clusters. (The two-node view appears only in a 5.6 or later version of the Witness.) To view the Witness dashboard, do the following:

1. Open a browser to https://witness_vm_ip_address.

2. Log in to the Witness VM. (The user name is **admin** and the password **Nutanix/4u** by default.)



Figure 57: Witness Login Screen

The Witness dashboard appears. In the **Witness Entities** pane on the left, click **2-Node Clusters** to display a table of registered two-node clusters in the right pane. A line appears for each cluster that includes the following fields:

- **Cluster Name:** Lists the name of the cluster.
- **Node IPs at Creation:** Displays the IP addresses of the two nodes in the cluster.
- **Cluster UUID:** Displays the cluster UUID value.
- **Status:** Indicates the current cluster status, either **OK** or a problem statement such as "one node down".
- **Last Keep Alive:** Lists the date and time when the current connection was established (continuous monitoring since that time).

To log out from the Witness VM, go to the user icon menu (upper right of the screen) and select **Logout** from the pull-down list.

CLUSTER NAME	NODE IP(S) AT CREATION	CLUSTER UUID	STATUS	LAST KEEP ALIVE
Fault6	10.4.77.210, 10.4.77.210	00056621-5803-59e6-0000-000000006ad2	One node down, current leader: b726e56a-a22c-439...	Mar 16, 2018, 4:42 PM
Unnamed	10.4.161.226, 10.4.161.227	00056622-943d-e1b0-0000-000000000963	One node down, current leader: 19a52814-1acc-47b...	Mar 6, 2018, 1:26 PM
robot-5	10.1.85.220, 10.1.85.99	00056804-b479-e054-4b2f-cc38c04283f	One node down, current leader: 5d8f6977-e021-495...	Mar 30, 2018, 3:38 PM
Robot_ArchiNode	10.1.86.237, 10.1.86.253	000568a4-35e1-faa3-42f3-ac1f6a042633	OK	Apr 2, 2018, 5:31 PM

Figure 58: Witness Dashboard (2-node cluster view)

Failure and Recovery Scenarios

There are several potential failure scenarios between the nodes and the Witness VM. The steps for recovering from a failure depend on the nature of the failure. This section describes the steps needed (or not needed) when a failure occurs. Each failure scenario generates one or more alerts that you can review. See [Alerts Dashboard](#) on page 327 for information about viewing alert messages.

Node Failure

When a node goes down, the live node sends a leadership request to the Witness VM and goes into single-node mode. In this mode RF2 is still retained at the disk level, meaning data is copied to two disks. (Normally, RF2 is maintained at the node level normally meaning data is copied to each node.) If one of the two metadata SSDs fails while in single-node mode, the cluster (node) goes into read-only mode until a new SSD is picked for metadata service. When the node that was down is back up and stable again, the system automatically returns to the previous state (RF2 at the node level). No user intervention is necessary during this transition.




Figure 59: Failure Diagram: Node Down

Network Failure Between The Nodes

When the network connection between the nodes fails, both nodes send a leadership request to the Witness VM. Whichever node gets the leadership lock stays active and goes into single-node mode. All operations and services on the other node are shut down, and the node goes into a waiting state. When the connection is re-established, the same recovery process as in the node failure scenario begins.




Figure 60: Failure Diagram: Node-to-Node Communication Down

Network Failure Between Node and Witness VM

When the network connection between a single node (Node A in this example) and the Witness fails, an alert is generated that Node A is not able to reach the Witness. The cluster is otherwise unaffected, and no administrator intervention is required.




Figure 61: Failure Diagram: Node-to-Witness Communication Down

Witness VM Failure

When the Witness goes down (or the network connections to both nodes and the Witness fail), an alert is generated but the cluster is otherwise unaffected. When connection to the Witness is re-established, the Witness process resumes automatically. No administrator intervention is required.

If the Witness VM goes down permanently (unrecoverable), follow the steps for configuring a new Witness through the **Configure Witness** option of the Prism web console (see [Configuring a Witness \(Two-Node Cluster\)](#) on page 84).




Figure 62: Failure Diagram: Witness Down

Increasing the Cluster Fault Tolerance Level

If you have a cluster set to Fault Tolerance one (FT1), you have the option to increase the cluster to FT2. Redundancy factor 3 is only supported in clusters that have FT2. This topic describes how to increase the cluster FT level.


About this task

Note:

- If you enabling block awareness in addition to increasing the FT level, you need a minimum of 5 blocks.
- Increasing the cluster FT level might require at least 30% of your disk space.
- Clusters must have a minimum of five nodes.
- Changes to Fault Tolerance cannot be reverted. You also cannot reduce the cluster fault tolerance level. For example, you cannot change the redundancy factor from 3 to 2 or 2 to 1. If you attempt to reduce this setting, the web console displays an error.

Procedure

1. To view the number of host failures that Zookeeper can tolerate, open a web browser and log on to the Prism web console. Click **Data Resiliency Status** on the main dashboard to display the Data Resiliency Status window.



Fault Domain Type: Host		
COMPONENT	FAILURES TOLERABLE	MESSAGE
Static Configuration	1	
Extent Groups	1	
Erasure Code Strip Size	1	
Metadata	1	
ZooKeeper	1	
Stargate Health	1	
Oplog	1	
Free Space	1	

Fault Domain Type: Block		
COMPONENT	FAILURES TOLERABLE	MESSAGE

OK

Figure 63: Data Resiliency Status window (RF 2)

2. Click the gear icon in the main menu and then select **Redundancy State** in the **Settings** page. The Redundancy Factor Readiness dialog box appears. This displays the existing Redundancy Factor of the cluster.





Figure 64: Redundancy Factor Readiness

3. Select **3** in the **Desired Redundancy Factor** drop-down list to set the cluster to redundancy factor 3.
4. Click **Data Resiliency Status** on the main dashboard to verify that the redundancy factor is now 3.

The process of increasing the cluster redundancy factor can take hours to complete because it must replicate metadata across the ring. When the value of **Failures Tolerable**

for Zookeeper increases to 2 as displayed in the following example, it confirms that the redundancy factor is now 3.



The image shows a 'Data Resiliency Status' window with a dark header bar containing a question mark and a close button. The main content is divided into two sections: 'Fault Domain Type: Host' and 'Fault Domain Type: Block'. The 'Host' section contains a table with three columns: 'COMPONENT', 'FAILURES TOLERABLE', and 'MESSAGE'. The 'ZooKeeper' row is highlighted with a red border. The 'Block' section is currently empty. An 'OK' button is located at the bottom right of the window.

COMPONENT	FAILURES TOLERABLE	MESSAGE
Static Configuration	2	
ZooKeeper	2	
Stargate Health	2	
Oplog	2	
Erasure Code Strip Size	2	
Extent Groups	1	Based on placement of extent group replicas the cluster can tolerate a maximum of 1 node failure(s)
Metadata	2	
Free Space	2	

Figure 65: Data Resiliency Status window (RF 3)

5. Set the replication factor to 3 for every storage container you want to have three copies of the data.

Increasing the redundancy factor for the cluster does not automatically increase the replication factor for any container. This gives you granular control in case you have a container where you do not want to incur the overhead of a third copy. However, this also

means no container has a third copy of the data until you explicitly increase the replication factor for that container to 3.

Note: Be sure to change the replication factor to 3 for all desired containers, including any created by the system such as the SelfServiceContainer or NutanixManagementShare.

- a. Log in to the Controller VM and type the following command.

```
nutanix@cvm$ ncli
```

- b. Locate the storage container(s) to update.

```
ncli> ctr list
```

Output similar to the following is displayed.

```
Id : 00052e5a-bc71-2112-0000-00000000261a::11
Uuid : c49eb9af-2eba-41b1-bae5-08227f7cfff13
Name : storage_container_name
Storage Pool Id : 00052e5a-bc71-2112-0000-00000000261a::10
Storage Pool Uuid : b785ff57-9d53-4f05-bad7-3ac3587f2960
Free Space (Logical) : 10.29 TiB (11,316,325,361,581 bytes)
Used Space (Logical) : 0 bytes
Allowed Max Capacity : 10.29 TiB (11,316,325,361,581 bytes)
Used by other Containers : 0 bytes
Explicit Reservation : 0 bytes
Thick Provisioned : 0 bytes
Replication Factor : 2
Oplog Replication Factor : 2
NFS Whitelist Inherited : true
Container NFS Whitelist :
VStore Name(s) : default-container-9754
Random I/O Pri Order : SSD-PCIE, SSD-SATA, DAS-SATA
Sequential I/O Pri Order : SSD-PCIE, SSD-SATA, DAS-SATA
Compression : off
Fingerprint On Write : off
On-Disk Dedup : none
Erasure Code : off
```

- c. For each target storage container, set the replication factor to 3.

```
ncli> ctr edit name=storage_container_name rf=3
```

Replace `storage_container_name` with the name of the storage container.

Output similar to the following is displayed. This shows that the replication factor is now at the correct state (3) for the storage container.

```
Id : 00052e5a-bc71-2112-0000-00000000261a::1381
Uuid : a567b66d-b356-4883-90c5-c37b2f3e0fad
Name : storage_container_name
Storage Pool Id : 00052e5a-bc71-2112-0000-00000000261a::10
Storage Pool Uuid : b785ff57-9d53-4f05-bad7-3ac3587f2960
Free Space (Logical) : 6.86 TiB (7,544,216,907,720 bytes)
Used Space (Logical) : 0 bytes
Allowed Max Capacity : 6.86 TiB (7,544,216,907,720 bytes)
Used by other Containers : 0 bytes
Explicit Reservation : 0 bytes
Thick Provisioned : 0 bytes
Replication Factor : 3
Oplog Replication Factor : 3
NFS Whitelist Inherited : true
Container NFS Whitelist :
```



```

VStore Name(s)      : aaa
Random I/O Pri Order : SSD-PCIe, SSD-SATA, DAS-SATA
Sequential I/O Pri Order : SSD-PCIe, SSD-SATA, DAS-SATA
Compression          : off
Fingerprint On Write : off
On-Disk Dedup         : none

Erasure Code         : off

```

As with increasing the redundancy factor, increasing the replication factor can take some time to complete. You can verify the status by again going to the Data Resiliency Status window in Prism. When the value of **Failures Tolerable** for Extent Groups (which reflects the container replication level) increases to 2, it confirms that the replication factor is now 3.

Increasing the Controller VM Memory Size

Foundation provisions the default memory and vCPUs to each Controller VM according to your platform model. See the "CVM vCPU and vRAM Allocation" topic in the "Field Installation Guide" for information about how Foundation allocates memory and vCPUs to your platform model. After Foundation, you can increase the memory reserved for each Controller VM in your cluster by using the 1-click Controller VM Memory Upgrade feature available in the Prism Element web console. You might need to increase the Controller VM memory if your cluster configuration is as one or more of the configurations described in this document.

About this task

Note: Nutanix does not support decreasing the Controller VM memory below the recommended minimum requirements.

The following tables provide a guideline on optimal memory sizes for various cluster configurations.

Table 13: Minimum Controller VM Memory Requirements for Remote Direct Memory Access (RDMA) Clusters

Two-socket nodes	Each CVM requires 32 GB of minimum memory.
Four-socket nodes	Each CVM requires 48 GB of minimum memory.

Table 14: Minimum Memory Required for Nodes With Both HDDs and SSDs

HDD capacity is > 80 TB	Minimum memory of 40 GB is required if hourly snapshots are enabled. Otherwise, 36 GB.
HDD capacity is > 32 TB and <= 80 TB	Minimum memory of 40 GB is required if hourly snapshots are enabled. Otherwise, 28 GB.



HDD capacity is <= 32 TB	Minimum memory of (20 GB + values from Table 4) or 32 GB, whichever is less, is required regardless of whether hourly snapshots are enabled or not.
Note: For nodes with both HDDs and SSDs, SSD capacity plays no role in memory requirement. SSD capacity is considered only on all-flash nodes.	

Table 15: Minimum Memory Required for All-Flash Nodes

SSD capacity is > 92 TB	Minimum memory of 40 GB is required if hourly snapshots are enabled. Otherwise, 36 GB.
SSD capacity is > 48 TB and <= 92 TB	Minimum memory of 40 GB is required if hourly snapshots are enabled. Otherwise, 28 GB.
SSD capacity is <= 48 TB	Minimum memory of (20 GB + values from Table 4) or 32 GB, whichever is less, is required regardless of whether hourly snapshots are enabled or not.

Table 16: Additional Memory Required for Cluster Features

Capacity tier deduplication (includes performance tier deduplication)	Add 12 GB
Performance tier deduplication	Add 8 GB
Redundancy factor 3	Add 8 GB

Note: Add-ins such as Nutanix Files and Nutanix Objects have their own memory requirements. See the documentation specific to the add-in (Nutanix Files, Nutanix Objects, or others) for any specific requirements.

The following procedure updates the memory of all the CVMs in the cluster. For add-ins such as Nutanix Files and Nutanix Objects, you might have to update the memory of the CVM of only one node. To update the memory of the CVM of only one node, depending on the hypervisor, see the following:

- For ESXi, see the *Changing CVM Memory Configuration (ESXi)* topic in the *vSphere Administration Guide for Acropolis (using vSphere HTML 5 Client)* guide.
- For Hyper-V, see the *Changing CVM Memory Configuration (Hyper-V)* topic in the *Hyper-V Administration for Acropolis* guide.
- For AHV, contact Nutanix Support.

Perform the following procedure to increase the Controller VM memory in the Prism Element web console.

Procedure

1. Run NCC as described in [Run NCC Checks](#).
2. Log on to the Prism Element web console.
3. Click the gear icon in the main menu and then select **Configure CVM** in the **Settings** page. The Configure CVM dialog box is displayed.




Figure 66: Configure CVM Window

4. If your cluster is running the ESXi hypervisor, you are managing your cluster by using VMware vCenter Server, and you have not authenticated the vCenter Server instance, perform the following steps:

- a. In the **Configure CVM** dialog box, click **Add vCenter**.
- b. Enter the vCenter Server IP address and administrator credentials.
- c. Click **Add**.

5. Select the **Target CVM Memory Allocation** memory size and click **Apply**.

The values available from the drop-down menu can range from 16 GB to the maximum available memory in GB.

AOS applies memory to each Controller VM that is below the amount you choose. Resizing memory for a Controller VM requires a reboot as part of the process.

If a Controller VM was already allocated more memory than your choice, it remains at the allocated amount. For example, if a Controller VM is at 20 GB memory and you select 28 GB, the Controller VM memory is upgraded to 28 GB. However, if a Controller VM is at 48 GB memory and you select 28 GB, the Controller VM memory remains unchanged at 48 GB.

Restarting the Nodes in the Cluster (AHV and XenServer)

About this task

The Request Reboot operation in the Prism web console gracefully restarts the selected nodes one after the other.

You might have to restart the nodes in the cluster for reasons such as successfully applying XenServer patches and so on.

Perform the following procedure to restart the nodes in the cluster.

Procedure

1. Click the gear icon in the main menu and then select **Reboot** in the **Settings** page.
2. In the Request Reboot window, select the nodes you want to restart, and click **Reboot**.




Figure 67: Request Reboot

A progress bar is displayed that indicates the progress of the restart of each node.

STORAGE MANAGEMENT

Storage in a Nutanix cluster is organized hierarchically into several components that allow you to manage capacity and performance. Following is a brief outline of the storage features of Nutanix and how to navigate these storage features.

- Nutanix clusters contain storage pool, storage container, volume group, and virtual disk components to organize storage. (see [Storage Components](#) on page 97).
- The web console allows you to monitor storage usage across the cluster (see [Storage Dashboard](#) on page 103).
- You can create storage pools (see [Creating a Storage Pool](#) on page 125), storage containers (see [Creating a Storage Container](#) on page 126), and volume groups (see [Creating a Volume Group](#) on page 133) through the web console.

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 (see [Storage Table View](#) on page 114).

Storage Pools

Storage pools are groups of physical disks from one or more tiers (see [Creating a Storage Pool](#) on page 125). Storage pools provide physical separation because a storage device can only be assigned to a single storage pool at a time. Nutanix recommends creating a single storage pool for each cluster. This configuration allows the cluster to dynamically optimize capacity and performance. Isolating disks into separate storage pools provides physical separation, but can create an imbalance of these resources if the disks are not actively used. When you expand your cluster by adding new nodes, the new disks can also be added to the existing storage pool. This scale-out architecture allows you to build a cluster that grows with your needs.

Storage Containers

A storage container is a subset of available storage within a storage pool (see [Creating a Storage Container](#) on page 126). Storage containers are created within a storage pool to hold virtual disks (vDisks) used by virtual machines. By default, storage is thinly provisioned, which means that the physical storage is allocated to the 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.

Volume Groups

A volume group is a collection of logically related virtual disks (or volumes). A volume group is attached to 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 of 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 a SCSI 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 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 provides the choice by supporting 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 it is the recommended datastore type.

In Hyper-V environments, storage containers are mounted as an SMB share.

Note: Using a Nutanix storage container as a general-purpose NFS or SMB share is not recommended. For NFS and SMB file service, use Nutanix Files.

NFS Datastores. The Distributed Storage Fabric (DSF) 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 ESX 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 68: vSphere Configuration of NFS Datastore

SMB Library Share. The Nutanix SMB share implementation is the Hyper-V equivalent of an NFS Datastore with feature and performance parity with a vSphere configuration. The registration of a Nutanix storage container as an SMB Library share can be accomplished through a single powershell script, or through the Virtual Machine Manager GUI.




Figure 69: Hyper-V Configuration of an SMB Share

Compression

You can enable compression on a storage container. Compression can save physical storage space and improve I/O bandwidth and memory usage—which may have a positive impact on overall system performance.

Note: If the metadata usage is high, compression is automatically disabled. If compression is automatically disabled, an alert is generated.

The following types of compression are available.

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.

When you create a new storage container, post-process compression is enabled by default for all license levels.

Inline compression

Data is compressed as it is written.

Compression Ratios

You can view compression ratios and usage savings in the Prism Web Console.

- Cluster

In the Storage dashboard, under **Capacity Optimization**, click the **After** bar, and hover your mouse over **Compression**.

- Storage container

In the Storage dashboard **Table** view, on the **Storage Container** tab, click the storage container for which you want to view the compression ratio. You can see the compression ratio for the selected storage container under **Storage Container Details**.

Deduplication

Deduplication reduces space usage by consolidating duplicate data blocks on Nutanix storage. You can enable either cache deduplication only or both cache and capacity deduplication on a storage container.

Cache Deduplication

You can enable cache deduplication of read cache to optimize performance. Cache deduplication means deduplication performed on the data in memory.

Note: Cache deduplication is not enabled by default.

Capacity Deduplication

You can enable capacity deduplication of persistent data to reduce storage usage. Capacity deduplication means deduplication performed on the data in hard disk storage (HDD). Capacity deduplication is not enabled by default and is available if you have purchased a Pro or higher license. Note that you can enable capacity deduplication only if cache deduplication is enabled.

Enabling Deduplication

Deduplication is enabled at the storage container level (the **Cache** property for cache deduplication and the **Capacity** property for capacity deduplication). These storage container properties can be set in the web console or nCLI.

Note: Deduplication enabled on storage containers having protected VMs lowers the replication speed.

In addition, Controller VMs in clusters with deduplication enabled need to be configured with more RAM:

- Cache deduplication: 24 GB RAM
- Capacity deduplication: 32 GB RAM

For more information about enabling deduplication, see [Creating a Storage Container](#) on page 126.

Deduplication Best Practices

This table shows when deduplication is recommended and when it is not.



Enable deduplication	Do not enable deduplication
• Full clones • Physical-to-virtual (P2V) migration • Persistent desktops	• Linked clones or Nutanix VAAI clones: Duplicate data is managed efficiently by DSF so deduplication has no additional benefit • Server workloads: Redundant data is minimal so may not see significant benefit from deduplication

Erasure Coding

Erasure coding increases the usable capacity on a cluster. Instead of replicating data, erasure coding uses a parity information to rebuild data in the event of a disk failure. The capacity savings of erasure coding is in addition to deduplication and compression savings.

If you have configured redundancy factor 2, 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 redundancy factor 2. In the following image, the white text represents the data blocks and the green text represents the copies.




Figure 70: Data copies before Erasure Coding

When the data becomes cold, the erasure code engine computes parity “P” for the data by performing an exclusive OR operation.



Figure 71: Computing Parity

Once 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 72: After Computation of Parity

If the node containing 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 image.

Block c is then placed on a node that does not have any other members of this erasure coded stripe.




Figure 73: Post Node Failure

Note: When the cluster is configured for the redundancy factor 3, 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 survive the failure of two nodes.

Example of Savings from Erasure Coding

The space savings from the erasure coding depends on the cluster size, redundancy setting, and percentage of cold data.

In a 6-node cluster configured with redundancy factor 2, erasure coding uses a stripe size of 5: 4 nodes for data and 1 node for parity. The sixth node in the cluster ensures that if a node fails, a node is available for rebuild. With a stripe of 4 data to 1 parity, the overhead is 25%. Without erasure coding, the overhead is 100%.

Erasure coding stripe size adapts to the size of the cluster starting with the minimum 4 nodes with a maximum of 5 node stripe width. The following image displays the various configurations of cluster size, possible stripe widths, and approximate savings that might occur when erasure coding is enabled.



Figure 74: Example of Space Saving from Erasure Coding on 20 TiB Nodes

Erasure Coding Best Practices and Requirements

- A cluster must have at least four nodes populated with each storage tier (SSD/HDD) represented to enable erasure coding.
- Avoid strips greater than (4, 1) because capacity savings provide diminishing returns and the larger strip size increases the cost of rebuild.

- Erasure coding effectiveness (data reduction savings) might be reduced on workloads that have many overwrites outside of the erasure coding window, which by default is 7 days.
- Read performance can degrade during failure scenarios.
- Erasure coding is an asynchronous process, so space savings might not appear for some time.
- Multiple node removal operations can break the erasure coded strip. If it is necessary to remove multiple nodes from a cluster that uses erasure coding turn off erasure coding before removing the nodes.
- If erasure coding is enabled on any storage container, a minimum of four blocks for RF2 or six blocks for RF3 is required to maintain block awareness.
- Erasure coding mandates that data and parity strips must be placed on separate failure domains (node) and that there is an additional node available for recovery. For example, a strip size of (4, 1) requires you to have at least six nodes.

Capacity Reservation Best Practices

By default, each storage container has access to all of the unused storage in the storage pool. If a storage pool has multiple storage containers, one storage container may take all the remaining storage space and leave others with no available space. To make sure that there is space for a storage container, you can enable capacity reservation.

Capacity reservation allows you to guarantee that a storage container has a minimum amount reserved that is unavailable to other storage containers.

Following are best practices for capacity reservation.

- Reserve capacity for a storage container only if the storage pool has multiple storage containers. Unless there is a specific reason to have multiple storage containers, Nutanix recommends having a single storage pool with a single storage container.
- In total, reserve no more than 90% of the space in the storage pool.
- When setting an advertised capacity for a storage container, be aware that some extra space should be allocated beyond the projected size of any VMs placed in the container to allow room for data that has not yet been garbage collected, which can be substantial depending on the workload (10% or more of the storage capacity in some cases).

Storage Dashboard

The Storage dashboard displays dynamically updated information about the storage configuration in a cluster. To view the Storage dashboard, select **Storage** from the pull-down list on the far left of the main menu.


Menu Options

In addition to the main menu (see [Main Menu](#) on page 46), the Storage screen includes a menu bar with the following options:



- View selector. The Storage dashboard provides three viewing modes.
 - Click the **Overview** button on the left to display storage information in a summary view (see [Storage Overview View](#) on page 105).
 - Click the **Diagram** button to display a diagram of the storage pools and storage containers in the cluster nodes from which you get detailed information by clicking on a storage pool or storage container of interest (see [Storage Diagram View](#) on page 107).
 - Click the **Table** button to display hardware information in a tabular form. The Table screen is further divided into Volume Group, Storage Pool, and Storage Container views; click the **Volume Group** tab to view volume group information, the **Storage Pool** tab to view storage pool information, and the **Storage Container** tab to view storage container information (see [Storage Table View](#) on page 114).
- Action buttons. Click the **Volume Group** button on the right to add a volume group to the cluster (see [Creating a Volume Group](#) on page 133) in a storage container. Click the **Storage Container** button to add a storage container to a storage pool (see [Creating a Storage Container](#) on page 126). Click the **Storage Pool** button to add a storage pool to the cluster (see [Creating a Storage Pool](#) on page 125).
- Page selector. In the Table view, hosts and disks are listed 10 per page. When there are more than 10 items in the list, left and right paging arrows appear on the right, along with the total count and the count for the current page.
- Export table content. In the Table view, you can export the table information to a file in either CSV or JSON format by clicking the gear icon





on the right and selecting either **Export CSV** or **Export JSON** from the pull-down menu. (The browser must allow a dialog box for export to work.) Chrome, Internet Explorer, and Firefox download the data into a file; Safari opens the data in the current window.

Storage Overview View

The Storage Overview view displays storage-specific performance and usage statistics on the left plus the most recent storage-specific alert and event messages on the right. Several fields include a slide bar on the right to view additional information in that field. The displayed information is dynamically updated to remain current. The following figure is a sample view, and the table describes each field in this view.




Figure 75: Storage Overview View

Note: See [Understanding Displayed Statistics](#) on page 56 for information about how the statistics are derived.

Table 17: Storage Overview Screen Fields

Name	Description
Storage Summary	Displays the amount of used, free, and total storage space (in GB or TB) in the cluster.
Storage Containers	Displays the number of storage containers, number of VMs, and number of hosts on which the storage containers are mounted in the cluster.
Capacity Optimization	Displays the data reduction ratio (compression, deduplication, and erasure coding), data reduction savings (compression, deduplication, and erasure coding), and the current gained by enabling compression, deduplication, and erasure coding features.

Capacity Optimization

1.61 : 1
Data Reduction

Data Reduction Savings	1.84 TiB
Overall Efficiency	1.85 : 1

Name	Description
Cluster-wide Controller IOPS	Displays I/O operations per second (IOPS) in the cluster. The displayed time period is a rolling interval 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. (These display features also apply to the I/O bandwidth and I/O latency monitors.)
Cluster-wide Controller IO B/W	Displays I/O bandwidth used per second in the cluster. The value is displayed in an appropriate metric (MBps, KBps, and so on) depending on traffic volume. For more in depth analysis, you can add this chart (and any other charts on the page) to the analysis page by clicking the blue link in the upper right of the chart (see Analysis Dashboard on page 317).
Cluster-wide Controller Latency	Displays the average I/O latency (in milliseconds) in the cluster.
Cache Deduplication	Displays summary statistics (example below) for memory and flash (SSD) storage used and saved through deduplication (fingerprint-on-write enabled). Deduplication  The diagram illustrates the deduplication process. On the left, under 'LOGICAL', it shows 'Memory 64.72 GB' and 'SSD 82.99 GB'. An arrow points to the right, with '2.93 : 1' indicating the deduplication ratio. On the right, under 'PHYSICAL', it shows 'Saved 98.86 GB', 'Memory 22.07 GB', and 'SSD 26.78 GB'.
Storage Critical Alerts	Displays the five most recent unresolved storage-specific critical alert messages. Click a message to open the Alert screen at that message. You can also open the Alert screen by clicking the view all alerts button at the bottom of the list (see Alerts Dashboard).
Storage Warning Alerts	Displays the five most recent unresolved storage-specific warning alert messages. Click a message to open the Alert screen at that message. You can also open the Alert screen by clicking the view all alerts button at the bottom of the list.
Storage Events	Displays the ten most recent storage-specific event messages. Click a message to open the Event screen at that message. You can also open the Event screen by clicking the view all events button at the bottom of the list.

Storage Diagram View

The Storage Diagram view displays information about storage pools and storage containers. The displayed information is dynamically updated to remain current.

The Storage Diagram view screen is divided into two sections:

- The top section is a cascading diagram of the storage units. Initially, a cluster bar appears with storage information about the cluster (used, provisioned, and available storage) and colored blocks in the bar for each storage pool. Clicking on a storage pool block displays storage information about the storage pool and a bar with blocks for each storage container in that storage pool. Clicking on a storage container block displays storage information about the storage container and a bar for that storage container. The storage container bar includes **Physical** and **Logical** buttons on the right that you can click to display information about the physical capacity, which is the total addressable space, or logical capacity, which

is the usable space after accounting for the data replication factor. Because the default replication factor is set to two, the logical space is typically half the physical space. You can edit a storage pool or storage container by clicking the pencil (edit) or X (delete) icon to the right of the name. Clicking the **close** link at the far right hides that storage pool or storage container bar from the display.

- The bottom **Summary** section provides additional information. It includes a details column on the left and a set of tabs on the right. The details column and tab content varies depending on what has been selected.

Note: See [Understanding Displayed Statistics](#) on page 56 for information about how the statistics are derived.




Figure 76: Storage Diagram View

Storage Container Details

Selecting a storage container in the diagram displays information about that storage container in the lower section of the screen (see [Creating a Storage Container](#) on page 126).

- When a storage container is selected, **Summary: *storage_container_name*** appears below the diagram, and action links appear on the right of this line (see [Modifying a Storage Container](#) on page 131 for more information about these actions):
 - Click the **Update** link to update the settings for this storage container.
 - Click the **Delete** link to delete this storage container configuration.
- In the **Summary: *storage_container_name*** table, hover your mouse over the value to see additional details of that parameter.
- Four tabs appear that display information about the selected storage container (see following sections for details about each tab): **Storage Container Usage**, **Storage Container Performance**, **Storage Container Alerts**, **Storage Container Events**.




Figure 77: Storage Diagram View: Storage Container

Table 18: Storage Container Details Fields

Parameter	Description	Values
Name	Displays the name of the storage container.	(name)
Replication Factor	Displays the replication factor, which is the number of maintained data copies. The replication factor is specified (2 or 3) when the storage container is created.	[2-3]
Protection Domain	Displays the name of the protection domain if the you have a configured a protection domain for that storage container.	(name)
Datastore	Displays the name of the datastore.	(name)
VMs	Displays the number of VMs	(number)
Free Space (Logical)	Displays the amount of free storage space available to the storage container that is unreserved.	xxx [GB TB]
Used	Displays the amount of used storage space for the storage container.	xxx [GB TB]

Parameter	Description	Values
Max Capacity	Displays the total amount of storage capacity available to the storage container. Nutanix employs a "thin" provisioning model when allocating storage space, which means space is assigned to a storage container only when it is actually needed. The maximum capacity value reflects the total available storage regardless of how many storage containers are defined. Therefore, when you have two storage containers, it can appear that you have twice as much capacity because maximum capacity for both storage containers show the full amount. Maximum capacity is calculated as the total logical capacity in the storage pool, minus any reserved capacity, minus space used by other storage containers.	xxx [TB]
Reserved	Displays the total reserved storage capacity in the storage container.	xxx [GB TB]
Data Reduction Ratio	Displays the ratio of how much the data size is reduced by enabling compression, deduplication, and erasure coding.	
Data Reduction Savings	Displays the data reduction savings by enabling compression, deduplication, and erasure coding.	
Effective Free	Displays the amount of usable free space after data reduction.	
Overall Efficiency	Displays the overall efficiency that you can achieve by enabling compression, deduplication, and erasure coding features.	
Compression	Displays whether compression is enabled.	[Off On]
Capacity Deduplication	Displays whether on disk deduplication is enabled on hard disks (HDD). Cache deduplication is a prerequisite for Capacity deduplication.	[Off On]
Cache Deduplication	Displays whether "fingerprint on write" is enabled, which allows data duplication when data is read. Data duplication (commonly referred to as dedup) is a specialized data reduction technique for eliminating duplicate copies of data. Setting this parameter to On causes dedup compression to be applied to data both in memory and in solid state storage (SSD).	[None, Off, On]
Erasure Coding	Displays whether erasure coding is enabled.	[On, Off]
Filesystem Whitelists	Displays whether you have configured filesystem whitelist for this storage container.	[None, On, Off]



Storage Pool Details

Selecting a storage pool in the diagram displays information about that storage pool in the lower section of the screen (see [Creating a Storage Pool](#) on page 125).

- When a storage pool is selected, **Summary:** `storage_pool_name` appears below the diagram, and action links appear on the right of this line (see [Modifying a Storage Pool](#) on page 126 for more information about these actions):
 - Click the **Update** link to update the settings for this storage pool.
 - Click the **Delete** link to delete this storage pool configuration.
- Four tabs appear that display information about the selected storage pool (see following sections for details about each tab): **Storage Pool Usage**, **Storage Pool Performance**, **Storage Pool Alerts**, **Storage Pool Events**.




Figure 78: Storage Diagram View: Storage Pool

Table 19: Storage Pool Details Fields

Parameter	Description	Values
Name	Displays the name of the storage pool.	(name)
Free (Physical)	Displays the total amount of physical storage space that is available.	xxx [GB TB]
Used (Physical)	Displays the total amount of physical storage space used in the storage pool.	xxx [GB TB]
Capacity (Physical)	Displays the total physical storage space capacity in the storage pool.	xxx [TB]
Disk Count	Displays the number of disks in the storage pool.	(number)

Cluster Summary Information

When a storage container or storage pool is not selected in the table (or when the word **Summary** is clicked), cluster-wide summary information appears in the lower part of the screen.

- The **Storage Summary** column (on the left) includes five fields:
 - **Free (Physical)**. Displays the amount of physical storage space still available in the cluster.
 - **Used (Physical)**. Displays the amount of physical storage space used currently in the cluster.
 - **Capacity (Physical)**. Displays the total physical storage capacity in the cluster.
 - **Storage Pool(s)**. Displays the names of the storage pool. Clicking on a name displays the detail information for that storage pool in this section.
 - **Storage Container(s)**. Displays the names of the storage containers. Clicking on a name displays the detail information for that storage container in this section.
- Four tabs appear that display cluster-wide information (see following sections for details about each tab): **Usage Summary**, **Performance Summary**, **Storage Alerts**, **Storage Events**.

Usage Tab

The Usage tab displays graphs of storage usage. The tab label varies depending on what is selected in the table:

- **Usage Summary** (no storage pool or storage container selected). Displays usage statistics across the cluster.
- **Storage Container Usage** (storage container selected). Displays usage statistics for the selected storage container.
- **Storage Pool Usage** (storage pool selected). Displays usage statistics for the selected storage pool.

The Usage tab displays the following two graphs:

- **Cluster-wide Usage Summary**: Displays a rolling time interval usage monitor 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 (see [Analysis Dashboard](#) on page 317).
- **Tier-wise Usage**: Displays a pie chart divided into the percentage of storage space used by each disk tier in the cluster, storage pool, or storage container. Disk tiers can include DAS-SATA, SSD-SATA, and SSD-PCIe depending on the Nutanix model type.




Figure 79: Storage Table View: Usage Tab

Performance Tab

The Performance tab displays graphs of performance metrics. The tab label varies depending on what is selected in the table:

- **Performance Summary** (no storage pool or storage container selected). Displays storage performance statistics across the cluster.
- **Storage Container Performance** (storage container selected). Displays storage performance statistics for the selected storage container.
- **Storage Pool Performance** (storage pool selected). Displays storage performance statistics for the selected storage pool.

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. For more in depth analysis, you can add a monitor to the analysis page by clicking the blue link in the upper right of the graph (see [Analysis Dashboard](#) on page 317). The Performance tab includes the following three graphs:

- **[Cluster-wide Hypervisor|Controller|Disk] IOPS:** Displays I/O operations per second (IOPS) for the cluster, selected storage container, or selected storage pool.
- **[Cluster-wide Hypervisor|Controller|Disk] I/O Bandwidth:** Displays I/O bandwidth used per second (MBps or KBps) for physical disk requests in the cluster, selected storage container, or selected storage pool.
- **[Cluster-wide Hypervisor|Controller|Disk] I/O Latency:** Displays the average I/O latency (in milliseconds) for physical disk requests in the cluster, selected storage container, or selected storage pool.




Figure 80: Storage Table View: Performance Tab

Storage Alerts Tab

The Alerts tab displays the unresolved alert messages about storage pools or storage containers in the same form as the Alerts page (see [Alerts Summary View](#) on page 327). Click the **Unresolved X** button in the filter field to also display resolved alerts.

Usage Summary		Performance Summary		Storage Alerts	Storage Events	
Unresolved x				1 Alert		
☐	SEVERITY	ISSUE	TIMESTAMP	ENTITIES	DOCUMENTATION	
☐	Warning	Disk space usage for one or more disks on Controller VM 10.3.17715 has exceeded 75%.	06-12-14, 06:44:03pm	Host	Cause Resolution	

Figure 81: Storage Table View: Alerts Tab

Storage Events Tab

The Events tab displays the unacknowledged event messages about storage pools or storage containers in the same form as the Events page (see [Events Summary View](#) on page 333). Click the **Include Acknowledged** button to also display acknowledged events.

Storage Container Usage		Storage Container Performance		Storage Container Alerts	Storage Container Events
				Include Acknowledged 1 Event	
☐	MESSAGE	ENTITIES	MODIFIED BY	ACKNOWLEDGED	TIMESTAMP
☐	Added Storage Container 'sgvnc_01_container' to storage pool 'sgvnc_01_storage_pool'	Storage Container Storage Pool	system		10/12/16, 3:52:28 PM

Figure 82: Storage Table View: Events Tab

Storage Table View

The Storage Table view displays information about storage pools and storage containers in a tabular form. Click the **Volume Group** tab to display volume group information; click the **Storage Pool** tab in the screen menu bar to display storage pool information; click the **Storage Container** tab to display storage container information. The displayed information is dynamically updated to remain current.

The Storage Table view is divided into two sections:

- The top section is a table. Each row represents a single volume group, storage pool, or storage container and includes basic information about that entity. Click a column header to order the rows by that column value (alphabetically or numerically as appropriate).
- The bottom **Summary** section provides additional information. It includes a details column on the left and a set of tabs on the right. The details column and tab content varies depending on what has been selected.

Note: See [Understanding Displayed Statistics](#) on page 56 for information about how the statistics are derived.



Figure 83: Storage Table View

Volume Group tab

Clicking the Volume Group tab displays information about volume groups in the cluster (see [Creating a Volume Group](#) on page 133).

- The table at the top of the screen displays information about all the configured volume groups, and the details column (lower left) displays additional information when a volume group is selected in the table. The following table describes the fields in the volume group table and detail column.

- When a volume group is selected, Summary: `volume_group_name` appears below the table, and action links appear on the right of this line (see [Modifying a Volume Group](#) on page 135 for more information about these actions):

- Click the **Update** link to update the settings for this volume group.
- Click the **Delete** link to delete this volume group.

Five tabs appear that display information about the selected volume group (see following sections for details about each tab): Performance Metrics, Virtual Disks, Volume Group Tasks, Volume Group Alerts, and Volume Group Events.




Figure 84: Storage Table View: Volume Group

Table 20: Volume Group Table and Detail Fields

Parameter	Description	Values
Volume Group Table Fields (upper screen)		
Name	Displays the name of the volume group.	(name)
Disks	Displays the number of disks in the volume group.	[0-256]
Controller 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]
Controller IO B/W	Displays I/O bandwidth used per second for Controller VM-serviced requests in this volume group.	xxx [MBps KBps]

Parameter	Description	Values
Controller IO Latency	Displays the average I/O latency for Controller VM-serviced requests in this volume group.	xxx [ms]
Volume Group Details Fields (lower screen)		
Name	Displays the name of the volume group.	(name)
Number of Virtual Disks	Displays the number of virtual disks in the volume group.	[0-256]
Total Size	Displays the total size of the volume group.	xxx [GB TB]
Flash Mode	Displays whether the flash mode is enabled.	
Shared	Indicates whether the volume group is shared across iSCSI initiators.	[Yes No]
Initiators	Displays the iSCSI initiators to which the volume group is attached.	(None List of names)
Target IQN Prefix	Displays the IQN prefix of the target iSCSI.	

Storage Container Tab

Clicking the **Storage Container** tab displays information about storage containers in the cluster (see [Creating a Storage Container](#) on page 126).

- The table at the top of the screen displays information about all the configured storage containers, and the details column (lower left) displays additional information when a storage container is selected in the table. The following table describes the fields in the storage container table and detail column.
- When a storage container is selected, **Summary: *storage_container_name*** appears below the table, and action links appear on the right of this line (see [Modifying a Storage Container](#) on page 131 for more information about these actions):
 - Click the **Update** link to update the settings for this storage container.
 - Click the **Delete** link to delete this storage container configuration.
- Four tabs appear that display information about the selected storage container (see following sections for details about each tab): **Storage Container Usage**, **Storage Container Performance**, **Storage Container Alerts**, **Storage Container Events**.






Figure 85: Storage Table View: Storage Container

Table 21: Storage Container Table and Detail Fields

Parameter	Description	Values
Storage Container Table Fields (upper screen)		
Name	Displays the name of the storage container.	(name)
RF	Displays the replication factor, which is the number of maintained data copies. The replication factor is specified (2 or 3) when the storage container is created.	[2-3]
Compression	Displays whether compression is enabled.	[Off On]
Cache Deduplication	Displays whether "fingerprint on write" is enabled, which allows data duplication compression when data is read. Data duplication (commonly referred to as dedup) is a data reduction technique for eliminating duplicate copies of data. Setting this parameter to On causes dedup to be applied to data both in memory and in solid state storage (SSD).	[None, On, Off]
Capacity Deduplication	Displays whether on disk deduplication is enabled on hard disks (HDD). Cache deduplication is a prerequisite for Capacity deduplication.	[On, Off]
Erasure Coding	Displays whether erasure coding is enabled for the storage container or not	[On, Off]
Free (logical)	Displays the amount of free storage space in the storage container.	xxx [GB TB]
Used	Displays the amount of used storage space in the storage container.	xxx [GB TB]

Parameter	Description	Values
Max Capacity	Displays the total amount of storage capacity available to the storage container. Nutanix employs a "thin" provisioning model when allocating storage space, which means space is assigned to a storage container only when it is actually needed. The maximum capacity value reflects total available storage regardless of how many storage containers are defined. Therefore, when you have two storage containers, it can appear you have twice as much capacity because maximum capacity for both storage containers show the full amount. Maximum capacity is calculated as the total logical capacity in the storage pool, minus any reserved capacity, minus space used by other storage containers.	xxx [TB]
Controller 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. The I/O can be served from memory, cache (SSD), or disk.	[0 - unlimited]
Controller IO B/W	Displays I/O bandwidth used per second for Controller VM-serviced requests in this storage container.	xxx [MBps KBps]
Controller IO Latency	Displays the average I/O latency for Controller VM-serviced requests in this storage container.	xxx [ms]
Storage Container Details Fields (lower screen)		
Name	Displays the name of the storage container.	(name)
Free Space (Logical)	Displays the amount of free storage space available to the storage container that is unreserved.	xxx [GB TB]
Used	Displays the amount of used storage space for the storage container.	xxx [GB TB]
Max Capacity	Displays the total amount of storage capacity available to the storage container (see the Reserved Capacity description).	xxx [TB]
Reserved	Displays the total reserved storage capacity in the storage container.	xxx [GB TB]
Replication Factor	Displays the replication factor, which is the number of maintained data copies. The replication factor is specified (2 or 3) when the storage container is created.	[2-3]
Compression	Displays whether compression is enabled.	[Off On]



Parameter	Description	Values
Data Reduction Ratio	Displays the ratio of how much the data size is reduced by enabling compression, deduplication, and erasure coding.	x.xx : 1
Data Reduction Savings	Displays the data reduction savings by enabling compression, deduplication, and erasure coding.	xxx [GB TB]
Effective Free	Displays the amount of usable free space after data reduction.	xxx [GB TB]
Overall Efficiency	Displays the overall efficiency that you can achieve by enabling compression, deduplication, and erasure coding features.	xxx [GB TB]
Cache Deduplication	Displays whether "fingerprint on write" is enabled, which allows data duplication compression when data is read. Data duplication (commonly referred to as dedup) is a specialized data compression technique for eliminating duplicate copies of repeating data. Setting this parameter to On causes dedup compression to be applied to data both in memory and in solid state storage (SSD).	[None, On, Off]
Capacity Deduplication	Displays whether on disk deduplication is enabled, that is dedup compression applied to data on hard disks (HDD). Cache deduplication is a prerequisite for Capacity deduplication.	[On, Off]
Filesystem Whitelists	Displays whether you have configured filesystem whitelist for this storage container.	[None, On, Off]
Erasure Coding	Displays whether erasure coding is enabled or not.	[On, Off]

Storage Pool Tab

Clicking the **Storage Pool** tab displays information about storage pools in the cluster (see [Creating a Storage Pool](#) on page 125).

- The table at the top of the screen displays information about all the configured storage pools, and the details column (lower left) displays additional information when a storage pool is selected in the table. The following table describes the fields in the storage pool table and detail column.
- When a storage pool is selected, **Summary:** `storage_pool_name` appears below the table, and action links appear on the right of this line (see [Modifying a Storage Pool](#) on page 126 for more information about these actions):
 - Click the **Update** link to update the settings for this storage pool.
 - Click the **Delete** link to delete this storage pool configuration.
- Four tabs appear that display information about the selected storage pool (see following sections for details about each tab): **Storage Pool Usage**, **Storage Pool Performance**, **Storage Pool Alerts**, **Storage Pool Events**.





Figure 86: Storage Table View: Storage Pool

Table 22: Storage Pool Table and Detail Fields

Parameter	Description	Values
Storage Pool Table Fields (upper screen)		
Name	Displays the name of the storage pool.	(name)
Disks	Displays the number of disks in the storage pool.	(number)
Free (Physical)	Displays the total amount of physical storage space that is available.	xxx [GB TB]
Used (Physical)	Displays the total amount of physical storage space used in the storage pool.	xxx [GB TB]
Max Capacity (Physical)	Displays the total physical storage space capacity in the storage pool.	xxx [TB]
Disk IOPS	Displays the current I/O operations per second (IOPS) for the storage pool. The IOPS, I/O bandwidth, and I/O latency fields record the I/O requests serviced by physical disks across the storage pool.	[0 - unlimited]
Disk IO B/W	Displays the I/O bandwidth used per second for physical disk requests in this storage pool.	xxx [MBps KBps]
Disk Avg IO Latency	Displays the average I/O latency for physical disk requests in this storage pool.	xxx [ms]
Storage Pool Details Fields (lower screen)		
Name	Displays the name of the storage pool.	(name)
Free (Physical)	Displays the total amount of physical storage space that is available.	xxx [GB TB]
Used (Physical)	Displays the total amount of physical storage space used in the storage pool.	xxx [GB TB]
Capacity (Physical)	Displays the total physical storage space capacity in the storage pool.	xxx [TB]

Parameter	Description	Values
Disk Count	Displays the number of disks in the storage pool.	(number)

Cluster Summary Information

When a storage pool, storage container, or volume group is not selected in the table (or when the word **Summary** is clicked), cluster-wide summary information appears in the lower part of the screen.

- The **Storage Summary** column (on the left) includes five fields:
 - **Available (Physical)**. Displays the amount of physical storage space still available in the cluster.
 - **Used (Physical)**. Displays the amount of physical storage space used currently in the cluster.
 - **Capacity (Physical)**. Displays the total physical storage capacity in the cluster.
 - **Storage Pool(s)**. Displays the names of the storage pools in the cluster. Clicking a name displays detailed information about that storage pool in this section.
 - **Storage Container(s)**. Displays the names of the storage containers. Clicking a name displays detailed information about that storage container in this section.
- Four tabs appear that display cluster-wide information (see following sections for details about each tab): **Usage Summary**, **Performance Summary**, **Storage Alerts**, **Storage Events**.

Usage Tab

The Usage tab displays graphs of storage usage. The tab label varies depending on what is selected in the table:

- **Usage Summary** (no storage pool, storage container, or volume group selected). Displays usage statistics across the cluster.
- **Storage Container Usage** (storage container selected). Displays usage statistics for the selected storage container.
- **Storage Pool Usage** (storage pool selected). Displays usage statistics for the selected storage pool.
- **Volume Group Usage** (volume group selected). Displays usage statistics for the selected volume group.

The Usage tab displays the following two graphs:

- **Usage Summary**: Displays a rolling time interval usage monitor 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 (see [Analysis Dashboard](#) on page 317).
- **Tier-wise Usage**: Displays a pie chart divided into the percentage of storage space used by each disk tier in the cluster, storage pool, or storage container. Disk tiers can include DAS-SATA, SSD-SATA, and SSD-PCIe depending on the Nutanix model type.




Figure 87: Storage Table View: Usage Tab

Performance Tab

The Performance tab displays graphs of performance metrics. The tab label varies depending on what is selected in the table:

- **Performance Summary** (no storage pool, storage container, or volume group selected). Displays storage performance statistics across the cluster.
- **Storage Container Performance** (storage container selected). Displays storage performance statistics for the selected storage container.
- **Storage Pool Performance** (storage pool selected). Displays storage performance statistics for the selected storage pool.
- **Volume Group Performance** (volume group selected). Displays storage performance statistics for the selected volume group.

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. For more in depth analysis, you can add a monitor to the analysis page by clicking the blue link in the upper right of the graph (see [Analysis Dashboard](#) on page 317). The Performance tab includes the following three graphs:

- **[Cluster-wide Hypervisor|Controller|Disk] IOPS**: Displays I/O operations per second (IOPS) for the cluster, selected storage container, selected storage pool, or selected volume group.
- **[Cluster-wide Hypervisor|Controller|Disk] I/O Bandwidth**: Displays I/O bandwidth used per second (MBps or KBps) for physical disk requests in the cluster, selected storage container, selected storage pool, or selected volume group.
- **[Cluster-wide Hypervisor|Controller|Disk] I/O Latency**: Displays the average I/O latency (in milliseconds) for physical disk requests in the cluster, selected storage container, selected storage pool, or selected volume group.




Figure 88: Storage Table View: Performance Tab

Alerts Tab

The Alerts tab displays the unresolved alert messages about storage pools, storage containers, or volume groups in the same form as the Alerts page (see [Alerts Summary View](#) on page 327). Click the **Unresolved X** button in the filter field to also display resolved alerts.

	SEVERITY	ISSUE	TIMESTAMP	ENTITIES	DOCUMENTATION
☐	Warning	Disk space usage for one or more disks on Controller VM 10.3.17715 has exceeded 75%.	06-12-14, 06:44:03pm	Host	Cause Resolution

Figure 89: Storage Table View: Alerts Tab

Events Tab

The Events tab displays the unacknowledged event messages about storage pools, storage containers, or volume groups in the same form as the Events page (see [Events Summary View](#) on page 333). Click the **Include Acknowledged** button to also display acknowledged events.

MESSAGE	ENTITY	MODIFIED BY	TIMESTAMP
Added Storage Container: <path>_st_container to storage pool: <path>_st_storage_pool	Storage Container: StoragePool	system	06/12/14 3:52:28 PM

Figure 90: Storage Table View: Events Tab

Creating a Storage Pool

A storage pool is a collection of physical disks.

About this task

A storage pool is created automatically when the cluster is created. It is recommended that you have one storage pool within the cluster. This configuration allows the cluster to optimize storage capacity and performance. Isolating disks into separate storage pools provides physical separation between VMs, but it can create a performance imbalance and less efficient use of space.

To create a storage pool, do the following:

Procedure

1. In the Storage dashboard (see [Storage Dashboard](#) on page 103), click the **Storage Pool** button.

The Create Storage Pool dialog box appears.




Figure 91: Create Storage Pool Window

2. In the **Name** field, enter a name for the storage pool.

Note: This entity has the following naming restrictions:

- The maximum length is 75 characters.
- Allowed characters are uppercase and lowercase standard Latin letters (A-Z and a-z), Simplified Chinese, decimal digits (0-9), dots (.), hyphens (-), and underscores (_).

3. In the **Capacity** field, check the **Use unallocated capacity** box (if present) to add the available unallocated capacity to this storage pool.

The **Capacity** field indicates the amount of space in the cluster that is available to this storage pool. The listed storage size (0 GB initially) is the assigned capacity for this storage pool. When there is unallocated capacity in the cluster, that means there are storage resources that have not yet been assigned to a storage pool, and the **Use unallocated capacity** box appears. Checking this box adds that space to this storage pool. The check box does not appear if there is no unallocated storage capacity.

4. When the field entries are correct, click the **Save** button.

Modifying a Storage Pool

A storage pool is a defined group of physical disks that can be modified as the cluster changes.

About this task

The main reason to modify a storage pool is to add the storage capacity from newly installed nodes into the storage pool. To modify a storage pool, do the following. You cannot delete a storage pool.

Procedure

1. In the web console select **Storage** from the pull-down main menu (upper left of screen) and then select the **Table** and **Storage Pool** tabs.
2. To update the storage pool, select the target storage pool and then click the **Update** link. The Update Storage Pool window appears, which includes the same fields as the Create Storage Pool window (see [Creating a Storage Pool](#) on page 125). Through this window you can change the storage pool name, add unallocated storage capacity, or change the information life cycle management (ILM) threshold.

Creating a Storage Container

A storage container is a subset of available storage within a storage pool. Storage efficiency features including compression, deduplication and erasure coding are enabled at the container level.

Before you begin

Confirm that the cluster is configured to synchronize time with NTP servers (see [Configuring NTP Servers](#) on page 597) and that the time on the Controller VMs is synchronized and current. If the time on the Controller VMs is ahead of the current time, cluster services may fail to start. Files within the storage containers may also have timestamps ahead the current time when viewed from the hypervisor.

About this task

AOS automatically creates the correct type of access to the storage container for each hypervisor.

- Hyper-V: The storage container is accessible as an SMB share.
- vSphere: The storage container is accessible as an NFS datastore. This requires access to the vSphere APIs. Ensure that you have appropriate license of vSphere to access the APIs.
- AHV: The storage container is accessible transparently.

Note:

- A storage pool and one storage container are created automatically when the cluster is created.
- A storage container is not created if you have not configured the Controller VMs with enough memory. Controller VM memory allocation requirements differ depending on the models and features that are being used. For more information, see *Controller VM Memory Configurations* in *Acropolis Advanced Administration Guide*.
- 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. To ensure proper operation of these features, do not delete this storage container. Nutanix also recommends that you do not delete this storage container even if you are not using these features. The NutanixManagementShare storage container is not intended to be used as storage for vDisks, including Nutanix Volumes.

To create a storage container, do the following:

Procedure


1. In the Storage dashboard (see [Storage Dashboard](#) on page 103), click the **Storage Container** button.

The Create Storage Container dialog box appears.




The dialog box is titled "Create Storage Container". It contains a text input field for "NAME" and a dropdown menu for "STORAGE POOL" with "default-storage-pool-30323" selected. Below this, it shows "MAX CAPACITY" as "12.95 TiB (Physical) Based on storage pool free unreserved capacity". Under "NFS DATASTORE", there are two radio buttons: "Mount on all ESXi hosts" (selected) and "Mount on the following ESXi hosts". At the bottom, there are buttons for "Advanced Settings", "Cancel", and "Save".

Figure 92: Create Storage Container (vSphere)



The dialog box is titled "Create Storage Container". It contains a text input field for "NAME" and a dropdown menu for "STORAGE POOL" with "agave_all_storage_pool" selected. Below this, it shows "MAX CAPACITY" as "13.77 TiB (Physical) Based on storage pool free unreserved capacity". At the bottom, there are buttons for "Advanced Settings", "Cancel", and "Save".

Figure 93: Create Storage Container (AHV)



The dialog box is titled "Create Storage Container". It contains a text input field for "NAME" and a dropdown menu for "STORAGE POOL" with "agave_all_storage_pool" selected. Below this, it shows "MAX CAPACITY" as "10.27 TiB (Physical) Based on storage pool free unreserved capacity". Under "MAKE THIS STORAGE CONTAINER THE DEFAULT STORE FOR VMS ON HYPER V HOSTS", there are three radio buttons: "None" (selected), "Make default on all Hyper V hosts", and "Make default on particular Hyper V hosts". At the bottom, there are buttons for "Advanced Settings", "Cancel", and "Save".

Figure 94: Create Storage Container (Hyper-V)

2. Do the following in the indicated fields:

a. **Name:** Enter a name for the storage container.

Note: This entity has the following naming restrictions:

- The maximum length is 75 characters.
- Allowed characters are uppercase and lowercase standard Latin letters (A-Z and a-z), Simplified Chinese, decimal digits (0-9), dots (.), hyphens (-), and underscores (_).

b. **Storage Pool:** Select a storage pool from the drop-down list.

The following field, **Max Capacity**, displays the amount of free space available in the selected storage pool.

c. (vSphere only) **NFS Datastore:** Select the **Mount on all ESXi hosts** button to mount the storage container on all hosts. Select the **Mount on the following ESXi hosts** button to mount the storage container on a subset of hosts, which displays a list of host IP addresses below this field. Check the boxes of the hosts to be included.

Note: Nutanix does not support the usage of third-party storage on the hosts that are part of Nutanix clusters. Normal cluster operations might be affected if there are connectivity issues with the third-party storage that you attach to the hosts in a Nutanix cluster.

d. (Hyper-V only) Set this storage container as default store for VMs on Hyper-V hosts. Depending on your selection the Create Virtual Machine Wizard of Hyper-V automatically populates the storage location with the relevant storage container.

Option	Description
Make default on all Hyper-V hosts	Makes this storage container a default location for storing virtual machine configuration and virtual hard disk files on all the Hyper-V hosts.
Make default on particular Hyper-V hosts	Provides you with an option to select the hosts on which you want to make this storage container a default location for storing virtual machine configuration and virtual hard disk files on all the Hyper-V hosts.

3. To configure additional parameters, click the **Advanced Settings** button.

The screenshot shows a web interface titled "Advanced Settings". It contains three input fields:

- REDUNDANCY FACTOR**: A dropdown menu with the value "2" selected. A help icon (?) is to the right of the label.
- RESERVED CAPACITY**: A text input field containing "0" and a "GiB" unit selector on the right.
- ADVERTISED CAPACITY**: A text input field containing "Total GiB" and a "GiB" unit selector on the right.

Figure 95: Create Storage Container Advanced Settings (1)

Do the following in the indicated fields:

- a. **Replication Factor:** Displays the number of data copies to maintain in the cluster. Nutanix supports a replication factor (RF) of 2 or 3. Setting the replication factor to 3 adds an extra layer of data protection at the cost of storing an additional copy of the data.

Note: To change the storage container level setting to replication factor 3, the cluster must be set to fault tolerance level 2. See [Increasing the Cluster Fault Tolerance Level](#) on page 88.

- b. **Reserved Capacity:** To reserve storage space for this storage container, enter the amount (in GiB) to reserve in this field.

You can reserve space for a storage container to ensure a minimum storage capacity is available. Reserving space for a storage container means that space is no longer available to other storage containers even if the reserved space is unused. See [Capacity Reservation Best Practices](#) on page 103 for more information.

- c. **Advertised Capacity:** Sets a maximum storage space for this storage container, enter the amount (in GiB) to reserve in this field.

This sets an "advertised" capacity, which is the maximum storage size that the storage container can use. This can be set to any value, but if a reserved capacity is configured, it must be set greater than or equal to the reservation on the storage container. The hypervisor ensures that the storage container storage does not go beyond the advertised

capacity. (When a storage container reaches a threshold percentage of the actual storage pool size, an alert is issued.)

☐ COMPRESSION

DEDUPLICATION

☐ CACHE ?
Perform inline deduplication of read caches to optimize performance.

☐ CAPACITY ?
Perform post-process deduplication of persistent data.

Figure 96: Create Storage Container Advanced Settings (2)

- d. **Compression:** Select the check box to enable compression. A **Delay (In Minutes)** field appears after checking the box. Enter a zero to enable inline compression or a value (number of minutes) to enable post-write compression, which can begin (up to) that number of minutes after the initial write.
All data in the storage container is compressed when this box is checked. See [Compression](#) on page 99 for guidelines about using compression.
- e. **Deduplication:** Select the **CACHE** check box to perform inline deduplication of read caches to optimize performance. If you enable this option, the Controller VMs must be configured to have at least 24 GB of RAM. This feature is primarily recommended for full-clone, persistent desktops, and physical to virtual migration use cases. Turning deduplication on for VAAI clone or linked clone environments is not recommended.
Select the **CAPACITY** check box to perform post-process deduplication of persistent data. This option is recommended primarily for full clone, persistent desktops, and physical to virtual migration use cases that need storage capacity savings (not just performance

savings from deduplication). It is further recommended that the Controller VMs have at least 32GB of RAM and 300GB SSDs for the metadata disk to use this option.

Figure 97: Create Storage Container Advanced Settings (3)

- f. **Erasure Coding:** Select the check box 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 101.
- g. **Filesystem Whitelists:** Enter the comma-separated IP address and netmask value (in the form `ip_address/ netmask`).
A whitelist is a set of addresses that are allowed access to this storage container. Whitelists are used to allow appropriate traffic when unauthorized access from other sources is denied.

Note: Setting a storage container level whitelist overrides any global whitelist for this storage container.

- 4. When all the field entries are correct, click the **Save** button.

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 to apply compression.

Note: The web console does not allow you to rename a storage container in an AHV cluster when modifying container details through the Update Storage Container dialog box. You cannot rename a storage container if it contains vdisks.

To modify (or delete) a storage container, do the following:

Procedure

1. In the web console select **Storage** from the pull-down main menu (upper left of screen) and then select the **Table** and **Storage Container** tabs.
2. To update the storage container, select the target storage container and then click the **Update** link.
The Update Storage Container window appears, which includes the same fields as the Create Storage Container window (see [Creating a Storage Container](#) on page 126). Through this

window you can specify NFS datastore mounts, reserve storage capacity, enable (or disable) compression, enable or disable erasure coding, select or deselect fingerprinting on writes which enables deduplication on reads, and configure filesystem whitelist.

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.
- The Prism web console does not provide an option to change the container replication factor. That can be done only through the nCLI (see [Increasing the Cluster Fault Tolerance Level](#) on page 88).

3. To delete the storage container, select the target storage container and then click the **Delete** link.

Volume Group Configuration

A volume group is a collection of virtual disks (vDisks). Each volume group is identified by a name and UUID. Each vDisk has a UUID and a SCSI index to specify the ordering of disks within the volume group.

You create a volume group on a storage container and enable access to the volume group in one of the following ways:

iSCSI

iSCSI access to a volume group is provided through Nutanix Volumes. For more information, see [Nutanix Volumes](#) on page 207.

Attach Directly to a VM

If the hypervisor is AHV, you can create a volume group and attach it to one or more VMs as a SCSI disk by using the Prism Element web console, the Nutanix REST API, or the nCLI or aCLI commands.

Once the volume group is attached to the VM by using the Prism Element web console, vDisks appear as SCSI devices to the guest operating system. AOS transparently handles load balancing and path resiliency on the guest VM. Directly attached volume groups support Controller VM failover and do not impact VM migration.

You can attach a volume group to a VM at any time. This option is available only to VMs deployed on AHV.

This method of attaching a volume group to a VM has the following limitations:

- The VM must run on the same cluster as the volume group.
- Directly attached volume groups do not support SCSI-3 persistent reservations (SCSI PR). SCSI PR is required for some Microsoft and Linux guest clustering software.

Volume groups are managed independently of the VMs to which they are attached. You can configure a directly attached volume group for access by a single VM (exclusive access) or by multiple VMs (shared access).

CAUTION: Multiple VMs writing to a vDisk that belongs to a shared volume group without additional software to manage the access can lead to data corruption. Use shared access when you are configuring VMs for use with cluster aware software.



Concurrent Access from Multiple Clients

Products, features, or solutions might require concurrent access to volume groups, either by multiple iSCSI initiators or multiple VM attachments.

Multiple iSCSI Initiators

The following products, features, or solutions are supported for concurrent access to volume groups:

- Oracle RAC (bare-metal and virtualized environments)
- Linux VMs
- Windows Failover Clustering
- IBM Spectrum Scale (GPFS)
- Veritas InfoScale Storage

Multiple Directly Attached VMs

The following products, features, or solutions are supported for concurrent access to volume groups:

- Oracle RAC
- Linux VMs
- IBM Spectrum Scale (GPFS)
- Veritas InfoScale Storage

The following products, features, or solutions are not supported for concurrent access to volume groups since direct attach does not support SCSI-3 persistent reserve.

- Windows Failover Clustering
- Linux guest VM clustering by using STONITH with SCSI-3 PR.

Creating a Volume Group

About this task

To create a volume group, do the following:

Procedure

1. In the **Storage** dashboard, click the **Volume Group** button.
The Create Volume Group dialog box is displayed.
2. In **Name**, enter a name for the volume group.

Note: This entity has the following naming restrictions:

- The maximum length is 75 characters.
- Allowed characters are uppercase and lowercase standard Latin letters (A-Z and a-z), decimal digits (0-9), dots (.), and hyphens (-).

3. The **iSCSI Target Name Prefix** is auto-filled with the volume group **Name**. You can accept this prefix or enter your own target name prefix for the volume group. This entity has the same naming restrictions as **Name**.
4. In **Description**, enter a description for the volume group.
5. To add a disk to the volume group, do the following:
 - a. In the **Storage** section, click **Add New Disk**.
 - b. In the **Add Disk** dialog box, in **Storage Container**, select the storage container to use from the pull-down list.
 - c. In **Size**, enter the disk size in GiBs.
 - d. Click **Add**.
 - e. Repeat these steps to add another disk for this volume group, if desired.
6. Do the following in the **Access Settings** section:
 - a. Select **Enable external client access** if you are whitelisting clients that are external to or not residing in this cluster. Otherwise, leave this clear. If you select this check box, it remains selected the next time you create a volume group.
 - b. If you are using one-way CHAP security, select **CHAP Authentication** and type a 12-character to 16-character password (also known as a CHAP secret) in the **Target Password** field.
Initiators must use the same password to authenticate to the AOS cluster.
 - c. Click **Add New Client** to configure the iSCSI initiators, and then enter the client Initiator iSCSI Qualified Name (IQN) in the **Client IQN/IP Address** field to create the whitelist. If you have configured Mutual CHAP authentication on the client, select **CHAP Authentication** and enter the iSCSI client password (secret). Click **Add**.

Note: Ensure that you enter the client IQN in the **Client IQN/IP Address** field and not the IP address. AOS does not support a whitelist containing IP addresses in a volume group.

Access Control displays any configured clients. This list includes any clients attached to volume groups in the cluster. Repeat this step to add more initiators allowed to access this storage. For information about which products, features, or solutions are supported for concurrent access to a volume group, see [Concurrent Access from Multiple Clients](#) on page 133.

7. To enable the Flash Mode feature, select **Enable Flash Mode**.

Note: Individual virtual disks of a volume group cannot be excluded from Flash Mode by using the Prism Element web console. However, you can exclude individual virtual disks from flash mode by using aCLI. For more information, see [Removing Flash Mode for Virtual Disks of a Volume Group](#) on page 137.

8. Click **Save**.

What to do next

If the hypervisor is AHV, you can now attach the volume group to the VM and start using the vDisks. If you want to use iSCSI and you already whitelisted the host IP addresses, log on to the VMs, and configure iSCSI.



Modifying a Volume Group

About this task

To modify a volume group, do the following:

Procedure

1. In the web console, select **Storage** from the pull-down main menu (upper left of screen), and then select the **Table** and **Volume Group** tabs.

2. To update a volume group, select the volume group, and then click the **Update** link.

The Update Volume Group window appears, which includes the same fields as the Create Volume Group window (see [Creating a Volume Group](#) on page 133). In this window you can change the volume group name, add and remove disks, configure the volume group for sharing, and add or remove (by unchecking) entries from the initiator whitelist. On AHV clusters, you can attach the volume group to a VM as a SCSI disk (described later in this procedure). If you attach a volume group to a VM that is part of a protection domain, the volume group is not protected automatically. Add the volume group to the same consistency group manually.

You can increase the size of the volume group. Reducing the size of a volume group is not supported.

Any unchecked initiator that remains unused by this or any other volume group ultimately disappears from this list.

3. To attach a volume group to a VM, click **Attach to a VM**, and then select the VM from the **Available VMs** list.
4. To delete a volume group, select the volume group, and then click the **Delete** link.
5. To manage the iSCSI client list, click the **Summary** link to deselect any volume groups, and then click the **Manage Initiators** link. (This link appears only when no volume groups are selected.)

The Manage iSCSI Clients window appears, which includes a list of the clients and available actions. To modify a client (enable/disable CHAP authentication), click the pencil icon for that client, which displays the Edit iSCSI client window. See [Nutanix Volumes](#) on page 207 for more information.

Flash Mode for Virtual Machines and Volume Groups

The flash mode for VM and allows you to set the storage tier preference to SSD for a virtual machine or volume group. Without flash mode, data for a mission critical application such as Relational database can run out of the room in the SSD tier because of other workloads running on the same cluster. When this happens the database could potentially migrate to the HDD tier. For extremely latency-sensitive workloads, this migration to the HDD tier could negatively affect the read and write performance.

By default, you can use up to 25% of the cluster-wide SSD tier as flash mode space for VMs or VGs. The data for the VMs and VGs with flash mode enabled is placed in the SSD tier and is not down migrated regardless of the amount of usage or degree of hotness of data belonging to other VMs or VGs with flash mode disabled. If the data size for flash mode enabled VMs or VGs exceeds 25% of the SSD capacity, the system may down migrate the data. Before performing down-migration, the flash mode feature tries to preserve the excess data on the SSD tier for some reasonable amount of time so that you can take corrective actions on the cluster and bring back to stable state. To reduce flash mode usage, you can disable flash mode on some VMs or VGs or add SSDs.



If you enable this feature on a VM, all the virtual disks that are attached to the VM are automatically placed on the SSD tier. Also, virtual disks added to this VM are automatically placed on SSD. However, you can update the VM configuration to remove the flash mode from any virtual disks.

CAUTION: The VM with flash mode enabled may increase the performance of the virtual disks that are attached to the VM, but enabling this feature can lower the performance of other VMs on which this feature is not enabled. Hence, ensure that you perform a thorough analysis of the performance impact of enabling this feature. To overcome this issue, you can update the VM configuration and remove the flash mode on individual virtual disks. For example, you can enable flash mode on the applications data disks and disable it on the log disks.

A node failure may cause a reduction in SSD tier capacity that can drive the flash mode usage above the 25% of the tier capacity. In this event, alerts will report the flash usage exceeding 25%.

Note:

- This feature is supported on ESXi and AHV for VMs and on all the hypervisors for VGs.
- For the cluster created using ESXi hosts, you must register your cluster with the vCenter Server. For more information, see [Registering a vCenter Server](#) on page 617.

You can enable the feature on the VM only during the VM update workflow from the Prism Element. However, for VGs, you can enable the feature during the creation of VGs. For more information on enabling the feature on VMs, see [VM Management](#) on page 268 and for enabling the feature on VGs, see [Creating a Volume Group](#) on page 133.

Note: You cannot enable this feature by using Prism Central.

Alerts are raised in the following scenarios.

- When the flash usage exceeds 25% of the SSD tier.
- When the VM has the flash mode feature enabled but is in the powered off state.

Guidelines for the Flash Mode feature

- If a VM or a VG with flash mode feature enabled is cloned, then the flash mode policies are not automatically applied to the cloned VM or VG.
- The flash mode policies are not automatically applied if a VM or VG is restored from the DR snapshot on the local or the remote site.
- The flash mode policies can be used for controlling down migration from the hot tier to the cold tier. The policies do not restrict the usage for the VM or VG on the hot tier. If space is available on the hot tier, more data from the VM or VG can reside on the hot tier.
- If you enable erasure coding on the cluster, the parity always reside on the HDD.
- If you replicate a VM or VG to a remote site (by using Metro Availability or Async DR), the remote site does not have flash mode feature enabled automatically. You need to manually enable the flash mode feature again.
- Enabling flash mode feature does not affect performance tier deduplication. Do not use flash mode feature with capacity tier deduplication.



- If you perform a storage migration or storage vMotion, the new virtual disk created on the target datastore will not have the flash mode feature enabled. You need to manually enable the flash mode feature again.
- If you try to upgrade the AOS with VM pinning feature enabled from the previous release, the upgrade process fails. Hence, ensure that you have disabled the VM pinning feature before performing an upgrade. Also, do not enable the VM pinning feature during the upgrade process.
- If you delete and then create the storage pool in your cluster, you cannot configure flash mode feature for a small duration of time because for some time there will be two storage pools in your cluster and flash mode feature is not supported if you have two storage pools in your cluster.

Removing Flash Mode for Virtual Disks of a Volume Group

Individual virtual disks of a volume group cannot be excluded from flash mode by using Prism. However, you can exclude individual virtual disks from flash mode by using aCLI.

Procedure

1. Log in to the Controller VM in your cluster through an SSH session and access the Acropolis command line.
2. (Optional) If you have not enabled the flash mode feature for the VG by using Prism, you can enable it by running the following command.

```
accli> vg.update vg_name flash_mode=true
```

Replace *vg_name* with the name of the VG on which you want to enable the flash mode.

3. Remove the flash mode feature on a particular disk.

```
accli> vg.disk_update vg_name index_value flash_mode=false
```

Replace *vg_name* with the name of the VG and *index_value* with the index value of the VG. For example to disable flash mode for disk at index 0 with name example_vg, use the following command.

```
accli> vg.disk_update example_vg 0 flash_mode=false
```



NETWORK MANAGEMENT

Nutanix provides several features to manage and monitor network settings for the cluster.

- In clusters with Nutanix virtualization management (such as those running AHV as the hypervisor), you can configure network connections through the web console (see [Network Configuration for VM Interfaces](#) on page 138).
- To track and record networking statistics for a cluster, the cluster requires information about the first-hop network switches and the switch ports being used. You can configure one or more network switches for statistics collection (see [Configuring Network Switch Information](#) on page 144).
- A network visualizer is provided that presents a consolidated graphical representation of the network formed by the VMs and hosts in a Nutanix cluster and first-hop switches. You can use the visualizer to monitor the network and to obtain information that helps you troubleshoot network issues (see [Network Visualization](#) on page 151).

Network Configuration for VM Interfaces

In clusters with Nutanix virtualization management (such as those running AHV), you can configure network connections through the web console. Each VM network interface is bound to a virtual network, and each virtual network is bound to a single VLAN. Additionally, on AHV, Hyper-V, and ESXi, you can secure intra-cluster communications by segmenting the virtual network on the cluster.

About this task

To configure a virtual network for VM interfaces, do one of the following:

Procedure

- To configure virtual networks for user VM interfaces, see [Configuring a Virtual Network for Guest VM Interfaces](#) on page 138.
- To secure intra-cluster communications by configuring network segmentation, see "" () .

Configuring a Virtual Network for Guest VM Interfaces

Follow this procedure to configure a virtual network for guest VMs. You can associate the network with a VLAN and configure IP address management for interfaces on the network.

About this task

Note: Do not add any other device, including guest VMs, to the VLAN to which the Controller VM and hypervisor host are assigned. Isolate guest VMs on one or more separate VLANs.

To configure a virtual network for guest VM interfaces, do the following:

Procedure

1. Click the gear icon in the main menu and then select Network Configuration in the **Setting** page.


Note: This option only appears in clusters that support this feature.

The Network Configuration window appears.



Figure 98: Network Configuration Window

2. Click the **Virtual Networks** tab and then click the **Create Network** button.
The **Create Network** dialog box appears. Do the following in the indicated fields:



The 'Create Network' dialog box is a vertical form with a title bar at the top containing the text 'Create Network' and a question mark icon. The form is divided into several sections by horizontal lines. The first section contains a 'Name' label and a text input field. The second section contains a 'VLAN ID' label with a help icon and a text input field. The third section starts with a checked checkbox labeled 'Enable IP address management', followed by a descriptive text: 'This gives AHV control of IP address assignments within the network.' Below this are two text input fields labeled 'Network IP Address / Prefix Length' and 'Gateway IP Address'. The fourth section starts with a checked checkbox labeled 'Configure Domain Settings', followed by four text input fields labeled 'Domain Name Servers (Comma Separated)', 'Domain Search (Comma Separated)', 'Domain Name', and 'TFTP Server Name'. The fifth section contains a text input field labeled 'Boot File Name'. The sixth section is titled 'IP Address Pools' with a help icon, and contains a button labeled '+ Create Pool'. Below this is the text 'None defined.'. The seventh section contains an unchecked checkbox labeled 'Override DHCP server' with a help icon. At the bottom right of the dialog are two buttons: 'Cancel' and 'Save'.

Figure 99: Create Network Dialog Box

- a. **Name:** Enter a name for the network.

- b. **VLAN ID:** 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 appears as **vlan.1** or **vlan.27** in displays.

- c. **Enable IP Address Management:** Check the box to have the cluster control IP addressing in the network.

Checking this box displays more fields. If you do not check the box, Files does not attempt network management. In this case, it is assumed management for this virtual LAN is handled outside the cluster.

Note: If you do not enable this option while creating a network, you cannot enable IP address management.

- d. **Network IP Address/Prefix Length:** Enter the IP address of the gateway for the network and prefix with the network prefix (CIDR notation, for example, 10.1.1.0/24).
- e. **Gateway IP Address:** Enter the VLAN default gateway IP address.
- f. **Configure Domain Settings:** Check this box to display fields for defining a domain. Checking this box displays fields to specify DNS servers and domains. Unchecking this box hides those fields.
- g. **Domain Name Servers (comma separated):** Enter a comma-delimited list of DNS servers.
- h. **Domain Search (comma separated):** Enter a comma-delimited list of domains.
- i. **Domain Name:** Enter the VLAN domain name.
- j. **TFTP Server Name:** Enter the host name or IP address of the TFTP server from which virtual machines can download a boot file. Required in a Pre-boot eXecution Environment (PXE).
- k. **Boot File Name:** Name of the boot file to download from the TFTP server.



3. To define a range of addresses for automatic assignment to virtual NICs, click the **Create Pool** button (under **IP Address Pools**) and enter the following in the Add IP Pool dialog box:
If you do not assign a pool, assign IP addresses to VMs manually.


The image shows a dialog box titled "Add IP Pool" with a question mark icon in the top right corner. Inside the dialog, there are two text input fields. The first field is labeled "Start Address" and the second field is labeled "End Address". At the bottom of the dialog, there are two buttons: "Cancel" and "Submit".

Figure 100: Add IP Pool Dialog Box

- a. Enter the starting IP address of the range in the **Start Address** field.
 - b. Enter the ending IP address of the range in the **End Address** field.
 - c. Click the **Submit** button to close the window and return to the Create Network dialog box.
4. To configure a DHCP server, click the **Override DHCP server** box and enter an IP address in the **DHCP Server IP Address** field.
This address (reserved IP address for the Acropolis DHCP server) is visible only to VMs on this network and responds only to DHCP requests. If you do not check the box, the **DHCP Server IP Address** field does not display, and the DHCP server IP address generates 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`.

The image shows a section of a form. At the top, there is a checkbox that is checked, followed by the text "OVERRIDE DHCP SERVER". To the right of this text is a circular information icon containing the letter 'i'. Below this, there is a label "DHCP SERVER IP ADDRESS" followed by a large, empty text input field.

Figure 101: DHCP Server IP Field

- When all the information is correct, click the **Save** button to configure the network connection, close the Create Network dialog box, and return to the Network Configuration window.

Note:

- You can also specify network mapping to control network configuration for the VMs when they start on the remote site. For more information about configuring networking mapping on remote site, see "Configuring a Remote Site (Physical Cluster)" in the *Prism Web Console Guide*.
- Verify (or update as needed) that the physical switch configuration allows traffic for the same VLANs that are configured for the virtual networks.

Modifying Network Connections

Follow this procedure to modify network connections defined through the Nutanix virtualization management.

About this task

To modify or delete a network configuration, do the following:

Procedure

- Click the gear icon in the main menu and then select Network Configuration in the **Settings** page.

Note: This option only appears when running a hypervisor that supports this feature.

The Network Configuration window appears. It lists the following information for each defined network configuration:

- NAME:** Displays the name of the network.
- VLAN ID:** Displays the VLAN identification number for the network in the form **vlan.#**, for example **vlan.27** for virtual LAN number 27.



Figure 102: Network Configuration Window

2. To modify a network configuration, select the target line and click the pencil icon (on the right).

The Update Network Configuration dialog box appears, which contains the same fields as the Create Network Configuration dialog box (see [Configuring a Virtual Network for Guest VM Interfaces](#) on page 138). Do the following:

- a. Update the field values as desired.
 - b. When the configuration is correct, click the **Save** button to close the dialog box, update the network configuration, and return to the Network Configuration window.
3. To delete a network configuration, select the target line and click the X icon (on the right). A window prompt appears to verify the action; click the **OK** button. The network configuration is removed from the list.

Network Segmentation

You can segment the network on a Nutanix cluster by using Prism Element. For more information, see “” () .

Configuring Network Switch Information

About this task

To track and record networking statistics for a cluster, the cluster requires information about the first-hop network switches and the switch ports being used. Switch port discovery is supported with switches that are RFC 2674 compliant. Switch port discovery involves obtaining statistics from the Q-BRIDGE-MIB on the switch and then identifying the MAC address that corresponds to the host. Such discovery is currently best-effort, so it is possible that, at times, an interface might not be discovered.

Before you begin

Before you configure network switch information in the Prism web console, configure the corresponding SNMP settings on the first-hop network switch.

About this task

To configure one or more network switches for statistics collection, do the following:



Procedure

1. Click the gear icon in the main menu and then select **Network Switch** in the **Settings** page.

Note: Network switch configuration is supported only for AHV clusters.

The Network Switch Configuration dialog box appears.




Figure 103: Network Switch Configuration Window

2. Click the **Switch Configuration** tab and then click the **Add Switch Configuration** button.

3. Do the following in the indicated fields:

- a. **Switch Management IP Address:** Enter the Management IP address of the switch or the fully qualified switch name.
- b. **Host IP Addresses or Host Name (separated by commas):** Enter the IP address or fully qualified host name of each host in the cluster that uses this switch to route traffic.
When there are multiple hosts, enter the addresses in a comma separated list. Failing to add the host list might result in issues with switch port discovery.
- c. **SNMP Profile:** Select the SNMP profile to apply (or **None**) from the pull-down list.
Any SNMP profiles you have created appear in the list (see [Creating SNMP Profiles](#) on page 148).

Note: Selecting a profile populates the remaining fields automatically with the profile values. If you have not created a profile (or select **None** from the list), you must enter the values in the remaining fields manually.

- d. **SNMP Version:** Select the SNMP version to use from the pull-down list.
The options are **SNMPv2c** and **SNMPv3**.
- e. **SNMP Security Level (SNMPv3 only):** Select the security level to enforce from the pull-down list.
The options are **No Authorization No Privacy**, **Authorization But No Privacy**, and **Authorization and Privacy**. This field appears only when SNMPv3 is selected as the version.
- f. **SNMP Community Name:** Enter the SNMP community name to use.
- g. **SNMP Username:** Enter the SNMP user name.
- h. **SNMP Authentication Type:** Select the authentication protocol to use from the pull-down list.
The options are **None** and **SHA**.

Note: This field and the following three fields are set to **None** or left blank (and read only) when the version is SNMPv2c or the security level is set to no authorization.

- i. **SNMP Authentication Pass Phrase:** Enter the appropriate authentication pass phrase.
- j. **SNMP Privacy Type:** Select the privacy protocol to use from the pull-down list.
The options are **None**, **AES**, and **DES**.
- k. **SNMP Privacy Pass Phrase:** Enter the appropriate privacy pass phrase.
- l. When all the fields are correct, click the **Save** button.
This saves the profile and displays the new entry in the Switch Configuration tab.

Note: As a security protection, the **SNMP Authentication Pass Phrase** and **SNMP Privacy Pass Phrase** fields appear blank after saving (but the entered phrases are saved).



Network Switch Configuration

?

Switch Management IP Address

Host IP Addresses Or Host Names (Separated By Commas)

SNMP Profile

None

SNMP Version

SNMP v3

SNMP Security Level

No Authorization No Privacy

SNMP Community Name

SNMP Username

SNMP Authentication Type

None

SNMP Authentication Pass Phrase

SNMP Privacy Type

None

SNMP Privacy Pass Phrase

< Cancel

Save

Figure 104: Add Network Switch Configuration

Creating SNMP Profiles

About this task

If you need to configure multiple network switches, it may be useful to create one or more SNMP profiles that can be applied when configuring the switches (see [Configuring Network Switch Information](#) on page 144). To create an SNMP profile, do the following:

Procedure

1. Click the gear icon in the main menu and then select **Network Switch** in the **Settings** page.

Note: Network switch configuration is supported only for AHV clusters.

The Network Switch Configuration dialog box appears.




Figure 105: Network Switch Configuration Window

2. Click the **SNMP Profile** tab and then click the **Add SNMP Profile** button.

3. Do the following in the indicated fields:

- a. **Profile Name:** Enter a name for the profile.
The name can be anything as it is for display purposes only.
- b. **SNMP Version:** Select the SNMP version to use from the pull-down list.
The options are **SNMPv2c** and **SNMPv3**.
- c. **SNMP Security Level** (SNMPv3 only): Select the security level to enforce from the pull-down list.
The options are **No Authorization No Privacy**, **Authorization But No Privacy**, and **Authorization and Privacy**. This field appears only when SNMPv3 is selected as the version.
- d. **SNMP Community Name:** Enter the SNMP community name to use.
- e. **SNMP Username:** Enter the SNMP user name.
- f. **SNMP Authentication Type:** Select the authentication protocol to use from the pull-down list.
The options are **None** and **SHA**.

Note: This field and the following three fields are set to **None** or left blank (and read only) when the version is SNMPv2c or the security level is set to no authorization.

- g. **SNMP Authentication Pass Phrase:** Enter the appropriate authentication pass phrase.
- h. **SNMP Privacy Type:** Select the privacy protocol to use from the pull-down list.
The options are **None**, **AES**, and **DES**.
- i. **SNMP Privacy Pass Phrase:** Enter the appropriate privacy pass phrase.
- j. When all the fields are correct, click the **Save** button.
This saves the profile and displays the new entry in the SNMP Profile tab.



Network Switch Configuration?

Profile Name

SNMP Version

SNMP v3

SNMP Security Level

No Authorization No Privacy

SNMP Community Name

SNMP Username

SNMP Authentication Type

None

SNMP Authentication Pass Phrase

SNMP Privacy Type

None

SNMP Privacy Pass Phrase

SNMP PRIVACY PROTOCOL

None

< Cancel

Save

Figure 106: Network Switch Configuration Window (add SNMP profile)

Modifying Switch Information

About this task

To modify the information about a network switch (or the SNMP profile used in the switch information), do the following:

Procedure

1. Click the gear icon in the main menu and then select **Network Switch** in the **Settings** page.

Note: Network switch configuration is supported only for AHV clusters.

The Network Switch Configuration dialog box appears.




Figure 107: Network Switch Configuration Window

2. To modify a switch configuration (or SNMP profile), select the **Switch Configuration** (or **SNMP Profile**) tab, go to the target switch (or profile) entry in the table, and click the pencil icon.

This displays the configuration fields (see [Configuring Network Switch Information](#) on page 144 or [Creating SNMP Profiles](#) on page 148). Edit the entries as desired and then click the **Save** button.

3. To delete a switch configuration (or SNMP profile), select the **Switch Configuration** (or **SNMP Profile**) tab, go to the target switch (or profile) entry in the table, and click the X icon. This deletes the entry and removes it from the table.

Network Visualization

The network visualizer is a consolidated graphical representation of the network formed by the VMs and hosts in a Nutanix cluster and first-hop switches. You can use the visualizer to monitor the network and to obtain information that helps you troubleshoot network issues.

You can use the visualizer to view the following:

- Physical and logical network topology.
- Summary of the number and types of devices in the network.
- Network configuration of the devices in the topology and of components such as physical and virtual NICs.
- Real-time usage graphs of physical and virtual interfaces.

CAUTION: You cannot use the visualizer to configure the network. The network visualizer is available only on AHV clusters.

Prerequisites

Before you use the network visualizer, do the following:

- Configure network switch information on the Nutanix cluster. See [Configuring Network Switch Information](#) on page 144.
- Configure SNMP v3 or SNMP v2c on the first-hop switches. The visualizer uses SNMP for discovery and to obtain real-time usage statistics from the switches. For information about configuring SNMP, see the switch manufacturer's documentation.
- Enable LLDP or CDP on the first-hop switches. The visualizer uses LLDP or CDP to determine which switch port is connected to a given host interface. If LLDP or CDP is unavailable, SNMP data is used on a best-effort basis. For information about configuring LLDP or CDP, see the switch manufacturer's documentation.

Note: Network visualization depends on a functional internal DNS system to map switch hostnames and IP addresses based on the LLDP responses. Incorrect or partial DNS configuration displays inaccurate network details in the UI. To troubleshoot the network visualization issues, see [KB-4085](#) in the Nutanix Support portal.

Network Visualizer

The network visualizer displays interactive visual elements for the networked devices and for network components such as physical and logical interfaces. It also provides filtering and grouping capabilities that you can use to limit the display to a specific set of networked devices and connections.




Figure 108: Network Visualizer

The network visualizer includes the virtual networks pane and the topology view.

Virtual Networks Pane

Lists the virtual networks (VLANs) configured on the Nutanix cluster. Selecting a VLAN includes the VMs on that VLAN in the topology view. Conversely, clearing a check box excludes the VMs on that VLAN from the topology view. You can show up to five VLANs at a time.

The check box titled Other corresponds to VMs that are not on any VLAN. At a minimum, this check box is associated with the Controller VMs in the cluster.

Topology View

Displays the network topology. The topology view shows the following entities:

VLANs

VLANs configured on the cluster. The visualizer displays a different color for each VLAN. It shows the VLANs to which a VM or the VMs in a VM group belong. It also shows which VLANs are configured on a first-hop switch.

VMs

VMs on the VLANs that are selected in the virtual networks pane. Filter and group-by options enable you to customize the topology view.

Hosts

Hosts in the Nutanix cluster. The filter above the hosts enables you to specify which hosts you want to show in the topology view.

Switches

First-hop switches and the VLANs configured on each of them. The filter above the switches enables you to specify which switches you want to show in the topology view.

Viewing the Network Visualizer

About this task

To view the network visualizer, do the following:

Procedure

1. Log on to the web console.
2. Click **Network** in the main menu.

Customizing the Topology View

You can group VMs by a VM property or use a search filter to specify which network devices you want to show or exclude.

About this task

To customize the topology view, do the following:



Procedure

1. Specify which VLANs you want to show in the topology view:

- In the virtual networks pane, select the VLANs that you want to show in the topology view and clear those that you do not want to show.
- Select or clear **Other** if you want to include or exclude, respectively, the VMs that are not on any VLAN.

You can show a maximum of five VLANs at a time.

2. Specify which VMs you want to show in the topology view:

- Select a group-by option from the menu at the top of the VMs. The following group-by options are available:
 - **Power State.** Groups VMs by states such as On and Off. By default, the VMs are grouped by power state.
 - **Host.** Groups VMs by the host on which they are running.
 - **VM Type.** Groups VMs into guest VMs and Controller VMs.
- Enter a string in the search filter field to filter VMs that match the search string.
- If the group-by and filter operations result in a VM group, click the VM group to show the VMs in the group.

When you click a VM group, the visualizer displays ten VMs at first. To load ten more VMs, click **Load More VMs**.

To group the VMs again or to clear the filter, click **Back** beside the group-by menu.

3. Specify which Nutanix hosts you want to show in the topology view:

- a. Click the menu above the Nutanix hosts.
- b. Select the hosts that you want to show and clear those that you want to exclude.

4. Specify which switches you want to show in the topology view:

- a. Click the menu above the switches.
- b. Select the switches that you want to show and clear those that you want to exclude.

Viewing VM Information

In the visualizer, you can view the settings and real-time statistics of a virtual NIC.

About this task

To view VM information, do the following:

Procedure

1. Use the group-by and filtering capabilities of the visualizer (see [Customizing the Topology View](#) on page 153) to show the VM you want, and then click the name of the VM. The VM network information window appears.



2. From the virtual NIC list at the top of the dialog box, select the NIC for which you want show settings and statistics.

For information about the statistics that are displayed for the virtual NICs, see [VM NICs Tab](#) on page 263.

3. Optionally, point to a location on a graph to view the value at that point in time.
4. If you want additional information about the VM, click **Go to VM Details**.
The **VM** table view on the **VM** page is displayed.

Tip: You can return to the visualizer by pressing the back button in your browser.

Viewing Host Information

In the visualizer, you can view the internal network diagram of a Nutanix host. You can click a network component, such as a bridge or bond, to view the settings and statistics of that component.

About this task

To view host information, do the following:

Procedure

1. In the topology view, click the name of a host.
The host network information window appears. The window shows a diagram of the host's internal network. The right pane displays host network information.
2. Click any network component in the diagram and view its network settings and statistics in the right pane.

- » Click the Controller VM, and then, in the right pane, select a virtual NIC from the list to view settings and statistics of that virtual interface. Optionally, point to a location on a graph to view the value at that point in time.

For information about the statistics that are displayed for the virtual NICs, see [VM NICs Tab](#) on page 263.

- » Click a bridge to view the settings of the bridge.

A solid rectangle indicates an external bridge. A dotted rectangle indicates an internal bridge.

- » Click a bond to view the settings of the bond.

- » Click a host interface to view the settings of the interface. For a host interface, statistics are shown in addition to interface settings. Point to a location on a graph to view the value at that point in time. For information about the statistics that are displayed for the host NICs, see [Host NICs Tab](#) on page 184.

A solid line leading from a bond to a host interface indicates an active connection. A dotted line indicates a backup connection.

3. If you want additional information about the host, click **Go to VM Details**.
The **Host** table view on the **Hardware** page is displayed.

Tip: You can return to the visualizer by pressing the back button in your browser.



Viewing Switch Information

In the visualizer, you can view both switch-level information and interface-level information.

Viewing Switch-Level Information

Switch-level information includes details such as switch name, vendor, and management IP address.

About this task

To view switch-level information, do the following:

Procedure

1. In the topology view, click the name of a switch.
The switch information window is displayed.
2. To view additional switch information, click **Go to Switch Details**.
The **Switch** view on the **Hardware** page is displayed.

Tip: You can return to the visualizer by pressing the back button in your browser.

Viewing Interface-Level Information

Interface-level information includes details such as interface name, physical address, and interface type. The visualizer also displays statistics for each interface.

About this task

To view interface-level information, do the following:

Procedure

1. Click a switch interface.
The switch port information window is displayed. The window shows both network settings and interface statistics.

For information about the statistics that are displayed for switch interfaces, see [Switch Tab](#) on page 178.
2. Optionally, to show the value at any given point in time in a graph, point to the location on the graph.

Uplink Configuration (AHV Only)

The default NIC-teaming policy of the bond br0-up of the default bridge br0 in your AHV cluster is **Active-Backup**. You can change the NIC-teaming policy of br0-up to **Active-Active**, **Active-Active with MAC pinning**, or retain the default **Active-Backup** policy. If you select the **Active-Active** policy, you must manually enable Link Aggregation (LAG) and LACP on the corresponding top-of-rack (ToR) switch for each node in the cluster.

Note: The uplink configuration feature is supported only on AHV.

By default, the bond br0-up aggregates all the physical interfaces available on the node. You can modify this default NIC configuration by selecting the interfaces that must belong to br0-up. You can either choose to have only 10G interfaces, only 1G interfaces, or retain the default setting, that is all the available interfaces aggregated into br0-up.

See the *Host Network Management* chapter in the *AHV Administration Guide* for more information about the default bridge br0.



If you change the uplink configuration of br0-up, AOS applies the updated settings to all the nodes in the cluster one after the other (rolling update). To update the settings in a cluster, AOS performs the following tasks.

1. Puts the node in maintenance mode (migrates VMs out of the node)
2. Applies the updated settings
3. Checks connectivity with the default gateway
4. Exits maintenance mode
5. Proceeds to apply the updated settings to the next node

AOS takes a backup of the existing network settings before applying the updated settings on a randomly chosen first node. If the application of the updated settings fails on the first node, the previous settings are restored to the node and the updated settings are not applied to the rest of the nodes in the cluster. However, if the updated settings are successfully applied to the first node but fail on any subsequent node(s), alerts and warnings for the failed node(s) are generated and AOS proceeds to apply the updated settings on the remaining nodes in the cluster.

Note the following points about the uplink configuration.

- If the cluster has one or more compute-only nodes, AOS does not apply the **Active-Active with MAC pinning** and **Active-Active** policies.
- If you select the **Active-Active** policy, you must manually enable LAG and LACP on the corresponding ToR switch for each node in the cluster.
- If you reimage a cluster with the **Active-Active** policy enabled, the default NIC-teaming policy on the reimaged cluster is once again the **Active-Backup** policy.
- Nutanix recommends configuring LACP with fallback to active-backup or individual mode on the ToR switches. The configuration and behavior varies based on the switch vendor. Use a switch configuration that allows both switch interfaces to pass traffic after LACP negotiation fails. Otherwise the **eth** interface that the AHV node chooses as active might be different than the switch port that the switch chooses as active for some switch models. Therefore, approximately only 50% of the nodes may have connectivity and might result in connectivity failures once the cluster comes back up after it is reimaged.

To avoid any connectivity failures, Nutanix recommends that you keep one switch port powered on and shut down all the other switch ports on the switches connecting to br0-up for the same node. This setting ensures that both the node and the switch choose the same interface as active.

- The **Select All** option, which allows you to aggregate all the available physical interfaces into br0-up, is available only for the **Active-Backup** policy.

Updating the Uplink Configuration (AHV Only)

You can change the NIC-teaming policy and grouping of interfaces in br0-up.

About this task

Perform the following procedure to update the uplink configuration.

CAUTION: Changing the uplink configuration is not reversible. Be sure the planned updates are correct before attempting this procedure.

Procedure

1. Log on to the Prism Element web console.



- Click **Network** in the main menu and, in the network visualizer screen, click **Uplink Configuration**.




Figure 109: Uplink Configuration

- Select one of the following NIC-teaming policies.
 - » **Active-Backup.** A single interface in the bond is chosen randomly at boot to be active. The other interfaces act as a backup until the active interface fails. Do not configure LAG on the connected ToR switch. This is the recommended and default policy.
 - » **Active-Active with MAC pinning.** All the interfaces in the bond are active. Source MAC addresses are periodically rebalanced among interfaces based on the network load. Do not configure link aggregation on the connected ToR switch.
 - » **Active-Active.** All the interfaces in the bond are active. Traffic is balanced among adapters based on source and destination addresses and TCP and UDP ports.If you select this option, you must manually enable LAG and LACP on the corresponding ToR switch for each node in the cluster one after the other. See [Enabling LAG and LACP on the ToR Switch \(AHV Only\)](#) on page 159 for more information.

4. Select one of the following NIC configurations.

- » **Select All.** Aggregates all the physical interfaces available on the node. This is the default option.

Note: This option is available only for the active-backup policy.

- » **10G.** Aggregates only the 10G interfaces available in the node into br0-up.
- » **1G.** Aggregates only the 1G interfaces available in the node into br0-up.

Uplink Configuration

This configuration will apply to every host - any OVS configurations will be changed.

Uplink Configuration Without LAG

☒ Active-Backup

Show Details

☐ Active-Active with MAC pinning

Show Details

Uplink Configuration Requiring LAG

☐ Active-Active

Show Details

NIC Configuration

Select All NICs

Cancel

Save

Figure 110: Update Uplink Configuration

5. Click **Save**.

Note: AOS applies this configuration to every node in the cluster, one node at a time.

You can click **Tasks** in the main menu to monitor the progress of this operation.

Enabling LAG and LACP on the ToR Switch (AHV Only)

Active-Active with MAC pinning does not require link aggregation since each source MAC address is balanced to only one adapter at a time. If you select the **Active-Active** NIC-teaming policy, you must enable LAG and LACP on the corresponding ToR switch for each node in the cluster one after the other.

About this task

In an **Active-Active** NIC-teaming configuration, network traffic is balanced among the members of the team based on TCP and UDP ports. With link aggregation negotiated by LACP, the uplinks appear as a single layer-2 link so a VM MAC address can appear on multiple links and use the bandwidth of all uplinks. If you do not enable LAG and LACP in this **Active-Active** configuration, network traffic might be dropped by the switch.

Perform the following procedure to enable LAG and LACP on a TOR switch.

Procedure

1. Put the node in maintenance mode. This is in addition to the previous maintenance mode that enabled Active-Active on the node.
2. Enable LAG and LACP on the ToR switch connected to that node.
3. Exit maintenance mode after LAG and LACP is successfully enabled.
4. Repeat steps 1 to 3 for every node in the cluster.

HARDWARE MANAGEMENT

A Nutanix block is a 1U or 2U chassis that contains one to four independent nodes, each optimized for high-performance compute, memory, and storage. Each node includes CPUs, DIMMs, SSD and HDD storage, fans, network connectors (multiple 10 and 1GbE ports), and other typical hardware components. Each block also includes dual power supplies. A Nutanix cluster can contain an unlimited number of nodes (and blocks), and you can expand the cluster size at any time.

- You can monitor hardware configurations and status across the cluster through the web console (see [Hardware Dashboard](#) on page 161).
- You can expand the cluster through the web console (see [Expanding a Cluster](#) on page 185).



Figure 111: Nutanix Block

Hardware Dashboard


The Hardware dashboard displays dynamically updated information about the hardware configuration in a cluster. To view the Hardware dashboard, select **Hardware** from the pull-down list on the far left of the main menu.

Menu Options

In addition to the main menu (see [Main Menu](#) on page 46), the Hardware screen includes a menu bar with the following options:

- View selector. The Hardware dashboard provides three viewing modes.
 - Click the **Overview** button on the left to display hardware information in a summary view (see [Hardware Overview View](#) on page 163).
 - Click the **Diagram** button to display a diagram of the cluster nodes from which you get detailed hardware information by clicking on a component of interest (see [Hardware Diagram View](#) on page 165).
 - Click the **Table** button to display hardware information in a tabular form. The Table screen is further divided into host, disk, and switch views; click the **Host** tab to view host information, the **Disk** tab to view disk information, or the **Switch** tab to view switch information (see [Hardware Table View](#) on page 173).
- Expand cluster button. Click the **Expand cluster** button on the right to add nodes to the cluster (see [Expanding a Cluster](#) on page 185).
- Repair Host Boot Disk. Click the Repair Host Boot Disk button on the right to repair the boot drive of your hosts (see [Starting Host Boot Disk Repair \(Proactive or Graceful Replacement\)](#) on page 194).
- Page selector. In the Table view, hosts and disks are listed 10 per page. When there are more than 10 items in the list, left and right paging arrows appear on the right, along with the total count and the count for the current page.
- Export table content. In the Table view, you can export the table information to a file in either CSV or JSON format by clicking the gear icon





on the right and selecting either **Export CSV** or **Export JSON** from the pull-down menu. (The browser must allow a dialog box for export to work.) Chrome, Internet Explorer, and Firefox download the data into a file; Safari opens the data in the current window.



Figure 112: Hardware Dashboard Menu

Hardware Overview View

The Hardware Overview view displays hardware-specific performance and usage statistics on the left plus the most recent hardware-specific alert and event messages on the right. Several fields include a slide bar on the right to view additional information in that field. The displayed information is dynamically updated to remain current.

The following figure is a sample view, and the table describes each field in this view.




Figure 113: Hardware Overview View

Note: See [Understanding Displayed Statistics](#) on page 56 for information about how the statistics are derived.

Table 23: Hardware Overview View Fields

Name	Description
Hardware Summary	Displays the number of hosts and blocks in the cluster. It also displays the Nutanix model number.
Hosts	Displays the number of hosts in the cluster broken down by on, off, and suspended states. It also displays the number of discovered nodes that have not yet been added to the cluster.
Disks	Displays the total number of disks in the cluster broken down by tier type (SSD-PCIe, SSD-SATA, DAS-SATA).
CPU	Displays the total amount of CPU capacity (in GHz) in the cluster.
Memory	Displays the total amount of memory (in GBs) in the cluster.
Top Hosts by Disk IOPS	Displays I/O operations per host for the 10 most active hosts.
Top Hosts by Disk IO Bandwidth	Displays I/O bandwidth used per host for the 10 most active hosts. The value is displayed in an appropriate metric (MBps, KBps, and so on) depending on traffic volume.
Top Hosts by Memory Usage	Displays the percentage of memory capacity used per host for the 10 most active hosts.
Top Hosts by CPU Usage	Displays the percentage of CPU capacity used per host for the 10 most active hosts.

Name	Description
Hardware Critical Alerts	Displays the five most recent unresolved hardware-specific critical alert messages. Click a message to open the Alert screen at that message. You can also open the Alert screen by clicking the view all alerts button at the bottom of the list (see Alerts Dashboard).
Hardware Warning Alerts	Displays the five most recent unresolved hardware-specific warning alert messages. Click a message to open the Alert screen at that message. You can also open the Alert screen by clicking the view all alerts button at the bottom of the list.
Hardware Events	Displays the ten most recent hardware-specific event messages. Click a message to open the Event screen at that message. You can also open the Event screen by clicking the view all events button at the bottom of the list.

Hardware Diagram View

The Hardware Diagram view displays information about hosts and disks. The displayed information is dynamically updated to remain current.

The Hardware Diagram view screen is divided into two sections:

- The top section is an interactive diagram of the cluster blocks. Clicking on a disk or host (node) in the cluster diagram displays information about that disk or host in the summary section.
- The bottom **Summary** section provides additional information. It includes a details column on the left and a set of tabs on the right. The details column and tab content varies depending on what has been selected.

Note: See [Understanding Displayed Statistics](#) on page 56 for information about how the statistics are derived.




Figure 114: Hardware Diagram View

Host Details

Selecting a host in the diagram displays information about that host in the lower section of the screen.

- When a host is selected, **Summary: *host_name*** appears below the diagram, and action links appear on the right of this line:
 - Click the **Turn On LED** link to light up the host LED light on the chassis.
 - Click the **Turn Off LED** link to turn off the host LED light on the chassis.
 - Click the **Remove Host** link to remove this host from the cluster.
- Five tabs appear that display information about the selected host (see following sections for details about each tab): **Host Performance**, **Host Usage**, **Host Alerts**, **Host Events**, **Host NICs**.



Figure 115: Hardware Diagram View: Hosts

Table 24: Host Detail Fields

Parameter	Description	Values
Host Name	Displays the name of the host.	(host name)
Hypervisor IP	Displays the IP address assigned to the hypervisor running on the host.	(IP address)
Controller VM IP	Displays the IP address assigned to the Controller VM.	(IP address)
IPMI IP	Displays the IP address of the Intelligent Platform Management Interface (IPMI) port. An IPMI port is used for the hypervisor host console. This field does not appear in Prism Central.	(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)
Block Serial	Displays the block serial number.	(block serial number)

Parameter	Description	Values
Block Model	Displays the block model number.	(model series number)
Storage Capacity	Displays the total amount of storage capacity on this host.	xxx [GB TB]
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)
Memory	Displays the total memory capacity for this host.	xxx [MB GB]
CPU Capacity	Displays the total CPU capacity for this host.	xxx [GHz]
CPU Model	Displays the CPU model name	(CPU model name)
No. of CPU Cores	Displays the number of CPU cores on this host.	(number of CPU cores)
No. of Sockets	Displays the number of sockets.	(number of sockets)
No. of VMs	Displays the number of VMs running on this host.	(number)
Oplog Disk %	Displays the percentage of the operations log (oplog) capacity currently being used. The oplog resides on the metadata disk.	[0 - 100%]
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]
Monitored	Displays whether the host is high availability (HA) protected. A true value means HA is active for this host. A false 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 true . A false value is likely a sign of a problem situation that should be investigated.	[true false]
Hypervisor	Displays the name and version number of the hypervisor running on this host.	(name and version #)
Datastores	Displays the names of any datastores. This field does not appear in Prism Central.	(name)

Disk Details

Selecting a disk in the diagram displays information about that disk in the lower section of the screen.



- When a disk is selected, **Summary:** `disk_name` appears below the diagram, and action links appear on the right of this line:
 - Click the **Turn On LED** link to light up the LED light on the disk.
 - Click the **Turn Off LED** link to turn off the LED light on the disk.
 - Click the **Remove Disk** link to remove this disk from the cluster.
- Four tabs appear that display information about the selected storage container (see following sections for details about each tab): **Disk Usage**, **Disk Performance**, **Disk Alerts**, **Disk Events**.



Figure 116: Storage Diagram View: Disks

Table 25: Disk Detail Fields

Parameter	Description	Values
ID	Displays the disk identification number.	(ID number)
Serial Number	Displays the disk serial number.	(serial number)
Model	Displays the disk model number.	(model number)
Storage 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]
Used (Physical)	Displays the amount of used space on the drive.	xxx [GB TB]
Capacity (Physical)	Displays the total physical space on the drive.	xxx [GB TB]
Hypervisor	Displays the IP address of the hypervisor controlling the disk.	(IP address)
Storage Pool	Displays the name of the storage pool in which the disk resides.	(name)

Parameter	Description	Values
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
Mode	Displays whether the disk is currently online or offline.	[online offline]
Self Encryption Drive	Displays whether this is a self encrypting drive (SED).	[present not present]
Password Protection Mode [SED only]	Displays whether data-at-rest encryption is enabled for the cluster. When it is enabled (see Configuring Data-at-Rest Encryption (SEDs)), a key is required to access (read or write) data on the drive. This field appears only when the drive is a SED.	[protected not protected]

Cluster Summary Information

When a host or disk is not selected in the diagram (or when the word **Summary** is clicked), cluster-wide summary information appears in the lower part of the screen.

- The **Hardware Summary** column (on the left) includes the following fields:
 - **Blocks.** Displays the number of blocks in the cluster.
 - **Hosts.** Displays the number of hosts in the cluster.
 - **Total Memory.** Displays the total memory capacity (GBs) in the cluster.
 - **Total CPU Capacity.** Displays the total CPU capacity (GHz) in the cluster.
 - **Disks.** Displays the number of disks in each storage tier (DAS-SATA, SSD-SATA, and SSD-PCIe) in the cluster. Tier types vary depending on the Nutanix model type.
 - **Network Switches.** Displays the number of network switches being used in the cluster.
 - **GPUs.** (AHV only) Comma-separated list of GPUs installed on the host. GPU information includes the model name and a count in parentheses if multiple GPUs of the same type are installed on the host. If the firmware on the GPU is in compute mode, the string “compute” is appended to the model name. No string is appended if the GPU is in graphics mode.
The field is hidden if no GPUs are configured or if the hypervisor is not AHV.
- Three tabs appear that display cluster-wide information (see following sections for details about each tab): **Performance Summary**, **Hardware Alerts**, **Hardware Events**.



Performance Tab

The Performance tab displays graphs of performance metrics. The tab label and number of graphs varies depending on what is selected in the diagram:

- **Performance Summary** (no host or disk selected). Displays resource performance statistics (CPU, memory, and disk) across the cluster.
- **Host Performance** (host selected). Displays resource performance statistics (CPU, memory, and disk) for the selected host.
- **Disk Performance** (disk selected). Displays disk performance statistics for the selected disk.

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. For more in depth analysis, you can add a monitor to the analysis page by clicking the blue link in the upper right of the graph (see [Analysis Dashboard](#) on page 317). The Performance tab includes the following graphs:

- **[Cluster-wide] CPU Usage**: Displays the percentage of CPU capacity currently being used (0 - 100%) either across the cluster or for the selected host. (This graph does not appear when a disk is selected.)
- **[Cluster-wide] Memory Usage**: Displays the percentage of memory capacity currently being used (0 - 100%) either across the cluster or for the selected host. (This graph does not appear when a disk is selected.)
- **[Cluster-wide] IOPS**: Displays I/O operations per second (IOPS) for the cluster, selected host, or selected disk.
- **[Cluster-wide] I/O Bandwidth**: Displays I/O bandwidth used per second (MBps or KBps) for physical disk requests in the cluster, selected host, or selected disk.
- **[Cluster-wide] I/O Latency**: Displays the average I/O latency (in milliseconds) for physical disk requests in the cluster, selected host, or selected disk.




Figure 117: Performance Tab

Usage Tab

The Usage tab displays graphs of storage usage. This tab appears only when a host or disk is selected. The tab label varies depending on what is selected in the diagram:

- **Host Usage** (host selected). Displays usage statistics for the selected host.
- **Disk Usage** (disk selected). Displays usage statistics for the selected disk.

The Usage tab displays one or both of the following graphs:

- **Usage Summary:** Displays a rolling time interval usage monitor 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 (see [Analysis Dashboard](#) on page 317).
- **Tier-wise Usage** (host only): Displays a pie chart divided into the percentage of storage space used by each disk tier on the host. Disk tiers can include DAS-SATA, SSD-SATA, and SSD-PCIe depending on the Nutanix model type.




Figure 118: Usage Tab

Alerts Tab

The Alerts tab displays the unresolved alert messages about hosts, disks, and other hardware in the same form as the Alerts page (see [Alerts Summary View](#) on page 327). Click the **Unresolved X** button in the filter field to also display resolved alerts.

The screenshot shows the 'Host Alerts' tab with a table of unresolved alerts. The table has columns for SEVERITY, ISSUE, TIMESTAMP, ENTITIES, and DOCUMENTATION. The filter field shows 'Unresolved X' and '6 Alerts'.

SEVERITY	ISSUE	TIMESTAMP	ENTITIES	DOCUMENTATION
Info	NIC eth0 in host 10.4.60.203 has encountered more than 200 rx_errors in one hour.	01-16-15, 11:00:07pm	Host	Cause Resolution
Critical	There have been 10 or more cluster services restarts within 15 minutes in the Controller VM 10.4.60.207.	01-16-15, 04:28:07pm	Host	Cause Resolution
Critical	Controller VM 10.4.60.207 has been rebooted on Fri Jan 16 16:27:00 2015.	01-16-15, 04:29:04pm	Host	Cause Resolution
Critical	Stargate on Controller VM 10.4.60.207 is down for 331 seconds.	01-16-15, 04:22:45pm	Host	Cause Resolution
Critical	Hypervisor 10.4.60.203 is not reachable from Controller VM 10.4.60.208 in the last six attempts.	01-16-15, 04:22:57pm	Host	Cause Resolution
Critical	Hypervisor 10.4.60.203 is not reachable from Controller VM 10.4.60.208 in the last three attempts.	01-16-15, 04:19:06pm	Host	Cause Resolution

Figure 119: Alerts Tab

Events Tab

The Events tab displays the unacknowledged event messages about hosts, disks, and other hardware in the same form as the Events page (see [Events Summary View](#) on page 333). Click the **Include Acknowledged** button to also display acknowledged events.




Figure 120: Events Tab

Host NICs Tab

The Host NICs tab displays information in tabular form about the host NICs used to support traffic through the virtual NICs. (This tab appears only when a host is selected and the hypervisor is ESXi.) Each line represent a host NIC, and the following information is displayed for each NIC:

- **Host NIC.** Displays the host NIC name.
- **Speed (in KBps).** Displays the host NIC transmission speed.
- **MAC Address.** Displays the host NIC MAC address.
- **Received Packets.** Displays the number of packets received by the host NIC.
- **Transmitted Packets.** Displays the number of packets transmitted by the host NIC.
- **Dropped Rx Packets.** Displays the number of received packets dropped by the host NIC.
- **Dropped Tx Packets.** Displays the number of transmitted packets dropped by the host NIC.
- **Rx Rate.** Displays the host NIC received packets rate.
- **Tx Rate.** Displays the host NIC transmitted packets rate.
- **Network Usage Rate.** Displays the host NIC total network usage rate.

When you click a host NIC entry, a set of usage graphs about that NIC appear below the table:

- **Total Packets Received.** Displays a monitor of the total packets received over time. Place the cursor anywhere on the line to see the rate for that point in time. (This applies to all the monitors in this tab.)
- **Total Packets Transmitted.** Displays a monitor of the total packets that were transmitted.
- **Dropped Packets Received.** Displays a monitor of received packets that were dropped.
- **Dropped Packets Transmitted.** Displays a monitor of transmitted packets that were dropped.
- **Error Packets Received.** Displays a monitor of error packets that were received.



Figure 121: Host NICs Tab

Hardware Table View

The Hardware Table view displays information about hosts and disks in a tabular form. Click the **Host** tab in the screen menu bar to display host information; click the **Disk** tab to display disk information; click the **Switch** tab to display switch information. The displayed information is dynamically updated to remain current.

The Hardware table view is divided into two sections:

- The top section is a table. Each row represents a single host or disk and includes basic information about that host or disk. Click a column header to order the rows by that column value (alphabetically or numerically as appropriate).
- The bottom **Summary** section provides additional information. It includes a details column on the left and a set of tabs on the right. The details column and tab content varies depending on what has been selected.

Note: See [Understanding Displayed Statistics](#) on page 56 for information about how the statistics are derived.




Figure 122: Hardware Table Screen

Host Tab

Clicking the **Host** tab displays information about hosts in the cluster.

- The table at the top of the screen displays information about all the hosts, and the details column (lower left) displays additional information when a host is selected in the table. The following table describes the fields in the host table and detail column.
- When a host is selected, **Summary: *host_name*** appears below the table, and action links appear on the right of this line:
 - Click the **Turn On LED** link to light up the host LED light on the chassis.
 - Click the **Turn Off LED** link to turn off the host LED light on the chassis.
 - Click the **Repair Host Boot Device** link to repair the boot drive of the selected host.
 - Click the **Remove Host** link to remove this host from the cluster.
- Five tabs appear that display information about the selected host (see following sections for details about each tab): **Host Performance**, **Host Usage**, **Host Alerts**, **Host Events**, **Host NICs**.



Figure 123: Hardware Table View: Host Tab

Table 26: Host Table and Detail Fields

Parameter	Description	Values
Host Table Fields (upper screen)		
Host Name	Displays the name of the host.	(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.	[ESXi AHV Hyper-V]
CPU Usage	Displays the percentage of CPU capacity currently being used.	0 - 100%
CPU Capacity	Displays the CPU capacity of this host.	xxx [GHz]
Memory Usage	Displays the percentage of memory capacity currently being used by this host.	0 -100%

Parameter	Description	Values
Memory Capacity	Displays the memory capacity of this host.	xxx [MB GB]
Total Disk Usage	Displays the percentage of storage space used and total disk capacity of this host.	[0 - 100%] of xxx [TB]
Disk IOPS	Displays I/O operations per second (IOPS) for this host.	[0 - unlimited]
Disk IO B/W	Displays I/O bandwidth used per second for this host.	xxx [MBps KBps]
Disk IO Latency	Displays the average I/O latency (in milliseconds) for this host.	xxx [ms]
Host Detail Fields (lower screen)		
Host Name	Displays the name of the host.	(host name)
Hypervisor IP	Displays the IP address assigned to the hypervisor running on the host.	(IP address)
Controller VM IP	Displays the IP address assigned to the Controller VM.	(IP address)
IPMI IP	Displays the IP address of the Intelligent Platform Management Interface (IPMI) port. An IPMI port is used for the hypervisor host console. This field does not appear in Prism Central.	(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)
Block Serial	Displays the block serial number.	(block serial number)
Block Model	Displays the block model number.	(model series number)
Storage Capacity	Displays the total amount of storage capacity on this host.	xxx [GB TB]
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)
Memory	Displays the total memory capacity for this host.	xxx [MB GB]
CPU Capacity	Displays the total CPU capacity for this host.	xxx [GHz]
CPU Model	Displays the CPU model name	(CPU model name)
No. of CPU Cores	Displays the number of CPU cores on this host.	(number of CPU cores)
No. of Sockets	Displays the number of sockets.	(number of sockets)
No. of VMs	Displays the number of VMs running on this host.	(number)



Parameter	Description	Values
Oplog Disk %	Displays the percentage of the operations log (oplog) capacity currently being used. The oplog resides on every SSD.	[0 - 100%]
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 every SSD is reserved for the oplog.	xxx [GB]
Monitored	Displays whether the host is high availability (HA) protected. A true value means HA is active for this host. A false 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 true . A false value is likely a sign of a problem situation that should be investigated.	[true false]
Hypervisor	Displays the name and version number of the hypervisor running on this host.	(name and version #)
Datastores	Displays the names of any datastores.	(name)

Disk Tab

Clicking the **Disk** tab displays information about disks in the cluster.

- The table at the top of the screen displays information about all the disks, and the details column (lower left) displays additional information when a disk is selected in the table. The following table describes the fields in the disk table and detail column.
- When a disk is selected, **Summary: *disk_name*** appears below the table, and action links appear on the right of this line:
 - Click the **Turn On LED** link to light up the LED light on the disk.
 - Click the **Turn Off LED** link to turn off the LED light on the disk.
 - Click the **Remove Disk** link to remove this disk from the cluster.
- Four tabs appear that display information about the selected storage container (see following sections for details about each tab): **Disk Usage**, **Disk Performance**, **Disk Alerts**, **Disk Events**.



Figure 124: Hardware Table View: Disk Tab

Table 27: Disk Table and Detail Fields

Parameter	Description	Values
Disk Table Fields (upper screen)		
Disk ID	Displays the disk identification number.	(ID number)
Serial Number	Displays the disk serial number.	(serial number)
Hypervisor IP	Displays the IP address assigned to the hypervisor running on 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 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]
Disk IOPS	Displays I/O operations per second (IOPS) for this disk.	[0 - unlimited]
Disk IO B/W	Displays I/O bandwidth used per second for this disk.	xxx [MBps KBps]
Disk Avg IO Latency	Displays the average I/O latency for this disk.	xxx [ms]
Disk Detail Fields (lower screen)		
ID	Displays the disk identification number.	(ID number)
Serial Number	Displays the disk serial number.	(serial number)
Model	Displays the disk model number.	

Parameter	Description	Values
Storage 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]
Used (Physical)	Displays the amount of used space on the drive.	xxx [GB TB]
Capacity (Logical)	Displays the total physical space on the drive.	xxx [GB TB]
Hypervisor	Displays the IP address of the hypervisor controlling the disk.	(IP address)
Storage Pool	Displays the name of the storage pool in which the disk resides.	(name)
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
Mode	Displays whether the disk is currently online or offline.	[online offline]
Self Encryption Drive	Displays whether this is a self encrypting drive (SED).	[present not present]
Password Protection Mode [SED only]	Displays whether data-at-rest encryption is enabled for the cluster. When it is enabled (see Configuring Data-at-Rest Encryption (SEDs)), a key is required to access (read or write) data on the drive. This field appears only when the drive is a SED.	[protected not protected]

Switch Tab

Clicking the **Switch** tab displays information about the physical switches used by the host NICs to support traffic through the virtual NICs. The table at the top of the screen displays information about the switches, and the lower portion of the screen displays additional information when a switch is selected in the table. You can configure any number of switches (see [Configuring Network Switch Information](#) on page 144), but only the switches that are actually being used for virtual NIC traffic appear in this table. The following table describes the fields in the switch table, in the detail column (lower left), and in the Physical Switch Interfaces tab (lower right).



OverviewDiagramTab

HostCivSwitch

2 Servers

search in table

ID	Switch Name	Management Addresses	Vendor Name	Location Info	Contact Info	Description
dbr-h3-sw1	dbr-h3-sw1	10.4.8.99	Arista Networks			Arista Networks EOS version 4.25.8P running on an Arista Networks DCS-7950R-32
dbr-h3-sw2	dbr-h3-sw2	10.4.8.90	Cisco			Cisco IOS Software, C2960X Software (C2960X-UNIVERSALK9-M), Version 15.0(2)E1, RELEASE SOFTWARE (fc) Technical Support: http://www.cisco.com/techsupport Copyright (c) 1999-2013 by Cisco Systems, Inc. Compiled Wed 11-Sep-13 02:04 by 0151_RL29887

table

Summarydbr-h3-sw1

Switch Details

Physical Switch Interfaces

switch details

Physical Switch Interface	Switch ID	ID	IP	MAC Address	Linkset Tx Pkts	Linkset Rx Pkts	Error Tx Pkts	Error Rx Pkts	Discard Tx Pkts	Discard Rx Pkts
Ethernet2/1	dbr-h3-sw1	2008	9214	01:07:35:4b:0e:0e	10027	9887	0	0	0	0
Ethernet2/11	dbr-h3-sw1	22001	9216	01:07:35:4b:0e:0e	10349	10706	0	0	0	0
Ethernet2/2	dbr-h3-sw1	22002	9218	01:07:35:4b:0e:0e	33330	29609	0	0	0	0
Ethernet2/3	dbr-h3-sw1	22003	9216	01:07:35:4b:0e:0e	28990	28998	0	0	0	0

interface details

Figure 125: Hardware Table View: Switch Tab

Table 28: Switch Table and Detail Fields

Parameter	Description	Values
Switch Table Fields (upper screen)		
Switch ID	Displays the switch identification number.	(ID value)
Switch Name	Displays the switch name.	(name)
Management Addresses	Displays the switch management IP address(es).	(IP address)
Vendor Name	Displays the name of the switch vendor.	(company name)
Location Info	Displays the switch vendor location.	(company address)
Contact Info	Displays the switch vendor contact information.	(company contact)
Description	Describes the switch model and type.	(switch description)
Switch Detail Fields (lower left screen)		
Name	Displays the switch name.	(name)
Vendor Name	Displays the name of the switch vendor.	(company name)
Management Addresses	Displays the IP address(es) for the switch management ports.	(IP address)
Services	Displays the number of services being used.	(number)
Physical Switch Interfaces Fields (lower right screen)		
Physical Switch Interface	Displays the interface name.	(name)

Parameter	Description	Values
Switch ID	Displays the switch identification number.	(ID value)
Index	Displays the index value.	(number)
MTU (in bytes)	Displays the size in bytes of the largest protocol data unit (maximum transmission unit) that the layer can pass onwards.	(number)
MAC Address	Displays the interface MAC address	(mac address)
Unicast Rx Pkts	Displays the number of unicast packets received.	(number)
Unicast Tx Pkts	Displays the number of unicast packets transmitted.	(number)
Error Rx Pkts	Displays the number of received packets with an error.	(number)
Error Tx Pkts	Displays the number of transmitted packets with an error.	(number)
Discard Rx Pkts	Displays the number of received packets that were discarded.	(number)
Discard Tx Pkts	Displays the number of transmitted packets that were discarded.	(number)

When you click a physical switch interface entry, usage graphs about that interface appear below the table:

- **Unicast Packets Received.** Displays a monitor of the received unicast packets over time. Place the cursor anywhere on the line to see the value for that point in time. (This applies to all the monitors in this tab.)
- **Unicast Packets Transmitted.** Displays a monitor of the transmitted unicast packets.
- **Error Packets Received.** Displays a monitor of error packets received.
- **Dropped Packets Received.** Displays a monitor of received packets that were dropped.
- **Dropped Packets Transmitted.** Displays a monitor of transmitted packets that were dropped.




Figure 126: Physical Switch Interface Monitors

Cluster Summary Information

When a host or disk is not selected in the table (or when the word **Summary** is clicked), cluster-wide summary information appears in the lower part of the screen.

- The **Hardware Summary** column (on the left) includes the following fields:
 - Blocks.** Displays the number of blocks in the cluster.
 - Hosts.** Displays the number of hosts in the cluster.
 - Total Memory.** Displays the total memory capacity (GBs) in the cluster.
 - Total CPU Capacity.** Displays the total CPU capacity (GHz) in the cluster.
 - Disks.** Displays the number of disks in each storage tier (DAS-SATA, SSD-SATA, and SSD-PCIe) in the cluster. Tier types vary depending on the Nutanix model type.
 - Network Switches.** Displays the number of network switches being used in the cluster.
 - GPUs.** (AHV only) Comma-separated list of GPUs installed on the host. GPU information includes the model name and a count in parentheses if multiple GPUs of the same type are installed on the host. If the firmware on the GPU is in compute mode, the string “compute” is appended to the model name. No string is appended if the GPU is in graphics mode.

The field is hidden if no GPUs are configured or if the hypervisor is not AHV.

- Three tabs appear that display cluster-wide information (see following sections for details about each tab): **Performance Summary**, **Hardware Alerts**, **Hardware Events**.

Performance Tab

The Performance tab displays graphs of performance metrics. The tab label and number of graphs varies depending on what is selected in the table:

- Performance Summary** (no host or disk selected). Displays resource performance statistics (CPU, memory, and disk) across the cluster.
- Host Performance** (host selected). Displays resource performance statistics (CPU, memory, and disk) for the selected host.
- Disk Performance** (disk selected). Displays disk performance statistics for the selected disk.

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. For more in depth analysis, you can add a monitor to the analysis page by clicking the blue link in the upper right of the graph (see [Analysis Dashboard](#) on page 317). The Performance tab includes the following graphs:

- **[Cluster-wide] CPU Usage:** Displays the percentage of CPU capacity currently being used (0 - 100%) either across the cluster or for the selected host. (This graph does not appear when a disk is selected.)
- **[Cluster-wide] Memory Usage:** Displays the percentage of memory capacity currently being used (0 - 100%) either across the cluster or for the selected host. (This graph does not appear when a disk is selected.)
- **[Cluster-wide] IOPS:** Displays I/O operations per second (IOPS) for the cluster, selected host, or selected disk.
- **[Cluster-wide] I/O Bandwidth:** Displays I/O bandwidth used per second (MBps or KBps) for physical disk requests in the cluster, selected host, or selected disk.
- **[Cluster-wide] I/O Latency:** Displays the average I/O latency (in milliseconds) for physical disk requests in the cluster, selected host, or selected disk.




Figure 127: Hardware Table View: Performance Tab

Usage Tab

The Usage tab displays graphs of storage usage. This tab appears only when a host or disk is selected. The tab label varies depending on what is selected in the table:

- **Host Usage** (host selected). Displays usage statistics for the selected host.
- **Disk Usage** (disk selected). Displays usage statistics for the selected disk.

The Usage tab displays one or both of the following graphs:

- **Usage Summary:** Displays a rolling time interval usage monitor 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 (see [Analysis Dashboard](#) on page 317).
- **Tier-wise Usage (host only):** Displays a pie chart divided into the percentage of storage space used by each disk tier on the host. Disk tiers can include DAS-SATA, SSD-SATA, and SSD-PCIe depending on the Nutanix model type.




Figure 128: Hardware Table View: Usage Tab

Alerts Tab

The Alerts tab displays the unresolved alert messages about hosts, disks, and other hardware in the same form as the Alerts page (see [Alerts Summary View](#) on page 327). Click the **Unresolved X** button in the filter field to also display resolved alerts.

SEVERITY	ISSUE	TIME STAMP	ENTITIES	DOCUMENTATION
Info	NIC eth0 in host 10.4.60.203 has encountered more than 200 rx_errors in one hour	01-16-15, 11:00:07pm	Host	Cause Resolution
Critical	There have been 10 or more cluster services restarts within 15 minutes in the Controller VM 10.4.60.207	01-16-15, 04:29:07pm	Host	Cause Resolution
Critical	Controller VM 10.4.60.207 has been rebooted on Fri Jan 16 16:27:00 2015	01-16-15, 04:29:04pm	Host	Cause Resolution
Critical	Stargator on Controller VM 10.4.60.207 is down for 331 seconds.	01-16-15, 04:22:45pm	Host	Cause Resolution
Critical	Hypervisor 10.4.60.203 is not reachable from Controller VM 10.4.60.208 in the last six attempts.	01-16-15, 04:22:07pm	Host	Cause Resolution
Critical	Hypervisor 10.4.60.203 is not reachable from Controller VM 10.4.60.208 in the last three attempts.	01-16-15, 04:19:06pm	Host	Cause Resolution

Figure 129: Hardware Table View: Alerts Tab

Events Tab

The Events tab displays the unacknowledged event messages about hosts, disks, and other hardware in the same form as the Events page (see [Events Summary View](#) on page 333). Click the **Include Acknowledged** button to also display acknowledged events.

Performance Summary	Hardware Alerts	Hardware Events	
☐ Include Acknowledged · No entities found · < > ·  			
MESSAGE	ENTITIES	MODIFIED BY	TIMESTAMP

Figure 130: Hardware Table View: Events Tab

Host NICs Tab

The Host NICs tab displays information in tabular form about the host NICs used to support traffic through the virtual NICs. (This tab appears only when a host is selected.) Each line represents a host NIC, and the following information is displayed for each NIC:

- **Host NIC.** Displays the host NIC name.
- **Speed (in KBps).** Displays the host NIC transmission speed.
- **MAC Address.** Displays the host NIC MAC address.
- **Received Packets.** Displays the number of packets received by the host NIC.
- **Transmitted Packets.** Displays the number of packets transmitted by the host NIC.
- **Dropped Rx Packets.** Displays the number of received packets dropped by the host NIC.
- **Dropped Tx Packets.** Displays the number of transmitted packets dropped by the host NIC.
- **Rx Packet Errors.** Displays the number of error packets received by the host NIC.
- **Tx Packet Errors.** Displays the number of error packets transmitted by the host NIC.

When you click a host NIC entry, a set of usage graphs about that NIC appear below the table:

- **Total Packets Received.** Displays a monitor of the total packets received by the host NIC (in KBs or MBs) over time. Place the cursor anywhere on the line to see the value for that point in time. (This applies to all the monitors in this tab.)
- **Total Packets Transmitted.** Displays a monitor of the total packets transmitted by the host NIC (in KBs or MBs) .
- **Dropped Packets Received.** Displays a monitor of received packets that were dropped.
- **Dropped Packets Transmitted.** Displays a monitor of transmitted packets that were dropped.
- **Error Packets Received.** Displays a monitor for error packets received.



Figure 131: Host NICs Tab

Expanding a Cluster

A cluster is a collection of nodes. You can add new nodes to a cluster at any time after physically installing and 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 nodes to have the same AOS version.

Before you begin

- Review the relevant sections in [Cluster Expansion Guidelines](#) on page 190 before attempting to add a node to the cluster. The process for adding a node varies depending on several factors, and this section covers specific considerations based on your AOS, hypervisor, encryption, and hardware configuration.
- Check the Health Dashboard (see [Health Dashboard](#) on page 239). If any health checks are failing, resolve them to ensure that the cluster is healthy before adding any nodes.
- Allow any current add node operations to complete.
- Ensure that all nodes are in the correct metadata state by checking the Hardware dashboard (see [Hardware Dashboard](#) on page 161). 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**.

About this task

To add one or more nodes to an existing cluster (you can add multiple nodes at the same time), do the following:

Procedure

- Either click the gear icon in the main menu and then select **Expand Cluster** in the **Settings** page or go to the Hardware dashboard and click the **Expand Cluster** button.

The network is searched for Nutanix nodes and then the Expand Cluster dialog box appears (on the "Select Host" screen) with a graphical list of the discovered blocks and nodes. Discovered blocks are blocks with one or more unassigned factory-prepared nodes (hypervisor and Controller VM installed) residing on the same subnet as the cluster. Discovery requires that IPv6 multicast packets are allowed through the physical switch.

A lack of IPv6 multicast support might prevent node discovery and successful cluster expansion.

2. Select the check box for each block to be added to the cluster. All nodes within a checked block are also checked automatically; uncheck any nodes you do not want added to the cluster.

When a block is checked, more fields appear below the block diagram. A separate line for each node (host) in the block appears under each field name.

3. Do the following in the indicated fields for each checked block:

- a. **Host Name:** Enter the name of the host.

Enter just the host name, not the fully qualified domain name. The host name is required for Hyper-V clusters, but it is ignored for ESXi and AHV clusters.

- b. **Controller VM IP:** Review the Controller VM IP address assigned to each host and do one of the following:

- If the address is correct, do nothing in this field.
- If the address is not correct, either change the incorrect address or enter a starting address on the top line (for multiple hosts). The entered address is assigned to the Controller VM of the first host, and consecutive IP addresses (sequentially from the entered address) are assigned automatically to the remaining hosts.

- c. **Hypervisor IP:** Repeat the previous step for this field.

This field sets the hypervisor IP addresses for all the hosts to be added.

- d. **IPMI IP:** Repeat the previous step for this field.

This field sets the IPMI port IP addresses for all the hosts to be added. An IPMI port is used for the hypervisor host console.

- e. When all the node values are correct, click the **Next** button (lower right).

The network addresses are validated before continuing. If an issue is discovered, the problem addresses are highlighted in red. If there are no issues, the process moves to the Assign Rack screen with a message at the top when the hypervisor, AOS, or other relevant feature is incompatible with the cluster version.

4. [rack fault tolerance only] In the **Assign Rack** screen, select the block and the rack (from the **Assign to Rack** drop-down menu) to map the block placement in that rack. To add a rack to the list, click **Create Rack**.

This tab is available only when rack fault tolerance is enabled. See [Configuring Rack Fault Tolerance](#) on page 18 for details.



5. [data-at-rest encryption only] On the encryption page, do the following.

This step applies when data-at-rest encryption is enabled (either SEDs or software only). See "" () for more information.

NODE ADDRESS	ACTION
10.46.241.157	Generate and Download
10.46.241.154	Generate and Download
10.46.241.156	Generate and Download

Name	Address	Action
kms_safenet	10.1133.3 : 5696	Select Files

Figure 132: Expand Cluster Window: Encryption Page

- a. In the Certificate Signing Request Information field, click the **generate and download** link for each node to be added.

Clicking the link generates a certificate signing request (CSR) named **csr_for_discovered_node** for the node, which you download to your workstation.

- b. Get the CSRs signed by a certificate authority (CA).
- c. Click the **select files** link for each key management server and upload the signed certificates for the nodes to be added.
- d. Click the **Next** button.

The Configure Host dialog box opens.

6. If the node must be imaged, do the following in the **Hypervisor: <type>** field:

- » If a hypervisor image is listed in the **Hypervisor: <type>** field and it is the desired one, skip to the next step. If you uploaded a hypervisor image when adding nodes previously (and it is still compatible for imaging the new nodes), that image appears here. You can use that image or upload a different one.
- » If no hypervisor image is listed or the listed one is not the desired one, click the **Choose File** button. In the search window, find the image file on your workstation, and then click the **Open** button (in the search window) to upload that image file.

You must provide an ISO file for imaging ESXi or Hyper-V. You can get a compatible AHV image from the **Downloads > Hypervisor Details** page of the Nutanix Support portal (see [Accessing the Nutanix Support Portal](#) on page 648).

The AHV image to upload is an installation bundle (not an ISO file) such as **kvm_host_bundle_version#.tar.gz**. AOS includes a default AHV installation bundle named **kvm_host_bundle_version#.tar.gz**. If this version is appropriate, use it instead of uploading an installation bundle from the support portal.

- » If a message appears that the hypervisor image is not compatible, either select (choose and upload) a hypervisor image that is compatible or update the hypervisor ISO whitelist.

The cluster includes a hypervisor ISO whitelist, which lists the approved hypervisor images. You can only use an image that is in the approved list. See "Hypervisor ISO Images" in the *Field Installation Guide* for more information about the whitelist. If your hypervisor image does not appear in the whitelist (because you created the cluster before the image was approved), you can update the whitelist as follows:

1. Download the latest hypervisor ISO whitelist from the **Downloads > Foundation** page of the Nutanix Support portal.
 2. Click the **Update** link in the **Hypervisor ISO Whitelist** field.
 3. Click the **Choose File** button (which opens a search window), find and select the whitelist file on your workstation, and then click the **Open** button.
 4. Click the **Upload** button. The incompatible message disappears and you can continue provided the hypervisor image is on the uploaded whitelist.
7. [Hyper-V only] Specify the credentials to join the new nodes to Active Directory and to a failover cluster.
 - a. Specify the name of the Hyper-V failover cluster in the **Failover Cluster Name** text box.
 - b. Specify the user name and password of the domain account that has the privileges to create a new or modify an existing computer account in the Active Directory domain. The user name must be in the DOMAIN\USERNAME format.
 8. [Citrix Hypervisor only] Specify the credentials of the master node of the Citrix Hypervisor pool.
 9. When all the fields are correct, click the **Expand Cluster** button.

The Expand Cluster dialog box closes and the add node process begins. As the nodes are added, messages appear on the dashboard. A blue bar indicates that the task is progressing normally. Nodes are processed (upgraded or reimaged as needed) and added in parallel. Adding nodes can take some time. Imaging a node typically takes a half hour or more depending on the hypervisor.

Note: A red bar indicates a problem; hovering the cursor over the red bar text displays more information about the problem.




Figure 133: Progress Messages

10. If the cluster has multiple storage pools, assign the new storage capacity to a storage pool after the nodes are added successfully:

Note: When the cluster has only one storage pool, this step not needed because the new storage is added to that storage pool automatically.

- a. Select **Storage** from the pull-down main menu (upper left of screen) and then select the **Table** and **Storage Pool** tabs.
- b. Select the target storage pool (upper display) and then click the **update** link.

A storage pool must exist already to do this step (see [Creating a Storage Pool](#) on page 125).




Figure 134: Storage Pool Screen

- c. In the Update Storage Pool window, check the **Use unallocated capacity** box in the **Capacity** line (see [Creating a Storage Pool](#) on page 125) and then click the **Save** button.

This step adds all the unallocated capacity to the selected storage pool.

What to do next

One or more of the following items might apply to the added nodes (see also [Cluster Expansion Guidelines](#) on page 190):

- Non-default timezones are not updated on added nodes and must be reconfigured manually.
- If the Controller VM password for the cluster was changed from the default, the password on any new nodes gets updated automatically to match the cluster password.
- After cluster expansion, you can verify that the Controller VM memory on the new nodes is the same as the other Controller VMs in the cluster and sufficient for the intended workload. Run the `cvm_same_mem_level_check` and `cvm_memory_check` NCC checks (see [Running Checks by Using Web Console](#) on page 247).
- If the `cvm_same_mem_level_check` result is FAIL, the memory size is not the same as the other Controller VMs in the cluster. If the memory is less than the common size, increase the memory to match.
- If the `cvm_memory_check` result is FAIL, the Controller VM memory is less than the minimum required for the workload. Increase the memory to (at least) the minimum size.

To increase the memory, see [Increasing the Controller VM Memory Size](#) on page 93. Also, the hypervisor-specific administration guides include the procedure to change

(increase) Controller VM memory (*AHV Administration Guide*, *vSphere Administration Guide for Acropolis*, *Hyper-V Administration for Acropolis*, or *Citrix Hypervisor on Nutanix Administration Guide*). For Controller VM memory size recommendations, see the "Controller VM Memory Configurations" section of the *Acropolis Advanced Administration Guide*.

Cluster Expansion Guidelines

The process for adding a node varies depending on the AOS version, hypervisor type (AHV, ESXi, Hyper-V, or Citrix Hypervisor), data-at-rest encryption status, and certain hardware configuration factors.

AOS Considerations

The following apply to any cluster:

- The cluster expansion process checks the memory allocation of the existing Controller VMs in the cluster and increases the base Controller VM memory allocation on the new node if necessary. The new Controller VM is upgraded to a maximum 32 GB.

The Controller VM is upgraded to a maximum 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 node is reimaged before being added under certain conditions. The following table describes those conditions.

Table 29: 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 reimaged before it is added. Note: If the AOS version on the node is different (lower) but the hypervisor version is the same, you can upgrade just AOS from the command line. To upgrade AOS, log on to a Controller VM in the cluster and run the following command: <pre>nutanix@cvm\$ /home/nutanix/cluster/bin/cluster -u new_node_cvm_ip_address upgrade_node</pre> After the upgrade is complete, you can add the node to the cluster without reimaging it. Alternately, if the AOS version on the node is higher than the cluster, you must either upgrade the cluster to that version or reimage the node. To upgrade your cluster, see the <i>Acropolis Upgrade Guide</i>.
same AOS version but different hypervisor version	You have the option to reimage the node before adding it. Reimaging is appropriate in many such cases, but sometimes it may not be necessary such as for a minor version difference.

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).

ESXi Considerations

The following apply to clusters running ESXi:

- If the Controller VMs in the cluster reside in a VLAN configured network, you must first configure the nodes in the VLAN before attempting to add them. Otherwise, the discovery process does not find these nodes. See the "Configuring the Cluster in a VLAN-Segmented Network" section in the *Acropolis Advanced Setup Guide* for VLAN configuration instructions.
- To expand a cluster 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 target storage containers must be set to mount on the new hosts. You can check the mount status from the Storage dashboard (see [Storage Table View](#) on page 114). 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 (see [Hardware Diagram View](#) on page 165 or [Hardware Table View](#) on page 173). For instructions on enabling EVC, see the vSphere documentation.

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 CVM 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 need to enable APD and APD timeout on the new host.
- If you are adding new nodes to vCenter EVC configured cluster, ensure the following requirements.
 - ESXi version should be version 6.0 or above for enabling EVC.
 - All ESXi nodes should be on same version and build.

Hyper-V Considerations

The following apply to clusters running Hyper-V:

- If the Controller VMs in the cluster reside in a VLAN configured network, you must first configure the nodes in the VLAN before attempting to add them. Otherwise, the discovery process does not find these nodes. See the "Configuring the Cluster in a VLAN-Segmented Network" section in the *Acropolis Advanced Setup Guide* for VLAN configuration instructions.
- 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 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.

Citrix Hypervisor Considerations

When expanding a Citrix Hypervisor cluster, ensure the following:

- The nodes that you want to add are not part of any other Citrix Hypervisor pool.
- The nodes do not have any shared repositories.
- The nodes do not have any VMs running on them.



Data-At-Rest Encryption Considerations

The following apply to clusters with data-at-rest encryption enabled (see "" ()):

- Configure data-at-rest encryption for the new nodes. The new nodes must have self-encrypting drives or software-only encryption.
- If the cluster uses an external key manager, either with self-encrypting drives or software only, then reimaging cannot be done through the expand cluster work flow. In this case, any nodes to add must already have the correct hypervisor and AOS version. Use Foundation to image the nodes (if needed) before attempting to add them to the cluster.
- Adding a node to a cluster with self-encrypting drives where the added node is running a different hypervisor is not supported. In this case, image the node to the same hypervisor by using Foundation before adding it to the cluster. For more information, see [KB-4098](#).

Hardware Considerations

Note the following when it applies to the nodes you are adding:

- 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 G4 processor changes.

Power cycle guest VMs from the Prism 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.



Starting Host Boot Disk Repair (Proactive or Graceful Replacement)

This procedure describes how to proactively install a hypervisor on a boot device that has not completely failed (for example, a SATA DOM, or an M.2 device). A master node helps clone the boot device image onto the target node. This procedure works for all hypervisors. This procedure can clone and restore a complete image before the boot device fails. If this procedure fails, perform the "ungraceful" procedure for a boot device that has completely failed.

Before you begin

This process runs validation checks and clones the boot device to a hard drive. To begin this procedure the host must be running. Later, the node will be placed in maintenance mode and will be powered off.

CAUTION: Verify that it is safe to remove the node from the cluster before you start this procedure.

If you are running Hyper-V you must be using the Microsoft SCVMM to perform these procedures.

About this task

Procedure

1. From the Prism web console, **Hardware>Diagram** view, click the **Repair Boot Disk** button at the top right part of the window.




Figure 135: Repair Host Boot Device Menu (NX-3050 shown)

2. In the new window, select the host (node) that you will repair and click **Next**.




Figure 136: Select Node to Repair Window

Note: If the host IP cannot be reached, you must follow the procedure "Hypervisor Boot Drive Replacement (Failed Boot Device)" which is a separate document on the Nutanix service documentation portal. You will have to replace the failed boot device before you continue. (Some document titles might read "Failed SATA DOM" or "Failed M.2 Device".) If the IP address of the host can be reached, you can perform a graceful (proactive) boot device replacement by continuing this procedure.

3. Select **Restore boot disk data from snapshot** and click **Next**.





Figure 137: Repair Host Boot Disk Window (Graceful or proactive Replacement)

In AHV, the image is cloned immediately when you click **Next**. In ESXi, you must enter the IP credentials for vCenter.

4. (vSphere) Enter the IP credentials for vCenter.



Once the snapshot of the current disk is complete please replace the boot disk with a new one. Are you sure you wish to proceed?

Please enter the server details and credentials

vCenter IP Address:

Username:


Password:

No Yes

Figure 138: vCenter Login window

- a. Enter the vCenter IP address.
- b. Type the administrator username and password.
- c. Click **Yes**.

The Start host bootdisk repair precheck window is displayed.



Repair Host Boot Disk

Start host bootdisk repair precheck

Close

Figure 139: Start host bootdisk repair precheck

The Validation checks passed window is displayed.

The Host boot disk repair preprocess window appears. (You can also monitor the installation progress from the tasks menu or a console window.)




Figure 140: Tasks menu: Host boot disk repair preprocess window

The node is automatically placed in maintenance mode.

The window Waiting for the Node to boot up in phoenix is displayed.

The window Rebooting to phoenix is displayed.

The Cloning Bootdisk status window is displayed.




Figure 141: Rebooting to Phoenix window




Figure 142: Cloning bootdisk window

Note: You can check the status of this procedure at any time. Click the **Tasks** tab, or at the BOTTOM of the window, click the **Repair Boot Disk** button.

When the cloning process is complete, the node under repair is automatically powered-off and the Download Phoenix Image window is displayed.

5. Click **Download** in the Download Phoenix Image window.




Figure 143: Download Phoenix Image window

6. When the download is complete, click **Next**.




Figure 144: Download Phoenix Image window




Figure 145: Waiting for Image to Come Up Window

The Waiting for phoenix window is displayed. (You can also check the Prism Web Console **Tasks** menu: it displays the host status as **Waiting for phoenix to come up**). The wait time is approximately 10 hours.

7. Complete the Physically Replacing the Hypervisor Boot Drive topic and Booting the Node from the Phoenix Image section as soon as possible. The repair process waits for approximately 10 hours.

Modifying a Cluster

Hardware components (nodes and disks) can be removed from a cluster or reconfigured in other ways when conditions warrant it.

Before you begin

If the Data-at-Rest Encryption is enabled then before removing a drive or node from a cluster, test the certificates again by clicking **Test all nodes** and ensure that testing is successful and the status is **Verified**. For a detailed procedure, refer to the [Configuring Data-at-Rest Encryption \(SEDs\)](#). In case of an SED drive or node, if the drive or node is not removed as recommended then the drive or node will be locked.

About this task

To reconfigure one or more hardware components in a cluster, do the following:

Procedure

1. In the Hardware dashboard (see [Hardware Dashboard](#) on page 161), click the **Diagram** or **Table** tab.

The following actions can be performed from either the **Diagram** or **Table** view.




Figure 146: Cluster Hardware (Diagram View)




Figure 147: Cluster Hardware (Table View)

2. To remove a disk from the cluster (which is necessary when replacing a failed disk), either select the target disk in the diagram (Diagram view) or click the **Disk** tab and select that disk in the table (Table view), and click the **Remove Disk** link on the right of the **Summary** line. A dialog box appears to verify the action; click the **OK** button.

Removing a disk can take some time because data on that disk must be migrated to other disks before it can be removed from a node. You can monitor progress through the dashboard messages.

CAUTION: Do not physically remove a disk until that disk appears red in the diagram. The status message might indicate the data migration is complete, but the disk is not ready for removal until that disk turns red in the diagram.

3. To remove a host (node) from the cluster, either select the target host in the diagram (Diagram view) or click the **Host** tab and select that host in the table (Table view), and click the **Remove Host** link on the right of the **Summary** line. A dialog box appears to verify the action; click the **OK** button.
 - The Prism web console displays a warning message that you need to reclaim the license after you have removed the node. See the topics *Reclaiming Licenses* or *Reclaiming Licenses (Portal Connection)* in the [Nutanix Licensing Guide](#).
 - Removing a host takes some time because data on that host must be migrated to other hosts before it can be removed from the cluster. You can monitor progress through the dashboard messages. Removing a host automatically removes all the disks in that host. Only one host can be removed at a time. If you want to remove multiple hosts, you must wait until the first host is removed completely before attempting to remove the next host.
 - (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 is displayed.

Node node id is a part of a Hyper-V failover cluster . 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, cluster administrators must use the management tools provided by Microsoft such as Failover Cluster Manager to drain all the highly-available roles off the host and then remove the host from the failover cluster followed by removing the host from the AOS cluster.

- (ESXi only) VMware/cluster administrators must use the management tools provided by VMware such as vCenter to migrate all the VMs off the host and then remove the host from the vCenter cluster followed by removing the host from the AOS cluster.

CAUTION: Be sure to migrate VMs in an ESXi cluster prior to removing a host. Verify you have enough available compute capacity in the cluster (ESXi or Hyper-V) before actually migrating the VMs. If you click **Remove Host** without first migrating the VMs, the VM's data may be migrated without any notice and those VMs will lose service.

After a node is removed, it goes into an unconfigured state. You can add such a node back into the cluster through the expand cluster workflow (see [Expanding a Cluster](#) on page 185).

CAUTION: If the removed node is the first node (lowest IP address in the cluster) and you reuse this node in another cluster, the same cluster ID is applied to both clusters if this node again becomes the first node (lowest IP address in the new cluster). To prevent this, reimage the node (through Foundation) before adding it to the new cluster. See the preparing bare metal nodes chapter in the *Field Installation Guide* for details on how to image a node.



4. To add a host into the metadata store, either select the target host in the diagram (Diagram view) or click the **Host** tab and select that host in the table (Table view), and click the **Enable Metadata Store** link on the right of the **Summary** line.

Each node includes a disk used for metadata storage, and AOS maintains a metadata store across these disks to ensure uninterrupted resiliency should a metadata disk fail. After such a failure, that node is taken out of the metadata store group and the cluster continues to operate seamlessly without it. Normally, the node is brought back into the metadata store automatically after the failed metadata disk is replaced. However, under certain (rare) circumstances this might not happen. If the node is ready but was not added back automatically, the alert message **Node ready to be added to metadata store** is displayed and the **Enable Metadata Store** link appears indicating you can add the node back into the metadata store manually.

Compute-Only Node Configuration (AHV Only)

A compute-only (CO) node allows you to seamlessly and efficiently expand the computing capacity (CPU and memory) of your AHV cluster. The Nutanix cluster uses the resources (CPUs and memory) of a CO node exclusively for computing purposes.

You can use a supported server or an existing hyperconverged (HC) node as a CO node. To use a node as CO, you must first image the node as CO by using Foundation and then add that node to the cluster by using the Prism Element web console. See *Field Installation Guide* for information about how to image a node as a CO node.

Note: If you want an existing HC node that is already a part of the cluster to work as a CO node, you must first remove that node from the cluster, image that node as CO by using Foundation, and then add that node back to the cluster. See the [Modifying a Cluster](#) on page 198 topic for information about how to remove a node.

Key Features of Compute-Only Node

Following are the key features of CO nodes.

- CO nodes do not have a Controller VM (CVM) and local storage.
- AOS sources the storage for vDisks associated with VMs running on CO nodes from the hyperconverged (HC) nodes in the cluster.
- You can seamlessly manage your VMs (CRUD operations, ADS, and HA) by using the Prism Element web console.
- AHV runs on the local storage media of the CO node.
- To update AHV on a cluster that contains a compute-only node, use the Life Cycle Manager. See *Life Cycle Manager Guide: LCM Updates*.

Use Case of Compute-Only Node

CO nodes enable you to achieve more control and value from restrictive licenses such as Oracle. A CO node is part of a Nutanix HC cluster, and there is no CVM running on the CO node (VMs use CVMs running on the HC nodes to access disks). As a result, licensed cores on the CO node are used only for the application VMs.

Applications or databases that are licensed on a per CPU core basis require the entire node to be licensed and that also includes the cores on which the CVM runs. With CO nodes, you get a much higher ROI on the purchase of your database licenses (such as Oracle and Microsoft SQL Server) since the CVM does not consume any compute resources.



Minimum Cluster Requirements

Following are the minimum cluster requirements for compute-only nodes.

- The Nutanix cluster must be at least a four-node cluster before you add a compute-only node.
- The ratio of CO to HC nodes in a cluster must not exceed the following:
1 CO : 2 HC
- All the hyper-converged nodes in the cluster must be all-flash nodes.
- The number of vCPUs assigned to CVMs on the hyper-converged nodes must be greater than or equal to the total number of available cores on all the compute-only nodes in the cluster.
- The total amount of NIC bandwidth allocated to all the hyperconverged nodes must be twice the amount of the total NIC bandwidth allocated to all the compute-only nodes in the cluster.
- The AHV version of the compute-only node must be the same as the other nodes in the cluster.

See the [Expanding a Cluster](#) on page 185 topic for general requirements about adding a node to a Nutanix cluster.

Restrictions

Nutanix does not support the following features or tasks on a CO node in this release:

1. Host boot disk replacement
2. Network segmentation

Supported AOS Versions

Nutanix supports compute-only nodes on AOS releases 5.11 or later.

Supported Hardware Platforms

Compute-only nodes are supported on the following hardware platforms.

- All the NX series hardware
- Dell XC Core
- Cisco UCS

Networking Configuration

To perform network tasks on a compute-only node such as creating an Open vSwitch bridge or changing uplink load balancing, you must add the **--host** flag to the **manage_ovs** commands as shown in the following example:

```
nutanix@cvm$ manage_ovs --host IP_address_of_co_node --bridge_name bridge_name  
create_single_bridge
```

Replace **IP_address_of_co_node** with the IP address of the CO node and **bridge_name** with the name of bridge you want to create.

Note: Run the **manage_ovs** commands for a CO from any CVM running on a hyperconverged node.

You must perform the networking tasks for each CO node in the cluster individually.



See the *Host Network Management* section in the *AHV Administration Guide* for more information about networking configuration of the AHV hosts.

Adding a Compute-Only Node to an AHV Cluster

About this task

Perform the following procedure to add a compute-only node to a Nutanix cluster.

Procedure

1. Either click the gear icon in the main menu and then select **Expand Cluster** in the **Settings** page or go to the hardware dashboard (see [Hardware Dashboard](#) on page 161) and click **Expand Cluster**.
2. In the **Select Host** screen, scroll down and, under **Manual Host Discovery**, click **Discover Hosts Manually**.

Expand Cluster

1. Select Host 2. Configure Host

☐ null (Serial Number: 10-40-170-46)

☐ A ⓘ

☐ NX-8150 (Serial Number: 15SM85010042)

☐ A ⓘ

Manual Host Discovery ⓘ

+ Discover Hosts Manually

Cancel Next

Figure 148: Discover Hosts Manually

3. Click **Add Host**.

Expand Cluster

1. Select Host 2. Configure Host

Manual Host Discovery

Add host or CVM IPs to discover hosts within the same network.

+ Add Host

HOST OR CVM IP	ACTIONS
No Data	

Back Discover and Add Hosts

Figure 149: Add Host

4. Under **Host or CVM IP**, type the IP address of the AHV host and click **Save**.
This node does not have a Controller VM and you must therefore provide the IP address of the AHV host.
5. Click **Discover and Add Hosts**.
Prism Element discovers this node and the node appears in the list of nodes in the Select Host screen.
6. Select the node to display the details of the compute-only node.
7. Click **Next**.
8. In the **Configure Host** screen, click **Expand Cluster**.
The add node process begins and Prism Element performs a set of checks before the node is added to the cluster.
9. Check the **Hardware Diagram** view to verify if the node is added to the cluster.

Storage-only Node Configuration

A storage-only node allows you to seamlessly expand the storage capacity in your cluster. A storage-only node always runs the AHV hypervisor. If you add a storage-only node to any cluster (regardless of the hypervisor), the hypervisor on the storage-only node is always AHV. Therefore, if you want to scale up only the storage capacity in your cluster, you do not need to purchase additional hypervisor licenses.

A storage-only node in an AHV cluster (all AHV nodes) is called an AHV never-schedulable node. See the *Adding a Never-Schedulable Node* topic in the *AHV Administration Guide* for more information.

Key Features of Storage-only nodes

Following are the key features of storage-only nodes:

- A storage-only node includes an AHV hypervisor, a Controller VM (CVM), and memory and CPU resources enough to run only the CVM.
- A storage-only node always runs the AHV hypervisor. Therefore, if you add a storage-only node to an ESXi or Hyper-V cluster, the hypervisor on the storage-only node is always AHV.
- A storage-only node is supported on all NX hardware models. For other hardware models, check with the manufacturer if storage-only nodes are supported on that hardware.
- A storage-only node is supported on ESXi, Hyper-V, and AHV clusters.

If you add storage-only nodes to ESXi or Hyper-V clusters, you cannot run guest VMs on the storage-only nodes.

To ensure that no guest VMs run on a storage-only node in an AHV cluster, you can either configure AHV affinity rules or add the storage-only node as a never-schedulable node. See the *Adding a Never-Schedulable Node* topic in *AHV Administration Guide* for more information about never-schedulable nodes. You can run guest VMs on a storage-only node that is part of an AHV cluster (all AHV nodes). However, ensure that the storage-only node has enough memory and CPU resources.

Note: A storage-only node is not the same as a storage-heavy node. A storage-heavy node is a regular Nutanix hyperconverged node, but with a greater storage capacity. A storage-heavy node can run any hypervisor (AHV, ESXi, Hyper-V, or XenServer) and can run guest VMs.

Storage-only Node Best Practices

Following are the best practice recommendations if you plan to include storage-only nodes in your cluster.

- Ensure that your cluster has sufficient capacity if a storage-only node fails.

Your cluster must have sufficient capacity to handle the rolling-upgrade operations during which nodes are restarted.

The minimum number of storage-only nodes in your cluster must be equal to the replication factor (RF) of your cluster. For example, if the RF of your cluster is two, add at least two storage-only nodes. If the RF of your cluster is three, add at least three storage-only nodes.

- Ensure that the SSD capacity of the node is at least 10% of the HDD capacity of the node. For example, if your cluster has 20 TB HDD capacity, the SSD capacity of the node is at least 2 TB.

This recommendation does not apply to all-flash clusters.

- Ensure that you configure all CVMs in the cluster with the same resources. For example, if the CVMs on nodes other than the storage-only nodes have 32 GB RAM, configure the CVMs on the storage-only nodes with 32 GB RAM.
- Ensure that a cluster that has all storage-only nodes has at least 3 nodes.

The following table shows the supported and unsupported configurations for a cluster with storage-only nodes.



Configuration	Supported	Description
A four-node cluster with 4 TB capacity on each node. One large capacity storage-only node of 30 TB capacity.	X	If the only storage-only node in the cluster fails, the cluster does not have the sufficient capacity of 30 TB. Resolution: Add the minimum number of storage-only nodes equal to the replication factor of your cluster.
A four-node cluster with 4 TB capacity on each node. Two large capacity storage-only nodes of 30 TB capacity each.	#	If one storage-only node fails, the other can handle the required capacity.
A four-node cluster with 4 TB capacity on each node. One same-capacity storage-only node as the rest of the cluster.	#	The cluster is able to handle the failure of any node in the cluster, because all nodes in the cluster are of the same capacity.
Nutanix cluster with 2 TB SSD capacity and 20 TB HDD capacity.	#	The total SSD capacity is minimum 10% of the total HDD capacity.
Nutanix cluster with 2 TB SSD capacity and 40 TB HDD capacity.	X	The total SSD capacity is not minimum 10% of the total HDD capacity.

= Supported configuration

X = Unsupported configuration

Imaging and Adding Storage-only Nodes

You can image storage-only nodes when you are creating a cluster by using Foundation. You can also expand your existing cluster to add a storage-only node.

See the *Field Installation Guide* for instructions about to image storage-only nodes when you are creating a cluster.

If you want to add a node as a storage-only node to an existing cluster, you must first image the node as a storage-only node by using the Foundation Standalone VM (see the *Field Installation Guide* for more information) and use the Prism Element web console to add the storage-only node to the cluster (see the [Expanding a Cluster](#) on page 185 topic for more information).

NUTANIX VOLUMES

The Nutanix Volumes topics describe basic configuration information about this feature. For best practice information, Nutanix provides the [Nutanix Volumes Best Practice Recommendations](#). Volumes is not supported for ESXi clients or initiators.

AOS supports two methods for client iSCSI connectivity to volume groups.

Table 30: Available Methods for iSCSI Connectivity to Volume Groups

Method	Description
iSCSI initiators with MPIO	In this scenario, the operating system's native Multipath I/O (MPIO) capability controls path management for vDisk load balancing and path resiliency.
iSCSI initiators with Nutanix Volumes (Volumes)	Nutanix recommends that you use Nutanix Volumes instead of MPIO. In this scenario, Volumes provides highly available and high performance access to the iSCSI LUNs to clients with simple client configuration. Clients can be non-Nutanix servers external to the cluster or guest VMs internal or external to the cluster.

Volumes acts 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.

To provide access to the iSCSI storage and simplify client configuration management, Volumes exposes a single *iSCSI data services IP address* to clients for target discovery, which also simplifies external iSCSI configuration on clients. The iSCSI data services IP acts as an iSCSI target discovery portal and initial connection point. The client only needs this single IP address to help load balance storage requests and provide path optimization in the cluster, preventing bottlenecks.

Volumes does not require use of MPIO on the client. It is compatible with clients that are currently using or configured with MPIO.

A significant advantage of Volumes is that the cluster-wide iSCSI data services IP address enables you to expand your cluster without requiring you to reconfigure clients using Volumes. Volumes also makes it easy to scale storage capacity and performance without the need for downtime.

Volumes provides high availability by design. Because all nodes in a Nutanix cluster support Volumes, storage access and requests are redirected to surviving nodes if any node fails or becomes unavailable. Failover in this case is non-disruptive, with an initial minimal delay of 10 seconds or less possibly experienced by applications.

Securing Initiators and The Nutanix Cluster Target with CHAP

For additional security, you can use Challenge-Handshake Authentication Protocol (CHAP) authentication for iSCSI as part of [Creating a Volume Group for Use with Volumes](#) on page 211.

- Use one-way CHAP for basic security between the initiator and Nutanix cluster target. Nutanix recommends using one-way CHAP authentication.
- Use Mutual CHAP for additional security, where the client and target authenticate each other. When you configure a volume group, you can set a shared initiator CHAP secret, known only to the authenticator and peer. AOS manages the secrets, with the AOS `iscsi_adapter` implementing CHAP.

Example Use Cases Supported by Volumes

Volumes can support use cases including but not limited to:

- iSCSI for Microsoft Exchange Server. Volumes enables Microsoft Exchange Server environments to use iSCSI as the primary storage protocol.
- Shared storage for Windows Server Failover Clustering (WSFC). Volumes supports SCSI-3 persistent reservations for shared storage-based Windows clusters, commonly used with Microsoft SQL Server and clustered file servers.
- Bare-metal environments. Volumes enables non-Nutanix server hardware to consume the Acropolis Distributed Storage Fabric (DSF) resources. Workloads not targeted for virtualization can also use the DSF.
- Boot over iSCSI. Volumes enables and supports the ability to boot an operating system over iSCSI for physical servers.

Note: Linux guest VM clustering is supported only for Oracle RAC with Oracle Clusterware and Microsoft Windows Failover Clusters.

If You are Already Using Volume Groups

Volumes does not require multipath I/O (MPIO) configuration on the client but it is compatible with clients that are currently using MPIO.

If you are already using volume groups with MPIO and individual Controller VM IP addresses as targets, upgrading AOS will not disable or affect your existing configuration or functionality.

Nutanix recommends that you convert to Nutanix Volumes where applicable to take advantage of load balancing and other capabilities provided by the iSCSI data services IP address. See [Converting Volume Groups and Updating Clients to Use Volumes](#) on page 231.

Volumes Requirements and Supported Clients

Supported clients and associated requirements for using Nutanix Volumes. Unless otherwise specified, Nutanix has qualified 64-bit guest OS versions.

Volumes Requirements

- Volumes is not supported for ESXi clients or initiators.
- Synchronous Replication or Metro Availability are not currently supported for volume groups.



- You must configure an iSCSI data services IP address in **Cluster Details** available from the Prism web console. See [Modifying Cluster Details](#) on page 58.
- Ensure that ports 3260 and 3205 are open on any clients accessing the cluster where Nutanix Volumes is enabled.

Supported Client Operating Systems

- Microsoft Windows Server 2019
- Microsoft Windows Server 2016
- Microsoft Windows Server 2012 R2
- Microsoft Windows Server 2008 R2
- Red Hat Enterprise Linux 6.7
- Red Hat Enterprise Linux 6.8
- Red Hat Enterprise Linux 6.9
- Red Hat Enterprise Linux 7.2
- Red Hat Enterprise Linux 7.3
- Red Hat Enterprise Linux 7.4
- Red Hat Enterprise Linux 7.6
- Red Hat Enterprise Linux 7.7
- Oracle Linux 6.x
- Oracle Linux 7.x
- CentOS 7.x
- Oracle Solaris 11.3 on SPARC
- SuSE Linux Enterprise Server 11 / 12 (x86 servers)
- IBM AIX 7.1 / 7.2 on POWER¹

1. PowerHA cluster configurations are not supported as AIX requires a shared drive for cluster configuration information and that drive cannot be connected over iSCSI.

Supported Network Cards and Client Operating Systems, Booting An Operating System Over iSCSI Feature

Volumes supports the ability to boot an operating system over iSCSI for physical servers. In this configuration, a host can start a supported operating system from a LUN instead of a local disk instance.



Table 31: Client OS Support, Boot Over iSCSI

Network Device	Intel 10 GbE Network Adapter X520 / I350	QLogic QLE 8442 Converged Network Adapter	PCIe2 4-Port (10Gb +1GbE) SR+RJ45 Adapter (FC ENOS)
Client OS	-	-	-
Microsoft Windows Server 2008 R2	¹ #	#	
Microsoft Windows Server 2012 R2	#		
Red Hat Enterprise Linux 6.7 / 6.8	#	#	
Red Hat Enterprise Linux 7.2	#	#	
Oracle Linux 6.x			
Oracle Linux 7.x			
Oracle Solaris 11.3 on SPARC	#		
SuSE Linux Enterprise Server 11 / 12 (x86 servers)			
IBM AIX 7.1 / 7.2 on POWER			² #

1. # = Supported

2. To boot AIX over iSCSI, set the **use_redirection** property to **false** for the new or existing volume group where the boot LUN resides. See [Modifying a Volume Group \(AIX Boot Over iSCSI\)](#) on page 231.

Enabling Nutanix Volumes

About this task

This topic describes the first-use workflow for enabling and implementing Nutanix Volumes. To convert existing AOS volume groups and clients, see [Converting Volume Groups and Updating Clients to Use Volumes](#) on page 231.

Note: A valid client is a server running an operating system listed in [Volumes Requirements and Supported Clients](#) on page 208, or guest VMs using in-guest iSCSI.

Procedure

- Get the client IP address that you will add to a volume group client whitelist.
You can also use unique iSCSI initiator names (Initiator iSCSI Qualified Name [IQN]) from each client. See [Obtaining the Windows Client iSCSI Initiator Name](#) on page 220 or [Obtaining the Linux Client iSCSI Initiator Name](#) on page 224.



2. Create an iSCSI data services IP address for the Nutanix cluster.

This address cannot be the same as the cluster virtual IP address. See [Modifying Cluster Details](#) on page 58 and [iSCSI Data Services IP Address Impact](#) on page 59.

3. Provision storage on the Nutanix cluster
 - a. Create a volume group consisting of one or more vDisks.
 - b. Create a client whitelist to enable access to the volume group by clients, and specify the client IP addresses or client initiator IQNs in a whitelist (as part of the volume group configuration).
 - c. Create a secret for the volume group if you are using CHAP authentication.
See [Creating a Volume Group for Use with Volumes](#) on page 211.
4. Perform an iSCSI target discovery of the Nutanix cluster from the clients.
5. [Optional] Configure CHAP or Mutual CHAP authentication on the initiators and target Nutanix clusters.

Creating a Volume Group for Use with Volumes


About this task

Create a volume group to serve storage to one or more client initiators. Use the Create Volume Group window scrollbar to see all fields and buttons.



Procedure

1. In the **Storage** dashboard (see [Storage Dashboard](#) on page 103), click the **Volume Group** button.
The Create Volume Group dialog box is displayed.



The image shows a 'Create Volume Group' dialog box. It has a dark header bar with the title 'Create Volume Group', a question mark icon, and a close 'X' icon. Below the header, there are three input fields: 'NAME' (with a blue border and a cursor), 'ISCSI TARGET NAME PREFIX', and 'DESCRIPTION'. Below these fields is a section titled 'Storage' which contains a message 'You haven't added any disks yet.' and a blue button labeled '+ Add New Disk'.

Figure 150: Create Volume Group Dialog Box

2. Name and describe the volume group.

a. In **Name**, enter a name for the volume group.

Note: This entity has the following naming restrictions:

- The maximum length is 75 characters.
- Allowed characters are uppercase and lowercase standard Latin letters (A-Z and a-z), decimal digits (0-9), dots (.), and hyphens (-).

b. The **iSCSI Target Name Prefix** is auto-filled with the volume group **Name**. You can accept this prefix or enter your own target name prefix for the volume group. This entity has the same naming restrictions as **Name**.

You also use **iSCSI Target Name Prefix** as the user name (`username_in`) if you configure Mutual CHAP authentication for Linux clients.

c. In **Description**, provide a description for the volume group.

3. To add one or more disks to the volume group, do the following:

a. In the **Storage** section, click **Add New Disk**.

b. In the **Add Disk** dialog box, select the storage container to use from the **Storage Container** pull-down list. The list includes all storage containers created on this cluster.

c. In **Size**, enter the disk size in GiBs.

d. Click **Add**.

e. Repeat these steps to add another disk for this volume group, if desired.

What to do next

Add the iSCSI initiators to the volume group and enable client access to the cluster storage.



Add the Client iSCSI Initiators to the Volume Group

Procedure

1. Scroll down to the **Access Settings** section.

Create Volume Group

+ Add New Disk

Access Settings

☐ Enable external client access

Grant access to external clients such as bare-metal servers or VMs running outside the Nutanix cluster. This setting applies to all volume groups on this cluster.

☐ CHAP Authentication

TARGET PASSWORD

Client

You haven't added any clients yet.

+ Add New Client

Advanced Settings

☐ Enable Flash Mode ?

Cancel Save

Figure 151: Enable External Client Access

2. Select **Enable external client access** if you are whitelisting clients that are external to this cluster. Otherwise, leave this clear. If you select this checkbox, it remains automatically selected if you create more volume groups.
3. If you are using one-way CHAP security, select **CHAP Authentication** and type a 12 to 16 character password (also known as a CHAP secret) in the **Target Password** field. Initiators must use the same password to authenticate to the AOS cluster.
4. Click **Add New Client** to configure the initiators.
 - a. Enter the client Initiator iSCSI Qualified Name (IQN) in the **Client IQN/IP Address** field to create the whitelist.

Note: Ensure that you enter the client IQN in the **Client IQN/IP Address** field and not the IP address. AOS does not support a whitelist containing IP addresses in a volume group.

See [Obtaining the Windows Client iSCSI Initiator Name](#) on page 220 or [Obtaining the Linux Client iSCSI Initiator Name](#) on page 224.

- b. [Option] Select **CHAP Authentication** and enter the iSCSI client password (secret) used when you configure Mutual CHAP authentication on the client. See [Discovering the Volumes Target from the Windows Client](#) on page 215, [Configuring Mutual CHAP](#)

[Authentication \(Linux\)](#) on page 223, or [Adding iSCSI Targets on the AIX Host](#) on page 225.

c. Click **Add**.

Access Control displays any configured clients. This list includes any clients attached to volume groups in the cluster.

5. Click **Add New Client** in the **Create Volume Group** window to add more initiators allowed to access this storage.
6. Repeat these steps until you have added all the initiators that you want to whitelist.
7. Click **Save**.

What to do next

For the whitelisted initiators, perform an iSCSI discovery and target connection of the data services IP address for the Nutanix cluster on the client (Windows or Linux). See:

- [Discovering the Volumes Target from the Windows Client](#) on page 215
- [Discovering the Volumes Target from the Linux Client](#) on page 222
- [Changing the AIX Initiator Name](#) on page 225

Configuring Windows Clients

Discovering the Volumes Target from the Windows Client

Before you begin

1. You will need the Nutanix cluster data services IP address you created when you modified cluster details. See [Modifying Cluster Details](#) on page 58 and [iSCSI Data Services IP Address Impact](#) on page 59.
2. Make sure you have already created a volume group with the client IP address or initiator IQNs specified in the **Access Control** whitelist. Any LUNs associated with the volume group will be automatically discovered. See [Creating a Volume Group for Use with Volumes](#) on page 211.
3. From the Windows client, make sure that you have:
 - Started the *Microsoft iSCSI Initiator Service* and that the **Startup Type** is set to **Automatic**.
 - Set your firewall rules to allow iSCSI service traffic.

About this task

This procedure describes how to perform a Windows client iSCSI discovery and connection of the data services IP address for the Nutanix cluster (the Volumes target). This allows the cluster and its volume groups to be discoverable by the initiators. Do the following from Windows to discover and connect the Volumes target. You can also configure CHAP authentication here.

Procedure

1. Open the iSCSI Initiator Properties window.
If iSCSI Initiator is not available from **Administrative Tools**, you can open it by clicking **Start**, typing **iSCSI** in the search box, and clicking **iSCSI Initiator** from Programs.
2. In the iSCSI Initiator Properties window, click the **Discovery** tab.



3. Click **Discover Portal**, add the iSCSI data services IP address for the Nutanix cluster.




Figure 152: iSCSI Initiator Properties

4. Click **OK**.

5. To connect to the Volumes target, go to the **Targets** tab.
 - a. Click **Refresh**. The Volumes target is shown as Inactive.
 - b. Select the target and click **Connect**.
 - c. Select Add this connection to the list of Favorite Targets to ensure this connection is persistent. Ensure that **Enable multi-path** is disabled (not selected).
 - d. Click **OK**.




Figure 153: iSCSI Volumes Target

The Volumes target is now shown as **Connected**.

What to do next

See [Configuring CHAP Authentication \(Windows\)](#) on page 217.

Configuring CHAP Authentication (Windows)

Before you begin

Nutanix recommends setting one-way or Mutual CHAP authentication. Set the CHAP secret in the Nutanix cluster from the Prism web console Updating Volume Group dialog box first. Then configure CHAP authentication from the client side on the Windows server hosting your initiators.

About this task

This procedure describes how to configure CHAP authentication.

Procedure

1. Open the iSCSI Initiator Properties window.
2. In the iSCSI Initiator Properties window, click the **Targets** tab and select the Volumes target.

3. Click **Properties** to configure CHAP Authentication.
 - a. Ensure that **Enable Chap log on** is cleared (not selected).
 - b. Use the default **Name** as the initiator name
 - c. Type the same password in the **Target Secret** field as the one you used in the **Target Password** field in the Prism web console **Create Volume Group** dialog box.
 - d. For mutual CHAP authentication, select **Perform mutual authentication**. Otherwise, leave this setting cleared for one-way CHAP authentication.
 - e. Click **OK**.



The image shows a Windows-style dialog box titled "Advanced Settings". It has two tabs: "General" and "IPsec", with "IPsec" currently selected. The dialog is divided into several sections:

- Connect using:** A group box containing three dropdown menus: "Local adapter:" (set to "Default"), "Initiator IP:" (set to "Default"), and "Target portal IP:" (empty).
- CRC / Checksum:** A group box containing two checkboxes: "Data digest" and "Header digest", both of which are unchecked.
- Enable CHAP log on:** An unchecked checkbox.
- CHAP Log on information:** A text box explaining that CHAP ensures connection security by providing authentication between a target and an initiator. It instructs the user to specify the same name and CHAP secret configured on the target. The name will default to the Initiator Name of the system unless another name is specified.
- Name:** A text field containing the value "iqn.1991-05.com.microsoft:abs-pc.corp.nutanix.com".
- Target secret:** An empty text field.
- Perform mutual authentication:** An unchecked checkbox. Below it is a text box explaining that to use mutual CHAP, either specify an initiator secret on the Configuration page or use RADIUS.
- Use RADIUS to generate user authentication credentials:** An unchecked checkbox.
- Use RADIUS to authenticate target credentials:** An unchecked checkbox.

At the bottom right of the dialog are three buttons: "OK", "Cancel", and "Apply".

Figure 154: Configuring CHAP Authentication

- [Mutual CHAP configuration only; otherwise skip this step.] In the iSCSI Initiator Properties window, click the **Configuration** tab.
- Click **CHAP...**, type a password (or secret) in the **Initiator CHAP Secret** field, and click **OK**.




Figure 155: Windows Initiator Name

Obtaining the Windows Client iSCSI Initiator Name

About this task

This task is optional.

If you are configuring a volume group for use with Nutanix Volumes (Volumes), you can optionally use the iSCSI initiator name (Initiator iSCSI Qualified Name [IQN]) instead of the IP address from the Windows clients. For supported clients, see [Volumes Requirements and Supported Clients](#) on page 208.

Note: Each initiator IQN must be unique. Connections from two clients having same initiator IQN name might result in undesirable behavior or results.

Procedure

1. Open the iSCSI Initiator Properties window.

If iSCSI Initiator is not available from Administrative Tools, you can open it by clicking **Start**, typing iSCSI in the search box, and clicking **iSCSI Initiator** under Programs.

2. In the iSCSI Initiator Properties window, click the **Configuration** tab.
The **Initiator Name** field contains the initiator IQN name. Copy this name for use with the procedures in this section.




Figure 156: Windows Initiator Name

What to do next

Create a volume group. See [Creating a Volume Group for Use with Volumes](#) on page 211.

Configuring Linux Clients

Discovering the Volumes Target from the Linux Client

Before you begin

1. You will need the Nutanix cluster data services IP address you created when you modified cluster details. See [Modifying Cluster Details](#) on page 58 and [iSCSI Data Services IP Address Impact](#) on page 59.
2. Make sure you have already created a volume group with the initiators specified in the whitelist. Any vDisks associated with the volume group will be automatically discovered. See [Creating a Volume Group for Use with Volumes](#) on page 211.
3. For additional security, see [Configuring Mutual CHAP Authentication \(Linux\)](#) on page 223.

About this task

This procedure describes how to perform a Linux client iSCSI discovery and connection of the data services IP address for the Nutanix cluster (the Volumes target). This allows the cluster and its volume groups to be discoverable by the initiators. Do the following from a Linux client terminal window to discover and connect the Volumes target.

Procedure

1. On the linux client, ensure that the iSCSI service is started.
 - a. For Red Hat Enterprise Linux 6.0:

```
$ sudo /etc/init.d/iscsi status
```
 - b. For Red Hat Enterprise Linux 6.7:

```
$ sudo service iscsid status
```
2. If the **iscsi status** command returns any status other than **running**, start the service.

- a. For Red Hat Enterprise Linux:

```
$ sudo /etc/init.d/iscsi start
```

- b. For SuSE Linux Enterprise Server:

```
$ sudo rcopen-iscsi restart
```

3. Discover the Volumes target by specifying the iSCSI data services IP address on the default port 3260.

```
$ sudo /sbin/iscsiadm --mode discovery --type sendtargets \
--portal data_services_IP_address:3260
```

The command output will appear similar to `external_data_services_IP_address:3260,1 iqn_name`, where *iqn_name* is the Volumes target IQN.

If your volume group contains more than 32 vDisks, the Linux initiator connects to all vDisks using the virtual targets. That is, once you log in to the 32 virtual disks, running the **lsscsi** command on the Linux host displays all vDisks you have defined.

4. Connect the Volumes target by specifying *iqn_name* from the previous command.

```
$ sudo /sbin/iscsiadm --mode node --targetname iqn_name \
--portal data_services_IP_address:3260,1 --login
```



What to do next

Nutanix recommends that you configure CHAP authentication. See [Configuring Mutual CHAP Authentication \(Linux\)](#) on page 223.

Configuring Mutual CHAP Authentication (Linux)

For additional security, Volumes also enables you to use Challenge-Handshake Authentication Protocol (CHAP) authentication with Linux clients. Set the Mutual CHAP user names and passwords (secrets) from the client side on the Linux server hosting your initiators.

Procedure

1. Log in to the linux client and get the name of the initiator(s).

```
$ sudo cat /etc/iscsi/initiatorname.iscsi
```

For example, the command displays:

```
InitiatorName=iqn.1991-05.com.redhat:8ef967b5b8f
```

Copy this name for use with the procedures in this section.

2. Open the `/etc/iscsi/iscsid.conf` file with a text editor.
Search for the **CHAP SETTINGS** section in the file. You might see text similar to the following.

```
#####
#CHAP Settings
#####

#To enable CHAP authentication set node.session.auth.authmethod
#to CHAP. The default is None.
#node.session.auth.authmethod = CHAP

#To set a CHAP username and password for initiator
#authentication by the target(s), uncomment the following lines:
#node.session.auth.username = username
#node.session.auth.password = password
node.session.auth.username = iqn.somename.somedomain.com
node.session.auth.password = xxxxxxxxxxxxxx

#To set a CHAP username and password for target(s)
#authentication by the initiator, uncomment the following lines:
#node.session.auth.username_in = username_in
#node.session.auth.password_in = password_in
node.session.auth.password_in = yyyyyyyyyyyyyy

#To enable CHAP authentication for a discovery session to the target
#set discovery.sendtargets.auth.authmethod to CHAP. The default is None.
#discovery.sendtargets.auth.authmethod = CHAP
#discovery.sendtargets.auth.authmethod = CHAP

#To set a discovery session CHAP username and password for the initiator
#authentication by the target(s), uncomment the following lines:
#discovery.sendtargets.auth.username = username
#discovery.sendtargets.auth.password = password

#To set a discovery session CHAP username and password for target(s)
#authentication by the initiator, uncomment the following lines:
#discovery.sendtargets.auth.username_in = username_in
#discovery.sendtargets.auth.password_in = password_in
```

3. Be sure to uncomment any commented lines by removing the hash character [#] in the following steps.



4. Enable CHAP authentication.

```
#To enable CHAP authentication set node.session.auth.authmethod
#to CHAP. The default is None.
node.session.auth.authmethod = CHAP
```

5. Configure the initiator user name (initiator IQN) and initiator password.

```
node.session.auth.username = initiator_IQN
node.session.auth.password = initiator_password
```

6. Configure the Volumes target IQN user name and password. From the Prism web console, use the volume group **Target IQN Prefix** as the `username_in`.

```
node.session.auth.username_in = username_in
node.session.auth.password_in = password
```

- a. Log in to the Prism web console, click **Storage**, then select the **Table** view.
- b. Click **Volume Group** and select the volume group.
- c. Get the **Target IQN Prefix** in the Summary for the volume group. All disks in the volume group will have the same target IQN prefix.
For example:

```
iqn.2010-06.com.nutanix:vg1-5ff34411-080e-4b95-97c8-c2e34d9e1a82
```

- a. Create a `password`. Ensure that this password matches the **Add New Client > Chap Authentication** client password you specify in the [Add the Client iSCSI Initiators to the Volume Group](#) on page 214 procedure.
7. Save and exit the file. Log out of any connected sessions and perform a discovery. See [Discovering the Volumes Target from the Linux Client](#) on page 222.
 8. Repeat for any additional initiators used with Volumes.

Obtaining the Linux Client iSCSI Initiator Name

About this task

This task is optional.

If you are configuring a volume group for use with Nutanix Volumes, you can optionally use the iSCSI initiator name (Initiator iSCSI Qualified Name [IQN]) instead of the IP address from the Linux clients. For supported clients, see [Volumes Requirements and Supported Clients](#) on page 208.

Note: Each initiator IQN must be unique. Connections from two clients having same initiator IQN name might result in undesirable behavior or results.

Procedure

On the Linux client, open a terminal window and type:

```
$ sudo cat /etc/iscsi/initiatorname.iscsi
```

For example, the command displays:

```
InitiatorName=iqn.1991-05.com.redhat:8ef967b5b8f
```

Copy this name for use with the procedures in this section.



Note: If you change the initiator names, you must restart the **iscsi** service by typing **sudo /etc/init.d/iscsi restart**.

What to do next

Create a volume group. See [Creating a Volume Group for Use with Volumes](#) on page 211.

Configuring AIX Clients

Changing the AIX Initiator Name

The default iSCSI initiator name (IQN) of an AIX host does not strictly conform to the iSCSI RFC standard. Change the IQN of each AIX host before connecting to the Nutanix cluster.

About this task

Change the IQN of each AIX host to include the date (year and month) before adding the initiators to the cluster volume group or connecting to the Nutanix cluster.

Procedure

1. Log on to the AIX host through SSH.
2. Display the iSCSI initiator name (IQN) of the AIX host. For example, note the absence of the year and month in the name.

```
# lsattr -E -l iscsi0
```

```
nodename.iqn.aixhost1.hostid.0a4cac48
```

3. In this example, change the IQN to include the date (format yyyy-mm).

```
# chdev -l iscsi0 -a initiator_name=iqn.2016-04.aixhost1.hostid.0a4cac48
```

A successful message of **iscsi0 changed** is displayed.

4. Verify the change.

```
# lsattr -E -l iscsi0
```

```
nodename.iqn.2016-04.aixhost1.hostid.0a4cac48
```

5. Repeat these steps for each AIX host.

What to do next

See [Creating a Volume Group for Use with Volumes](#) on page 211 and [Adding iSCSI Targets on the AIX Host](#) on page 225.

Adding iSCSI Targets on the AIX Host

Before you begin

The AIX software initiator supports static discovery only. As a result, add the cluster iSCSI targets to the AIX host by editing the **/etc/iscsi/targets** file on the AIX host so the host can discover the cluster targets. You can obtain the volume group base target name from the Prism web console.



Procedure

1. Log in to the Prism web console and get the volume group virtual target IQN prefix names
 - a. Click **Storage**, then select the **Table** view.
 - b. Click **Volume Group** and select the volume group.
 - c. Obtain the **Target IQN Prefix** in the Summary for the volume group. All disks in the volume group will have the same target IQN prefix.
For example:

```
iqn.2010-06.com.nutanix:vg1-5ff34411-080e-4b95-97c8-c2e34d9e1a82
```

2. Log on to the AIX host through SSH.
3. Using a text editor, open the `/etc/iscsi/targets` file and add the cluster target information. Include the Controller VM IP address or iSCSI data services IP address and port 3260, then append `-tgtx` to the target IQN entry, where `x` starts at 0. The number of entries depends of the number of disks in the volume group. The maximum number of targets is 32 (that is, `-tgt0` through `-tgt31`). If there are more than 32 disks in the volume group, you do not have to enter additional targets beyond the 32nd disk.

Example: For a volume group consisting of a single disk:

```
10.1.216.192 3260 iqn.2010-06.com.nutanix:vg1-5ff34411-080e-4b95-97c8-c2e34d9e1a82-tgt0
```

Example: For a volume group consisting of 5 disks:

```
10.1.216.192 3260 iqn.2010-06.com.nutanix:vg1-5ff34411-080e-4b95-97c8-c2e34d9e1a82-tgt0
10.1.216.192 3260 iqn.2010-06.com.nutanix:vg1-5ff34411-080e-4b95-97c8-c2e34d9e1a82-tgt1
10.1.216.192 3260 iqn.2010-06.com.nutanix:vg1-5ff34411-080e-4b95-97c8-c2e34d9e1a82-tgt2
10.1.216.192 3260 iqn.2010-06.com.nutanix:vg1-5ff34411-080e-4b95-97c8-c2e34d9e1a82-tgt3
10.1.216.192 3260 iqn.2010-06.com.nutanix:vg1-5ff34411-080e-4b95-97c8-c2e34d9e1a82-tgt4
```

4. If you are using CHAP authentication, append the CHAP secret to the end of each entry, enclosed in double-quote characters.
For example:

```
10.1.216.192 3260 iqn.2010-06.com.nutanix:vg1-5ff34411-080e-4b95-97c8-c2e34d9e1a82-tgt0
"AIXforyou@123"
```

5. Save and close the file.

What to do next

See [Discovering the Volumes Target from the AIX Client](#) on page 226.

Discovering the Volumes Target from the AIX Client

Before you begin

1. Modify the AIX IQN name to make it compliant and useable by Volumes. See [Changing the AIX Initiator Name](#) on page 225.
2. Make sure you have already created a volume group with the initiators specified in the whitelist. Any LUNs associated with the volume group will be automatically discovered. See [Creating a Volume Group for Use with Volumes](#) on page 211.
3. Add iSCSI targets to the AIX host. See [Adding iSCSI Targets on the AIX Host](#) on page 225.



About this task

This procedure describes how to perform an AIX client iSCSI discovery and connection of the data services IP address for the Nutanix cluster (the Volumes target). This allows the cluster and its volume groups to be discoverable by the initiators. Do the following from an AIX client terminal window to discover and connect the Volumes target.

Procedure

1. Log on to the AIX host through SSH (if you are not already logged on).
2. Verify that the AIX host can contact the cluster targets.

```
# cfgmgr -v -l iscsi0
```

3. If the command is successful, AIX defines a new hard disk for each discovered LUN. Display them as follows.

```
# lsdev -c disk
```

The command result shows iSCSI disk drives.

Booting Over iSCSI

Configuring the Intel NIC

Volumes supports the ability to boot an operating system over iSCSI for physical servers. In this configuration, a host can start a supported operating system from a LUN instead of a local disk instance. This procedure describes how to configure the IBM PCIe BIOS settings to enable this feature.

Before you begin

- See the Intel Ethernet Adapter vendor documentation for more details about *Intel iSCSI Boot* configuration.
- According to Intel, you might be able to configure these settings through the adapter's **Properties > Data Options** tab in Microsoft Windows Device Manager.
- See [Volumes Requirements and Supported Clients](#) on page 208 for a list of the supported network hardware and clients.
- See [Enabling Nutanix Volumes](#) on page 210 and read about the procedures to perform on the Nutanix cluster. Before performing an iSCSI target discovery of the Nutanix cluster, configure BIOS boot settings for the network adapter as described in these steps.

About this task

Before performing an iSCSI target discovery of the Nutanix cluster, configure BIOS boot settings for the network adapter as described in these steps.

Procedure

1. See [Enabling Nutanix Volumes](#) on page 210 and read about the procedures to perform on the Nutanix cluster. Before performing an iSCSI target discovery of the Nutanix cluster, configure BIOS boot settings for the network adapter as described in these steps.
2. Open the **iSCSI Port Selection** page to configure the network adapter.
Start or restart the host machine where the network adapter is installed and press the **Ctrl+D** keys when you see the **Press Ctrl-D to run setup...** message during the boot sequence.



3. Select the adapter you are configuring and press **P** to select the adapter as the primary boot device, then press **Enter**.
4. Select **iSCSI Boot Configuration**, press **Enter**, then complete these fields .
Initiator network information

- **Use DHCP**. Clear (not selected).
- **Initiator Name**. iSCSI RFC 3720 compliant name.
- **Initiator IP**. Client IP address.
- **Subnet Mask**. IP subnet mask address
- **Gateway**. Network gateway IP address.

Nutanix iSCSI target information

- **Use DHCP for iSCSI Target Information**. Clear (not selected).
- **Target Name**. Use **iSCSI Target Name Prefix** Nutanix cluster target name from the volume group you created in [Creating a Volume Group for Use with Volumes](#) on page 211 appended with the disk index number **-tgtx** as shown in the volume group dialog box. You can select any disk in the volume group.

For example, for a volume group consisting of a single disk:

```
iqn.2010-06.com.nutanix:vg1-5ff34411-080e-4b95-97c8-c2e34d9e1a82-tgt0
```

- **Target IP**. Use the iSCSI data services IP address of the Nutanix cluster. See [Modifying Cluster Details](#) on page 58.
 - **Target Port**. 3260
 - **Boot LUN**. 0
5. Select **OK**.
 6. [Optional] Select **iSCSI CHAP Configuration** to configure CHAP/Mutual CHAP authentication.
 - » Select **Use CHAP** and enter a user name and target secret (from the procedures described in [Creating a Volume Group for Use with Volumes](#) on page 211).
 - » Select **Use Mutual CHAP** and enter an initiator secret.
 7. Select **OK**.
 8. Select **Save changes and Exit**.
 9. Press **Esc** to exit the **iSCSI Port Selection** page.
 10. Install the supported operating systems listed in [Volumes Requirements and Supported Clients](#) on page 208 on the iSCSI LUN.
 11. Perform an iSCSI target discovery of the Nutanix cluster from the clients.
 12. [Optional] Configure CHAP or Mutual CHAP authentication on the initiators and target Nutanix clusters.
See [Configuring CHAP Authentication \(Windows\)](#) on page 217 or [Configuring Mutual CHAP Authentication \(Linux\)](#) on page 223.



Configuring the QLogic HBA

Volumes supports the ability to boot an operating system over iSCSI for physical servers. In this configuration, a host can start a supported operating system from a LUN instead of a local disk instance. This procedure describes how to configure the IBM PCIe BIOS settings to enable this feature.

Before you begin

- See the QLogic 10GbE Converged Network Adapter vendor documentation for more details.
- See [Volumes Requirements and Supported Clients](#) on page 208 for a list of the supported network hardware and clients.
- See [Enabling Nutanix Volumes](#) on page 210 and read about the procedures to perform on the Nutanix cluster. Before performing an iSCSI target discovery of the Nutanix cluster, configure BIOS boot settings for the network adapter as described in these steps.

About this task

Before performing an iSCSI target discovery of the Nutanix cluster, configure BIOS boot settings for the network adapter as described in these steps.

Procedure

1. Open the **Comprehensive Configuration Management** page to configure the network adapter.
Start or restart the host machine where the network adapter is installed and press the **Ctrl +S** keys when you see the Ethernet Boot Agent banner during the boot sequence.
2. Select the adapter you are configuring and press **Enter**, then select **MBA Configuration**.
3. Go to **Boot Protocol** and select **iSCSI**.
4. Press **Esc** to go back to the Main Menu, select **iSCSI Boot Configuration** and press **Enter**.
5. Select **1st Target Parameters**, enable **Connect**, then complete these fields:
 - **IP Address.** Use the iSCSI data services IP address of the Nutanix cluster. See [Modifying Cluster Details](#) on page 58.
 - **TCP Port.** 3260
 - **Boot LUN.** 0
 - **iSCSI Name.** Use **iSCSI Target Name Prefix** Nutanix cluster target name from the volume group you created in [Creating a Volume Group for Use with Volumes](#) on page 211 appended with the disk index number **-tgtx** as shown in the volume group dialog box. You can select any disk in the volume group.
For example, for a volume group consisting of a single disk:
`iqn.2010-06.com.nutanix:vg1-5ff34411-080e-4b95-97c8-c2e34d9e1a82-tgt0`
 - **CHAP ID.** Use for CHAP / Mutual CHAP authentication .
 - **CHAP Secret.** Use for Mutual CHAP authentication. (from the procedures described in [Creating a Volume Group for Use with Volumes](#) on page 211).
6. Press **Esc** twice to exit the menu and save the configuration.



7. Press **CTRL+ALT+DEL** to apply adapter changes and exit the **Comprehensive Configuration Management** page.
8. See the QLogic 10GbE Converged Network Adapter vendor documentation to install the supported operating systems listed in [Volumes Requirements and Supported Clients](#) on page 208 on the iSCSI LUN.

Note: As described in the vendor documentation, Windows installations require you to create a slipstream Windows Server image that includes the network adapter drivers. During Windows installation, the iSCSI LUN will then be displayed in the list of disks where Windows Server will be installed.

9. Perform an iSCSI target discovery of the Nutanix cluster from the clients.
10. [Optional] Configure CHAP or Mutual CHAP authentication on the client initiators and target Nutanix clusters.
See [Configuring CHAP Authentication \(Windows\)](#) on page 217 or [Configuring Mutual CHAP Authentication \(Linux\)](#) on page 223.

Configuring the IBM PCIe Adapter

Volumes supports the ability to boot an operating system over iSCSI for physical servers. In this configuration, a host can start a supported operating system from a LUN instead of a local disk instance. This procedure describes how to configure the IBM PCIe BIOS settings to enable this feature.

Before you begin

- See the IBM Adapter vendor documentation for more details about [installing the base operating system on the iSCSI disk](#) and configuring iSCSI Boot.
- See [Volumes Requirements and Supported Clients](#) on page 208 for a list of the supported network hardware and clients.
- See [Enabling Nutanix Volumes](#) on page 210 and read about the procedures to perform on the Nutanix cluster. Before performing an iSCSI target discovery of the Nutanix cluster, configure BIOS boot settings for the network adapter as described in these steps.

About this task

Before performing an iSCSI target discovery of the Nutanix cluster, configure BIOS boot settings for the network adapter as described in these steps.

Procedure

1. Installing the AIX base operating system on the iSCSI disk.
2. After installing the base operating system on the iSCSI disk, power cycle the AIX machine to restart it.
3. As the machine is starting and icons appear at the bottom of your display, press F1 to open the System Management Services (SMS) menu.
4. Choose **Select Boot Options**.
5. Choose **Select Install/Boot Device**.
6. Choose **Network**, then select **iSCSI**.
The list shows the IBM adapter.



7. Select the IBM adapter and press **Enter**.
Your AIX machine restarts.

What to do next

See also [Creating a Volume Group for Use with Volumes](#) on page 211 and [Modifying a Volume Group \(AIX Boot Over iSCSI\)](#) on page 231.

Modifying a Volume Group (AIX Boot Over iSCSI)

Before you begin

Modify the default iSCSI initiator name (IQN) for AIX hosts to make it compliant and useable by Volumes. See [Changing the AIX Initiator Name](#) on page 225.

About this task

To boot AIX over iSCSI, use the Acropolis command line **accli** to set the **use_redirection** property to **false** for the new or existing volume group where the boot LUN resides.

Procedure

1. To access the Acropolis command line, log on to a Controller VM in the cluster with SSH and type **accli** at the shell prompt.
2. Set the volume group **use_redirection** property to **false**. Here, the example volume group name is **vg1** and example initiator name is **iqn.2016-04.aixhost1.hostid.0a4cac48**.

» Do this step for a new volume group.

```
vg.attach_external vg1 initiator_name=iqn.2016-04.aixhost1.hostid.0a4cac48:aeou \  
use_redirection=false
```

» Do this step for an existing volume group where the initiator is already attached by adding **use_redirection=false**.

```
vg.update_external vg1 initiator_name=iqn.2016-04.aixhost1.hostid.0a4cac48:aeou \  
use_redirection=false
```

3. Exit the command line by typing **exit**, then log out of the SSH session.

Converting Volume Groups and Updating Clients to Use Volumes

About this task

This topic describes the workflow to use Nutanix Volumes if you want to convert an existing volume group that includes attached external initiators with previously-discovered Controller VM IP address as target portal addresses. Use these procedures also if you are using clients currently configured to use multipath I/O (MPIO) and want to use Nutanix Volumes instead.

Procedure

1. From the Windows or Linux client: Disconnect and remove any existing targets configured on the client.
 - » [Modifying Linux Client Settings to Use Volumes](#) on page 232
 - » [Modifying Windows Client Settings to Use Volumes](#) on page 232



2. From the Prism web console:
 - a. Modify the existing AOS volume group to delete (detach) existing whitelisted initiator IQNs.
 - b. Add (re-attach) the initiators to the volume group.See [Modifying a Volume Group for Use with Volumes](#) on page 233.
3. Create an iSCSI data services IP address for the Nutanix cluster. This address cannot be the same as the cluster virtual IP address. See [Modifying Cluster Details](#) on page 58.
4. From the Windows or Linux client: Discover and connect to the new targets. Do not use MPIO. (For Windows, ensure that **Enable multi-path** is disabled (not selected)).

Modifying Windows Client Settings to Use Volumes

About this task

For all Controller VM IP addresses previously discovered as targets, disconnect each session and delete the target portal IP addresses. Then discover and connect to the new target.

Procedure

1. From Windows Server, open the iSCSI Initiator Properties window to disconnect and remove any existing targets.

If iSCSI Initiator is not available from **Administrative Tools**, you can open it by clicking **Start**, typing **iSCSI** in the search box, and clicking **iSCSI Initiator** from Programs.

 - a. In the iSCSI Initiator Properties window, click the **Targets** tab.
 - b. Select a discovered target, then click **Disconnect**.
 - c. In the iSCSI Initiator Properties window, click the **Discovery** tab.
 - d. In **Target portals**, select each Controller VM IP address, then click **Remove**.
 - e. Click the **Favorite Targets** tab, select the discovered target, then click **Remove**.

This step is needed because Windows clients will connect to **Favorites Targets** even if you remove them from the **Discovery** tab.
2. From the Prism web console, delete (detach), then add (re-attach) initiator IQNs in the AOS volume group as described in [Modifying a Volume Group for Use with Volumes](#) on page 233.
3. Create an iSCSI data services IP address for the Nutanix cluster. This address cannot be the same as the cluster virtual IP address. See [Modifying Cluster Details](#) on page 58.
4. Discover and connect to the new target as described in [Discovering the Volumes Target from the Windows Client](#) on page 215.

Modifying Linux Client Settings to Use Volumes

About this task

For all Controller VM IP addresses previously discovered as targets, disconnect by logging out of each session and delete the target portal IP addresses. Then discover and connect to the new target.



Procedure

1. Open a Linux terminal window and log out of the iSCSI targets to disconnect the iSCSI session by doing one of the following.

- a. Log out of all targets.

```
$ sudo /sbin/iscsiadm -m node -u
```

- b. Log out of a specific target.

```
$ sudo /sbin/iscsiadm --mode node --targetname iqn_name \  
--portal target_IP_address:3260,1 --logout
```

2. Remove the now-disconnected target records from **discoverydb**.

```
$ sudo /sbin/iscsiadm -m node -o delete
```

3. From the Prism web console, delete (detach), then add (re-attach) initiator IQNs in the AOS volume group as described in [Modifying a Volume Group for Use with Volumes](#) on page 233.
4. Create an iSCSI data services IP address for the Nutanix cluster. This address cannot be the same as the cluster virtual IP address. See [Modifying Cluster Details](#) on page 58.
5. Discover and connect to the new target as described in [Discovering the Volumes Target from the Linux Client](#) on page 222.

Modifying a Volume Group for Use with Volumes

Before you begin

If you choose to use iSCSI initiator names instead of whitelisting IP addresses, get the iSCSI initiator name(s) from your client server (Windows Server or Linux). See [Obtaining the Windows Client iSCSI Initiator Name](#) on page 220 or [Obtaining the Linux Client iSCSI Initiator Name](#) on page 224.

To boot AIX over iSCSI, set the **use_redirection** property to **false** for the new or existing volume group where the boot LUN resides. See [Modifying a Volume Group \(AIX Boot Over iSCSI\)](#) on page 231.

About this task

Modify an existing AOS volume group to delete (detach) existing whitelisted initiator IQNs from it, then add (re-attach) the initiators to the volume group. Use the scrollbar to see all fields and buttons.

Procedure

1. In the Prism web console, select **Storage** from the pull-down main menu, and then select the **Table** and **Volume Group** tabs.
2. To update a volume group, select the volume group, and then click the **Update** link. The Update Volume Group dialog box is displayed.
3. Scroll down to the **Access Settings** section and clear (uncheck) any IQN initiators listed. Any unchecked IQN initiators that remains unused by this or any other volume group will ultimately disappear from this list.
4. Click **Save**.



5. Select the volume group again and click **Update**.
6. To whitelist the initiators that must access the volume group, in the **Initiators** section, see [Add the Client iSCSI Initiators to the Volume Group](#) on page 214.

Note: If the initiator is a VM on AHV and you want to attach the volume group as a SCSI disk, skip this step and save the volume group. After it is saved, modify the volume group to attach VMs.

7. Click **Save**.

What to do next

For the whitelisted initiators, perform an iSCSI discovery and connection of the data services IP address for the Nutanix cluster on the client (Windows or Linux). See [Discovering the Volumes Target from the Windows Client](#) on page 215 or [Discovering the Volumes Target from the Linux Client](#) on page 222.

Creating a Linux Guest Cluster

Linux guest clustering provides high-availability for cluster-aware workloads by failing over applications to VMs on a designated host. Supported client distributions include CentOS, Red Hat Enterprise Linux (RHEL), and SUSE Enterprise Linux (SLES). See [Volumes Requirements and Supported Clients](#) on page 208 for the specific CentOS, RHEL, and SLES versions.

Before you begin

- Linux guest clustering is only available for Volumes. Use a command-line interface to configure the failover environment on a guest VM.
- AOS clusters running the following hypervisor hosts support Linux guest clustering.
 - AHV
 - Hyper-V
 - ESXi
- When using SUSE Enterprise Linux, ensure that volume used by target VMs is configured to automatically mount when these VMs start or restart. You can configure this setting by using the YaST tool available in SLES.

Set startup mode to automatic for the corresponding targets for the iSCSI Initiator: **YAST > Network Services > iSCSI Initiator > Connected Targets > Edit > Startup > automatic**

- As an admin user, perform the steps described in this topic for every VM on the Linux guest cluster.

Procedure

1. Log on to the web console for your AOS cluster as the admin user and create a volume group (VG) with whitelisted external clients, see [Creating a Volume Group for Use with Volumes](#) on page 211.
2. Install the iSCSI initiator package on the client (Linux guest VM).

```
user@host$ sudo yum install iscsi-initiator-utils
```



3. Check that the iSCSI daemon is running on the guest VM. For example:

» For RHEL 6.

```
user@host$ sudo /etc/init.d/iscsi status
```

» For RHEL 7.

```
user@host$ sudo service iscsid status
```

» For SLES 12.

```
user@host$ sudo systemctl status iscsid
```

If the returned status is **running**, continue to step 5. Otherwise, continue to the next step.

4. If the returned status is not **running**, start the iSCSI daemon. For example:

» For RHEL 6.

```
user@host$ sudo /etc/init.d/iscsi start
```

» For RHEL 7.

```
user@host$ sudo service iscsid start
```

» For SLES 12.

```
user@host$ sudo systemctl start iscsid
```

5. Attach the initiator iSCSI Qualified Name (IQN) to the VG:

a. On the **Home** page in Prism Element, click the name of the cluster to launch the Cluster Details window.

b. Locate the **ISCSI Data Services IP**.

c. Now, use the iSCSI data services IP to discover the available targets on the host.

```
user@host$ sudo iscsiadm -m discovery -t st -p <external-data-services-ip-address>
```

6. Connect to the discovered targets.

```
user@host$ sudo iscsiadm -m node --login
```

7. Verify that the client successfully attached the vDisks.

```
user@host$ ls --l /dev/disk/by-path
```

What to do next

- For RHEL VMs: Install Red Hat High Availability Add-On packages and use the newly created vDisks to set up a guest cluster. Refer to the configuration steps in [Red Hat documentation](#) Clustering administration topics about configuring Red Hat high availability add-on with Pacemaker.



- For SLES VMs:
 - Refer to the [SUSE high availability extension documentation](#).
 - In rare cases when using SLES configured with target storage in an AHV cluster, after presenting the disk as iSCSI to the SLES OS, the OS might not start or restart. Log on to a Controller VM in the cluster with SSH and type the following, where `vm_name` is the guest VM.

```
nutanix@cvm$ acli vm.update_boot_device vm_name boot_device_order=kDisk disk_addr=scsi.0
```

FILE SERVER MANAGEMENT

Nutanix Files (Files) lets you share files across user work stations in a centralized and protected location to eliminate the requirement for a third-party file server. The File Server dashboard displays dynamically updated information about the file servers and shares/exports in a cluster. To view the File Server dashboard, select **File Server** from the pull-down list on the left of the main menu. See the [Nutanix Files Guide](#) for more information.

DATA PROTECTION

Nutanix offers data protection solutions for virtual datacenters. Nutanix provides data protection functions at the VM, file, and volume group level, so VMs and data remain safe in a crash-consistent environment.

Note: The following links are to topics in the *Data Protection and Recovery with Prism Element* manual. See the [Leap Administration Guide](#) for additional data protection solutions available from Nutanix.

- Nutanix supports several types of protection strategies including one-to-one or one-to-many replication (see “” ()).
- You can configure asynchronous disaster recovery (see “” ()).
 - The cloud connect feature gives you the option to asynchronously back up data on the Amazon Web Service (AWS) or Microsoft Azure cloud (see “” ()).
- You can also configure synchronous disaster recovery. There are two options:
 - For environments that support metro availability, you can configure a metro availability protection domain (see “” ()).
 - For other environments, you can configure synchronous replication (see “” ()).

HEALTH MONITORING

Nutanix provides a range of status checks to monitor the health of a cluster.

- Summary health status information for VMs, hosts, and disks appears on the home dashboard (see [Home Dashboard](#) on page 53).
- In depth health status information for VMs, hosts, and disks is available through the Health dashboard (see [Health Dashboard](#) on page 239).
- You can customize the frequency of the scheduled health checks and how frequently to run them. (see [Configuring Health Checks](#) on page 244).
- You can run NCC health checks directly from the Prism. (see [Running Checks by Using Web Console](#) on page 247).
- You can collect logs for all the nodes and components. (see [Collecting Logs by Using Web Console](#) on page 248).
- For a description of each available health check, see [Alerts/Health checks](#) on page 340.

Note: If the Cluster Health service status is DOWN for more than 15 minutes, an alert email is sent by the AOS cluster to configured addresses and Nutanix support (if selected). In this case, no alert is generated in the Prism web console. The email is sent once per 24 hours. You can run the NCC check `cluster_services_down_check` to see the service status.

Health Dashboard

The Health dashboard displays dynamically updated health information about VMs, hosts, and disks in the cluster. To view the Health dashboard, select **Health** from the pull-down list on the left of the main menu.

Menu Options

The Health dashboard does not include menu options other than those available from the main menu (see [Main Menu](#) on page 46).

Note: When you first visit the Health dashboard, a tutorial opens that takes you through a guided tour of the health analysis features. Read the message and then click the **Next** button in the text box to continue the tutorial (or click the **X** in the upper right to close the text box and end the tutorial.) You can view this tutorial at any time by selecting the **Health Tutorial** option in the user menu (see [Main Menu](#) on page 46).

Screen Details

The Health dashboard is divided into three columns:

- The left column displays tabs for each entity type (VMs, hosts, disks, storage pools, storage containers, cluster services, and [when configured] protection domains and remote sites). Each tab displays the entity total for the cluster (such as the total number of disks) and the

number in each health state. Clicking a tab expands the displayed information (see following section).

- The middle column displays more detailed information about whatever is selected in the left column.
- The right column displays summary of all the health checks. You also have an option to view individual checks from the **Checks** button (success, warning, failure, or disabled).
- The **Summary** tab provides summarized view of all the health checks according to check status and check type.
- The **Checks** tab provides information about individual checks. Hovering the cursor over an entry displays more information about that health check (see following figure). You can filter the checks by clicking appropriate field type and clicking **Apply**. The checks are categorized as follows.

Filter by Status

Passed, Failed, Warning, Error, Off, or All

Filter by Type

Scheduled, Not Scheduled, Event Triggered, or All

Filter by Entity Type

VM, Host, Disk, Storage Pool, Storage Container, Cluster Service, or All




Figure 157: Hover Information

For example, if you want to see only the failed checks, filter the checks by selecting the **Failed** option. If you click on the specific check, the middle column will provide the detailed history of when the checks failed and what is the percentage of the check failure. If you click the bar, a detailed graph of the pass and fail history is displayed (as shown

below). Hovering the mouse along the graph line displays information about that point in time.





Figure 158: Filter Categorization

Note: For the checks with status as error, follow the similar process as described above to get detailed information about the errors.

You can also search for specific checks by clicking the



icon and then entering a string in the search box.

- The **Actions** tab provides you with an option to manage checks, run checks, and collect logs.




Figure 159: Health Dashboard

Focus and Filtering Options

The Health dashboard allows you to display entity health information through various views. Clicking a left column tab expands that tab to display grouping categories for that entity type (VMs, hosts, or disks). The middle section also expands with additional detail. The **Checks** tab of the right column displays the health checks that are relevant for that entity type.




Figure 160: Health Dashboard Expanded

Clicking on a grouping category displays detailed information about that grouping:

- The left column expands to display a set of grouping and filter options. The selected grouping is highlighted. You can select a different grouping by clicking on that grouping. Each grouping entry lists how many categories are in that grouping, and the middle section displays information about those categories. In the following example, the disks storage tier is selected, and there are two categories (SSD and HDD) in that grouping. By default, all entities (in this example, all disks) are included in the category information. You can narrow the included list by clicking one or more of the filters.
- The middle column displays a field for each category in the selected grouping. Each field provides details about that category. You can see additional information by hovering the cursor over a specific entry. There is a drop-down select list for filtering (same as the grouping filters) and a drop-down sort by list for ordering the information.
- The right column continues to display the health checks that are relevant for that entity type.




Figure 161: Health Dashboard Grouping Screen

The middle column provides two viewing options: a diagram view (see figure above) and a table view (see figure below). The table view provides more detailed information in tabular form. You can sort the entries by clicking a column header.

DISK ID	HYPERVERSOR ID	TIER	DISK USAGE	DISK IOPS	DISK I/O BW	DISK AVG I/O LATENCY	GROUP
00053e42-3a718c3b-0000-00000000444f-39	10.4.45.80	HDD	0%	0	0 Kbps	0 ms	HDD
00053e42-3a718c3b-0000-00000000444f-41	10.4.45.79	HDD	0%	0	0 Kbps	0 ms	HDD
00053e42-3a718c3b-0000-00000000444f-42	10.4.45.79	HDD	0%	0	0 Kbps	0 ms	HDD
00053e42-3a718c3b-0000-00000000444f-46	10.4.45.78	HDD	0%	0	0 Kbps	0 ms	HDD
00053e42-3a718c3b-0000-00000000444f-47	10.4.45.78	HDD	0%	0	0 Kbps	0 ms	HDD
00053e42-3a718c3b-0000-00000000444f-40	10.4.45.80	SSD	0%	0	0 Kbps	0 ms	SSD
00053e42-3a718c3b-0000-00000000444f-43	10.4.45.79	SSD	0%	0	0 Kbps	0 ms	SSD
00053e42-3a718c3b-0000-00000000444f-48	10.4.45.78	SSD	0%	0	0 Kbps	0 ms	SSD

Figure 162: Health Dashboard Table View (middle column)

The middle column also includes watch list information at the top ("Currently watching **xx** entities" or "Currently watching **x** / **xx** total **entity_type**"). The Health dashboard is dynamically adjusted to reflect information about the entities in the current watch list. In this example, all disks are currently selected for the watch list ("Currently watching 18 / 18 total disks"), so the status information (middle column) and relevant health checks (right column) reflect the 18 disks. When you change the watch list to a subset of the current entity type (such as a single disk) or a different entity type (such as hosts), the status information and relevant health checks are customized accordingly for the new watch list.

Configuring Health Checks

A set of health checks are run regularly that provide a range of clusters health indicators. You can specify which checks to run and configure the schedulable checks and other parameters for each health check.

About this task

The cluster health checks cover a range of entities including AOS, hypervisor, and hardware components. A set of checks are enabled by default, but you can run, disable, or reconfigure any of the checks at any time. To reconfigure one or more health checks, do the following:

Procedure

1. In the Health dashboard (see [Health Dashboard](#) on page 239), click **Actions > Manage Checks**.

The Health dashboard redisplay with information about the health checks. If you clicked on a specific health check, that check is highlighted. Either you are prompted to select a health check (first time) or the previously selected health check is highlighted. The following information is displayed:

- The left column lists the health checks with the selected check highlighted. Click any of the entries to select and highlight that health check.
- The middle column describes what this health check does, and it provides the run schedule and history across affected entities (hosts, disks, or VMs).
- The right column describes what failing this health check means (cause, resolution, and impact).




Figure 163: Health Check Configuration Screen

2. To run a particular check, click **Run Check**.
3. To turn off (or turn on) a health check, click the **Turn Check Off** (or **Turn Check On**) link at the top of the middle column and then click the **Yes** button in the dialog box.
4. To change a parameter setting (for those health checks that have configurable parameters), click the **Parameters** link at the top of the middle column, change one or more of the parameter values in the drop-down window, and then click the **Update** button.

This link appears only when the health check includes configurable parameters. The configurable parameters are specific to that health check. For example, the CPU Utilization

health check includes parameters to specify the host average CPU utilization threshold percentage and host peak CPU utilization threshold percentage.



Figure 164: Health Check Parameters

- To change the schedule for running the health check, click the **Schedule** link for the schedulable checks at the top of the middle column, and select an interval from the drop-down list.

You can choose an interval from 1 minute to 48 hours. Each check has a default interval, which varies from as short as once a minute to as long as once a day depending on the health check. The default intervals are optimal in most cases, and changing the interval is not recommended (unless requested to do so by Nutanix customer support).




Figure 165: Health Check Schedule

Configuring NCC Frequency

Perform the following procedure to configure NCC to run automatically after the specified period of time from Prism.

Procedure

- In the Health dashboard, from the **Actions** drop-down menu select **Set NCC Frequency**.
- Select the configuration schedule.

» **Every 4 hours:** Select this option to run the NCC checks at four hours interval.

» **Every Day:** Select this option to run the NCC checks on a daily basis.

Select the time of the day when you want to run the checks from **Start Time** field.

» **Every Week:** Select this option to run the NCC checks on a weekly basis.

Select the day and time of the week when you want to run the checks from the **On** and **Start Time** fields. For example, if you select Sunday and Monday from the **On** field and

select 3:00 p.m. from the **Start Time** field, every Sunday and Monday at 3 p.m. the NCC checks are run automatically.

The Email address that you have configured by using [Configuring Alert Emails](#) on page 335 is also displayed. A report will be sent as an email to all the recipients.

3. Click **Save**.

Running Checks by Using Web Console

You can now run the NCC checks from the **Health** dashboard of the Prism web console. You can select to run all the checks at once, the checks that have failed or displayed some warning, or even specific checks of your choice.

About this task

Note: If you are running checks by using web console, you will not be able to collect the logs at the same time.

Procedure


1. In the Health dashboard, from the **Actions** drop-down menu select **Run Checks**.
2. Select the checks that you want to run for the cluster.
 - a. **All checks:** Select this option to run all the checks at once.
 - b. **Only Failed and Warning Checks:** Select this option to run only the checks that were failed or gave warning during the health check runs.
 - c. **Specific Checks:** Select this option and type the check or checks name in the text box that appears that you want to run.

This field gets auto-populated once you start typing the name of the check. All the checks that you have selected for this run are listed in the **Added Checks** box.



3. Select the **Send the cluster check report in the email** option to receive the report after the cluster check.

To receive the email configuration ensure that you have configured email configuration for alerts. For more information, see [Configuring Alert Emails](#) on page 335.

A dark-themed dialog box titled "Run Checks" with a question mark icon and a close (X) icon in the top right corner. The main content area is light gray and contains the text "Select checks you want to run:". Below this text are three radio button options: "All Checks (249)" (which is selected), "Only Failed And Warning Checks (8)", and "Specific Checks". Below these options is a horizontal line, followed by a checkbox labeled "Send the cluster check report in the email" with a question mark icon to its right. Underneath the checkbox is the text "Recipients: none". At the bottom right of the dialog box are two buttons: "Cancel" and "Run".

Run Checks

Select checks you want to run:

☒ All Checks (249)

☐ Only Failed And Warning Checks (8)

☐ Specific Checks

☐ Send the cluster check report in the email ?

Recipients: none

Cancel Run

Figure 166: Run Health Checks

The status of the run (succeeded or aborted) is available in the **Tasks** dashboard. By default, all the event triggered checks are passed. Also, the **Summary** page of the **Health** dashboard will be updated with the status according to health check runs.

Collecting Logs by Using Web Console

You can collect logs for Controller VMs, file server, hardware, alerts, hypervisor, and for the system. After the task finishes, the log bundle is available for download purpose from the **Tasks** dashboard.

About this task

Log bundle includes logs and configuration information from one or more Controller VMs, configuration information for hypervisors, and information about alerts and so on. To collect the logs and download the log bundle, perform the following procedure.

Note: While this method works, the preferred method for collecting logs is through a CLI tool called **logbay**. See the *Nutanix Cluster Check (NCC) Guide* for information about **logbay**.

Procedure

1. In the **Health** dashboard, from the **Actions** drop-down menu, select **Collect Logs**.
2. In **Node Selection**, click **+ Select Nodes**. Select the nodes for which you want to collect the logs and click **Done**.
3. Click **Next**.

4. In **Log Settings**, select one of the following:

- » **All**. Select this option if you want to collect the logs for all the tags.
- » **Specific (by tags)**. Select this option, click **+ Select Tags** if you want to collect the logs only for the selected tags and then click **Done**.



5. In **Output Preferences**, do the following in the indicated fields.

- 1. **Select Duration.** Select the duration for which you want to collect the logs. You can collect the logs either in hours or days. Click the drop-down list to select the required option.
- 2. **Cluster Date.** Select the date from which you want to start the log collection operation. Click the drop-down list to select either **Before** or **After** to collect logs before or after a selected date.
- 3. **Cluster Time.** Select the time from when you want to start the log collection operation.
- 4. **Select Destination for the collected logs.** Click the drop-down list to select the server where you want the logs to be collected.
 - **Download Locally**
 - **Nutanix Support FTP.** If you select this option, enter the case number in the **Case Number** field.
 - **Nutanix Support SFTP.** If you select this option, enter the case number in the **Case Number** field.
 - **Custom Server.** Enter server name, port, username, password, and archive path if you select this option.
- 5. **Anonymize Output.** Select this option if you want to mask all the sensitive information like the IP addresses.

1 Node Selection

2 Log Settings

3 Output Preferences

Select Duration
Select Duration for the Logs to be collected

10

Hours ▾

Cluster Date

Cluster Time

After ▾

03/24/2019

📅

22:42:30

🕒

Corresponding Browser time : 11:12:30, 25/03/2019

Select Destination for the collected logs
Select Server where the logs can be collected

Download Locally

▾

Logs will be collected on respective nodes

☐ **Anonymize Output**

This masks all the sensitive information like the IP addresses. It is not recommended for effective debuggability.

◀ Back

Cancel

Collect

Figure 167: Collect Logs

6. To start the operation, click **Collect**.
7. After the operation completes, you can download the log bundle for the last two runs and (as needed) add it to a support case as follows:
 - a. Go to the **Task** dashboard (see [Task Status](#) on page 592), find the log bundle task entry, and click the **Succeeded** link for that task (in the Status column) to download the log bundle.

Note: If a pop-up blocker in your browser stops the download, turn off the pop-up blocker and try again.

- b. Log in the support portal, click on the target case in the **360 View** widget on the dashboard (or click the **Create a New Case** button to create a new case), and upload the log bundle to the case (click the **Choose Files** button in the **Attach Files** section to select the file to upload).

VIRTUAL MACHINE MANAGEMENT

Each node in a cluster includes local storage (flash and hard disk), a Controller VM, a hypervisor, and any number of host VMs running on that hypervisor.

- The web console allows you to monitor status of the VMs across the cluster (see [VM Dashboard](#) on page 252).
- In Acropolis managed clusters, the web console also allows you to do the following:
 - Create VMs (see [Creating a VM \(AHV\)](#) on page 268).
 - Manage VMs (see [Managing a VM \(AHV\)](#) on page 275).
 - Enable VM high availability (see [Enabling High Availability for the Cluster](#) on page 296).
 - Configure network connections (see [Network Configuration for VM Interfaces](#) on page 138).
- You can create and manage VMs directly from Prism Element when the hypervisor is ESXi.
 - Create VMs (see [Creating a VM \(ESXi\)](#) on page 283).
 - Manage VMs (see [Managing a VM \(ESXi\)](#) on page 285).



Figure 168: Node Architecture

VM Dashboard

The virtual machine (VM) dashboard displays dynamically updated information about virtual machines in the cluster. To view the VM dashboard, select **VM** from the pull-down list on the left of the main menu.

Menu Options


In addition to the main menu (see [Main Menu](#) on page 46), the VM screen includes a menu bar with the following options:

- View selector. Click the **Overview** button on the left to display the VM dashboard (see [VM Overview View](#) on page 254), or click the **Table** button to display VM information in a tabular form (see [VM Table View](#) on page 255).

- Action buttons. Click the **Create VM** button on the right to create a virtual machine (see [Creating a VM \(AHV\)](#) on page 268). Click the **Network Config** button to configure a network connection (see [Network Configuration for VM Interfaces](#) on page 138).

Note: The action buttons appear only in Acropolis managed clusters.

- CVM filter. In the Table view, the Controller VMs are not listed by default. To include them in the table list, check the **Include Controller VMs** box.
- Page selector. In the Table view, VMs are listed 10 per page. When there are more than 10 VMs, left and right paging arrows appear on the right, along with the total VM count and the VM numbers for the current page.
- Export table content. In the Table view, you can export the table information to a file in either CSV or JSON format by clicking the gear icon



on the right and selecting either **Export CSV** or **Export JSON** from the pull-down menu. (The browser must allow a dialog box for export to work.) Chrome, Internet Explorer, and Firefox download the data into a file; Safari opens the data in the current window.

- Search box. In the Table view, you can search for entries in the table by entering a search string in the box.



Figure 169: VM Dashboard Menu

VM Overview View

The VM Overview view displays VM-specific performance and usage statistics on the left plus the most recent VM-specific alert and event messages on the right. The following figure is a sample view, and the table describes each field in this view. Several fields include a slide bar on the right to view additional information in that field. The displayed information is dynamically updated to remain current.




Figure 170: VM Overview View

Note: See [Understanding Displayed Statistics](#) on page 56 for information about how the statistics are derived.

Table 32: VM Overview View Fields

Name	Description
Hypervisor Summary	Displays the name and version number of the hypervisor.
VM Summary	Displays the total number of VMs in the cluster broken down by on, off, and suspended states.
CPU	Displays the total number of provisioned virtual CPUs and the total amount of reserved CPU capacity in GHz for the VMs.

Name	Description
Memory	Displays the total amount of provisioned and reserved memory in GBs for the VMs.
Top User VMs by Controller IOPS	Displays I/O operations per VM for the 10 most active VMs.
Top User VMs by Controller IO Latency	Displays I/O bandwidth used per VM for the 10 most active VMs. The value is displayed in an appropriate metric (MBps, KBps, and so on) depending on traffic volume.
Top User VMs by Memory Usage	Displays the percentage of reserved memory capacity used per VM for the 10 most active VMs.
Top User VMs by CPU Usage	Displays the percentage of reserved CPU capacity used per VM for the 10 most active VMs.
VM Critical Alerts	Displays the five most recent unresolved VM-specific critical alert messages. Click a message to open the Alert screen at that message. You can also open the Alert screen by clicking the "view all alerts" button at the bottom of the list (see Alerts Dashboard).
VM Warning Alerts	Displays the five most recent unresolved VM-specific warning alert messages. Click a message to open the Alert screen at that message. You can also open the Alert screen by clicking the "view all alerts" button at the bottom of the list.
VM Events	Displays the ten most recent VM-specific event messages. Click a message to open the Event screen at that message. You can also open the Event screen by clicking the "view all events" button at the bottom of the list.

VM Table View

The VM Table view displays information about each VM in a tabular form. The displayed information is dynamically updated to remain current. In Acropolis managed clusters, you can both monitor and manage VMs through the VM Table view.

Table View Fields

The VM Table view is divided into two sections:

- The top section is a table. Each row represents a single VM and includes basic information about that VM. Click a column header to order the rows by that column value (alphabetically or numerically as appropriate).
- The bottom **Summary** section provides additional information. It includes a summary or details column on the left and a set of tabs on the right. The details column and tab content varies depending on what has been selected.

The following table describes each field in the table portion of the view. The details portion and tab contents are described in the subsequent sections.

Note: See [Understanding Displayed Statistics](#) on page 56 for information about how the statistics are derived.






Figure 171: VM Table View

Table 33: VM Table View: Table Fields

Parameter	Description	Values
VM Name	Displays the name given the VM.	(VM name)
Host	Displays the name of the host	(Host name)
IP Addresses	Displays the IP address assigned to the VM.	(IP address)
Cores	Displays the number of CPU cores being used by this VM.	(number)
Memory Capacity	Displays the total amount of memory available to this VM.	xxx [MB GB]
Storage	Displays the used capacity in relation to the total capacity (for example 1.9 GiB/ 5 GiB).	xxx / xxx [MiB GiB]
CPU Usage	Displays the percentage of allocated CPU capacity currently being used by this VM.	0 - 100%
Memory Usage	Displays the percentage of allocated memory capacity currently being used by this VM.	0 -100%

Parameter	Description	Values
[Controller] Read IOPS	Displays read I/O operations per second (IOPS) for this VM.	(number)
Note: In this and the following three fields, the column name includes the word Controller if the information comes from the Controller VM instead of the hypervisor. For ESXi, the information comes from the hypervisor; for Hyper-V and AHV, the information comes from the Controller VM.		
[Controller] Write IOPS	Displays write I/O operations per second for this VM.	(number)
[Controller] IO Bandwidth	Displays I/O bandwidth used per second for this VM.	xxx [MBps KBps]
[Controller] Avg IO Latency	Displays the average I/O latency for this VM.	xxx [ms]
Backup and Recovery Capable	Indicates (Yes or No) whether the VM can be protected (create backup snapshots) and recovered if needed. When the value is No, click the question mark icon for an explanation.	[Yes No]
Flash Mode	Displays whether flash mode feature is enabled or not for the VM	[Yes No]

VM Detail Information

When a VM is selected in the table, information about that VM appears in the lower part of the screen.

- **Summary:** **vm_name** appears below the table and **VM Details** fields appear in the lower left column. The following table describes the fields in this column.



- For VMs in Acropolis managed clusters, action links appear on the right of the **Summary:** `vm_name` line (see [Managing a VM \(AHV\)](#) on page 275 for more information about these actions):

- Click **Manage NGT** to enable and mount Nutanix guest tools for this VM.
- Click the **Launch Console** link to open a console window for this VM.
- Click the **Power on** (or **Power Off Actions**) link to start (or shut down) this VM.
- Click the **Take Snapshot** link to create a backup snapshot on demand.
- Click the **Migrate** link to migrate the VM onto a different host.
- Click the **Pause** (or **Resume**) link to pause (or resume) this VM.

Note: VM pause and resume feature is not supported on AHV.

- Click the **Clone** link to clone a copy of the VM.
- Click the **Update** link to update the VM configuration.
- Click the **Delete** link to delete the VM. (A VM must be powered off before it can be deleted.)

Note: If the VM is a Prism Central VM, all Actions except **Launch Console**, **Take Snapshot**, and **Migrate** are disabled as a protective measure.

- A set of tabs appear to the right of the details section that display information about the selected VM. The set of tabs varies depending on whether the VM is an Acropolis managed VM or not. The following sections describe each tab.

- Standard VM tabs: **VM Performance**, **Virtual Disks**, **VM NICs**, **VM Alerts**, **VM Events**, **I/O Metrics**, and **Console**.
- Acropolis managed VM tabs: **VM Performance**, **Virtual Disks**, **VM NICs**, **VM Snapshots**, **VM Tasks**, **I/O Metrics**, and **Console**.



Figure 172: VM Table View: VM Details

Table 34: VM Detail Fields

Parameter	Description	Values
Name	Displays the name given the VM.	(VM name)

Parameter	Description	Values
Host	Displays the host name on which this VM is running.	(IP address)
Host IP	Displays the host IP address for this VM.	(IP address)
Guest OS	Displays the operating system running on this VM, such as Windows 7 or Ubuntu Linux. (This information is not available when running AHV.)	(operating system name)
Memory	Displays the amount of memory available to this VM.	xxx [MB GB]
Reserved Memory	Displays the amount of memory reserved for this VM (by the hypervisor).	xxx [MB GB]
Assigned Memory (Hyper-V only)	Displays the amount of dynamic memory currently assigned to the VM by the hypervisor.	xxx [MB GB]
Cores	Displays the number of CPU cores being used by this VM.	(number)
Reserved CPU	Displays the amount of CPU power reserved for this VM (by the hypervisor).	xxx [GHz]
Disk Capacity	Displays the total disk capacity available to this VM.	xxx [GB TB]
Network Adapters	Displays the number of network adapters available to this VM.	(# of adapter ports)
IP Addresses	Displays the IP address assigned to the VM.	(IP address)
Storage Container	Displays the name of the storage container in which the VM resides.	(storage container name)
Virtual Disks	Displays the number of virtual disks in the VM.	(number)
NGT Enabled	Displays whether NGT is enabled or not for the VM.	[Yes No]
NGT Mounted	Displays whether NGT is mounted or not for the VM.	[Yes No]
GPU Configuration	(AHV only) Comma-separated list of GPUs configured for the VM. GPU information includes the model name and a count in parentheses if multiple GPUs of the same type are configured for the VM. If the firmware on the GPU is in compute mode, the string compute is appended to the model name. The field is hidden if no GPUs are configured or if the hypervisor is not AHV.	(list of GPUs)
GPUs in Use	(AHV only) Number of GPUs in use by a VM. The field is hidden if no GPUs are configured or if the hypervisor is not AHV.	(number)

Parameter	Description	Values
VMware Guest Tools Mounted	Displays whether VMware guest tools are mounted or not on the VM	[Yes No]
VMware Guest Tools Running Status	Displays whether VMware guest tools are running or not on the VM.	[Yes No]

Cluster Summary Information

When a VM is not selected in the table (or when the word **Summary** is clicked), summary information across all VMs in the cluster appears in the lower part of the screen.

- The **VM Summary** fields appear in the lower left column. The following table describes the fields in this column.
- Three tabs appear that display cluster-wide information (see following sections for details about each tab): **Performance Summary**, **All VM Alerts**, **All VM Events**.




Figure 173: VM Table View: VM Summary

Table 35: VM Summary Fields

Parameter	Description	Values
Total VMs	Displays the total number of VMs in the cluster.	(number)
VM State	Displays the number of powered on, powered off, and suspended VMs in the cluster.	[number] powered on, powered off, suspended
Total Provisioned vCPU	Displays the total number of provisioned virtual CPUs in the cluster.	(number)
Total Reserved CPU	Displays the total amount of CPU power reserved for the VMs (by the hypervisor).	xxx [GHz]
Total Provisioned Memory	Displays the total amount of memory provisioned for all VMs.	xxx [GB]
Total Reserved Memory	Displays the total amount of memory reserved for all VMs (by the hypervisor).	xxx [GB]

Performance Tab

The Performance tab displays graphs of performance metrics. The tab label varies depending on what is selected in the table:

- **Performance Summary** (no VM selected). Displays resource performance statistics (CPU, memory, and I/O) across all VMs in the cluster.
- **VM Performance** (VM selected). Displays resource performance statistics (CPU, memory, and I/O) for the selected VM.

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. For more in depth analysis, you can add a monitor to the analysis page by clicking the blue link in the upper right of the graph (see [Analysis Dashboard](#) on page 317). The Performance tab includes the following graphs:

- **[Cluster-wide] CPU Usage**: Displays the percentage of CPU capacity currently being used (0 - 100%) across all VMs or for the selected VM.
- **[Cluster-wide] Memory Usage**: Displays the percentage of memory capacity currently being used (0 - 100%) across all VMs or for the selected VM. (This field does not appear when the hypervisor is Hyper-V.)
- **[Cluster-wide] {Hypervisor|Controller} IOPS**: Displays I/O operations per second (IOPS) across all VMs or for the selected VM.

Note: In this and the following two fields, the field name is either Controller or Hypervisor to indicate where the information comes from. For ESXi, the information comes from the hypervisor; for Hyper-V and AHV, the information comes from the Controller VM.

- **[Cluster-wide] {Hypervisor|Controller} I/O Bandwidth**: Displays I/O bandwidth used per second (MBps or KBps) across all VMs or for the selected VM.
- **[Cluster-wide] {Hypervisor|Controller} Avg I/O Latency**: Displays the average I/O latency (in milliseconds) across all VMs or for the selected VM.




Figure 174: VM Table View: Performance Tab

Virtual Disks Tab

The Virtual Disks tab displays information in tabular form about the virtual disks in a selected VM. (This tab appears only when a VM is selected.) Each line represents a virtual disk, and the following information is displayed for each disk organized under "Default" and "Additional Stats" subtabs.

Default Tab:

Note: Clicking on a virtual disk (line) displays subtabs for total, read, and write IOPS, I/O bandwidth, and I/O latency performance graphs for the virtual disk (see the "Performance Tab" section for more information about the graphs).

- **Virtual Disk.** Displays the virtual disk identification number.
- **Total Capacity.** Displays the total capacity of the virtual disk (in GiBs).
- **Physical Usage.** Displays the used space of the virtual disks (in GiBs).
- **Read IOPS.** Displays the read IOPS for the virtual disk.
- **Read BW.** Displays the bandwidth used by the virtual disk for read operations.
- **Read Latency.** Displays the average I/O latency for read requests to this virtual disk.
- **Write IOPS.** Displays the write IOPS for the virtual disk.
- **Write BW.** Displays the bandwidth used by the virtual disk for write operations.
- **Write Latency.** Displays the average I/O latency for write requests to this virtual disk.
- **Flash Mode.** Displays whether flash mode is enabled for the virtual disk or not.

Additional Stats Tab:

- **Total IOPS.** Displays the total (both read and write) I/O operations per second (IOPS) for the virtual disk.
- **Random IO.** Displays the percentage of I/O that is random (not sequential).
- **Read Source Cache.** Displays the amount of cache data accessed for read requests.
- **Read Source SSD.** Displays the amount of SSD data accessed for read requests.
- **Read Source HDD.** Displays the amount of HDD data accessed for read requests.
- **Read Working Size Set.** Displays the amount of data actively being read by applications in the VM that are using this virtual disk.
- **Write Working Size Set.** Displays the amount of data actively being written by applications in the VM that are using this virtual disk.
- **Union Working Size Set.** Displays the total amount of data used by the VM for either reads or writes.

Virtual Disks										
Default										
VIRTUAL_DISK	TOTAL_CAPACITY	PHYSICAL_USAGE	READ-IOPS	READ-BW	READ-LATENCY	WRITE-IOPS	WRITE-BW	WRITE-LATENCY	FLASH-MODE	
default_vdi_1	10 GiB	4.14 GiB	0	0 KiB/s	0 ms	0	0 KiB/s	0 ms	Disabled	
default_vdi_2	10 GiB	-	0	0 KiB/s	0 ms	0	0 KiB/s	0 ms	Disabled	

Figure 175: Virtual Disks Tab

VM NICs Tab

The VM NICs tab displays information in tabular form about the virtual NICs in a selected VM. (This tab appears only when a VM is selected.) Each line represents a virtual NIC, and the following information is displayed for each NIC:

- **Virtual NIC.** Displays the virtual NIC identification number.
- **Adapter Type.** Displays the adaptor type defined for the virtual NIC.
- **MAC Address.** Displays the virtual NIC MAC address
- **IPv4 Addresses.** Displays the virtual NIC IPv4 address(es).
- **IPv6 Addresses.** Displays the virtual NIC IPv6 address(es).
- **Received Packets.** Displays the number of packets received by the virtual NIC.
- **Transmitted Packets.** Displays the number of packets transmitted by the virtual NIC.
- **Dropped Rx Packets.** Displays the number of received packets dropped by the virtual NIC.
- **Dropped Tx Packets.** Displays the number of transmitted packets dropped by the virtual NIC.

VM Performance		Virtual Disks		VM NICs		VM Events		VM Events	
Virtual NIC	Adapter Type	MAC Address	IPv4 Address	IPv6 Address	Rx Pkts	Tx Pkts	Dropped Rx Pkts	Dropped Tx Pkts	
Network adapter 1	vmxnet3	00:50:56:b0:60:c0	10.4.34.49/22	fe80::25c:56f:a6b:80c46a	332	80	251	0	
Network adapter 2	vmxnet3	00:50:56:b0:70:04		fe80::25c:56f:a6b:70a164	-	-	-	-	

Figure 176: VM NICs Tab

When you click a virtual NIC entry, three more tabs appear below the list of virtual NICs. Clicking the **Virtual NICs Stats** tab displays the following statistics for that virtual NIC:

- **Total Packets Received.** Displays a monitor of the total packets received (in KB or MB) over time. Place the cursor anywhere on the line to see the value for that point in time. (This applies to all the monitors on this tab.)
- **Total Packets Transmitted.** Displays a monitor of the transmitted data rate.
- **Dropped Packets Received.** Displays a monitor of received packets that were dropped.
- **Dropped Packets Transmitted.** Displays a monitor of transmitted packets that were dropped.
- **Error Packets Received.** Displays a monitor for error packets received.



Figure 177: VM NICs Tab: Virtual NIC Stats

Clicking the **Host NICs Stats** tab displays the following statistics for each host NIC (one per line) that is used by the selected virtual NIC to send the traffic:

- **Host NIC.** Displays the host NIC name.
- **Speed (in KBps).** Displays the host NIC transmission speed.
- **MAC Address.** Displays the host NIC MAC address.
- **Received Packets.** Displays the number of packets received by the host NIC.
- **Transmitted Packets.** Displays the number of packets transmitted by the host NIC.
- **Dropped Rx Packets.** Displays the number of received packets dropped by the host NIC.
- **Dropped Tx Packets.** Displays the number of transmitted packets dropped by the host NIC.
- **Rx Packet Errors.** Displays the number of error packets received by the host NIC.
- **Tx Packet Errors.** Displays the number of error packets transmitted by the host NIC.

Virtual NICs Stats		Host NICs Stats				Physical Switch Interface Stats			
HOST NIC	SPEED (in KBPS)	MAC ADDRESS	RX PKTS	TX PKTS	DROPPED RX PKTS	DROPPED TX PKTS	RX PKT ERRORS	TX PKT ERRORS	
vmnic0	1250000	00:25:90:e7:8e:0c	79344	140394	0	0	0	0	
vmnic1		00:25:90:e7:8e:0f	0	0	0	0	0	0	
vmnic3		00:25:90:e7:8e:0d	0	0	0	0	0	0	
vmnic2	125000	00:25:90:e7:8e:0b	334	0	0	0	0	0	

Figure 178: Virtual NICs Tab: Host NIC Stats

Clicking the **Physical Switch Interface Stats** tab displays the following statistics for each physical switch interface (one per line) used by the selected virtual NIC to send the traffic:

- **Physical Switch Interface.** Displays the switch interface name.
- **Switch ID.** Displays the switch interface ID value.
- **Index.** Displays the switch interface index number.
- **MTU (in Bytes).** Displays the size in bytes of the largest protocol data unit (maximum transmission unit) that the layer can pass onwards.
- **MAC Address.** Displays the interface MAC address.
- **Unicast Rx Pkts.** Displays the number of unicast packets received.
- **Unicast Tx Pkts.** Displays the number of unicast packets transmitted.
- **Error Rx Pkts.** Displays the number of received packets with an error.
- **Error Tx Pkts.** Displays the number of transmitted packets with an error.
- **Discard Rx Pkts.** Displays the number of received packets that were discarded.
- **Discard Tx Pkts.** Displays the number of transmitted packets that were discarded.

Virtual NICs Stats		Host NICs Stats				Physical Switch Interface Stats				
PHYSICAL SWITCH INTERFACE	SWITCH ID	INDEX	MTU (in BYTES)	MAC ADDRESS	UNICAST RX PKTS	UNICAST TX PKTS	ERROR RX PKTS	ERROR TX PKTS	DISCARD RX PKTS	DISCARD TX PKTS
Ethernet2/3	dit-r3-sw1	22003	9214	01:72:54:00:e4	23463	24199	0	0	0	0

Figure 179: Virtual NICs Tab: Physical Switch Interface Stats

VM Alerts Tab

The VM Alerts tab displays the unresolved alert messages about all VMs or the selected VM in the same form as the Alerts page (see [Alerts Summary View](#) on page 327). Click the **Unresolved X** button in the filter field to also display resolved alerts.


VM Events Tab

The VM Events tab displays the unacknowledged event messages about all VMs or the selected VM in the same form as the Events page (see [Events Summary View](#) on page 333). Click the **Include Acknowledged** button to also display acknowledged events.

VM Snapshots Tab (Acropolis only)

The VM Snapshots tab displays information in tabular form about backup snapshots of the VM. (This tab appears only when a VM is selected.) Each line represents a snapshot, and the following information is displayed for each snapshot:

- **Create Time.** Displays the time the backup snapshot was created (completed).
- **Name.** Displays a name for the backup if one was created.
- **Actions.** Displays four action links:
 - Click the **Details** link to open a window that displays the VM configuration (see [Creating a VM \(AHV\)](#) on page 268) plus a creation time stamp field.



The screenshot shows a window titled "Snapshot Details" with a dark header bar containing a question mark icon and a close button. The window is divided into sections: "General Configuration" with a "NAME" field containing "NTNX-lb-rc3-3"; "Compute Details" with "CREATE TIME" (12/16/16, 11:46:35 AM), "VCPUS" (4), "NUMBER OF CORES PER VCPU" (empty), and "MEMORY" (12 GiB); and "Disks" with a table header showing "TYPE", "ADDRESS", and "PARAMETERS".

Figure 180: Snapshot Details Window

- Click the **Clone** link to clone a VM from the snapshot (see [Managing a VM \(AHV\)](#) on page 275).
- Click the **Restore** link to restore the VM from the snapshot. This restores the VM back to the state of the selected snapshot.
- Click the **Delete** link to delete the snapshot.

VM Performance	Virtual Disks	VM NICs	VM Snapshots	VM Tasks	I/O Metrics	Console
1 Snapshot						
search in table						
CREATE TIME	NAME		ACTIONS			
12/01/16, 11:48:25 AM	100		Details - Clone - Revert - Delete			

Figure 181: VM Snapshots Tab

VM Tasks Tab (Acropolis only)

The VM Tasks (selected VM) or All VM Tasks (cluster-wide) tab displays a log in tabular form of the running and completed tasks for the selected VM or across the cluster. Each line represents a task, and the following information is displayed for each task:

- **Operation.** Describes the type of task.
- **Entities.** Lists the affected VM and node.
- **Percent Complete.** Displays the run status (0-100%) of the task.
- **Progress Status.** Displays the completion status of the task (succeeded, in progress, failed).
- **Create Time.** Displays when the task began.
- **Duration.** Displays how long the task took to complete.

VM Performance	Virtual Disks	VM NICs	VM Snapshots	VM Tasks	I/O Metrics	Console
10 VM Tasks						
search in table						
OPERATION	ENTITIES		PERCENT COMPLETE	PROGRESS STATUS	CREATE TIME	DURATION
Create Snapshot	VM: HTNKH-BL-c3-3	Cluster: m1	100	Succeeded	12/01/16, 11:48:36 AM	Under 1 second
VerifyPowerState	VM: HTNKH-BL-c3-3	Node: HTNKH-ES00005320000-B	100	Succeeded	12/01/16, 3:57:36 PM	4 seconds
Create Virtual Disk	VM: HTNKH-BL-c3-3		100	Succeeded	12/01/16, 3:57:39 PM	Under 1 second
Create Virtual Disk	VM: HTNKH-BL-c3-3		100	Succeeded	12/01/16, 3:57:38 PM	Under 1 second
Create Virtual Disk	VM: HTNKH-BL-c3-3		100	Succeeded	12/01/16, 3:57:38 PM	10 seconds
Create Virtual Disk	VM: HTNKH-BL-c3-3		100	Succeeded	12/01/16, 3:57:38 PM	Under 1 second
VerifyPowerState	VM: HTNKH-BL-c3-3	Node: HTNKH-ES00005320000-B	100	Succeeded	12/01/16, 3:57:43 PM	2 seconds
VerifyPowerState	VM: HTNKH-BL-c3-3	Node: HTNKH-ES00005320000-B	100	Succeeded	12/01/16, 3:58:39 PM	2 seconds
Update	VM: HTNKH-BL-c3-3		100	Succeeded	12/01/16, 3:58:28 PM	Under 1 second
Create	VM: HTNKH-BL-c3-3		100	Succeeded	12/01/16, 3:58:22 PM	Under 1 second

Figure 182: [All] VM Tasks Tab

I/O Metrics

The I/O Metrics tab displays information about different I/O metrics for the VM (latency and performance distribution).

- Average I/O latency for the VM. Displays a graph of average I/O latency (in milliseconds) for reads and writes over a period of time. If you hover the cursor over the graph, the read and write latency at that particular time is displayed.
- Performance distribution. Displays a bar chart for the performance distribution of read and write size (in bytes) and read and write latency (in milliseconds). This widget also provides information on the read source (HDD, SSD, or DRAM), random vs sequential read and writes in a pie chart format. This information changes according to the time that you select in the **Avg I/O Latency** widget.



Figure 183: I/O Metrics Tab

Console Tab

The Console tab displays a live console window. (This tab appears only when a VM is selected.) In addition to entering commands in the console window, you can invoke several options from this tab:

- Click the language (left-most) button and select the desired language from the pull-down list to set the language key mapping for the console keyboard.
- Click the **Send Ctrl+Alt+Del** button to send a Ctrl+Alt+Delete key signal. This is the same as pressing Ctrl+Alt+Delete from the console keyboard.
- Click the **Take Screenshot** button to take a screen shot of the console display that you can save for future reference.
- Click the **New Window** button to open a new console window. This is the same as clicking the **Launch Console** link on the Summary line.




Figure 184: Console Tab

VM Management

You can create and manage VMs directly from Prism Element when the hypervisor is either ESXi or AHV. Following topics provide more information on creating and managing VM configuration on ESXi and AHV.

- ESXi
 - To create a VM, see [Creating a VM \(ESXi\)](#) on page 283.
 - To manage guest tools, launch VM console, power actions, clone, update, or delete operations, see [Managing a VM \(ESXi\)](#) on page 285.
- AHV
 - To create a VM, see [Creating a VM \(AHV\)](#) on page 268.
 - To manage guest tools, launch console, power actions, take snapshot, migrate, power operations, clone, update, or delete operations, see [Managing a VM \(AHV\)](#) on page 275.

Creating a VM (AHV)

In Acropolis managed clusters, you can create a new virtual machine (VM) through the web console.

About this task

When creating a VM, you can configure all of its components, such as number of vCPUs and memory, but you cannot attach a volume group to the VM. Attaching a volume group is possible only when you are modifying a VM.

To create a VM, do the following:



Procedure

1. In the VM dashboard , click the **Create VM** button.

Note: This option does not appear in clusters that do not support this feature.

The Create VM dialog box appears.

Figure 185 shows the 'Create VM' dialog box. It contains the following fields and options:

- General Configuration**
 - Name:** Text input field with placeholder 'Name'.
 - Description:** Text input field with placeholder 'Optional'.
 - Timezone:** Dropdown menu showing '(UTC) UTC'.
 - ☐ Use this VM as an agent VM
- Compute Details**
 - vCPU(s):** Text input field with placeholder 'Value'.
 - Number Of Cores Per vCPU:** Text input field with value '1'.

Buttons: Cancel, Save

Figure 185: Create VM Dialog Box

2. Do the following in the indicated fields:
 - a. **Name:** Enter a name for the VM.
 - b. **Description** (optional): Enter a description for the VM.
 - c. **Timezone:** Select the timezone that you want the VM to use. If you are creating a Linux VM, select **(UTC) UTC**.

Note:

The RTC of Linux VMs must be in UTC, so select the UTC timezone if you are creating a Linux VM.

Windows VMs preserve the RTC in the local timezone, so set up the Windows VM with the hardware clock pointing to the desired timezone.

- 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 the VMs are powered on on the host) and must be powered off after the rest of the VMs are powered off (for example, during maintenance mode operations). Agent VMs are never migrated to any other host in the cluster. If an HA event occurs or the host is put in maintenance mode,

agent VMs are powered off and are powered on on the same host once that host comes back to a normal state.

If an agent VM is powered off, you can manually start that agent VM on another host and the agent VM now permanently resides on the new host. The agent VM is never migrated back to the original host. Note that you cannot migrate an agent VM to another host while the agent VM is powered on.

- 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 MiBs) to allocate to this VM.
3. (For GPU-enabled AHV clusters only) To configure GPU access, click **Add GPU** in the **Graphics** section, and then do the following in the **Add GPU** dialog box:

NAME	VIRTUAL SLICE	FRAMEBUFFER	VMS ASSIGNED	DETAILS
☐ GRID M10-0B	1/16 GPU	512 MiB	0	?
☐ GRID M10-1B	1/8 GPU	1 GiB	0	?
☐ GRID M60-0B	1/16 GPU	512 MiB	0	?
☐ GRID M60-1B	1/8 GPU	1 GiB	0	?

Figure 186: Add GPU Dialog Box

- a. To configure GPU pass-through, in **GPU Mode**, click **Passthrough**, select the GPU that you want to allocate, and then click **Add**.
If you want to allocate additional GPUs to the VM, repeat the procedure as many times as you need to. Make sure that all the allocated pass-through GPUs are on the same host. If all specified GPUs of the type that you want to allocate are in use, you can proceed to allocate the GPU to the VM, but you cannot power on the VM until a VM that is using the specified GPU type is powered off.
- b. To configure virtual GPU access, in **GPU Mode**, click **virtual GPU**, select a GRID license, and then select a virtual GPU profile from the list.
You can assign only one virtual GPU to a VM. A vGPU is assigned to the VM only if a vGPU is available when the VM is starting up.

4. To attach a disk to the VM, click the **Add New Disk** button.
The Add Disks dialog box appears.


The image shows a dialog box titled "Add Disk" with a question mark icon and a close button (X). It contains several dropdown menus and a text input field. The "TYPE" dropdown is set to "DISK". The "OPERATION" dropdown is set to "Allocate on Storage Container". The "BUS TYPE" dropdown is set to "SCSI". The "STORAGE CONTAINER" dropdown is set to "default-container-71260". Below these is a "SIZE (GiB)" text input field with a warning message: "Please note that reducing the size of a disk is not allowed." At the bottom right are "Cancel" and "Add" buttons.

Figure 187: Add Disk Dialog Box

Do the following in the indicated fields:

- a. **Type:** Select the type of storage device, **DISK** or **CD-ROM**, from the pull-down list.
The following fields and options vary depending on whether you choose **DISK** or **CD-ROM**.
- b. **Operation:** Specify the device contents from the pull-down list.
 - Select **Clone from ADSF file** to copy any file from the cluster that can be used as an image onto the disk.
 - Select **Empty CD-ROM** to create a blank CD-ROM device. (This option appears only when **CD-ROM** is selected in the previous field.) A CD-ROM device is needed when you intend to provide a system image from CD-ROM.
 - Select **Allocate on Storage Container** to allocate space without specifying an image. (This option appears only when **DISK** is selected in the previous field.) Selecting this option means you are allocating space only. You have to provide a system image later from a CD-ROM or other source.
 - Select **Clone from Image Service** to copy an image that you have imported by using image service feature onto the disk. For more information on the Image Service feature, see the *Configuring Images* section of *Migration Guide*.
- c. **Bus Type:** Select the bus type from the pull-down list. The choices are **IDE**, **SCSI**, or **SATA**.
- d. **ADSF Path:** Enter the path to the desired system image.
This field appears only when **Clone from ADSF file** is selected. It specifies the image to copy. Enter the path name as `/storage_container_name/iso_name.iso`. For example to clone an image from `myos.iso` in a storage container named `crt1`, enter `/crt1/myos.iso`. When a user types the storage container name (`/storage_container_name/`), a list appears of the

ISO files in that storage container (assuming one or more ISO files had previously been copied to that storage container).

- e. **Image:** Select the image that you have created by using the image service feature. This field appears only when **Clone from Image Service** is selected. It specifies the image to copy.
 - f. **Storage Container:** Select the storage container to use from the pull-down list. This field appears only when **Allocate on Storage Container** is selected. The list includes all storage containers created for this cluster.
 - g. **Size:** Enter the disk size in GiBs.
 - h. When all the field entries are correct, click the **Add** button to attach the disk to the VM and return to the Create VM dialog box.
 - i. Repeat this step to attach additional devices to the VM.
5. Select one of the following firmware to boot the VM.
- » **Legacy BIOS:** Select legacy BIOS to boot the VM with legacy BIOS firmware.
 - » **UEFI:** 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 *UEFI Support for VM* section in the *AHV Administration Guide*.

6. To create a network interface for the VM, click the **Add New NIC** button. The Create NIC dialog box appears.


The image shows a 'Create NIC' dialog box with a dark header bar containing a question mark and a close button. Below the header, there are several fields: 'VLAN NAME' is a pull-down menu showing 'Ackbar'; 'VLAN ID' is a read-only field showing 'vlan.124'; 'VLAN UUID' is a read-only field showing 'f530d179-16a1-43f0-bbfa-01fd64576a7c'; 'NETWORK ADDRESS / PREFIX' is a read-only field showing '10.5.124.0/22'; and 'IP ADDRESS' is an empty text input field with a help icon. At the bottom right, there are 'Cancel' and 'Add' buttons.

Figure 188: Create NIC Dialog Box

Do the following in the indicated fields:

- a. **VLAN Name:** Select the target virtual LAN from the pull-down list.
The list includes all defined networks (see [Network Configuration for VM Interfaces](#) on page 138).
 - b. **VLAN ID:** This is a read-only field that displays the VLAN ID.
 - c. **VLAN UUID:** This is a read-only field that displays the VLAN UUID.
 - d. **Network Address/Prefix:** This is a read-only field that displays the network IP address and prefix.
 - e. **IP Address:** Enter an IP address for the VLAN.
This field appears only if the NIC is placed in a managed network. Entering an IP address in this field is optional when the network configuration provides an IP pool. If the field is left blank, the NIC is assigned an IP address from the pool.
 - f. When all the field entries are correct, click the **Add** button to create a network interface for the VM and return to the Create VM dialog box.
 - g. Repeat this step to create additional network interfaces for the VM.
7. To configure affinity policy for this VM, click **Set Affinity**.
 - a. Select the host or hosts on which you want configure the affinity for this VM.
 - b. Click **Save**.
The selected host or hosts are listed. This configuration is permanent. The VM will not be moved from this host or hosts even in case of HA event and will take effect once the VM starts.

8. To customize the VM by using Cloud-init (for Linux VMs) or Sysprep (for Windows VMs), select the **Custom Script** check box. Fields required for configuring Cloud-init and Sysprep, such as options for specifying a configuration script or answer file and text boxes for specifying paths to required files, appear below the check box.

☒ **Custom Script**

Provide a Cloudinit or Sysprep script to customize this VM.

☒ **ADSF path**

Start typing for suggestions

☐ **Upload a file**

Choose File No file chosen

☐ **Type or paste script**

FILES TO COPY
Specify external files to copy inside of the guest VM.

Source File ADSF Path Destination Path in VM +

Cancel Save

Figure 189: Create VM Dialog Box (custom script fields)

9. To specify a user data file (Linux VMs) or answer file (Windows VMs) for unattended provisioning, do one of the following:
- » If you uploaded the file to a storage container on the cluster, click **ADSF path**, and then enter the path to the file.
Enter the ADSF prefix (**adsf://**) followed by the absolute path to the file. For example, if the user data is in **/home/my_dir/cloud.cfg**, enter **adsf:///home/my_dir/cloud.cfg**. Note the use of three slashes.
 - » If the file is available on your local computer, click **Upload a file**, click **Choose File**, and then upload the file.
 - » If you want to create or paste the contents of the file, click **Type or paste script**, and then use the text box that is provided.

10. To copy one or more files to a location on the VM (Linux VMs) or to a location in the ISO file (Windows VMs) during initialization, do the following:
 - a. In **Source File ADSF Path**, enter the ADSF prefix (**adsf://**) followed by the absolute path to the file.
For example, if the file is `/home/my_dir/myfile.txt`, enter **adsf:///home/my_dir/myfile.txt**. Note the use of three slashes.
 - b. In **Destination Path in VM**, enter the absolute path to the target directory and the file name. You do not have to enter a path prefix in this field.
For example, if the source file entry is **adsf:///home/my_dir/myfile.txt** then the entry for the **Destination Path in VM** should be `/<directory_name>/copy_desititation` i.e. `/mnt/myfile.txt`.
 - c. To add another file or directory, click the button beside the destination path field. In the new row that appears, specify the source and target details.
11. When all the field entries are correct, click the **Save** button to create the VM and close the Create VM dialog box.
The new VM appears in the VM table view (see [VM Table View](#) on page 255).

Managing a VM (AHV)

You can use the web console to manage virtual machines (VMs) in Acropolis managed clusters.

About this task

After creating a VM (see [Creating a VM \(AHV\)](#) on page 268), you can use the web console to start or shut down the VM, launch a console window, update the VM configuration, take a snapshot, attach a volume group, migrate the VM, clone the VM, or delete the VM.

Note: Your available options depend on the VM status, type, and permissions. Unavailable options are grayed out.

To accomplish one or more of these tasks, do the following:

Procedure

1. In the VM dashboard (see [VM Dashboard](#) on page 252), click the **Table** view.



- The Summary line (middle of screen) displays the VM name with a set of relevant action links on the right. The possible actions are **Manage Guest Tools**, **Launch Console**, **Power on** (or **Power off**), **Take Snapshot**, **Migrate**, **Clone**, **Update**, and **Delete**.

The following steps describe how to perform each action.

[illegible]

Figure 190: VM Action Links

- You can also enable NGT applications (self-service restore, Volume Snapshot Service and application-consistent snapshots) also as part of manage guest tools.

- b. Select **Mount Nutanix Guest Tools** to mount NGT on the selected VM.

The VM is registered with the NGT service. NGT is enabled and mounted on the selected virtual machine. A CD with volume label NUTANIX_TOOLS gets attached to the VM.

- The Self-Service Restore feature is enabled of the VM. The guest VM administrator can restore the desired file or files from the VM. For more information on self-service restore feature, see Self-Service Restore in the guide titled *Data Protection and Recovery with Prism Element*.

- After this feature is enabled, Nutanix native in-guest VmQuiesced Snapshot Service (VSS) agent is used 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 Nutanix platform in a hypervisor-agnostic manner. For more information, see [Implementation Guidelines for](#)

Asynchronous Disaster Recovery in the guide titled *Data Protection and Recovery with Prism Element*.

e. Click **Submit**.

The VM is registered with the NGT service. NGT is enabled and mounted on the selected virtual machine. A CD with volume label NUTANIX_TOOLS gets attached to the VM.

Note:

- If you clone a VM, by default NGT is not enabled on the cloned VM. If the cloned VM is powered off, enable NGT from the UI and power on the VM. If cloned VM is powered on, enable NGT from the UI and restart the nutanix guest agent service.
- If you want to enable NGT on multiple VMs simultaneously, see [Enabling NGT and Mounting the NGT Installer Simultaneously on Multiple Cloned VMs](#) on page 304.

If you eject the CD, you can mount the CD back again by logging into the Controller VM and running the following nCLI command.

```
ncli> ngt mount vm-id=virtual_machine_id
```

For example, to mount the NGT on the VM with VM_ID=00051a34-066f-72ed-0000-000000005400::38dc7bf2-a345-4e52-9af6-c1601e759987, type the following command.

```
ncli> ngt mount vm-id=00051a34-066f-72ed-0000-000000005400::38dc7bf2-a345-4e52-9af6-c1601e759987
```

CAUTION: In AOS 4.6, for the powered-on Linux VMs on AHV, ensure that the NGT ISO is ejected or unmounted within the guest VM before disabling NGT by using the web console. This issue is specific for 4.6 version and does not occur from AOS 4.6.x or later releases.

Note: If you have created the NGT ISO CD-ROMs prior to AOS 4.6 or later releases, the NGT functionality will not work even if you upgrade your cluster because REST APIs have been disabled. You need to unmount the ISO, remount the ISO, install the NGT software again, and then upgrade to 4.6 or later version.



4. To launch a console window, click the **Launch Console** action link.
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 console window includes four menu options (top right):
 - Clicking the **Mount ISO** button displays the following window that allows you to mount an ISO image to the VM. To mount an image, select the desired image and CD-ROM drive from the pull-down lists and then click the **Mount** button.

Mount Disk Image

Select the desired ISO along with the CD-ROM drive you wish to mount it to. Please note, any ISO currently attached to the selected drive will be ejected and replaced by your selection.

ISO File

centos7.4

CD-ROM Drive

ide.0 - Empty

Cancel

Mount

Figure 191: Mount Disk Image Window

- Clicking the **C-A-D** icon button sends a **CtrlAltDel** command to the VM.

- Clicking the camera icon button takes a screenshot of the console window.
- Clicking the power icon button allows you to power on/off the VM. These are the same options that you can access from the **Power On Actions** or **Power Off Actions** action link below the VM table (see next step).




Figure 192: Virtual Network Computing (VNC) Window

Note: A VNC client may not function properly on all browsers. Some keys are not recognized when the browser is Chrome. (Firefox typically works best.)

- To start or shut down the VM, click the **Power on** (or **Power off**) action link.
Power on begins immediately. If you want to power off the VMs, you are prompted to select one of the following options:
 - **Power Off.** Hypervisor performs a hard power off action on the VM.
 - **Power Cycle.** Hypervisor performs a hard restart action on the VM.
 - **Reset.** Hypervisor performs an ACPI reset action through the BIOS on the VM.
 - **Guest Shutdown.** Operating system of the VM performs a graceful shutdown.
 - **Guest Reboot.** Operating system of the VM performs a graceful restart.
- To make a backup of the VM, click the **Take Snapshot** action link.
This displays the Take Snapshot dialog box. Enter a name for the snapshot and then click the **Submit** button to start the backup.

Note: These snapshots (stored locally) cannot be replicated to other sites.

Figure 193: Take Snapshot Dialog Box

7. To migrate the VM to another host, click the **Migrate** action link.
This displays the Migrate VM dialog box. Select the target host from the pull-down list (or select the **System will automatically select a host** option to let the system choose the host) and then click the **Migrate** button to start the migration.


The image shows a 'Migrate VM' dialog box. It has a dark header bar with the title 'Migrate VM' and icons for help and close. Below the header, there is a 'HOST' section with a pull-down menu that currently displays 'System will automatically select a host.'. At the bottom of the dialog, there are two buttons: 'Cancel' and 'Migrate'.

Figure 194: Migrate VM Dialog Box

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.

8. To clone the VM, click the **Clone** action link.
This displays the Clone VM dialog box, which includes the same fields as the Create VM dialog box but with all fields (except the name) filled in with the current VM settings and number of clones needed. Enter a name for the clone and number of clones of the VM that are required and then click the **Save** button to create the clone. You can create a modified clone by changing some of the settings. You can also customize the VM during initialization by providing a custom script and specifying files needed during the customization process (see [Creating a VM \(AHV\)](#) on page 268).


The image shows a 'Clone VM' dialog box. It has a dark header bar with the title 'Clone VM' and icons for help and close. The main content area is divided into sections: 'General Configuration' with 'NUMBER OF CLONES' (set to 1) and 'NAME' (set to 'NTNX-acl-1-1'); 'Compute Details' with 'VCPU(S)' (set to 4), 'NUMBER OF CORES PER VCPU' (set to 1), and 'MEMORY' (set to 12 GIB); and 'Disks' with a table showing disk details and an '+ Add New Disk' link. At the bottom, there are 'Cancel' and 'Save' buttons.

Figure 195: Clone VM Dialog Box

9. To modify the VM configuration, click the **Update** action link.

The Update VM dialog box appears, which includes the same fields as the Create VM dialog box. Modify the configuration as needed (see [Creating a VM \(AHV\)](#) on page 268), and then save the configuration. In addition to modifying the configuration, you can attach a volume group to the VM and enable flash mode on the VM. If you attach a volume group to a VM that is part of a protection domain, the VM is not protected automatically. Add the VM to the same consistency group manually.

(For GPU-enabled AHV clusters only) 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, and change the specified vGPU license. However, you need to power off the VM before you perform these operations.

Update VM

Compute Details

VCPU(S)

2

NUMBER OF CORES PER VCPU

1

MEMORY

16 GiB

Disks [+ Add New Disk](#)

BOOT	DEVICE	TYPE	ADDRESS	PARAMETERS
☐	CD-ROM	ide.0	SIZE=0.01GiB; CONTAINER=d...	↑ ✎ ✕
☒	DISK	scsi.0	SIZE=32GiB; CONTAINER=de...	✎ ✕

Volume Groups

[Close](#) [Save](#)

Figure 196: VM Update Dialog Box

Note: If you delete a vDisk attached to a VM and snapshots associated with this VM exist, space associated with that vDisk is not reclaimed unless you also delete the VM snapshots.

To increase the memory allocation and the number of vCPUs on your VMs while the VMs are powered on (hot-pluggable), do the following:

- In the **vCPUs** field, you can increase the number of vCPUs on your VMs while the VMs are powered on.
- In the **Number of Cores Per vCPU** field, you can change the number of cores per vCPUs while the VMs are powered on.
- In the **Memory** field, you can increase the memory allocation on your VMs while the VMs are powered on.

For more information about hot-pluggable vCPUs and memory, see the *Virtual Machine Memory and CPU Hot-Plug Configurations* topic in the *AHV Administration Guide*.

To attach a volume group to the VM, do the following:

- In the **Volume Groups** section, click **Add volume group**, and then do one of the following:
 - From the **Available Volume Groups** list, select the volume group that you want to attach to the VM.
 - Click **Create new volume group**, and then, in the **Create Volume Group** dialog box, create a volume group (see [Creating a Volume Group](#) on page 133). After you create a volume group, select it from the **Available Volume Groups** list.

Repeat these steps until you have added all the volume groups that you want to attach to the VM.

b. Click **Add**.

a. To enable flash mode on the VM, click the **Enable Flash Mode** check box.

- » After you enable this feature on the VM, the status is updated in the VM table view. To view the status of individual virtual disks (disks that are flashed to the SSD), go the **Virtual Disks** tab in the VM table view.
- » You can disable the flash mode feature for individual virtual disks. 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.

10. To delete the VM, click the **Delete** action link. A window prompt appears; click the **OK** button to delete the VM.

The deleted VM disappears from the list of VMs in the table.

Creating a VM (ESXi)

In ESXi clusters, you can create a new virtual machine (VM) through the web console.

Before you begin

- See the requirements and limitations section in the [VM Management through Prism Element \(ESXi\)](#) on page 616 before proceeding.
- Register the vCenter Server with your cluster. For more information, see [Registering a vCenter Server](#) on page 617.

About this task

When creating a VM, you can configure all of its components, such as number of vCPUs and memory, but you cannot attach a volume group to the VM.

To create a VM, do the following:

Procedure

1. In the VM dashboard , click the **Create VM** button.
The Create VM dialog box appears.

2. Do the following in the indicated fields:

a. **Name**: Enter a name for the VM.

b. **Description** (optional): Enter a description for the VM.

c. **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. See the list of supported operating systems in [VM Management through Prism Element \(ESXi\)](#) on page 616 topic.

d. **vCPU(s)**: Enter the number of virtual CPUs to allocate to this VM.

e. **Number of Cores per vCPU**: Enter the number of cores assigned to each virtual CPU.

f. **Memory**: Enter the amount of memory (in GiBs) to allocate to this VM.



3. To attach a disk to the VM, click the **Add New Disk** button.
The Add Disks dialog box appears.


The image shows a dialog box titled "Add Disk" with a dark header bar containing a question mark and a close button. The dialog has several fields: "TYPE" with a dropdown menu showing "DISK"; "OPERATION" with a dropdown menu showing "Allocate on Storage Container"; "BUS TYPE" with a dropdown menu showing "SCSI"; "STORAGE CONTAINER" with a dropdown menu showing "default-container-54681"; and "SIZE (GiB)" with an empty text input field. At the bottom right, there are "Cancel" and "Add" buttons.

Figure 197: Add Disk Dialog Box

Do the following in the indicated fields:

- a. **Type:** Select the type of storage device, **DISK** or **CD-ROM**, from the pull-down list.
The following fields and options vary depending on whether you choose **DISK** or **CD-ROM**.
- b. **Operation:** Specify the device contents from the pull-down list.
 - Select **Clone from ADSF file** to copy any file from the cluster that can be used as an image onto the disk.
 - Select **Allocate on Storage Container** to allocate space without specifying an image. (This option appears only when **DISK** is selected in the previous field.) Selecting this option means you are allocating space only. You have to provide a system image later from a CD-ROM or other source.
- c. **Bus Type:** Select the bus type from the pull-down list. The choices are **IDE** or **SCSI**.
- d. **ADSF Path:** Enter the path to the desired system image.
This field appears only when **Clone from ADSF file** is selected. It specifies the image to copy. Enter the path name as `/storage_container_name/vmdk_name.vmdk`. For example to clone an image from `myvm-flat.vmdk` in a storage container named `crt1`, enter `/crt1/myvm-flat.vmdk`. When a user types the storage container name

([/storage_container_name/](#)), a list appears of the VMDK files in that storage container (assuming one or more VMDK files had previously been copied to that storage container).

Note: Make sure you're copying from a flat file.

- e. **Storage Container:** Select the storage container to use from the pull-down list.
This field appears only when **Allocate on Storage Container** is selected. The list includes all storage containers created for this cluster.
 - f. **Size:** Enter the disk size in GiBs.
 - g. When all the field entries are correct, click the **Add** button to attach the disk to the VM and return to the Create VM dialog box.
 - h. Repeat this step to attach additional devices to the VM.
4. To create a network interface for the VM, click the **Add New NIC** button.
The Create NIC dialog box appears. Do the following in the indicated fields:
 - a. **VLAN Name:** Select the target virtual LAN from the pull-down list.
The list includes all defined networks (see [Network Configuration for VM Interfaces](#) on page 138).
 - b. **Network Adapter Type:** Select the network adapter type from the pull-down list. See [VM Management through Prism Element \(ESXi\)](#) on page 616 for the list of supported adapter types
 - c. **Network UUID:** This is a read-only field that displays the network UUID.
 - d. **Network Address/Prefix:** This is a read-only field that displays the network IP address and prefix.
 - e. When all the field entries are correct, click the **Add** button to create a network interface for the VM and return to the Create VM dialog box.
 - f. Repeat this step to create additional network interfaces for the VM.
 5. When all the field entries are correct, click the **Save** button to create the VM and close the Create VM dialog box.
The new VM appears in the VM table view (see [VM Table View](#) on page 255).

Managing a VM (ESXi)

You can use the web console to manage virtual machines (VMs) in the ESXi clusters.

Before you begin

- See the requirements and limitations section in the [VM Management through Prism Element \(ESXi\)](#) on page 616 before proceeding.
- Ensure that you have registered the vCenter Server with your cluster. For more information, see [Registering a vCenter Server](#) on page 617.

About this task

After creating a VM, you can use the web console to manage guest tools, power operations, suspend, launch a VM console window, update the VM configuration, clone the VM, or delete the VM. To accomplish one or more of these tasks, do the following:



Note: Your available options depend on the VM status, type, and permissions. Unavailable options are grayed out.

Procedure

1. In the VM dashboard , click the **Table** view.
2. Select the target VM in the table (top section of screen).

The Summary line (middle of screen) displays the VM name with a set of relevant action links on the right. The possible actions are **Manage Guest Tools**, **Launch Console**, **Power on** (or **Power off actions**), **Suspend** (or **Resume**), **Clone**, **Update**, and **Delete**. The following steps describe how to perform each action.

The screenshot shows the VM dashboard in 'Table' view. A table lists several VMs with columns for VM Name, Host, IP Address, Config, Memory Capacity, Storage, CPU Usage, Memory Usage, Controller Status, Controller Write Ops, Controller Read Ops, Controller AVIO Latency, Backup Ops, and Disk Mode. The VM 'vmm001' is selected, highlighted in blue. Below the table, a summary line for 'vmm001' is shown, with a red arrow pointing to the 'selected VM' label and a yellow box labeled 'action links' pointing to the right. The action links include 'Manage Guest Tools', 'Launch Console', 'Power Off Actions', 'Suspend', 'Clone', 'Update', and 'Delete'.

Figure 198: VM Action Links

3. To manage guest tools as follows, click **Manage Guest Tools**.

You can also enable NGT applications (self-service restore, Volume Snapshot Service and application-consistent snapshots) as part of manage guest tools.

- a. Select **Enable Nutanix Guest Tools** check box to enable NGT on the selected VM.
- b. Select **Mount Nutanix Guest Tools** to mount NGT on the selected VM.

Ensure that VM must have at least one empty IDE CD-ROM slot to attach the ISO.

The VM is registered with the NGT service. NGT is enabled and mounted on the selected virtual machine. A CD with volume label NUTANIX_TOOLS gets attached to the VM.

- c. To enable self-service restore feature for Windows VMs, click **Self Service Restore (SSR)** check box.

The Self-Service Restore feature is enabled of the VM. The guest VM administrator can restore the desired file or files from the VM. For more information on self-service restore feature, see Self-Service Restore in the guide titled *Data Protection and Recovery with Prism Element*.

- d. After you select **Enable Nutanix Guest Tools** check box the VSS and application-consistent snapshot feature is enabled by default.

After this feature is enabled, Nutanix native in-guest VmQuiesced Snapshot Service (VSS) agent is used 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 Nutanix platform in a hypervisor-

agnostic manner. For more information, see Implementation Guidelines for Asynchronous Disaster Recovery in the guide titled *Data Protection and Recovery with Prism Element*.

- e. To mount VMware guest tools, click **Mount VMware Guest Tools** check box.
The VMware guest tools are mounted on the VM.

Note: You can mount both VMware guest tools and Nutanix guest tools at the same time on a particular VM provided the VM has sufficient empty CD-ROM slots.

- f. Click **Submit**.

The VM is registered with the NGT service. NGT is enabled and mounted on the selected virtual machine. A CD with volume label NUTANIX_TOOLS gets attached to the VM.

Note:

- If you clone a VM, by default NGT is not enabled on the cloned VM. If the cloned VM is powered off, enable NGT from the UI and power on the VM. If cloned VM is powered on, enable NGT from the UI and restart the nutanix guest agent service.
- If you want to enable NGT on multiple VMs simultaneously, see [Enabling NGT and Mounting the NGT Installer Simultaneously on Multiple Cloned VMs](#) on page 304.

If you eject the CD, you can mount the CD back again by logging into the Controller VM and running the following nCLI command.

```
ncli> ngmt mount vm-id=virtual_machine_id
```

For example, to mount the NGT on the VM with VM_ID=00051a34-066f-72ed-0000-000000005400::38dc7bf2-a345-4e52-9af6-c1601e759987, type the following command.

```
ncli> ngmt mount vm-id=00051a34-066f-72ed-0000-000000005400::38dc7bf2-a345-4e52-9af6-c1601e759987
```

CAUTION: In AOS 4.6, for the powered-on Linux VMs on AHV, ensure that the NGT ISO is ejected or unmounted within the guest VM before disabling NGT by using the web console. This issue is specific for 4.6 version and does not occur from AOS 4.6.x or later releases.

Note: If you have created the NGT ISO CD-ROMs prior to AOS 4.6 or later releases, the NGT functionality will not work even if you upgrade your cluster because REST APIs have been disabled. You need to unmount the ISO, remount the ISO, install the NGT software again, and then upgrade to 4.6 or later version.

4. To launch a VM console window, click the **Launch Console** action link.

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 Off Actions** action link below the VM table can also be accessed from the VNC console window. To access the VM power options, click the **Power** button at the top-right corner of the console window.

Note: A VNC client may not function properly on all browsers. Some keys are not recognized when the browser is Chrome. (Firefox typically works best.)



5. To start (or shut down) the VM, click the **Power on** (or **Power off**) action link.

Power on begins immediately. If you want to power off the VMs, you are prompted to select one of the following options:

- **Power Off.** Hypervisor performs a hard power off action on the VM.
- **Reset.** Hypervisor performs an ACPI reset action through the BIOS on the VM.
- **Guest Shutdown.** Operating system of the VM performs a graceful shutdown.
- **Guest Reboot.** Operating system of the VM performs a graceful restart.

Note: The **Guest Shutdown** and **Guest Reboot** options are available only when VMware guest tools are installed.

6. To pause (or resume) the VM, click the **Suspend** (or **Resume**) action link. This option is available only when the VM is powered on.

7. To clone the VM, click the **Clone** action link.

This displays the Clone VM dialog box, which includes the same fields as the Create VM dialog box but with all fields (except the name) filled in with the current VM settings and number of clones needed. Enter a name for the clone and number of clones of the VM that are required and then click the **Save** button to create the clone.

The screenshot shows the 'Clone VM' dialog box. It has a title bar with a question mark and a close button. The main content area is divided into sections: 'General Configuration' with 'NUMBER OF CLONES' set to 1 and 'NAME' set to 'NTNX-aci-1-1'; 'Compute Details' with 'VCPUS' set to 4, 'NUMBER OF CORES PER VCPU' set to 1, and 'MEMORY' set to 12 GiB; and 'Disks' with a table showing disk details and an 'Add New Disk' button. At the bottom are 'Cancel' and 'Save' buttons.

Figure 199: Clone VM Dialog Box

Note: In the Clone window, you cannot update the disks and network interfaces.

8. To modify the VM configuration, click the **Update** action link.

The Update VM dialog box appears, which includes the same fields as the Create VM dialog box. Modify the configuration as needed (see [Creating a VM \(ESXi\)](#) on page 283), and in addition you can enable flash mode for the VM.

Note: If you delete a vDisk attached to a VM and snapshots associated with this VM exist, space associated with that vDisk is not reclaimed unless you also delete the VM snapshots.

- a. Click the **Enable Flash Mode** check box.
 - » After you enable this feature on the VM, the status is updated in the VM table view. To view the status of individual virtual disks (disks that are flashed to the SSD), go the **Virtual Disks** tab in the VM table view.
 - » You can disable the flash mode feature for individual virtual disks. 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.
9. To delete the VM, click the **Delete** action link. A window prompt appears; click the **OK** button to delete the VM.

The deleted VM disappears from the list of VMs in the table. You can also delete a VM that is already powered on.

Configuring Images

In AHV clusters, you can import and configure operating system ISO and disk image files through the web console. You can also convert previously imported files to the format that AHV uses.

About this task

The image service feature allows you to build a store of imported files that you can use to create a CD-ROM from an ISO image or an operating system. Disk from a disk image when creating a VM. The image service supports raw, vhd, vhdx, vmdk, vdi, iso, and qcow2 disk formats. (The image service uses port 2007, so that port must be open.)

Image create, update, and delete (CUD) behavior depends on whether a cluster (also known as Prism Element) is registered to Prism Central.

- Images created on a Prism Element reside on Prism Element and can be managed from Prism Element.
- For better centralized management, you can migrate images manually to Prism Central by using the image import feature in Prism Central. An image migrated to Prism Central in this way remains on Prism Element, but you can manage the image only from Prism Central. Migrated images cannot be updated from the Prism Element.
- In the case of a local image upload, with more than one PE cluster managed by Prism Central, the image state is active on that PE cluster. All other PE clusters will show the image as inactive. If you create a VM from that image, the image bits are copied to the other PE clusters. The image then appears in an active state on all managed PE clusters.

To import and configure an image file, do the following:



Procedure

- 1. Click the gear icon in the main menu and then select **Image Configuration** in the **Settings** page.
The Image Configuration window appears.




Figure 200: Image Configuration Window

2. To upload an image file to the cluster, click the **Upload Image** button.
The Create Image window appears. Do the following in the indicated fields:

Create Image

Name

Annotation

Image Type

Storage Container

Image Source

☒ From URL

☐ Upload a file Choose File No file chosen

< Back Cancel Save

Figure 201: Create Image Window

- Name:** Enter a name for the image.
- Annotation** (optional): Enter a description for the image.
- Image Type** (optional): Select the image type, either **ISO** or **Disk**, from the pull-down list.
- Storage Container:** Select the storage container to use from the pull-down list.
The list includes all storage containers created for this cluster. If there are no storage containers currently, a Create Storage Container link appears to create a storage container (see [Creating a Storage Container](#) on page 126).
- Image Source:** Do one of the following:

- Click the **From URL** radio button to import the image from the Internet. 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](#) on page 595). A file uploaded through NFS must have 644 permissions.

If the image files have been copied to a container on the cluster, replace `IP_addr` with 127.0.0.1, the IP address of the loopback interface.

For example, enter `nfs://127.0.0.1/container_name/vm_name/vm_name.vmdk` or `nfs://127.0.0.1/container_name/vm_name/vm_name-flat.vmdk`.

Replace `container_name` with the name of the container where the image is placed and replace `vm_name` with the name of the VM where the image is placed.

- Click the **Upload a file** radio button to upload a file from your workstation. Click the **Choose File** button and then select the file to upload from the file search window.
 - f. When all the fields are correct, click the **Save** button.
The Create Image window closes and the Image Configuration window reappears with the new image appearing in the list.
3. To update the image information, click the pencil icon for that image.
The **Update Image** window appears. Update the fields as desired and then click the **Save** button.

Note: The pencil icon is unavailable for images imported to Prism Central. Use Prism Central to manage such images.

4. To delete an image file from the store, click the X icon for that image.
The image file is deleted and that entry disappears from the list.

Virtual Machine Customization

In an Acropolis cluster, you can use Cloud-init to customize Linux VMs and the System Preparation (Sysprep) tool to customize Windows VMs.

About Cloud-Init

Cloud-init is a utility that is used to customize Linux VMs during first-boot initialization. The utility must be pre-installed in the operating system image used to create VMs. Cloud-init runs early in the boot process and configures the operating system on the basis of data that you provide (user data). You can use Cloud-init to automate tasks such as setting a host name and locale, creating users and groups, generating and adding SSH keys so that users can log in, installing packages, copying files, and bootstrapping other configuration management tools such as Chef, Puppet, and Salt. For more information about Cloud-init, see <https://cloudinit.readthedocs.org/>.

About Sysprep

Sysprep is a utility that prepares a Windows installation for duplication (imaging) across multiple systems. Sysprep is most often used to generalize a Windows installation. During generalization, Sysprep removes system-specific information and settings such as the security identifier (SID) and leaves installed applications untouched. You can capture an image of the generalized installation and use the image with an answer file to customize the installation of Windows on other systems. The answer file contains the information that Sysprep needs to complete an unattended installation. For more information about Sysprep and answer files, see the Sysprep documentation at <https://technet.microsoft.com/>.



The Customization Process in a Nutanix Cluster

You can use Cloud-init or Sysprep both when creating and when cloning VMs in a Nutanix cluster. For unattended provisioning, you can specify a user data file for Cloud-init and an answer file for Sysprep. All Cloud-init user-data formats are supported. For example, you can use the Cloud Config format, which is written in YAML, or you can provide a multi-part archive. To enable Cloud-init or Sysprep to access the script, AOS creates a temporary ISO image that includes the script and attaches the ISO image to the VM when you power on the VM.

Note: The ISO image is mounted on bus IDE 3, so ensure that no other device is mounted on that bus.

You can also specify source paths to the files or directories that you want to copy to the VM, and you can specify the target directories for those files. This is particularly useful if you need to copy software that is needed at start time, such as software libraries and device drivers. For Linux VMs, AOS can copy files to the VM. For Windows VMs, AOS can copy files to the ISO image that it creates for the answer file.

After customizing a VM, you can copy the VDisk of the VM to Image Service for backup and duplication.

Customizing Linux Virtual Machines with Cloud-Init

Keep the user data file ready, either saved locally or uploaded to a storage container on the cluster. Alternatively, you can create or paste the script in the web console. If you want files copied to the VM during initialization, upload the files to a storage container on the cluster.

About this task

To customize a Linux VM by using Cloud-init, do the following:

Procedure

1. Log in to the web console by using the Nutanix credentials.
2. In the VM dashboard (see [VM Dashboard](#) on page 252), do one of the following:
 - » To create a VM, click **Create VM**.
 - » To clone a VM, click the VM that you want to clone, and then click **Clone**.
3. In the **Create VM** or **Clone VM** dialog box, specify a name for the VM and allocate resources such as vCPUs, memory, and storage. Select the **Custom Script** check box and specify how you want to customize the VM.

For information about creating a VM and specifying customization options, see [Creating a VM \(AHV\)](#) on page 268. For information about cloning a VM, see [Managing a VM \(AHV\)](#) on page 275.
4. In the VM dashboard, select the VM, and then click **Power On**.

The VM is powered on and initialized based on the directives in the user data file. To create a reference image from the VM, use Image Service. See Image Service in the VM Management chapter of the *Acropolis App Mobility Fabric Guide*.

Cloud-Init Limitations and Guidelines

- The maximum size of a VM guest customization script is 32 KB.
- Config drive version 2 data format is supported.



Customization of Windows Virtual Machines with System Preparation

To customize a Windows VM by using Sysprep, you need to perform the following tasks:

1. Create a reference image by using Sysprep.
2. Create a VM from the reference image.

You can also customize a VM when performing a fresh installation of Windows with an ISO file.

If you require unattended provisioning, keep the answer file ready, either saved locally or uploaded to a storage container on the cluster. Alternatively, you can create or paste the script in the web console. If you have files that need to be copied to the temporary ISO image, upload the files to a storage container on the cluster.

Preparing a VM to Create a Reference Image

Creating a reference image requires knowledge of Sysprep. For information about how to use Sysprep, see the Sysprep documentation on the Microsoft TechNet website.

About this task

To create a reference image, do the following:

Procedure

1. Log in to the web console by using the Nutanix credentials, and then browse to the VM dashboard (see [VM Dashboard](#) on page 252).
2. Select the VM that you want to clone, click **Launch Console**, and then log in to the VM with administrator credentials.
3. Configure Sysprep with the system cleanup action of your choice, specify whether or not you want to generalize the installation, and then choose to shut down the VM.

- Open the command prompt as an administrator and navigate to the sysprep folder.

```
cd C:\windows\system32\sysprep
```

- Generalize the installation and then shut down the VM.

```
C:\Windows\system32\sysprep>sysprep /generalize /shutdown /oobe
```

The VM shuts down automatically.

Note: Make sure to shut down the VM. Restarting the VM will result in the VM losing its generalized state and in Sysprep attempting to find an answer file that has not been provided yet. For the same reasons, until you have completed this procedure, do not start the VM.

4. Create a reference image from the VM by using Image Service. See [Configuring Images](#) on page 289 for more information.

Creating a Customized Virtual Machine from a Reference Image

About this task

To use a reference image, do the following:

Procedure

1. Log in to the web console by using the Nutanix credentials, and then browse to the VM dashboard (see [VM Dashboard](#) on page 252).



2. Click **Create VM**, and then, in the **Create VM** dialog box, do the following:
 - a. Specify a name for the VM and allocate resources such as vCPUs, memory, and storage.
 - b. Click **Add new disk**, select the **Clone from Image Service** operation, and select the Windows reference image that you copied to Image Service.
 - c. Click the **Custom Script** check box and specify how you want to customize the VM.For more information about creating a VM, see [Creating a VM \(AHV\)](#) on page 268.
3. In the VM dashboard, select the VM, and then click **Power On**.

The VM is powered on and initialized based on the directives in the answer file. To create a reference image from the VM, use Image Service. See Image Service in the VM Management chapter of the *Acropolis App Mobility Fabric Guide*.

Customizing a Fresh Installation

You can perform a fresh installation only if you attach an empty vDisk and an installation CD-ROM to the VM. If you specify an image from Image Service or ADSF, for use as a vDisk, the VM is created from that image, and the install is no longer a fresh install.

About this task

To customize a fresh installation of Windows by using Sysprep, do the following:

Procedure

1. Log in to the web console by using the Nutanix credentials, and then browse to the VM dashboard (see [VM Dashboard](#) on page 252).
2. Click **Create VM**, specify the details that are required for installing Windows on the new VM, and then do the following:
 - a. Specify a name for the VM and allocate resources such as vCPUs, memory, and storage.
 - b. In the **Disks** area, click the edit button that is provided against the default CD-ROM entry. In the **Update Disk** dialog box, select the operation (**Clone from ADSF File** or **Clone from Image Service**), and then specify the image that you want to use. Click **Update**.
 - c. Click **Add new disk**. Allocate space for a new disk on a storage container, and then click **Add**.
 - d. Click the **Custom Script** check box and specify how you want to customize the VM.For more information about creating a VM, see [Creating a VM \(AHV\)](#) on page 268.
3. In the VM dashboard, select the VM, and then click **Power On**.

The VM is powered on and initialized based on the directives in the answer file. To create a reference image from the VM, use Image Service. See Image Service in the VM Management chapter of the *Acropolis App Mobility Fabric Guide*.

VM High Availability in Acropolis

If you have not modified high availability configuration from previous version of AOS releases, best effort VM availability is by default implemented in Acropolis.

If you have not enabled high availability, in case of host failure VMs are restarted from the failed host to any available space on the other hosts in the cluster. Once the failed host joins the cluster again, VMs are migrated back to the host. This type of VM high availability is implemented without reserving any resources. Admission control is not enforced and hence there may not be sufficient capacity available to start all the VMs.



Note: Nutanix does not support VMs that are running with 100% remote storage for high availability. The VMs must have at least one local disk that is present on the cluster.

VM high availability can be configured for host based reservations using segment-based reservation method which is the default.

In segment-based reservation, the cluster is divided into segments to ensure enough space is reserved for any host failure. Each segment corresponds to the largest VM that is guaranteed to be restarted in case the failure occurs. The other factor is the number of host failures that can be tolerated. Using these inputs, the scheduler implements admission control to always have enough resources reserved so that the VMs can be restarted upon failure of any host in the cluster.

The segment size ensures that the largest VM can be powered on in HA failover when cluster is fully loaded (if the cluster is fully used except the reserved segments). The number of segments that is reserved is such a way that for each host enough resources are reserved to ensure any host failure in the cluster is tolerated. Multiple VMs may fit into a segment. If anything changes in the cluster, the reservation is computed again. The total resources reserved for segments can be more than the resources used by running VMs. This implementation guarantees successful failover even in the case of fragmentation of segments. The actual number of reserved resources depends on the current load of the cluster, but it is typically at 1 to 1.25 times the resource usage on the most loaded host.

If the host enters maintenance mode (in case of host upgrade), you might not be protected against further host failures. Maintenance mode uses reservations made for HA for migrating VMs from the host. Although, you are not protected against host failure if you have reservation for HA, hypervisor upgrade occurs without any difficulty because from the perspective of a user it is exactly the same as host failure except that the VMs are migrated (instead of restarted) and hence no runtime state is lost. The HA status goes through the same states as it goes when the host failure had occurred.

Segment-based reservation is the default method that Acropolis uses to enable VM high availability.

CAUTION: Do not change segment-based reservation method with the reserve host method.

Enabling High Availability for the Cluster

In Acropolis managed clusters, you can enable high availability for the cluster to ensure that VMs can be migrated and restarted on another host in case of failure.

About this task

After you enable high availability for the cluster, if a host failure occurs the cluster goes through following changes.

- **OK:** This state implies that the cluster is protected against a host failure.
- **Healing:** Healing period is the time that Acropolis brings the cluster to the protected state. There are two phases to this state. The first phase occurs when the host fails. The VMs are restarted on the available host. After restarting all the VMs if there are enough resources to protect the VM, the HA status of the cluster comes back to OK state. If this does not occur, the cluster goes into critical state. The second phase occurs when the host comes back from the failure. Once the host comes back from failure, no VMs are present on the host and hence during this healing phase restore locality task occurs (VMs are migrated back). Apart from restoring the locality of the VMs, the restore locality task ensures that the cluster is back to the same state before the HA failure. Once it is finished, the HA status is back to OK state.



- Critical: If the host is down, the HA status of the cluster goes into Critical state. This happens because the cluster cannot tolerate any more host failures. You have to ensure that you bring back the host so that your cluster is protected against any further host failures.

Note: On a less loaded cluster, it is possible for HA to go directly back to OK state if enough resources are reserved to protect another host failure. The start and migrate operations on the VMs are restricted in the Critical state because Acropolis continuously tries to ensure that the HA status is back to the OK state.

Procedure

1. Click the gear icon in the main menu and then select **Manage VM High Availability** in the **Settings** page.

Note: This option does not appear in clusters that do not support this feature.

The Manage VM High Availability dialog box appears.




Figure 202: Manage VM High Availability Window

2. Check the **Enable HA Reservation** box and then click the **Save** button to enable.

Nutanix Guest Tools

Nutanix guest tools (NGT) is a software package that you can install in a guest virtual machine (Microsoft Windows or Linux) to enable the advanced VM management features provided by Nutanix.

To use the NGT feature, you must do the following:

1. Enable the NGT feature for a VM in the Prism Element web console.

When you enable NGT for a VM, Prism Element prepares the VM so that NGT can be successfully installed and used in that VM.

2. Mount the NGT installer (ISO disk file) in a VM.
3. Install NGT in a VM.

If you install NGT in guest VMs, you can use the following advanced features (NGT applications):

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 by providing drivers for 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](#) on page 620.

VSS requestor and hardware provider for Windows VMs

Enables application-consistent snapshots of AHV or ESXi Windows VMs. For more information about Nutanix VSS-based snapshots for the Windows VMs, see the best practices section for taking application-consistent snapshots under Implementation Guidelines for Asynchronous Disaster Recovery in the *Data Protection and Recovery with Prism Element* guide.

Application-consistent snapshot for Linux VMs

Supports application-consistent snapshots for Linux VMs by running specific scripts on VM quiesce. For more information about Nutanix VSS-based snapshots for the Linux VMs, see the best practices section for taking application-consistent snapshots in Implementation Guidelines for Asynchronous Disaster Recovery in the *Data Protection and Recovery with Prism Element* guide.

The NGT software package includes the following:

NGT Installer

Allows you to install NGT in a guest VM.

Nutanix Guest Agent (NGA) Service

Maintains a communication channel between the Nutanix Controller VM and guest VMs.

Nutanix VirtIO Package

Includes the Nutanix VM mobility drivers that enable the following features:

- VM migration between ESXi and AHV
- In-place hypervisor conversion
- Cross-hypervisor disaster recovery (CHDR)

Nutanix Guest Tools Requirements and Limitations

Ensure that the following requirements are met, so that you can successfully use all the features of NGT.

General Requirements

- 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. See [Virtual IP Address Impact](#) on page 59 for more information.
- VMs must have at least one empty IDE CD-ROM or SATA slot to attach the ISO.
- The following ports must be open between the Controller VMs and guest VMs on which NGT is installed:
 - Controller VM: 2074. Port 2074 must be open in the Controller VM, so that user VMs can communicate with the Controller VM.
 - User VM: TCP port 23578 if you want to use the VSS service.



- Guest VMs must be connected to a network that can be accessed by using the virtual IP address of the cluster.
- Hypervisor versions must be as follows:
 - ESXi version 5.1 or later
 - AHV version 20160215 or later

Windows Requirements

- For Windows Server Edition VMs, ensure that Microsoft VSS is enabled before you install NGT.
- PowerShell path must be available in the environment for NGT installation. Otherwise, installation of NGT fails.
- For NGT installation on Windows VM, PowerShell path must be correctly set in the environment variables.
- For Windows to trust the device drivers for AHV, you must install the correct root CA in the domain. NGT installation fails if the root CA is not installed in the domain.

Linux Requirements

- For Linux operating systems, NGT is installed at **/usr/local**. Hence, ensure that the user has the write permissions for this location.
- For Linux operating systems, the lvm2 package version must be version 2.02.109 or later.

Limitations

- If a VM is connected to a volume group (VG), NGT captures the IQN of the VM and stores the information. If you change the VM IQN and you take a snapshot of the VM before the NGT refresh cycle (currently at 5 minutes) occurs, NGT will not be able to provide the auto restore capability because the snapshot operation will not be able to capture the VM-VG connection. As a workaround, restart the Nutanix Guest Agent service by running the **\$ sudo service ngt_guest_agent restart** command on Linux VMs. For Windows VMs, in the command prompt, run the **net stop "Nutanix Guest Tools Agent" && net start "Nutanix Guest Tools Agent"** command.

See the supported operating system information for the specific NGT features to verify if the operating system is supported for the particular feature or not.

Enabling NGT and Mounting the NGT Installer in a VM

By default, the NGT feature is disabled for a VM running in a Nutanix cluster. To install and use the NGT feature in a VM, you must first enable the NGT feature (allow the installation and usage of NGT) in a VM and then mount the NGT installer (ISO disk file) in that VM by using the Prism Element web console.

About this task

When you are enabling the NGT feature and mounting the NGT installer in a VM, you must also select the NGT applications (self-service restore, volume snapshot service, and application-consistent snapshots) that you want to use in that VM.

Perform the following procedure to enable the NGT feature, mount the NGT installer, and select the NGT applications that you want to use in a VM by using the Prism Element web console.



Procedure

1. Log on to the Prism Element web console.
2. In the drop-down list of the main menu, select **VM** and select the **Table** view.
3. Select the VM for which you want to enable NGT and click **Manage Guest Tools**.
The **Manage Guest Tools** dialog box appears. In the **Manage Guest Tools** dialog box, in addition to enabling the NGT feature and mounting the NGT installer, you must also select the NGT applications that you want to use in a VM.

Do the following in the indicated fields.

- a. **Enable Nutanix Guest Tools.** Select this check box to enable the NGT feature in the selected VM.
- b. **Mount Nutanix Guest Tools.** Select this check box to mount the NGT installer on the selected VM.

Note: Ensure that the VM has at least one empty IDE CD-ROM or SATA slot to attach the ISO.

- c. **Self Service Restore (SSR).** Select this check box to enable the self-service restore feature for Windows VMs.

For more information about the self-service restore feature, see *Self-Service Restore* in the *Data Protection and Recovery with Prism Element* guide.

- d. After you select the **Enable Nutanix Guest Tools** check box, the VSS and application-consistent snapshot applications are automatically selected.

With these applications enabled, the Nutanix native in-guest VmQuiesced Snapshot Service (VSS) agent is used 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, see

Implementation Guidelines for Asynchronous Disaster Recovery in the Data Protection and Recovery with Prism Element guide.

e. Click **Submit**.

Prism Element now enables the NGT feature, mounts the NGT installer, and attaches a CD (ISO disk file) with the volume label NUTANIX_TOOLS to the selected VM.




Figure 203: Manage NGT

Note:

If you clone a VM, NGT is not enabled on the cloned VM by default. If the cloned VM is powered off, enable NGT from the Prism Element web console and power on the VM. If the cloned VM is powered on, enable NGT from the Prism Element web console and restart the Nutanix Guest Agent service.

Restart the Nutanix guest agent service by running the **\$ sudo service ngt_guest_agent restart** command on Linux VMs. For Windows VMs, in the command prompt, run the **net stop "Nutanix Guest Tools Agent" && net start "Nutanix Guest Tools Agent"** command.

If you eject the CD, you can manually mount the CD again by running the following nCLI command.

```
ncli> ngt mount vm-id=virtual_machine_id
```

For example, to mount the CD on the VM with VM_ID=00051a34-066f-72ed-0000-000000005400::38dc7bf2-a345-4e52-9af6-c1601e759987, type the following command.

```
ncli> ngt mount vm-id=00051a34-066f-72ed-0000-000000005400::38dc7bf2-a345-4e52-9af6-c1601e759987
```

NGT Installation

After you enable NGT and mount the NGT installer in a VM, you must install NGT in a VM (Windows or Linux) by logging into that VM. You can also use the silent installer package to install NGT on Windows VMs without affecting the normal VM operations.



Installing NGT on a Windows VM

After you enable NGT and mount the NGT installer in a VM, log on to the Windows VM to install NGT.

About this task

Procedure

1. Log on to the Windows VM.
2. In the Windows explorer, in the left pane, click **CD Drive NUTANIX_TOOLS**, and, in the right pane, double-click **setup** (setup.exe file).




Figure 204: Start NGT Installation on Windows

As shown in the screenshot, the Nutanix icon appears next to the CD drive and setup.exe file.

Note: If you mount the NGT installer while the VM is powered off, the Nutanix icon is not displayed after you power on the VM. You must open the CD (right-click the CD and click **Open**) and double-click the setup.exe file.

3. Accept the license agreement and follow the prompts to install NGT in the virtual machine.

Note: The Nutanix VSS HW Provider does not get automatically registered after installation. The process of registration starts after you take the snapshot of the VM and it automatically unregisters after the operation is complete.

CAUTION: If the Nutanix VM Mobility driver version in AHV is of the latest version than the one installed in your VM, during the installation of NGT this driver is upgraded and a prompt that asks for a restart of the VMs is displayed. Restart the VMs to upgrade the mobility driver version successfully.

Installing NGT on Windows Machines (Silent Install)

The silent installer package allows you to install NGT on Windows VMs without affecting the normal VM operations. You can also use the silent installer package to install NGT on multiple Windows VMs simultaneously. You can achieve this by adding the commands described in this topic to a custom script to install NGT on multiple Windows VMs simultaneously.

About this task

The NGT silent installer package consists of the following components.

- Custom module that performs checks and fulfills all the prerequisites.
- Nutanix VM mobility drivers

- Python
- Microsoft Visual C++ 2008/2012 x64 redistributable setup for Windows Server operating systems
- Nutanix VSS package
- Nutanix Guest Agent service package

About this task

Perform the following procedure to install NGT by using the NGT silent installer package.

Procedure

1. Open the command prompt and go to the drive on which the NGT installer is mounted.
2. Install NGT by using the silent installer package.

CAUTION:

If the Nutanix VM Mobility driver version in AHV is of the latest version than the one installed in your VM, during the installation of NGT this driver is upgraded and you must restart the VMs to upgrade the mobility driver version successfully. See the Release Notes of your AHV and AOS versions to verify if the Nutanix VM Mobility driver version is different in AHV and your VM.

On AHV, depending on whether you need to restart the VMs or not, use one of the following commands to install NGT. You can use any of the following commands for ESXi, because VMs running on ESXi are not affected due to the mismatch of the Nutanix VM Mobility driver version.

» **drive:\> setup.exe /quiet ACCEPTEULA=yes /norestart**

Use this command to ensure that VMs are not restarted after installing NGT. Note that if you use this command, mobility drivers are not updated. You must manually restart the VMs to update the mobility drivers.

» **drive:\> setup.exe /quiet ACCEPTEULA=yes**

If you use this command, VMs are restarted and drivers are updated.

By default, NGT logs are generated in **Event Viewer** and in the **%TEMP%** directory starting with *Nutanix_Guest_Tools_timestamp*. Event Viewer log files are created for all the components as part of the NGT installation.

Note: If you get any warnings after running this command, these warnings are treated as errors and are captured in the logs. If you want to ignore the warnings and complete the installation, you can use the flag **IGNOREALLWARNINGS=yes**. For example, **drive:\> setup.exe /quiet ACCEPTEULA=yes IGNOREALLWARNINGS=yes**. If you enable this flag, all the warnings are ignored and the silent installation proceeds smoothly.



3. (Optional) Install NGT if you want the logs to be generated in a location other than the **%TEMP%** directory.

```
drive:\> setup.exe /quiet ACCEPTEULA=yes -log Log_path
```

Replace *Log_path* with the path where you want to create the log files.

Logs are generated in **Event Viewer** and in the path that you provide. Logs are not generated in the **%TEMP%** directory. Ensure that the path you provide has the necessary write permissions.

CAUTION: If the Nutanix VM Mobility driver version in AHV is of the latest version than the one installed in your VM, during the installation of NGT this driver is upgraded and you must restart the VMs to upgrade the mobility driver version successfully. See the Release Notes of your AHV and AOS versions to verify if the Nutanix VM Mobility driver version is different in AHV and your VM.

Installing NGT on Linux Machines

After you enable NGT and mount the NGT installer in a VM, log on to the Linux VM to install NGT. Graphical Linux deployments auto-discover the CD-ROM and mount the CD-ROM appropriately. For non-graphical Linux deployments, use the following procedure.

Procedure

1. Log on to the Linux VM.
2. Determine the directory in which the CD has been mounted.
For example, the CD might be mounted in */dev/sr0*, */dev/sr1*, and so on.
3. Create a temporary directory (*/mnt*) and mount the content of the CD in the temporary directory you created.

```
$ sudo mount /dev/directory-name /mnt
```

Replace *directory-name* with the directory you determined in step 2. For example, *sr0*, *sr1*, and so on.

4. Install NGT.

```
$ sudo /mnt/installer/linux/install_ngt.py
```

Enabling NGT and Mounting the NGT Installer Simultaneously on Multiple Cloned VMs

If you have cloned multiple VMs from a single VM (master VM), you can enable NGT and mount the NGT installer simultaneously on multiple VMs by using the master VM image.

Before you begin

Ensure the following before you perform this task:

1. Install NGT on the master VM.
2. Clone the required number of VMs from the master VM.
3. Shut down the cloned VMs.

Perform the following procedure to enable NGT and mount the NGT installer simultaneously on multiple VMs by using the master VM image.

Note: After you perform the following procedure, you do not need to separately install NGT on the cloned VMs.



Procedure

1. For every cloned VM, log on to the Controller VM and run the following command.

```
nccli> ngt mount vm-id=clone_vm_id
```

Replace *clone_vm_id* with the ID of the cloned VM.

To find the ID of the cloned VM, run `nccli> vm list name="<clone-vm-name>"` command. Note the value of the **Id** field as *clone_vm_id*.

```
<nccli> vm list name="<clone-vm-name>"
```

```
Id                : 00058a81-64bb-27af-69d2-ac1f6b602812::50132bb1-a07b-2e34-
a768-d6b0ffc940e7
Uuid              : 5522f8cd-1c11-456c-b287-485e5e9de123
Name              : abcdef asdf
VM IP Addresses   :
Hypervisor Host Id : 00058a81-64bb-27af-69d2-ac1f6b602812::3
Hypervisor Host Uuid : ae80deac-0224-43c9-9234-7a2071217e6d
Hypervisor Host Name : gologolo2-1
Memory            : 1 GiB (1,073,741,824 bytes)
Virtual CPUs      : 1
VDisk Count       : 0
Protection Domain :
Consistency Group  :
```

2. Power on the cloned VM.
3. (Optional) If you enable NGT and mount the NGT installer while the cloned VMs are powered on, do one of the following:
 - » Restart the Nutanix guest agent service by running the `$ sudo service ngt_guest_agent restart` command on Linux VMs. For Windows VMs, in the command prompt, run the `net stop "Nutanix Guest Tools Agent" && net start "Nutanix Guest Tools Agent"` command.
 - » Guest power cycle the cloned VMs.

NGA and Controller VM Communication

After NGT is installed in a VM, the Nutanix Guest Agent (NGA) service in the VM starts periodic communication with the Controller VM over SSL connections.

Each Nutanix cluster is configured as a Certificate Authority (CA). When NGT is enabled in a VM, a certificate pair is generated for that specific 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 running inside the VM initiates an SSL connection to the virtual IP port 2074 of the Controller VM to communicate with the Controller VM. No change in the firewall is required from the VM side. However, the VM must be able to reach port 2074 through SSL on the Controller VM.

NGT provides three levels of security for VM communication.

- SSL certificates
- Capability based authorization
- System or BIOS UUID of the VM must be same as seen by the hypervisor



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 will not be able to communicate with the Controller VM. You must separately enable NGT and the mount NGT installer on the cloned VM.

NGA also publishes information about the VM to the Controller VM, for example, guest OS type, status of VM mobility, and VSS services.

CD-ROM Eject Functionality of NGT

After NGT is installed in a VM, the CD-ROM is auto-ejected from the VM in the following two stages.

1. CD-ROM is automatically ejected from the guest VM immediately after the installation of NGT is successfully completed.
2. NGA then indicates to the Controller VM that the installation is complete, and the CD-ROM is detached from the VM through the hypervisor and destroyed. This operation depends on the periodicity of the VM and Controller VM communication and may take up to 10 minutes.

Note:

- There can be a time difference between these operations. Hence, there may be a possibility that the CD-ROM shows up as empty inside a VM, but from the NGT-Controller VM service Tools Mounted may be displayed as true. If this situation occurs, verify the hypervisor reported values.
- If the NGT software version in the ISO is greater than the installed version inside the VM, the CD-ROM does not get automatically ejected. This helps with the upgrade of the NGT software inside the VM.

Nutanix Guest Tools Usage in Disaster Recovery

If a VM that is protected by the Nutanix data protection feature and has NGT enabled for it or vice-versa, the relevant NGT information including capabilities is added as part of the disaster recovery snapshot record.

If you migrate a VM to a remote site, restore it in place, or clone it from a snapshot, the NGT information is preserved for the restored VM. A new NGT ISO image containing only the relevant configuration information (SSL certificates, Controller VM IP address, etc.) is created for the recovered 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 CD-ROM automatically.

If you restore a VM on the remote site or retrieve the snapshot back to the local or source cluster, NGT information or functionality is lost. Therefore, you must enable NGT again for the restored VM.

Upgrading NGT

If you have upgraded your AOS, you can also upgrade NGT. NGT is not automatically upgraded. Following are the general guidelines for upgrading NGT.

- After you upgrade AOS, you can choose to upgrade NGT. You must again enable NGT and mount the NGT installer in a VM to upgrade NGT to a new version. See [Enabling NGT and Mounting the NGT Installer in a VM](#) on page 299 for more information. You do not need to install NGT again if you just want to upgrade NGT to the latest version.
- If the Nutanix VM Mobility driver version in AHV is of the latest version than the one installed in your VM, during the installation of NGT this driver is upgraded and a prompt that asks



for a restart of the VMs is displayed. Restart the VMs to upgrade the mobility driver version successfully.

Reconfiguring NGT

If you reconfigure the cluster IP address, NGT loses connection with the CVM. You must then reconfigure NGT to establish the connection again.

About this task

Perform the following procedure to reconfigure NGT.

Procedure

1. Mount the NGT installer. For more information about mounting the NGT installer, see [Enabling NGT and Mounting the NGT Installer in a VM](#) on page 299.
2. Do one of the following.

- » Restart the VM.
- » If you do not want to restart the VM, restart the Nutanix Guest Agent (NGA) service. To restart the NGA service, use one of the following commands, depending on your operating system.

On Linux VMs, run the command: `$ sudo service ngt_guest_agent restart`.

On Windows VMs, in the command prompt, run the command: `net stop "Nutanix Guest Tools Agent" && net start "Nutanix Guest Tools Agent"`.

After you mount the NGT ISO and restart NGA, NGA fetches the latest configuration (new cluster IP address) mounted in the guest VM. The guest VM can now use the new IP address to communicate with the cluster.

Disabling, Uninstalling, and Removing Nutanix Guest Tools

You can use the Prism Element web console to disable and uninstall NGT in a VM.

About this task

If you disable NGT, only the communication between the VM and NGT-Controller VM service is stopped. To completely remove NGT from the database, you must use nCLI.

Procedure

1. Disable NGT.
 - a. Log on to the Prism Element web console.
 - b. Go to the **Table** view of the **VM** page.
 - c. Select the VM in which you want to disable NGT.
 - d. Click **Manage Guest Tools**.
 - e. Clear the **Enable Nutanix Guest Tools** check box.
 - f. Click **Submit**.

Disabling NGT rejects any form of communication between the VM and NGT-Controller VM service. The NGT software is not uninstalled from the VM.



2. Uninstall NGT.

- a. For Windows VMs, run the following command.

```
drive:\> setup.exe /uninstall /quiet ACCEPTEULA=yes /norestart
```

For Windows VMs, uninstalling the NGT software from a VM does not uninstall the Nutanix virtIO package. This ensures that the VM always has the required drivers for mobility purposes. For example, the VM might be migrated from one hypervisor to another and if drivers are uninstalled, the VM might not start properly.

Note: Uninstalling Nutanix VM Mobility drivers removes SCSI VirtIO drivers leading to VM boot failures.

- b. For Linux VMs, run the following command.

```
$ sudo python /usr/local/nutanix/ngt/python/bin/uninstall_ngt.py
```

3. Remove NGT completely.

- a. Log on to the Controller VM.

- b. Run the following command.

```
ncli> ngt delete vm-id=virtual_machine_id
```

Replace `virtual_machine_id` with the ID of the virtual machine. Deleting the NGT entity clears any relevant NGT information for the VM.

If you delete the VMs on which NGT is enabled, the NGT entity is cleaned up automatically and the ISO files are also deleted.

Affinity Policies for AHV

As an administrator in an Acropolis managed cluster, you can specify scheduling policies for virtual machines on an AHV cluster. By defining these policies, you can control placement of the virtual machines on the hosts within a cluster.

You can define two types of affinity policies.

VM-Host Affinity Policy

The VM-host affinity policy controls the placement of the VMs. You can use this policy to specify that a selected VM can only run on the members of the affinity host list. This policy checks and enforces where a VM can be hosted when you power on or migrate the VM.

Note:

- If you choose to apply the VM-host affinity policy, it limits Acropolis HA and Acropolis Dynamic Scheduling (ADS) in such a way that a virtual machine cannot be powered on or migrated to a host that does not conform to requirements of the affinity policy as this policy is enforced mandatorily.
- The VM-host anti-affinity policy is not supported.
- In previous AOS releases, VMs configured with host affinity settings retain these settings if the VM is migrated to a new cluster. Attempting to power on these VMs usually failed. Such VMs are now tagged as not CBR capable (Capable of Backup and Recovery), are not protected, and not migrated as part of protection domain operations. A web console message is displayed when selecting this VM type: **This**



VM is not capable of backup and recovery. Possible reason: VM-Host affinity is set.

You can define the VM-host affinity policies by using Prism Element during the VM create or update operation. For more information, see [Creating a VM \(AHV\)](#) on page 268.

VM-VM Anti-Affinity Policy

You can use this policy to specify anti-affinity between the virtual machines. The VM-VM anti-affinity policy keeps the specified virtual machines apart in such a way that when a problem occurs with one host, you should not lose both the virtual machines. However, this is a preferential policy. This policy does not limit the Acropolis Dynamic Scheduling (ADS) feature to take necessary action in case of resource constraints.

Note:

- Currently, you can only define VM-VM anti-affinity policy by using aCLI. For more information, see [Configuring VM-VM Anti-Affinity Policy](#) on page 309.
- The VM-VM affinity policy is not supported.

Note: If a VM is cloned that has the affinity policies configured, then the policies are not automatically applied to the cloned VM. However, if a VM is restored from a DR snapshot, the policies are automatically applied to the VM.

Limitations of Affinity Rules

- Even though if a host is removed from a cluster, the host UUID is not removed from the host-affinity list for a VM.
- The VM-host affinity cannot be configured on a cluster that has the HA configured by using reserved host method.
- You cannot remove the VM-host affinity for a powered on VM from Prism. You can use the `vm.affinity_unset vm_list` acli command to perform this operation.

Configuring VM-VM Anti-Affinity Policy

Perform the following procedure to configure the VM-VM anti-affinity policy.

About this task

To configure VM-VM anti-affinity policies, you must first define a group and then add all the VMs on which you want to define VM-VM anti-affinity policy.

Note: Currently, the VM-VM affinity policy is not supported.

Procedure

1. Log in to the Controller VM in your cluster through an SSH session and access the Acropolis command line.
2. Create a group.

```
accli> vm_group.create group_name
```

Replace `group_name` with the name of the group.



3. Add the VMs on which you want to define anti-affinity to the group.

```
accli> vm_group.add_vms group_name vm_list=vm_name
```

Replace `group_name` with the name of the group. Replace `vm_name` with the name of the VMs that you want to define anti-affinity on.

4. Configure VM-VM anti-affinity policy.

```
accli> vm_group.antiaffinity_set group_name
```

Replace `group_name` with the name of the group.

After you configure the group and then power on the VMs, the VMs that are part of the group are started (attempt to start) on the different hosts. However, this is a preferential policy. This policy does not limit the Acropolis Dynamic Scheduling (ADS) feature to take necessary action in case of resource constraints.

Removing VM-VM Anti-Affinity Policy

Perform the following procedure to remove the VM-VM anti-affinity policy.

Procedure

1. Log in to the Controller VM in your cluster through an SSH session and access the Acropolis command line.
2. Remove the VM-VM anti-affinity policy.

```
accli> vm_group.antiaffinity_unset group_name
```

Replace `group_name` with the name of the group.

The VM-VM anti-affinity policy is removed for the VMs that are present in the group, and they can start on any host during the next power on operation (as necessitated by the ADS feature).

Connect to Citrix Cloud

The **Connect to Citrix Cloud** feature leverages the automated installer functionality of the Citrix Cloud connector to establish a secure communication channel between Nutanix and Citrix Cloud. This feature provides single-click integration of the on-premise Nutanix clusters as a resource location with the Citrix Cloud environment.

Note: This feature is supported on AHV only.

This feature helps you configure the Citrix Cloud integration settings in the following way:

1. Establishes the connection to the Citrix Cloud workspace.
2. Configures Nutanix cluster as resource location in the Citrix cloud.
3. Configures the Citrix Cloud connector VM.
4. Registers the Citrix Cloud connector VM to the Active Directory (AD) domain on the Citrix cloud.

Once the integration is complete, VDIs can be created using the *Nutanix AHV MCS Plug-in for Citrix XenDesktop 1.1.1.0* or later. The AHV MCS Plug-in is designed to create and manage VDIs in a Nutanix Acropolis infrastructure environment. See the *Nutanix AHV MCS Plug-in for Citrix XenDesktop 7.9 or later install guide and release notes* for details.

Thus, to begin deploying your VMs and applications, you must perform the following:

- Configure the Citrix cloud integration settings using the **Connect to Citrix Cloud** feature.



- Install the *Nutanix CWA Plug-In for Citrix Cloud Connector*.

Refer to the *Citrix documentation* for details on Citrix Cloud connector.

Connecting to the Citrix Cloud

To integrate with the Citrix Cloud workspace, perform the following procedure.

Before you begin

Ensure that you have created a Sysprep (System Preparation) VM with Windows 2012 R2 server as the base image and the VM should be available on the Nutanix cluster.

CAUTION: The Sysprep VM must be in the powered-off state. If you power on the VM, you will lose the Sysprep state and the configuration will fail.

See *Microsoft documentation on Sysprep (Generalize) a Windows installation* for more information.

Procedure

1. Log on to the Prism Web Console using your Nutanix administrator credentials.

Note: The procedure does not work if Prism Element is launched from Prism Central web console.

2. Click the gear icon and select **Connect to Citrix Cloud**. Click the gear icon in the main menu and then select **Connect to Citrix Cloud** in the **Settings** page.




Figure 205: Connect to Citrix Cloud

3. Authenticate your connection to the Citrix cloud by using one of the following methods.

- **Enter Manually**

1. Enter your **Customer ID** for the Citrix Cloud.
2. Enter your secure **Client ID** for the Citrix Cloud.
3. Enter the downloaded secure client **Secret Key**.

Note: You can find the **Customer ID**, **Client ID**, and **Secret Key** from the *API Access* page in the Citrix Cloud console.

4. Click **Connect**.

- **Upload Credential Key**

1. Click **Upload Key File** to browse and select the key file in the CSV format.

You can create or download the key file from the *API Access* page in the Citrix Cloud console (within *Identity and Access Management*). This key file is used for the Citrix Cloud connector installation.

2. Enter your **Customer ID** for the Citrix Cloud.
3. Click **Connect**.

4. Select the **Resource Location** from the drop-down menu.

The drop-down menu displays the resource locations that are created on the Citrix Cloud. If you have not previously created any resource locations, enter a name for the resource location that you want to create in the **Resource location** field.

Connect to Citrix Cloud

?

✕

Citrix Cloud needs a connector agent to communicate with the Nutanix cluster. Select the Resource Location and the Connector VM details.

Citrix Cloud Details

Change

Client ID

Customer ID

Nutanix

☒ Ensure High Availability

On enabling, two connector nodes will be created for redundancy.

RESOURCE LOCATION

nutanix_test_1

VM MASTER IMAGE

CITRIXSYSP

CONNECTOR VM NAME

connector-1

SECONDARY CONNECTOR VM NAME (FOR HIGH AVAILABILITY)

connector-2

CONNECTOR VM PASSWORD

.....

Show

Domain Credentials

Enter the domain credentials to join the connector VMs to the domain.

DOMAIN USERNAME ?

user

DOMAIN PASSWORD

.....

Show

MACHINE DOMAIN NAME

nutanix.com

DNS GATEWAY ?

4.2.2.2

*Optional

Cancel

Save & Complete

Figure 206: Configuring Resource Location and Connector VM

5. (Optional) Review the **Citrix Cloud** details. If any change is required, click the **Change** hyperlink to edit the connection details.
6. (Optional) Select the **High Availability** check-box to enable or disable high availability for the connector nodes.
By default, high availability is enabled. On enabling high availability, two connector nodes are created for the redundancy purpose.

Note: Citrix Cloud recommends that you install two connectors for redundancy and high availability.

7. In the **VM Master Image** search box, start typing the initial letters of the previously created Sysprep VM image and select the auto-completed option.
8. Enter the **Connector VM Name**.
9. If high availability is enabled, enter the **Secondary Connector VM Name (for high availability)**.
10. Enter the **Connector VM Password**.
11. Enter the **Domain Credentials** to join your enterprise domain to the resource location.
 - a. Enter the **Domain Username** in the format `USERDOMAIN\USERNAME`.
 - b. Enter the **Domain Password**.
 - c. Enter the **Machine Domain Name**.
 - d. Enter the **DNS Gateway**.
12. Click **Save and Continue**.
The connector VM or VMs are configured successfully.

What to do next

- Verify the connection status, see [Viewing Citrix Connection Status](#) on page 315.
- Verify if the connector VM or VMs are created from the **VM Dashboard**.
- Wait for the state of the connector VM or VMs to change from powered off to powered on and then click **Launch Console**. The VM preparation process starts.
- Once the VM starts, install the *Nutanix CWA Plug-In for Citrix Cloud Connector* to start the application deployment.

Viewing Citrix Connection Status

You can view the current status of your Citrix Cloud connection. This helps to verify the registration of the connector VM to the Citrix cloud.

To view details of the Citrix connection, click the gear icon and select **Connect to Citrix Cloud**. If the connector VM is configured, the connector VM name and the status are listed under Connector VMs. It might take some time for the connection status to refresh.

The Connection Status shows the following statuses:

- **Connected** - The connector VM is successfully registered with Citrix cloud.
- **Connection in progress** - The connector VM is registration with Citrix cloud is in progress.



- Not Connected - The connector VM is not registered with the Citrix cloud or the Citrix registration has failed. In this case, wait for the connection status to refresh, it may take several minutes depending on your environment. Otherwise, click **Delete Connection** to delete the Citrix cloud connection and configure it again. The **Delete Connection** option deletes the connector VM and unregisters the account from the Citrix Cloud.

Note: For XenServer, the **Delete Connection** option does not delete the connector VMs. You must delete the connector VMs manually.

PERFORMANCE MONITORING

Nutanix provides several mechanisms to maximize performance in the cluster. The converged Distributed Storage Fabric (DSF) architecture is designed to service the VM requests locally on each node whenever possible. Each node employs data tiers so that frequently accessed ("hot") data is retained in memory or solid state disk (SSD) storage while seldom accessed ("cold") data is moved to hard disk drive (HDD) storage. Each Controller VM has an in-memory read cache to access highly requested data directly from memory.

- The web console allows you to monitor and analyze performance across the cluster (see [Analysis Dashboard](#) on page 317).

Analysis Dashboard

The Analysis dashboard allows you to create charts that can monitor dynamically a variety of performance measures. To view the Analysis dashboard, select **Analysis** from the pull-down list on the left of the main menu.

Menu Options

The Analysis dashboard does not include menu options other than those available from the main menu (see [Main Menu](#) on page 46).

Analysis Screen Details

The Analysis dashboard includes three sections.

- Chart definitions. The pane on the left lists the charts that can be run. No charts are provided by default, but you can create any number of charts. A chart defines the metrics to monitor. There are two types of charts, metric and entity. A metric chart monitors a single metric for one or more entities. An entity chart monitors one or more metrics for a single entity.

Note: You can change the color assigned to a metric or entity by clicking that color box in the chart (left pane) and then selecting a different color from the displayed palette.

- Chart monitors. When a chart definition is checked, the monitor appears in the middle pane. An Alerts & Events monitor always appears first. The remaining monitors are determined by which charts are checked in the left pane. You can customize the display by selecting a time interval (from 3 hours to a month) from the **Range** drop-down (above the charts) and then refining the monitored period by moving the time interval end points to the desired length.
- Alerts and events. Any alerts and events that occur during the interval specified by the time line in the middle pane appear in the pane on the right.

The following figure is a sample view, and the table describes each field in this view. Some fields can include a slide bar on the right to view additional information in that field. The displayed information is dynamically updated to remain current.

Note: See [Understanding Displayed Statistics](#) on page 56 for information about how the metrics are measured.






Figure 207: Analysis Dashboard

Table 36: Analysis Screen Fields

Name	Description
Charts	Displays the set of defined charts. Check the box next to a chart name to run that chart in the middle pane. The chart monitor appears in the middle pane shortly after checking the box. Uncheck the box to stop that monitor and remove it from the middle pane. To edit a chart definition, click the pencil icon to the right of the name. This opens the edit chart window, which is the same as the new chart window except for the title. To delete a chart, click the cross icon on the right.
New Metric Chart	Allows you to create a chart that tracks a single metric for one or more entities (see Creating a Metric Chart on page 320).
New Entity Chart	Allows you to create a chart that tracks one or more metrics for a single entity (see Creating an Entity Chart on page 319).
(range time line and monitor period)	Displays a time line that sets the duration for the monitor displays. To set the time interval, select the time period (3 hour, 6 hour, 1 day, 1 week, WTD [week to date], 1 month) from the Range field pull-down menu (far right of time line). To customize the monitor period, drag the time line end points to the desired times on the time line.

Name	Description
Alerts & Events Monitor	Displays a monitor of alert and event messages that were generated during the time interval. Alerts and events are tracked by a moving histogram with each bar indicating the number of messages generated during that time. The message types are color coded in the histogram bars (critical alert = red, warning alert = orange, informational alert = blue, event = gray).
(defined chart monitors)	Displays monitors for any enabled (checked) charts. In the figure above, three charts are enabled (memory usage, CPU/memory, and disk IOPS). You can export the chart data by clicking on the chart header. This displays a drop-down menu (below) to save the data in CSV or JSON format. It also includes a chart link option that displays the URL to that chart, which you can copy to a clipboard and use to import the chart.  Export Chart Data (CSV)...  Export Chart Data (JSON)...  Chart Link...
Alerts	Displays the alert messages that occurred during the time interval (see Alerts Dashboard). Clicking a message causes the monitor line to move to the time when that alert occurred.
Events	Displays the event messages that occurred during the time interval. Clicking a message causes the monitor line to move to the time when that event occurred.

Creating an Entity Chart

About this task

An entity chart monitors the performance of one or more metrics for a single entity. To create an entity chart definition, do the following:

Procedure

1. In the Analysis dashboard (see [Analysis Dashboard](#) on page 317), click **New > New Entity Chart**.

The New Entity Chart dialog box appears.

2. Do the following in the indicated fields:

a. **Chart Title:** Enter a title for this chart.

b. **Entity type:** Select an entity from the pull-down list.

The entity types include host, disk, storage pool, storage container, virtual machine, volume group, remote site, protection domain, replication link, virtual disk, and cluster.

c. **Entity:** Enter the name of the target entity.

As you enter characters in this field, it displays a list of matching entries of that entity type. Click the name when it appears in the search list.


Note: If you are creating this chart for Prism Central, the list spans the registered clusters. Otherwise, the list is limited to the current cluster.

- d. **Metric:** Select a metric from the pull-down list. (Repeat to include additional metrics.)
For descriptions of the available metrics, see [Chart Metrics](#) on page 321.



3. When all the field entries are correct, click the **Save** button.

The Analysis dashboard reappears with the new chart appearing in the list of charts on the left of the screen.



The image shows a 'New Entity Chart' dialog box. It has a dark header bar with the title 'New Entity Chart' and a close button (X). Below the header, there are four input fields: 'CHART TITLE' (a text box), 'ENTITY TYPE' (a dropdown menu with 'Host' selected), 'ENTITY' (a search box with a magnifying glass icon), and 'METRIC' (a dropdown menu with 'Select a metric' selected). At the bottom right of the dialog, there are two buttons: 'Cancel' and 'Save'.

Figure 208: New Entity Chart Window

Creating a Metric Chart

About this task

A metric chart monitors the performance of a single metric for one or more entities. To create a metric chart definition, do the following:

Procedure

1. In the Analysis dashboard (see [Analysis Dashboard](#) on page 317), click **New > New Metric Chart**.
The New Metric Chart dialog box appears.

2. Do the following in the indicated fields:

a. **Chart Title:** Enter a title for this chart.

b. **Metric:** Select a metric to monitor from the pull-down list.

For descriptions of the available metrics, see [Chart Metrics](#) on page 321.

c. **Entity Type:** Select an entity type from the pull-down list. (Repeat to include additional entities.)

The entity types include host and cluster.

d. **Entity:** Enter the name of the target entity.

As you enter characters in this field, it displays a list of matches of the entity type. Click the name when it appears in the search list. (Repeat to include additional names.)

Note: If you are creating this chart for Prism Central the list spans the registered clusters. Otherwise, the list is limited to the current cluster.

3. When all the field entries are correct, click the **Save** button.

The Analysis dashboard reappears with the new chart appearing in the list of charts on the left of the screen.




Figure 209: New Metric Chart Window

Chart Metrics

These metrics can be added to charts.

Metric	Description
Content Cache Hit Rate (%)	Content cache hits over all lookups. ID: <code>CONTENT_CACHE_HIT_PPM</code>
Content Cache Hits	Number of hits on the content cache. ID: <code>CONTENT_CACHE_NUM_HITS</code>
Content Cache Logical Memory Usage	Logical memory (in bytes) used to cache data without deduplication. ID: <code>CONTENT_CACHE_LOGICAL_MEMORY_USAGE_BYTES</code>
Content Cache Logical SSD Usage	Logical SSD memory (in bytes) used to cache data without deduplication. ID: <code>CONTENT_CACHE_LOGICAL_SSD_USAGE_BYTES</code>
Content Cache Lookups	Number of lookups on the content cache. ID: <code>CONTENT_CACHE_NUM_LOOKUPS</code>
Content Cache Memory Saved	Memory (in bytes) saved due to content cache deduplication. ID: <code>CONTENT_CACHE_SAVED_MEMORY_USAGE_BYTES</code>
Content Cache Physical Memory Usage	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	Average number of content cache references. ID: <code>CONTENT_CACHE_NUM_DEDUP_REF_COUNT_PPH</code>
Content Cache SSD Usage	Real SSD usage (in bytes) used to cache data by the content cache. ID: <code>CONTENT_CACHE_PHYSICAL_SSD_USAGE_BYTES</code>
Content Cache SSD Usage Saved	SSD usage (in bytes) saved due to content cache deduplication. ID: <code>CONTENT_CACHE_SAVED_SSD_USAGE_BYTES</code>
Deduplication Fingerprints Cleared	Number of written bytes for which fingerprints have been cleared. ID: <code>DEDUP_FINGERPRINT_CLEARED_BYTES</code>
Deduplication Fingerprints Written	Number of written bytes for which fingerprints have been added. ID: <code>DEDUP_FINGERPRINT_ADDED_BYTES</code>
Disk I/O Bandwidth	Data transferred per second in KB/second from disk. ID: <code>STATS_BANDWIDTH</code>



Metric	Description
Disk I/O Bandwidth - Read	Read data transferred per second in KB/second from disk. ID: <code>STATS_READ_BANDWIDTH</code>
Disk I/O Bandwidth - Write	Write data transferred per second in KB/second from disk. ID: <code>STATS_WRITE_BANDWIDTH</code>
Disk I/O Latency	I/O latency in milliseconds from disk. ID: <code>STATS_AVG_IO_LATENCY</code>
Disk IOPS	Input/Output operations per second from disk. ID: <code>STATS_NUM_IOPS</code>
Disk IOPS - Read	Input/Output read operations per second from disk. ID: <code>STATS_NUM_READ_IOPS</code>
Disk IOPS - Write	Input/Output write operations per second from disk. ID: <code>STATS_NUM_WRITE_IOPS</code>
Hypervisor CPU Usage (%)	Percent of CPU used by the hypervisor. ID: <code>STATS_HYP_CPU_USAGE</code>
Hypervisor Memory Usage (%)	Percent of memory used by the hypervisor. ID: <code>STATS_HYP_MEMORY_USAGE</code>
Hypervisor CPU Ready Time (%)	CPU ready time percentage. ID: <code>STATS_HYP_CPU_READY_TIME</code>
Hypervisor I/O Bandwidth	Data transferred per second in KB/second from Hypervisor. ID: <code>STATS_HYP_BANDWIDTH</code>
Hypervisor I/O Bandwidth - Read	Read data transferred per second in KB/second from Hypervisor. ID: <code>STATS_HYP_READ_BANDWIDTH</code>
Hypervisor I/O Bandwidth - Write	Write data transferred per second in KB/second from Hypervisor. ID: <code>STATS_HYP_WRITE_BANDWIDTH</code>
Hypervisor I/O Latency	I/O latency in milliseconds from Hypervisor. ID: <code>STATS_HYP_AVG_IO_LATENCY</code>
Hypervisor I/O Latency - Read	I/O read latency in milliseconds from Hypervisor. ID: <code>STATS_HYP_AVG_READ_IO_LATENCY</code>
Hypervisor I/O Latency - Write	I/O write latency in milliseconds from Hypervisor. ID: <code>STATS_HYP_AVG_WRITE_IO_LATENCY</code>

Metric	Description
Hypervisor IOPS	Input/Output operations per second from Hypervisor. ID: <code>STATS_HYP_NUM_IOPS</code>
Hypervisor IOPS - Read	Input/Output read operations per second from Hypervisor. ID: <code>STATS_HYP_NUM_READ_IOPS</code>
Hypervisor IOPS - Write	Input/Output write operations per second from Hypervisor. ID: <code>STATS_HYP_NUM_WRITE_IOPS</code>
Logical Usage	Logical usage of storage (physical usage divided by replication factor). ID: <code>STATS_UNTRANSFORMED_USAGE</code>
Physical Usage	Actual usage of storage. ID: <code>STATS_TRANSFORMED_USAGE</code>
Read IOPS (%)	Percent of IOPS that are reads. ID: <code>STATS_READ_IO_PPM</code>
Replication Bandwidth - Received	Replication data received per second in KB/second ID: <code>STATS_REP_BW_RECEIVED</code>
Replication Bandwidth - Transmitted	Replication data transferred per second in KB/second ID: <code>STATS_REP_BW_TRANSFERRED</code>
Replication Bytes - Received	Number of bytes received. ID: <code>STATS_REP_NUM_RECEIVED_BYTES</code>
Replication Bytes - Total Received	Total number of bytes received. ID: <code>STATS_REP_TOT_RECEIVED_BYTES</code>
Replication Bytes - Total Transmitted	Total number of bytes transmitted. ID: <code>STATS_REP_TOT_TRANSMITTED_BYTES</code>
Replication Bytes - Transmitted	Number of bytes transmitted. ID: <code>STATS_REP_NUM_TRANSMITTED_BYTES</code>
Storage Controller Bandwidth	Data transferred in KB/second from the Storage Controller. ID: <code>STATS_CONTROLLER_BANDWIDTH</code>
Storage Controller Bandwidth - Read	Read data transferred in KB/second from the Storage Controller. ID: <code>STATS_CONTROLLER_READ_BANDWIDTH</code>
Storage Controller Bandwidth - Write	Write data transferred in KB/second from the Storage Controller. ID: <code>STATS_CONTROLLER_WRITE_BANDWIDTH</code>



Metric	Description
Storage Controller IOPS	Input/Output operations per second from the Storage Controller ID: <code>STATS_CONTROLLER_NUM_IOPS</code>
Storage Controller IOPS - Read	Input/Output read operations per second from the Storage Controller ID: <code>STATS_CONTROLLER_NUM_READ_IOPS</code>
Storage Controller IOPS - Read (%)	Percent of Storage Controller IOPS that are reads. ID: <code>STATS_CONTROLLER_READ_IO_PPM</code>
Storage Controller IOPS - Write	Input/Output write operations per second from the Storage Controller ID: <code>STATS_CONTROLLER_NUM_WRITE_IOPS</code>
Storage Controller IOPS - Write (%)	Percent of Storage Controller IOPS that are writes. ID: <code>STATS_CONTROLLER_WRITE_IO_PPM</code>
Storage Controller Latency	I/O latency in milliseconds from the Storage Controller. ID: <code>STATS_CONTROLLER_AVG_IO_LATENCY</code>
Storage Controller Latency - Read	Storage Controller read latency in milliseconds. ID: <code>STATS_CONTROLLER_AVG_READ_IO_LATENCY</code>
Storage Controller Latency - Write	Storage Controller write latency in milliseconds. ID: <code>STATS_CONTROLLER_AVG_WRITE_IO_LATENCY</code>
Storage container own usage	Storage container's own usage + Reserved (not used). ID: <code>NEW_CONTAINER_OWN_USAGE_LOGICAL</code>
Swap In Rate	Rate of data being swapped in. ID: <code>STATS_HYP_SWAP_IN_RATE</code>
Swap Out Rate	Rate of data being swapped out. ID: <code>STATS_HYP_SWAP_OUT_RATE</code>
Virtual NIC bytes received packets with error	Virtual NIC bytes received packets with error. ID: <code>STATS_NETWORK_ERROR_RECEIVED_PACKETS</code>
Virtual NIC bytes received rate	Virtual NIC bytes received rate in kbps. ID: <code>STATS_NETWORK_RECEIVED_RATE</code>
Virtual NIC bytes transmitted rate	Virtual NIC bytes transmitted rate in kbps. ID: <code>STATS_NETWORK_TRANSMITTED_RATE</code>
Virtual NIC dropped transmitted packets	Number of dropped transmitted packets by the VNIC. ID: <code>STATS_NETWORK_DROPPED_TRANSMITTED_PACKETS</code>



Metric	Description
Virtual NIC received packets dropped	Number of receive packets dropped by the VNIC. ID: <code>STATS_NETWORK_DROPPED_RECEIVED_PACKETS</code>
Write IOPS (%)	Percent of IOPS that are writes. ID: <code>STATS_WRITE_IO_PPM</code>

Exporting Performance Data

This topic describes how to view and export the performance chart.

About this task

Performance data is stored for three months in the UI. You can export the performance data results in CSV or JSON format.

Note: Performance data is available for 3 months.

Procedure

1. Log into the Prism Web Console.
2. Navigate to the Analysis Dashboard. Click **Home** > **Analysis**.
The Analysis page displays.
3. Set the range to **1 Month**. The **1 Month** range shows the data in monthly segments.




Figure 210: Set Performance Data Range

4. Export the performance data into a CSV or JSON file. Click the drop-down arrow next to the cluster chart you want to export.



Figure 211: Export Performance Data

ALERT AND EVENT MONITORING

The Prism web console provides several mechanisms to monitor and configure cluster alerts and events.

- The [Alerts Dashboard](#) on page 327 allow you to view and monitor alert and event activity in the cluster.
- The [Alerts Summary View](#) on page 327 displays a list of alerts that you filter in various ways. You can also drill down for detailed information and corrective actions for an alert.
- The [Events Summary View](#) on page 333 displays a list of events that you filter in various ways. You can also drill down for detailed information for an event.
- You can configure rules for who should receive email alerts and customize the alert messages (see [Configuring Alert Emails](#) on page 335).
- You can customize the list of events that generate an alert (see [Configuring Alert Policies](#) on page 338).

When alerting is enabled, some potentially high impact events also generate a case automatically with Nutanix customer support. Events that might generate a Nutanix support case include hardware problems such as clock, power supply, or disk failures; monitoring problems such as Stargate outages or Curator scan failures; and data processing problems such as disabled compression or deduplication.

Note: These views reflect that Prism retains alerts and events for 90 days and hourly data for a year by default. (Some DR events are retained for just 7 days by default.)

Alerts Dashboard

The Alerts dashboard displays alert and event messages. To view the Alerts dashboard, select **Alerts** from the pull-down list on the far left of the main menu.

The Alerts dashboard provides two viewing modes.

- Click the **Alerts** tab (on the left) to display the alert messages (see [Alerts Summary View](#) on page 327).
- Click the **Events** tab to display the event messages (see [Events Summary View](#) on page 333).

Alerts Summary View

The Alerts summary view displays a list of alert messages. From this view you can customize the order and scope of listed alerts, display the details for a specific alert, and configure alert policies. The following figure is a sample view, and the table describes each field in this view.

Note: See [Alerts/Health checks](#) on page 340 for a list of alert messages and corrective actions.






Figure 212: Alerts View

- To order the alerts chronologically, click the **Create Time** column header, which toggles the alert list latest-to-earliest or earliest-to-latest.
- To download the table of alerts in CSV format, click the download icon.
- To display the details page for an alert, click the alert title (see [Alert Details](#) on page 330).
- To configure alert policies, do one or both of the following:
 - Select the **Configure > Alert Policy** button to specify what events should generate an alert (see [Configuring Alert Policies](#) on page 338).
 - Select the **Configure > Email Configuration** button to enable alert emails and specify email addresses to which alerts should be sent (see [Configuring Alert Emails](#) on page 335).
- To filter the list, click the **Filters** button (upper right). This displays a pane for selecting filter values. Check the box for each value to include in the filter. You can include multiple values. You can filter the search on the following alert parameters and values.
 - Severity: Critical, Warning, Info
 - Resolved: Yes (Manual), Yes (Auto), No
 - Acknowledged: Yes, No
 - Impact Type: Availability, Capacity, Configuration, Performance, System Indicator
 - Create Time: Last 1 hour, Last 24 hours, Last week, From **XXX** to **XXX**
- To select (or unselect) alerts in the list, do one of the following:
 - Click the check box on the far left for each alert to select. (Click the check box again to unselect that alert.)
 - Click the check box in the header row to select (unselect) all the alerts in the list.
 - In the check box pull-down menu (far left above the list of alerts), select **Select all** (or **Clear selections**) to select (unselect) all the alerts in the list.

- When one or more alerts are selected, **Resolve** and **Acknowledge** appear in the menu; click one or both to resolve and/or acknowledge the selected alerts. The user name and time appears in the appropriate field when an alert is acknowledged or resolved.

Table 37: Alerts View Fields

Parameter	Description	Values
(selection box)	Click this box to select the message for acknowledgement or resolution.	n/a
Title	Displays the title of the alert. Click the title to displays the details page for that alert (see Alert Details on page 330).	(alert title)
Source Entity	Displays the entity (such as a cluster, host, or VM name) to which the alert applies. A comma separate list appears if it applies to multiple entities. If there is an associated details page, the entity is a live link; clicking the link displays the details page.	(entity names)
Impact Type	Displays the category in which the alert is classified.	Availability, Capacity, Configuration, Performance, System Indicator
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 may have the potential to stop running, or it could run into irreparable issues. Warning An actionable issue has been detected, and user intervention is required. A more serious issue may develop if this is not resolved soon. Informational (Info) An actionable minor problem has been detected. It should be resolved relatively soon and not ignored.	Critical, Warning, Informational

Parameter	Description	Values
Resolved	Indicates whether the alert has been resolved. Resolving an error means you set that error as fixed. (The alert may 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 (see Configuring Alert Policies on page 338). - A blank value means the alert is not resolved. - An "Auto (<i>date_time</i>)" entry indicates this alert was resolved automatically at the specified date and time. This requires that auto alert is enabled for this alert type (see Configuring Alert Policies on page 338). - A "By <i>user</i> (<i>date_time</i>)" entry indicates this alert was resolved manually by the specified user at that date and time. To manually resolve an alert, check the box in the far left column and then select the Resolve button (left of screen above the search field). The field value changes from blank to the user name and date/time the alert was resolved.	(blank, Auto (<i>date_time</i>), By <i>user</i> (<i>date_time</i>))
Acknowledged	Indicates whether the alert has been acknowledged. Acknowledging an alert means you recognize the error exists (no more reminders for this condition), but the alert status remains. The options (blank, auto, or user) are the same as for Resolved. To manually acknowledge an alert, check the box in the far left column and then click the Acknowledge button (left of screen above the search field). The field value changes from blank to the user name and date/time the alert was acknowledged.	(blank, Auto (<i>date_time</i>), By <i>user</i> (<i>date_time</i>))
Create Time	Displays the date and time when the alert occurred.	(date and time)

Alert Details

Clicking on an alert message in the summary view (see [Alerts Summary View](#) on page 327) or anywhere else the alert title appears, such as in a search list, displays detailed information about that alert. The alert details appear in the left column. Possible causes for the alert appear to the right. The most likely cause appears first with other possible causes (if any) appearing below in the order of likelihood. Each cause includes a Recommendation section that describes the recommended corrective action and in many cases a Details section that provides additional context and instructions.






Figure 213: Alert Details Screen

- To switch among alert details pages, select the desired alert from the pull-down list in the upper left of the screen.
- To acknowledge or resolve the alert, click the **Resolve** or **Acknowledge** button in the upper right.
- To open a help page in a separate tab or window, click the question mark icon.
- To close the details screen, click the X icon in the upper right.

Table 38: Alert Details Fields

Parameter	Description	Values
Alert Title	Displays the title of the alert.	(alert title)
Source Entity	Displays a list of source entities or a dash if no specific entity is the source of the alert. Clicking an entity link opens the detail page for that entity.	(entity names)

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 may have the potential to stop running, or it could run into irreparable issues. Warning An actionable issue has been detected, and user intervention is required. A more serious issue may develop if this is not resolved soon. Informational (Info) An actionable minor problem has been detected. It should be resolved relatively soon and not ignored.	Critical, Warning, Informational
Created Time	Displays the date and time when the alert first occurred.	(time and date)
Last Occurred	Displays the date and time for the most recent occurrence of the alert.	(time and date)
Impact Type	Displays the impact type to which the alert is classified.	Availability, Capacity, Configuration, Performance, System Indicator
Policy	Displays the name of the alert policy.	(alert policy name)
Status (for manually resolved alerts)	Displays the current alert status: acknowledged, resolved, or neither (blank).	(blank), Acknowledged, Resolved
Status (for auto resolved alerts)	Displays the current alert status: acknowledged, resolved, auto resolved, or neither (blank).	(blank), Acknowledged, Auto Resolved

Parameter	Description	Values
Acknowledged By	Indicates whether the alert has been acknowledged. - A blank (dash) value means the alert is not acknowledged. - An "N/A <code>date_time</code>" entry indicates this alert was acknowledged automatically at the specified date and time. This requires that auto alert is enabled for this alert type (see Configuring Alert Policies on page 338). - A "<code>user date_time</code>" entry indicates this alert was acknowledged manually by the specified user at that date and time. To manually acknowledge an alert, click the Acknowledge button (upper right). The field value changes from blank to the user name and date/time the alert was acknowledged.	(blank), N/A <code>date_time</code> , <code>user date_time</code>
Resolved By	Indicates whether the alert has been resolved. The options (blank, N/A, or user) are the same as for Acknowledged. To manually resolve an alert, click the Resolve button (upper right). The field value changes from blank to the user name and date/time the alert was resolved.	(dash), N/A <code>date_time</code> , <code>user date_time</code>

Events Summary View

The Events summary view displays a list of event messages. Event messages describe cluster actions such as adding a storage pool or taking a snapshot. Unlike alerts, event messages are simply informational without the need to acknowledge or resolve. The following figure is a sample view, and the table describes each field in this view.

The screenshot shows the Prism interface with the 'Events' tab selected. The main content area displays a table of 29 events. The table has the following columns: Title, Source Entity, Event Type, and Create Time. The events listed are all 'User Action' events related to logging in and out from 10.149.76. Annotations highlight key features: 'display filter tab' points to the 'Filters' button in the top right; 'download table' points to the download icon; 'click to order by time' points to the sort icon; and 'click title for details' points to an event title.

Title	Source Entity	Event Type	Create Time
User admin has logged in from 10.149.76	Cluster	User Action	02/07/20, 03:39:21 PM
User admin has logged out from 10.149.76	Cluster	User Action	02/07/20, 02:45:02 ...
User admin has logged in from 10.149.76	Cluster	User Action	02/07/20, 02:26:37 P...
User admin has logged out from 10.149.76	Cluster	User Action	02/07/20, 02:17:31 PM
User admin has logged in from 10.149.76	Cluster	User Action	02/07/20, 01:55:08 PM

Figure 214: Events View

- To display the details page for an event, click the event title (see [Event Details](#) on page 334).
- To order the events chronologically, click the **Create Time** column header, which toggles the event list latest-to-earliest or earliest-to-latest.
- To filter the list, click the **Filters** button (upper right). This displays a pane for selecting filter values. Check the box for each value to include in the filter. You can include multiple values. You can filter the search on the following event parameters and values.
 - Event Type: Behavioral Anomaly, System Action, User Action
 - Time Range: Last 1 hour, Last 24 hours, Last week, From **xxx** to **xxx**

Table 39: Events View Fields

Parameter	Description	Values
Title	Displays the name of the event. Click the name to displays the details page for that event (see Event Details on page 334).	(event name)
Source Entity	Displays the entity (such as a cluster, host, or VM name) to which the event applies. A comma separated list appears if it applies to multiple entities. If there is an associated details page, the entity is a live link; clicking the link displays the details page.	(entity names)
Event Type	Displays the type of event. Three types are defined, and each event is assigned to one of these types.	Behavioral Anomaly, System Action, User Action
Create Time	Displays the date and time when the event occurred.	(time and date)

Event Details

Clicking on an event message in the summary view (see [Events Summary View](#) on page 333) displays detailed information about that event. The event details appear in the left column. Additional information, such as a description or performance graph, appears to the right.




Figure 215: Event Details Screen

- To switch among event details pages, select the desired event from the pull-down list in the upper left of the screen.
- To open a help page in a separate tab or window, click the question mark icon.
- To close the details screen, click the X icon in the upper right.

Table 40: Event Details Fields

Parameter	Description	Values
Event Title	Displays the title of the event.	(alert title)
Source Entity	Displays a list of source entities or a dash if no specific entity is the source of the event. Clicking an entity link opens the detail page for that entity.	(entity names)
Created Time	Displays the date and time when the event occurred.	(time and date)
Event Type	Displays the type of event. Three types are defined, and each event is assigned to one of these types.	Behavioral Anomaly, System Action, User Action

Configuring Alert Emails

About this task

E-mail notification of alerts is enabled by default and sends alert messages automatically to Nutanix customer support through customer-opened ports 80 or 8443. To automatically receive email notification alerts, ensure that **nos-alerts** and **nos-asup** recipients are added to the accepted domain of your SMTP server. To customize who should receive the alert e-mails (or to disable e-mail notification), do the following:

Note: An alert case is opened automatically with Nutanix customer support for any alert notifications sent to nos-alerts@nutanix.com.

Procedure

1. Do one of the following:
 - » Click the gear icon in the main menu and then select **Alert Email Configuration** in the **Settings** page.
 - » Go to the Alerts dashboard, click the **Configure** button, and select **Email Configuration** from the pull-down menu.

The Alert Email Configuration window appears.



2. Click the **Settings** tab and do the following:
 - a. Select the **Email Every Alert** box to send an email whenever the event occurs.
 - b. Select the **Email Daily Digest** box to send a cumulative (24 hour) list of alerts once a day. You can check one or both boxes. If neither box is checked, no alert emails are sent.
 - c. Select the **Nutanix Support** box to send alert notification to Nutanix customer support. Note that an alert notification to the Nutanix customer support is sent only if you have selected the **Email Every Alert** option.
 - d. To send alert notifications to others, enter their email addresses in a comma separated list in the field below the **Nutanix Support** box.
 - e. Click the **Apply** button to apply the changes and continue the configuration (or click the **Save** button to apply the changes and close the window).

The **Connection Status** section displays mail transport status information. In this example, default Nutanix tunnel is used for mail service. For more information on configuring the SMTP server settings window (see [Configuring an SMTP Server](#) on page 598).




Figure 216: Alert Email Configuration: Settings Tab

3. To create a custom alert email rule, click the **Rules** tab and the **New Rule** button, and then configure the rule as follows:
 - a. Specify the conditions for generating the alert:
 - **Severity**: Select one or more of the severities from the pull-down list (Critical, Warning, Info, All).
 - **Category**: Select one or more of the categories from the pull-down list (Availability, Capacity, Configuration, Performance, System Indicator, All).
 - b. Specify who should receive the alert email.
 - **Email Addresses**: Enter recipient email addresses as a comma separate list in the box.
 - **Global Recipient List**: Click this box to add everyone on the global recipient list. This is in addition to any users listed in the **Email Addresses** field.
 - c. Click the **Apply** button to apply the rule.
 - d. Repeat these steps to apply additional custom rules.

Alert Email Configuration ?

Settings **Rules** Template

When:

Severity Impact Type

Alert Contains

Send to:

Email Addresses

☐ Global Recipient List ?

< Back Apply

Figure 217: Alert Email Configuration: Settings Tab

4. To create a template for the email messages, click the **Template** tab and do the following:
 - a. In the **Prepend content to the email subject** field, enter the desired text.

This text will appear at the beginning of the subject field in each alert email. If the field is left blank, no prepended text will appear in the subject.
 - b. In the **Append content to the email body** field, enter the desired text.

This text will appear at the end of the message body in each alert email. If the field is left blank, no appended text will appear in the message body.
 - c. Click the **Save** button to apply the changes and close the window (or click the **Apply** button to apply the changes and continue the configuration).

The screenshot shows the 'Alert Email Configuration' window with the 'Template' tab selected. The 'Preview' section displays a subject line 'Data Disk Space Usage High' and a body line 'Disk space usage for {mount_path} on {entity} {ip_address} has exceeded {disk_usage_threshold}%'. Below the preview, there are two text input fields: 'Prepend Content To The Email Subject' and 'Append Content To The Email Body'. At the bottom right, there are 'Apply' and 'Save' buttons. The 'Template' tab is highlighted with a red box.

Figure 218: Alert Email Configuration: Settings Tab

Configuring Alert Policies

About this task

The system comes with default alert policies, but you can customize which events should trigger an alert. To customize the alert policies, do the following:

Procedure

1. Do one of the following:
 - » Click the gear icon in the main menu and then select **Alert Policies** in the **Settings** page.
 - » Go to the Alerts dashboard, click the **Configure** button, and select **Alert Policies** from the pull-down menu.

The Alert Policies page appears. This page displays the list of events that can trigger an alert. Each line includes the ID number (Axxxx), title, impact type (system indicator, availability, configuration, and so on), entity type (cluster, host, protection domain, and so on), rule setting (default or overridden), enabled and auto-resolve status (enabled or disabled), and

the last time this alert policy was updated. Enter a string in the search box at the top or use the scroll bar on the right to search through the list.




Figure 219: Alert Policies Page

2. To modify the alert policy for an event, click the title or pencil icon for that event to open the Update Policy window, and do the following in the indicated fields:




Figure 220: Update Policy Window

- a. **Critical, Warning, and/or Info** severity fields: Uncheck (or check) the box next to the severity to disable (or re-enable) this as an alert.

All the alerts are enabled by default (box checked). In most cases this field includes just a single box with the word Critical, Warning, or Info indicating the severity level. Checking the box means this event will trigger an alert of that severity. Unchecking the box means an alert will not be issued when the event occurs. In some cases, such as in the example figure about disk space usage, the event can trigger two alerts, a warning alert when one threshold is reached (in this example 90%) and a critical alert when a second threshold

is reached (95%). In these cases you can specify whether the alert should be triggered (check/uncheck the box) and at what threshold (enter a percentage in the box).

- b. **Auto Resolve These Alerts:** Uncheck (or check) the box to disable (or re-enable) automatic alert resolution.

Automatic alert resolution is enabled for all alert types (where applicable) by default. When this is enabled, the system will automatically resolve alerts under certain conditions such as when the system recognizes that the error has been resolved or when the initiating event has not reoccurred for 48 hours. (Automatic resolution is not allowed for some alert types, and this is noted in the policy window for those types.)

- c. Click the **Save** button to save the changes and return to the Alert Policies page.

3. Repeat the previous step for any additional alert policies that you want to update.

Alerts/Health checks

Availability Zone

Table 41: Unable to Retrieve the Availability Zone Endpoint [130173] [A130173]

Name	Unable to get Availability Zone Endpoint
Description	Unable to retrieve the endpoint information for the Availability Zone
Alert message	Unable to retrieve the endpoint information for the Availability Zone <i>availability_zone_physical_name</i>
Cause	Prism Central, Prism Element, or Xi - Availability Zone may be unreachable
Impact	VMs may not be fully protected
Resolution	Check if the Prism Central, Prism Element, or Xi - Availability Zone are reachable

Table 42: Unable to communicate with the Data Center Manager [130175] [A130175]

Name	Failure to communicate with DCM
Description	Unable to retrieve the endpoint information for the Availability Zone from the Data Center Manager
Alert message	Unable to retrieve the endpoint information for the Availability Zone <i>availability_zone_name</i> from the Data Center Manager
Cause	Data Center Manager may be unreachable or has provided an invalid response
Impact	VMs may not be fully protected
Resolution	Check if Data Center Manager is reachable from the Xi - Availability Zone

Cluster



Table 43: High vDisk count in the cluster [1032] [A1182]

Name	vDisk Count Check
Description	Checks for high vDisk count in the cluster
Alert message	Number of <code>vdisk_type</code> in the cluster is above the threshold (<code>vdisk_count/vdisk_threshold</code>)
Cause	Aggressive replication/snapshot schedules may generate large number of vDisks on the remote site
Impact	The cluster may become unstable due to increased load on core Nutanix services
Resolution	Immediately contact Nutanix support for assistance

Table 44: All flash nodes mixed with non-all-flash nodes [1054] [A1191]

Name	All-flash Node Intermixed Check
Description	Checks that all-flash nodes do not coexist with non-all-flash nodes in a cluster
Alert message	Cluster <code>cluster_id</code> has both all-flash nodes and non-all-flash nodes
Cause	All-flash nodes and non-all-flash nodes are put into the same cluster
Impact	The cluster may experience performance degradation or instability.
Resolution	Remove either all-flash nodes or non-all-flash nodes from the cluster

Table 45: Cluster In Read-Only Mode [1058] [A1195]

Name	Cluster In Read-Only Mode
Description	Single-node cluster is in read-only mode.
Alert message	Single node cluster with CVM <code>service_vm_external_ip</code> is running in read-only mode due to <code>reason</code>.
Cause	One of the metadata SSDs has failed.
Impact	Native backups are disabled.
Resolution	Replace the failed SSD. Refer to the Nutanix documentation for instructions.

Table 46: VM Memory Pressure [3024]

Name	VM Memory Pressure
Description	Checks balloon driver reclaimed memory is within target.
Cause	Ballooning driver could not reclaim enough memory.
Impact	VM and cluster performance may be degraded.
Resolution	Increase memory, reduce memory intensive VMs, or reduce the total number of VMs.



Table 47: CVM Time Difference High [3025] [A1017]

Name	Time Drift Check
Description	Checks whether CVM times match.
Alert message	Time differs between Controller VMs.
Cause	Timedrift between CVMs is high. NTP connectivity may be lost.
Impact	The cluster may experience downtime.
Resolution	Check NTP/network configuration or connectivity.
Time Drift Threshold	Warning Threshold: 3 seconds

Table 48: Upgrade bundle available [3031] [A3031]

Name	Upgrade Bundle Available
Description	Bundle available for upgrade
Alert message	Bundle <i>bundle_name</i> is available for <i>upgrade_type</i> upgrade.
Cause	Newer bundle available for upgrade
Impact	Cluster is running on older version.
Resolution	Upgrade the cluster with newer bundle. Refer to the Nutanix documentation for instructions.

Table 49: Cluster services are down [3032] [A3032]

Name	Cluster Services Down Check
Description	Check that no services are down in the cluster.
Alert message	Cluster Service: <i>service_name</i> is down on the Controller VM <i>ip_address</i>.
Cause	One or more services in the cluster are down.
Impact	Cluster performance may be significantly degraded. In the case of multiple services with the same condition, the cluster may become unable to serve I/O requests.
Resolution	[If the reason for the service(s) being down is unknown, run "ncc health_checks run_all", perform log collection covering the period of this alert timestamp, and provide the results to Nutanix Support via case attachment for further investigation. If due to planned maintenance or triage, please ensure the service is started again using "cluster start" as soon as possible. Report any reoccurring issues.]

Table 50: Alert Manager Service Check [3033]

Name	Alert Manager Service Check
Description	Check if Alert Manager service is available.



Cause	This check usually indicates that the Alert Manager service is not healthy, but there could be other causes.
Impact	Cluster issues may stay unnoticed.
Resolution	If this alert occurs once or infrequently, no action is necessary. If it is frequent, contact Nutanix support.

Table 51: Cluster Service Restarting Frequently [3034] [A3034]

Name	Cluster Services Status
Description	Check if services have restarted recently across the cluster.
Alert message	There have been <code>threshold</code> or more service restarts of <code>service</code> within one day across all Controller VM(s).
Cause	This alert indicates that one or more services in the cluster were restarted.
Impact	Cluster performance may be significantly degraded. In the case of multiple services with the same condition, the cluster may become unable to service I/O requests.
Resolution	If this alert occurs once or infrequently, no action is necessary. If it is frequent, contact Nutanix support.

Table 52: EXT4-fs Error Check [3038]

Name	EXT4-fs Error Check
Description	Captures EXT4-fs error messages
Cause	File system inconsistencies are present on the node.
Impact	Inability for the CVM to boot or for the upgrade pre-checks to run.
Resolution	Look for any problems in the file system. Review KB 8514.

Table 53: Host IP Not Reachable [3065] [A1006]

Name	Host IP Pingable
Description	Check that all host ips are pingable from local SVM.
Alert message	Hypervisor <code>target_ip</code> is not reachable from Controller VM <code>source_ip</code> in the last 6 attempts.
Cause	The hypervisor host is down or there is a network connectivity issue.
Impact	Cluster compute and storage capacity are reduced. Until data stored on this host is replicated to other hosts in the cluster, the cluster has one less copy of guest VM data.
Resolution	Ensure that the hypervisor host is running and that physical networking, VLANs, and virtual switches are configured correctly.



Table 54: NIC Link Down [3067] [A1082]

Name	NIC Link Status
Description	Checks whether any nic is down.
Alert message	Link on NIC <code>nic_name</code> of host <code>host_ip</code> is down. NIC description: <code>nic_description</code>
Cause	The NIC is disconnected from the switch, or the switch port is failing.
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Ensure that the NIC is connected to the switch and that the switch port is functioning properly.

Table 55: Remote Support Enabled [6206] [A1051]

Name	Remote Support Status
Description	Checks if remote support tunnel to Nutanix HQ is enabled on this cluster.
Alert message	Remote support tunnel to Nutanix HQ is enabled on this cluster.
Cause	Nutanix support staff are able to access the cluster to assist with any issue.
Impact	Nutanix support staff are able to access the cluster to assist with any issue.
Resolution	No action is necessary.

Table 56: CVM IP Address Mismatch [6207] [A1009]

Name	CVM IP Address Configuration
Description	Checks that cvm IP address is in sync with zeus configuration
Alert message	IP address of Controller VM <code>zeus_ip_address</code> has been updated to <code>invalid_ip_address</code>. The Controller VM will not be part of the cluster once the change comes into effect, unless zeus configuration is updated.
Cause	The IP address configured in the cluster does not match the actual setting of the Controller VM.
Impact	Cluster compute and storage capacity are reduced. Until data stored on this host is replicated to other hosts in the cluster, the cluster has one less copy of guest VM data.
Resolution	Follow the IP address change procedure in the Nutanix documentation.

Table 57: Zeus Config Mismatch [6208] [A1029]

Name	Hypervisor IP Address Configuration
Description	Checks that hypervisor IP address is in sync with zeus configuration



Alert message	Hypervisor IP address on Controller VM <code>svm_ip_address</code> was updated from <code>zeus_ip_address</code> to <code>invalid_ip_address</code> without following the Nutanix IP Reconfiguration procedure.
Cause	The IP address configured in the cluster does not match the actual setting of the hypervisor.
Impact	The hypervisor host is unreachable from other Controller VMs in the cluster.
Resolution	Follow the IP address change procedure in the Nutanix documentation.

Table 58: CVM Subnet Configuration [6212]

Name	CVM Subnet Configuration
Description	Checks that subnet is in sync with zeus configuration
Cause	The Subnet configured in the cluster does not match the actual subnet
Impact	The Controller VM will not be part of the cluster, unless zeus configuration is updated
Resolution	Follow the IP address change procedure in the Nutanix documentation.

Table 59: CVM using default password [6213] [A6213]

Name	CVM using default password
Description	CVM is using default password
Alert message	One or more CVMs using the default credentials
Cause	SSH access credentials have not changed from the shipped default.
Impact	Insecure configuration.
Resolution	Ensure that all CVMs are using a custom password for improved security.

Table 60: Host using default password [6214] [A6214]

Name	Host using default password
Description	Host is using default password
Alert message	One or more Hosts are still using the default credentials
Cause	SSH access credentials have not changed from the shipped default.
Impact	Insecure configuration.
Resolution	Ensure that all hosts are using a custom password for improved security.

Table 61: Multiple Vcenters discovered [6415] [A6415]

Name	Multiple vCenter Servers Discovered



Description	Multiple vCenter servers discovered
Alert message	Discovered more than one vCenter servers : '<code>vcenter_ips</code>'
Cause	Nodes in the current Nutanix cluster are managed by multiple vCenters.
Impact	You will not be able to perform VM Management related operations.
Resolution	Make sure all the nodes of this Nutanix cluster are moved to any one vCenter.
Resolution	Resolve the issue for the failure. If you cannot resolve the issue, contact Nutanix support.

Table 62: Common port group between ESXi hosts is absent [6416] [A6416]

Name	No common port group among ESXi hosts
Description	Common port group between ESXi hosts is not present
Alert message	'<code>message</code>'
Cause	ESXi hosts does not have any common port group.
Impact	You will not be able to perform network related operations.
Resolution	Ensure that all the nodes of this cluster have at least one common port group or all the nodes are in a single datacenter. If you are still not able to resolve the issue, contact Nutanix support.

Table 63: Number of datastores configured is higher than ESXiMaxConnPerIP. [6447] [A6447]

Name	Check ESXi stale worlds
Description	Checks if the number of datastores configured in the ESXi host is higher than ESXi MaxConnPerIP advanced setting.
Alert message	Number of datastores configured in the ESXi host is <code>no_of_datastore</code> and it is higher than ESXi MaxConnPerIP advanced setting value <code>#MaxConnPerIP@</code>. This may cause stale VM worlds leading to degraded IO performance on the host <code>host_ip</code>.
Cause	Number of datastores configured in the ESXi host is higher than ESXi MaxConnPerIP advanced setting. This may cause stale VM worlds leading to degraded IO performance on the host.
Impact	Degraded hypervisor IO performance
Resolution	Review KB-7833 to increase MaxConnPerIP parameters.

Table 64: ESXi SIOC status check [6448]

Name	ESXi SIOC status check
Description	Checks if the ESX Storage IO Controller (SIOC) feature is disabled
Cause	ESX SIOC feature may be enabled
Impact	ESX SIOC feature should be disabled on all Nutanix datastores



Resolution	ESX SIOC feature needs to be disabled

Table 65: vCenter Server Not Registered [6450] [A6450]

Name	vCenter Registration Check
Description	Check if vCenter server is registered.
Alert message	vCenter Server for this vSphere cluster is not registered in the Prism.
Cause	vCenter Sever is not registered in the Prism.
Impact	vSphere VM management operations on Prism will be disabled.
Impact	Some VMs might be missed from the snapshots created during upgrade.
Resolution	Register the vCenter server for this vSphere cluster in the Prism settings

Table 66: IPMI IP Address Mismatch [15013] [A1008]

Name	IPMI IP Address Configuration
Description	Checks that IPMI IP address is in sync with zeus configuration
Alert message	IPMI IP address on Controller VM <code>svm_ip_address</code> was updated from <code>zeus_ip_address</code> to <code>invalid_ip_address</code> without following the Nutanix IP Reconfiguration procedure.
Cause	The IP address configured in the cluster does not match the actual setting of the IPMI interface.
Impact	The host is unreachable through the IPMI interface.
Resolution	Follow the IP address change procedure in the Nutanix documentation.

Table 67: IPMI sensor values check [15021]

Name	IPMI sensor values check
Description	Check that all IPMI sensor values are within threshold
Cause	IPMI sensors are not within threshold.
Impact	Cannot monitor hardware health.
Resolution	Review KB article 1524.

Table 68: IPMI connectivity check [15022]

Name	IPMI connectivity check
Description	Check that IPMI is working on all nodes
Cause	Refer to KB 1502.



Impact	IPMI could be unavailable.
Resolution	Refer to KB 1502.

Table 69: IPMI sel assertions check [15023]

Name	IPMI sel assertions check
Description	Check for IPMI SEL assertions in the past 24 hours
Cause	IPMI events logged in the last 24 hours.
Impact	Host might have restarted.
Resolution	Refer to KB 1590.

Table 70: IPMI SEL Log Power Failure [15030] [A15030]

Name	IPMI sel power failure check
Description	Check for IPMI SEL Power failure in the past 24 hours
Alert message	Power failure Over-Current-Protection engaged - host <code>ip_address</code>: <code>alert_msg</code> do not remove and add back power - Immediately Contact Nutanix Support.
Cause	The power supply has detected an Over-Current condition for this block and its protection mechanism has activated.
Impact	Block will automatically shut down as part of the Over-Current protection mechanism.
Resolution	Do not replace or reseal the power supply. Immediately call Nutanix technical support for assistance.

Table 71: IPMI using default password [15037] [A15037]

Name	IPMI using default password
Description	IPMI is using default password
Alert message	One or more IPMI devices are still using the default credentials
Cause	SSH access credentials have not changed from the shipped default.
Impact	Insecure configuration.
Resolution	Ensure that all IPMI devices are using a custom password for improved security.

Table 72: Power supplies of different types detected on a node [15038] [A15038]

Name	Mixed Power Supply check
Description	Check for mixed power supplies on chassis



Alert message	per_block_msg
Cause	Power supplies of different types are installed on a chassis
Impact	Mixed power supplies is not a tested combination and may result in unanticipated results.
Resolution	Replace the power supply to match PSUs defined for the NX platform.

Table 73: Number of orphaned egroups is over the recommended threshold. [20013] [A20013]

Name	Garbage egroups check
Description	Check that the number of orphaned egroups are below the recommended threshold.
Alert message	Orphaned egroups are high
Cause	Number of orphaned groups is over the recommended threshold.
Impact	Space is underutilized.
Resolution	Review KB 1574.

Table 74: Snapshot chain height exceeds threshold [20014] [A20014]

Name	Snapshot chain height check
Description	Check that snapshot chain height is less than 25
Alert message	Snapshot chain height has exceeded the threshold limit.
Cause	Snapshot chain height exceeds threshold of 25.
Impact	Some VMs might experience increased latencies.
Resolution	Review KB 1732.

Table 75: Oplog Episode Count Check [20015] [A20015]

Name	Oplog episodes check
Description	Check that oplog episode count is within threshold
Alert message	Oplog episode count is high for some virtual disks for a sustained period of time.
Cause	Oplog episode count is high for one or more vDisks.
Impact	Oplog draining to extent store might get slower or stalled, impacting random writes latency.
Resolution	One or more VMs have an unusual burst of write IO, or there might be data inconsistency in oplog episode files.
Resolution	Review KB 1541.



Table 76: Automatic Dedup disabled check [20016]

Name	Automatic Dedup disabled check
Description	Check that Dedup is not automatically disabled.
Cause	Deduplication on the Storage Container is disabled by Stargate service.
Impact	Space savings may not be observed on the Storage Container.
Resolution	Review KB 1730.

Table 77: Compression disabled check [20017]

Name	Compression disabled check
Description	Check whether compressions is automatically disabled.
Cause	Metadata usage has exceeded the 60GB or 20 percent default threshold.
Impact	
Resolution	Contact Nutanix support to manually reenale compression.

Table 78: Linked clones on Dedup check [20018]

Name	Linked clones on Dedup check
Description	Verifies if there are shared disks on a Storage Container that has deduplication enabled that may be linked clones.
Cause	There are shared disks on a Storage Container that has deduplication enabled that may be linked clones.
Impact	Using both linked clones and deduplication on the same Storage Container may cause metadata bloat.
Resolution	For non-linked clone VMs that use deduplication, place the VMs in separate Storage Containers.

Table 79: Dedup and compression on Storage Container [20019]

Name	Dedup and compression on Storage Container
Description	Check that only one of dedup and compression is enabled on the Storage Containers
Cause	If dedup and compression enabled on the Storage Containers only compresion will take effect unless AOS is above 4.5.
Impact	Deduplication is not in affect.
Resolution	Enable either Dedup or compression on the Storage Container depending on the type of data.



Table 80: Cassandra service status check [21009]

Name	Cassandra service status check
Description	Check that all Cassandra nodes are online and and functioning normally
Cause	Cassandra service is not running in one or more hosts.
Impact	Cluster resilience might degrade.
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Contact Nutanix support.

Table 81: Cassandra service crashed [21010] [A21010]

Name	Cassandra service restarts check
Description	Check if Cassandra service is crashing continuously.
Alert message	Cassandra service has recently crashed. Contact Nutanix support for assistance.
Cause	Cassandra service has restarted frequently in the last 30 minutes.
Impact	Data resiliency is compromised.
Resolution	Contact Nutanix support.

Table 82: Cassandra service is running out of memory [21011] [A21011]

Name	Cassandra memory usage
Description	Check if the Cassandra service is running out of memory.
Alert message	Cassandra service is running out of memory. Contact Nutanix support for assistance.
Cause	Cassandra has crossed memory threshold more than 5 times in the last 3 minutes.
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Contact Nutanix support.

Table 83: Cassandra Keyspace/Column family check [21013]

Name	Cassandra Keyspace/Column family check
Description	Check that there are no undefined Keyspace or CF in Cassandra
Cause	
Impact	



Table 84: Cassandra Waiting For Disk Replacement [21014] [A21014]

Name	Cassandra Waiting For Disk Replacement
Description	Metadata disk not replaced for Disk Replace Op
Alert message	Disk <i>disk_id</i> for node <i>service_vm_external_ip</i> not replaced. Node performing disk replace: <i>ip_address_doing_ring_change</i>
Cause	Disk not replaced for Disk Replace Op.
Impact	Disk Replacement operation will not proceed to completion.
Resolution	Ensure that a replacement disk is present on the node undergoing disk replace. If it is and the alert persists, contact Nutanix support.

Table 85: Multiple Cassandra nodes have similar tokens [21015] [A21015]

Name	Cassandra Similar Token check
Description	Check for similar tokens in Cassandra
Alert message	Multiple Cassandra nodes have similar tokens
Cause	Multiple Cassandra nodes have similar tokens
Impact	Cluster may not tolerate a Cassandra failure.
Resolution	Refer to KB 4301.

Table 86: Domain Fault Tolerance is low for Metadata [21016] [A21016]

Name	Cassandra Domain Fault Tolerance Warning
Description	Domain fault tolerance is low for metadata
Alert message	<i>domain</i> fault tolerance is reduced from desired fault tolerance(<i>desired_fault_tolerance</i>) to <i>current_fault_tolerance</i> for metadata.
Cause	One or more partitions have lower domain fault tolerance than the desired fault tolerance.
Impact	Domain fault tolerance will be reduced.
Resolution	Cassandra-domain-aware fixer has been triggered to re-arrange metadata and improve the domain fault tolerance. No action is required from your side.

Table 87: Domain Fault Tolerance is low for Metadata [21017] [A21017]

Name	Cassandra Domain Fault Tolerance Error
Description	Domain fault tolerance is low for metadata



Alert message	<code>domain</code> fault tolerance is reduced from desired fault tolerance(<code>desired_fault_tolerance</code>) to <code>current_fault_tolerance</code> for metadata.
Cause	One or more partitions have lower domain fault tolerance than the desired fault tolerance.
Impact	Domain fault tolerance will be reduced.
Resolution	Please run Cassandra-domain-aware fixer operation using the nCLI command 'cluster cassandra-domain-aware-fixer-start'.

Table 88: Domain Fault Tolerance is low for Metadata [21018] [A21018]

Name	Cassandra Domain Fault Tolerance Critical
Description	Domain fault tolerance is low for metadata
Alert message	<code>domain</code> fault tolerance is reduced from desired fault tolerance(<code>desired_fault_tolerance</code>) to <code>current_fault_tolerance</code> for metadata.
Cause	One or more partitions have lower domain fault tolerance than the desired fault tolerance.
Impact	Domain fault tolerance will be reduced.
Resolution	Please add more nodes to reduce the skew in number of nodes in each domain. It is preferred to maintain almost equal number of nodes in all the domains to achieve best domain fault tolerance.

Table 89: Node detached from metadata ring [21019] [A21019]

Name	Metadata Drive Detached With Fixer Check
Description	Node detached from metadata ring
Alert message	Metadata drive on CVM <code>ip_address</code> is now detached from ring due to <code>reason</code>.
Cause	Node was removed from the metadata store as part of the ring fixer operation.
Impact	Cluster performance may be degraded till node gets added back.
Resolution	No action required. Node will get added back to the metadata store automatically.

Table 90: One or more Cassandra nodes have invalid tokens [21020] [A21020]

Name	Cassandra Invalid Token check
Description	Check for short tokens in Cassandra
Alert message	Nodes <code>node_list</code> have shorter tokens than recommended
Cause	One or more Cassandra nodes have short tokens
Impact	Cluster may not work properly after some time.



Resolution	Refer to KB 6656.

Table 91: Health warnings detected in metadata service [21021] [A21021]

Name	Cassandra SSTable Health Warning check
Description	Check for cassandra sstables with health warnings
Alert message	Health warnings detected in metadata service
Cause	Cassandra SSTable health warning
Impact	Data resiliency is critical. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Refer to KB 7077.

Table 92: ESXi RAM disk full check [101040]

Name	ESXi RAM disk full check
Description	Check for ramdisk full condition
Cause	pynfs server is using legacy setting to write log file on root partition of ESXi ramdisk.
Impact	ESXi host may be disconnected from vCenter.
Impact	Local datastore may go offline and will be unavailable.
Resolution	Clean up the growing file and change the pynfs server settings.
Resolution	Review KB 1718 for additional details.

Table 93: ESXi RAM disk root usage [101041]

Name	ESXi RAM disk root usage
Description	Check that RAM disk root directory usage below 85%
Cause	ESXi ramdisk root usage exceeds 85%.
Impact	Host might be inaccessible if ramdisk is full.
Resolution	Refer to KB 1846.

Table 94: Discoverable disks check [101042]

Name	Discoverable disks check
Description	Check if local disks are managed by the failover cluster
Cause	HyperV utilities in the host are unable to detect controllers and disks.
Cause	Local disks are managed by the failover cluster service, which makes them inaccessible for CVM



Impact	NA
Resolution	The host may need a reboot. Review KB 2712 for mitigation steps.
Resolution	Please remove local disks from failover cluster manager. Review KB 2172 for mitigation steps.

Table 95: Duplicate disk id check [101044]

Name	Duplicate disk id check
Description	Check for duplicate disk ids
Cause	Multiple disks with same disk id were found.
Impact	This can interfere with cluster functionality.
Resolution	Review KB 1876.

Table 96: Container available unreserved space check [101045]

Name	Container available unreserved space check
Description	Checks if available unreserved space on any Storage Container is below 10 percent of the maximum capacity.
Cause	Unreserved available space is below threshold.
Impact	Virtual Machines might going to a stunned state if the available unreserved capacity drops to zero.
Resolution	

Table 97: Descriptors to deleted files check [101046]

Name	Descriptors to deleted files check
Description	Check if there are file descriptors to deleted files
Cause	Running processes have file descriptors for deleted files that are over 1 GB.
Impact	Cluster upgrades may fail due to lack of space on CVM.
Resolution	Follow the instructions in KB 3644.

Table 98: Cluster can not tolerate node failure and guarantee available rebuild capacity [101047] [A101047]

Name	High disk space usage
Description	Check storage pool has sufficient capacity to sustain replication factor rebuild capacity in the event of a node failure, based on highest replication factor container requirements



Alert message	Cluster cannot tolerate <code>num_node_failure_limit</code> node failure. Current usage of <code>used_bytes</code> bytes exceeds max limit of <code>usable_bytes</code> bytes.
Cause	Inadequate free space on the cluster
Impact	Cluster can not tolerate single node(RF2) or two node(RF3) failure and guarantee available storage to provide data redundancy.
Resolution	Clean out unused data from guest VMs, unneeded reservation at container level, or VMs from the container to be below the threshold.

Table 99: Offline disk in a cluster [101048] [A101048]

Name	Disk online check
Description	Check for the offline disks
Alert message	Disk mounted at <code>mount_path</code> on Controller VM <code>cvm_ip</code> is marked offline.
Cause	A disk may have failed or was manually removed.
Impact	Cluster performance may be degraded.
Resolution	Please refer to KB 1536 for more information.

Table 100: Incomplete disk removal [101049]

Name	Incomplete disk removal
Description	Check for incomplete disk removal
Cause	Cluster Health detected an incomplete disk removal operation.
Impact	Data integrity might be compromised.
Resolution	If the disk removal has not been manually triggered contact Nutanix Support.

Table 101: SED protection consistency check [101050]

Name	SED protection consistency check
Description	Checks that all the drives in the system can authenticate to their internal security agents
Cause	
Impact	
Resolution	Contact Nutanix Support.

Table 102: Automatic disabling of Deduplication [101051]

Name	Automatic disabling of Deduplication
Description	Check that On disk Dedup is not automatically disabled



Cause	Deduplication was disabled due to high metadata size, CVM RAM or SSD requirements are not met.
Impact	Deduplication is disabled.
Resolution	Ensure the RAM and SSD requirement are met and enable dedup from prism or CLI.

Table 103: Sanity check on local.sh [101053]

Name	Sanity check on local.sh
Description	Validate local.sh file configuration.
Cause	pynfs server is using legacy setting to write log file on root partition of ESXi ramdisk.
Impact	ESXi host may be disconnected from vCenter.
Impact	Local datastore may go offline and will be unavailable.
Resolution	Clean up the growing file and change the pynfs server settings.
Resolution	Review KB 1718 for details.

Table 104: Cluster In Override Mode [101057] [A101057]

Name	Cluster In Override Mode
Description	Single-node cluster is in override mode.
Alert message	Single node cluster with CVM <code>service_vm_external_ip</code> is running in override mode due to <code>reason</code>.
Cause	One of the metadata SSDs has failed.
Impact	Cluster is not disk fault tolerant.
Resolution	Replace the failed SSD. Refer to the Nutanix documentation for instructions.

Table 105: Host FQDN resolution [103069]

Name	Host FQDN resolution
Description	Check FQDN resolution of host IPs
Cause	Unable to reach name server or name server doesn't have correct entry.
Impact	Unable to resolve "host FQDN to IP" or "IP to host FQDN".
Resolution	Check if hostname to IP and IP to hostname resolution is working.

Table 106: NTP server FQDN resolution [103070]

Name	NTP server FQDN resolution



Description	Check NTP Server FQDN resolution
Cause	Unable to reach name server or name server doesn't have correct entry.
Impact	Unable to resolve NTP server FQDN.
Resolution	Check if ntp server name resolution is working.

Table 107: SMTP server FQDN resolution [103071]

Name	SMTP server FQDN resolution
Description	Check SMTP server FQDN resolution
Cause	Unable to reach name server or name server doesn't have correct entry.
Impact	Unable to resolve SMTP server FQDN.
Resolution	Check if smtp server name resolution is working.

Table 108: NSC(Nutanix Service Center) server FQDN resolution [103072]

Name	NSC(Nutanix Service Center) server FQDN resolution
Description	Check NSC server FQDN resolution
Cause	Unable to reach name server or name server doesn't have correct entry.
Impact	Unable to resolve NSC(Nutanix Service Center) server FQDN.
Resolution	Check if cluster has Internet access and correct name server setting exists.

Table 109: CVM Not Uplinked to Active 10Gbps Link [103073] [A103073]

Name	CVM 10GBe uplink check
Description	Check that CVM uplinked to active 10Gbps link
Alert message	CVM is not uplinked to active 10Gbps link.
Cause	CVM is not uplinked to active 10Gbps link.
Impact	Performance may degrade.
Resolution	Verify bonds do not contain a 1Gbe link.
Resolution	Review KB 1584.

Table 110: Storage routed to alternate CVM check [103074]

Name	Storage routed to alternate CVM check
Description	Check if storage is routed to an alternate CVM
Cause	Stargate traffic reroute is engaged on host of the cluster.



Impact	Due to stargate traffic reroute, the data resilience and performance may get impacted
Resolution	Check the following cases: 1. CVM is completely down. 2. Stargate is not running or CVM is partly down. 3. Stargate is crashing.

Table 111: NIC driver and firmware version check [103075]

Name	NIC driver and firmware version check
Description	Check if Quad port NIC driver version is above min recommended version
Cause	Quad-port nic card on Nutanix Haswell platform is not running min recommended driver or firmware version.
Impact	NIC is not on the latest driver / firmware version.
Resolution	Upgrade to the latest version of driver and firmware version for the Quad d port nic card.

Table 112: Incorrect {vm_type} NTP Configuration [103076] [A103076]

Name	NTP configuration check
Description	Check that NTP is configured properly on the CVM and hypervisor
Alert message	vm_type NTP is not properly configured. <i>alert_msg</i>
Cause	Detected problems with NTP configuration.
Impact	Metadata operations or alerts might not work properly.
Resolution	Follow the instructions in KB 4519.

Table 113: VM IDE bus check [103079]

Name	VM IDE bus check
Description	Check UVMs with IDE disk
Cause	Detected VMs running in the cluster on IDE bus.
Impact	VMs Performance degradation.
Resolution	Review KB 2998.

Table 114: CVM memory check [103080]

Name	CVM memory check
Description	Check CVM Memory
Cause	CVM memory is less than the recommended threshold with the current feature set.

Impact	Cluster services may crash and fail to restart on CVMs with less than recommended memory.
Resolution	Increase the memory of the Controller VM as per platform requirements listed in KB1513.

Table 115: Same timezone check [103085] [A103085]

Name	CVM same timezone check
Description	Check that all CVMs are in the same timezone
Alert message	vm_types are not in the same timezone.
Cause	Some of the Controller VMs in the cluster are not in the same timezone.
Impact	It is recommended to have all the CVMs in the same timezone to avoid potential issues.
Resolution	Make sure all Controller VMs are in the same time zone. If they are not, then set the timezone using KB 1050

Table 116: Cluster NCC version check [103086]

Name	Cluster NCC version check
Description	Check that NCC version is consistent on the cluster
Cause	NCC version is different across the CVMs.
Impact	Some NCC checks can not be executed on all CVMs.
Resolution	Upgrade NCC to the desired version to keep it consistent across CVMs.

Table 117: CVM ip address is unreachable [103088] [A103088]

Name	Inter-CVM connectivity check
Description	Check that all CVMs are reachable via ping.
Alert message	CVM dest_ip is unreachable
Cause	The Controller VM is down or there is a network connectivity issue.
Impact	Storage may be unavailable and/or performance issues may be observed.
Resolution	If the Controller VM does not respond to ping, turn it on. Ensure that physical networking, VLANs, and virtual switches are configured correctly.

Table 118: Duplicate CVM IP address detected [103089] [A103089]

Name	Duplicate CVM IP check
Description	Check for duplicate CVM IP address
Alert message	Duplicate CVM IP address detected. Associated mac addresses are cvm_ip



Cause	One (or more) Nutanix cluster Controller VM external(eth0) interfaces conflicts with other hosts on the same network.
Impact	CVM unavailable from network
Resolution	Review KB 3025.

Table 119: Duplicate Hypervisor IP address check [103091]

Name	Duplicate Hypervisor IP address check
Description	Check for duplicate Hypervisor IP address
Cause	One or more of your Nutanix cluster hypervisor interfaces conflicts with other hosts on the same network.
Impact	Hypervisor host connectivity can become unstable or unavailable, leading to performance impact, redundancy concerns and and potential downtime.
Resolution	Reconfigure the conflicting hypervisor interfaces to have different IP addresses or VLAN settings

Table 120: Host same timezone check [103092]

Name	Host same timezone check
Description	Check that all hosts are in the same timezone
Cause	Not all hosts in the cluster are in the same timezone.
Impact	Ensure all hosts in the cluster are in the same timezone to help ensure all guest VMs share the same timezone and log messages are timestamped consistently.
Resolution	Ensure that all hosts are in the same timezone by following the instructions in KB 1050

Table 121: Network segmentation enabled on an AOS version lower than 5.5 [103093]

Name	Network segmentation enabled on an AOS version lower than 5.5
Description	Check if network segmentation enabled on a cluster with AOS version lower than 5.5
Cause	Network segmentation enabled on cluster running an AOS version lower than 5.5
Impact	Data may be lost during HA events
Resolution	Upgrade the cluster to AOS 5.5.0.1 or higher.

Table 122: CVM NIC Link Down [103094] [A103094]

Name	CVM NIC Link Status



Description	Checks whether any CVM NIC is down.
Alert message	Link on NIC <code>nic_name</code> of cvm <code>svm_ip</code> is down.
Cause	The rdma0 NIC is disconnected from the switch, or the switch port is failing.
Impact	Cluster performance may be significantly degraded. When rdma0 port is down, the CVM will switch to TCP mode via host vSwitch, and the RDMA mode of entire cluster is turned off. Critical large storage io service will suffer in non-RDMA mode.
Resolution	Ensure that the rdma0 NIC is connected to the switch and that the switch port is functioning properly.

Table 123: Network switch DCBX or NIC PFC is not configured correctly or enabled for RDMA NICs on CVM [103095] [A103095]

Name	RDMA network switch and NIC PFC config check
Description	Checks if network switch DCBX and NIC PFC are correctly configured and enabled.
Alert message	Network switch DCBX or NIC PFC is not configured correctly or enabled for RDMA NICs on CVM <code>cvm_ip</code>: <code>alert_msg</code>
Cause	Network switch DCBX is disabled or corresponding NIC PFC bit is not enabled.
Impact	RDMA traffic between CVMs is not enabled, which degrades cluster performance.
Resolution	Enable DCBX on network switch and configure PFC bit on RDMA NICs.

Table 124: NDP check [103100]

Name	NDP check
Description	Checks if the NDP(Nutanix Discovery Protocol) works in the cvm external network
Cause	IPv6 may be disabled at the CVM, or an upstream network switch may be blocking IPv6 multicast or has IPv6 MLD snooping enabled.
Impact	Node discovery will fail
Resolution	NDP sends an IPv6 UDP payload to multicast address ff02::1 DstPort 13000. Confirm IPv6 is enabled in your network and ensure multicast propagation is enabled on listeners (Controller VMs in the cluster) and senders (any node to be added) to avoid issues when expanding cluster.

Table 125: DVS ESXi compatibility check [103102]

Name	DVS ESXi compatibility check
Description	Check if DVS is being used on an affected ESXi version - ESXi 6.0GA,U1,U2,U3 and 6.5GA.



Cause	ESXi on the cluster is not of the recommended version - ESXi 6.0GA,U1,U2,U3 and 6.5GA.
Impact	CVMs might be disconnected after reboot and risk a hung upgrade.
Resolution	Upgrade ESXi on the cluster to ESXi 6.0 p07 Build 9239799 or ESXi 6.5 U1 or later.
Resolution	Please follow steps from KB 8149 to resolve this issue.

Table 126: Boot device connection check [106015]

Name	Boot device connection check
Description	Check connection to the boot device
Cause	Boot device may have lost connection.
Impact	Host OS may become unresponsive and host may not boot if restarted
Resolution	Check boot device status.

Table 127: Boot device status check [106016]

Name	Boot device status check
Description	Check status of the boot device
Cause	Boot device may be unhealthy and not responding correctly to status requests
Impact	Host responsiveness or availability may be affected
Resolution	Check boot device connection using vmkfstool -P bootbank / readbank on ESXi host.

Table 128: LSI Disk controller firmware status [106018]

Name	LSI Disk controller firmware status
Description	Check LSI Firmware Revision
Cause	LSI Disk Controller is running a superseded version of firmware.
Impact	LSI Disk Controller firmware issues can result in a node to hang and disk instability.
Resolution	Update the LSI Controller Firmware version. Perform inventory from Life Cycle Manager to check and update to the latest firmware version. Check KB 6937 for alternative update option.

Table 129: Sata Controller status is bad. [106020] [A106020]

Name	SATA controller



Description	Check SATA controller ports
Alert message	<code>controller_type controller_instance</code> on host <code>ip_address</code> has the following problems: <code>alert_msg</code>
Cause	The SATA controller is experiencing errors.
Impact	The system may experience issues while accessing physical storage within the node.
Resolution	Review KB 2312.
Resolution	Contact Nutanix support.

Table 130: Intel SSD S3610 on {ip_address} has configuration problems. [106021] [A106021]

Name	Intel SSD Configuration
Description	Check if node has supported configuration for Intel SSD S3610.
Alert message	Intel SSD S3610 on host <code>ip_address</code> has configuration problems: <code>alert_msg</code>.
Cause	Intel SSD S3610 is installed on an NX-1020 or NX-9040 node. Intel SSD S3610 is installed on a cluster whose AOS version is before 4.1.1.4. Intel SSD S3610 is not properly down-formatted.
Impact	Intel SSD S3610 drive performance and write endurance are degraded.
Resolution	Remove Intel SSD S3610 from the NX-1020 or NX-9040 nodes. Upgrade AOS to at least version 4.1.1.4. Contact Nutanix support to replace with a new drive. Refer to the Nutanix documentation for instructions.

Table 131: SSD Firmware Check [106022]

Name	SSD Firmware Check
Description	Check SSD firmware version
Cause	SSD firmware version is below the recommended version.
Impact	Intel S3610 SSDs running firmware version 0110 could be marked offline for upto a minute.
Resolution	Update SSD firmware to the recommended version or above. Perform inventory from Life Cycle Manager to check and update to the latest firmware version. Check the KB 2588 for more details.

Table 132: Hardware FRU information correctness [106025]

Name	Hardware FRU information correctness
Description	Check that FRU fields are programmed to correct value
Cause	Hardware information may be incorrectly configured.
Impact	No functional Impact.



Resolution	Contact Nutanix support to correctly configure the FRU fields.

Table 133: CVM virtual hardware version check [106027]

Name	CVM virtual hardware version check
Description	Check CVM virtual hardware version
Cause	One or more CVM is using old/different virtual hardware version
Impact	Virtual hardware version 7 might cause performance issue
Resolution	Use the recommended minimum Virtual machine hardware version of 8.
Resolution	All CVMs should be on same virtual hardware version.

Table 134: DIMMs of different manufacturers in one memory channel [106028] [A106028]

Name	DIMMs Interoperability Check
Description	DIMMs Interoperability Check
Alert message	Memory channels on host <code>ip_address</code> have DIMMs of multiple manufacturers: <code>err_msg</code>
Cause	DIMMs of different manufacturers are in the same memory channel.
Impact	Unsupported configuration.
Resolution	Ensure that all DIMMs in a memory channel are of the same manufacturer. Verify capacity, model, and manufacturer of DIMMs.

Table 135: DIMMs have invalid Part Number. [106038] [A106038]

Name	DIMM Part Number Check
Description	DIMMs Part Number Check
Alert message	DIMMs on host <code>ip_address</code> have invalid Part Number: <code>err_msg</code>
Cause	DIMMs have invalid Part Number.
Impact	Unsupported configuration.
Resolution	Replace DIMMs that have invalid Part Number. Ensure that all DIMMs have the valid Part Number M393A4K40BB2-CTD.

Table 136: DIMM Configuration is Wrong [106041] [A106041]

Name	DIMM Config Check
Description	DIMM Configuration Check
Alert message	<code>dim_id</code> on host <code>ip_address</code> has the following problem: <code>alert_msg</code>



Cause	The DIMM configuration is unsupported. Either the number of CPUs, the number of DIMMs in a memory channel of a CPU, the capacity of DIMMs, or the total number of DIMMs on a node is wrong or 128GB DIMMs installed on node with wrong type of CPUs.
Impact	Unsupported configuration. Performance may be degraded.
Resolution	Ensure that DIMM configuration is correct. Each node must have 1 or 2 CPUs; each memory channel of a CPU can have up to 2 DIMMs; the capacity of each DIMM is 32, 64 or 128 GB; each node can have 4, 6, 8, 12, 16 or 24 DIMMs.

Table 137: Samsung 863 or 863a on {ip_address} has configuration problems. [106046] [A106046]

Name	Samsung 863(a) Configuration
Description	Check if node has supported configuration for Samsung 863 or 863a.
Alert message	Samsung 863 or 863a on host <code>ip_address</code> has configuration problems: <code>alert_msg</code>.
Cause	Samsung 863 or 863a is not properly down-formatted.
Impact	Samsung 863 or 863a drive performance and write endurance are degraded.
Resolution	Contact Nutanix support to replace with a new drive. Refer to the Nutanix documentation for instructions.

Table 138: Supported virtual hardware version check [106047]

Name	Supported virtual hardware version check
Description	Supported virtual hardware version check
Cause	One or more CVM is using unsupported virtual hardware version
Impact	Virtual hardware version lesser than 7 might cause performance issue
Resolution	Use the recommended minimum Virtual machine hardware version

Table 139: USB boot device missing on node. [106048] [A106048]

Name	Klas USB boot device connection check
Description	USB boot device connection check on Klas platform
Alert message	USB boot device is not detected on node <code>ip_address</code> with Klas platform. <code>alert_msg</code>
Cause	USB boot device is not present on Klas platform.
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Replug USB boot device again into Klas node and restart the node. If restart fails, contact Klas for USB replacement.



Table 140: Secure boot feature status [106065] [A106065]

Name	Secure Boot Check
Description	Check if secure boot configuration is consistent across all nodes in the cluster.
Alert message	Secure boot is not enabled on following hosts: <i>host_ips</i>
Cause	Secure boot is not consistently configured across all nodes in the cluster.
Impact	Cluster is only partially protected by host secure boot.
Resolution	Enable secure boot on all hosts in the cluster to ensure intended cluster security.

Table 141: ESXi version compatibility [106415]

Name	ESXi version compatibility
Description	Check ESX version compatibility with AOS
Cause	ESXi on the cluster is not of the recommended version.
Impact	Identified ESXi version on the cluster is either incompatible or has issues with CBT/VADP APIs.
Resolution	Upgrade ESXi on the cluster.
Resolution	Refer to KB 1729 for more details.

Table 142: ESXi Driver compatibility check [106418]

Name	ESXi Driver compatibility check
Description	Check the compatibility of current installed ixgbe/ixgben driver versions with the current ESX version.
Cause	Current installed Intel ixgbe/ixgben driver version is not recommended on current ESX version.
Impact	Nutanix has seen instability in Network with the older driver version on later ESX versions.
Resolution	Nutanix recommends to keep the driver on the latest qualified version based on Intel 82599 and x540 10Gigabit Ethernet Controllers.

Table 143: ESXi services check [106419]

Name	ESXi services check
Description	Check if ESXi services are up on the host
Cause	One of ssh, ntpd, shell, vpxa, lbtd, DCUI, sfcdb services are not running.
Impact	Cluster connectivity, cluster synchronization, adding host to cluster, accessing console may be impacted.
Resolution	Check following services on the ESX host: SSH, ntpd, vpxa, lbtd, DCUI, sfcdb services.



Table 144: ESXi VAAI plugin installed [106421]

Name	ESXi VAAI plugin installed
Description	Check if the latest version of VAAI Plugin is installed
Cause	Incorrect ESXi VAAI plugin version for AOS version.
Impact	Suboptimal storage performance due to lack of supported VAAI primitives (VMware KB 1021976).
Resolution	Review KB 1868.

Table 145: ESXi VAAI plugin enabled [106422]

Name	ESXi VAAI plugin enabled
Description	Check if VAAI Plugin is enabled
Cause	ESXi VAAI plugin is not enabled.
Impact	Suboptimal storage performance due to lack of storage hardware acceleration (VMware KB 1021976).
Resolution	Review KB 1868.

Table 146: Some storage containers have a high number of NFS files [106426] [A106426]

Name	NFS file count check
Description	Check that NFS datastore file count below threshold 20000
Alert message	Some storage containers have a high number of NFS files
Cause	Number of files for respective storage containers has increased beyond 20K. This is expected with large VDI setups.
Impact	High number of NFS files may cause vpxa services on esxi hosts to restart.
Resolution	Reduce the number of files if you observe vpxa instability.

Table 147: hostd Service Not Running [106427] [A106427]

Name	VMware hostd service check
Description	Check if hostd access is available
Alert message	hostd service is not running.
Cause	VMware hostd service may not be running.
Impact	If hostd service is not running, AOS upgrade will fail.
Resolution	Check hostd status and restart it manually.



Table 148: CVM Port Group Renamed [106428] [A106428]

Name	CVM port group renamed
Description	Check that CVM port group name has not changed
Alert message	CVM port group has been renamed on the ESXi hosts
Cause	CVM port group has been renamed on the ESXi hosts.
Impact	AOS upgrade may fail.
Resolution	Refer to KB 1904

Table 149: ESXi Scratch Configuration [106429]

Name	ESXi Scratch Configuration
Description	Check ESXi Scratch settings
Cause	Host contains a scratch location that is not properly configured or the locker.conf cannot be read.
Impact	If scratch space is not available, temporary data is stored on a RAM disk with limited space, causing issues and log files not being updated.
Resolution	Manually configure the scratch location.
Resolution	Review KB 2046.

Table 150: ESXi configured VMK check [106430]

Name	ESXi configured VMK check
Description	Check that correct vmk is chosen
Cause	Host Management IP in esx.conf does not match the host ip in zeus config.
Impact	A cluster upgrade or node addition operation might fail
Impact	The cluster will choose the wrong ESXi interface as management server IP address when a node is added.
Resolution	Review KB 2049 and KB 2522

Table 151: Ivy Bridge performance check [106441]

Name	Ivy Bridge performance check
Description	Check Ivy Bridge performance degradation scenarios on ESXi
Cause	Multiple causes, review KB 2729
Impact	Performance issues may be observed at Applications or OS level.
Resolution	The performance degradation issue is resolved in ESXi 5.5 U2. Therefore upgrading to latest ESXi version will resolve the issue



Table 152: VMKNICs subnets check [106442]

Name	VMKNICs subnets check
Description	Check if vmknics have different subnets
Cause	vmknics have ip address configured in the same subnet
Impact	vmknics in the same ip subnet on the same esxi host is unsupported.
Resolution	Correct the IP addressing in the network.
Resolution	Review KB 2722 for more details.

Table 153: ESXi CPU model and UVM EVC mode check [106444]

Name	ESXi CPU model and UVM EVC mode check
Description	Check CPU model and UVM EVC mode on ESXi
Cause	Some SMEP enabled VMs detected.
Impact	Performance might be degraded for the affected VMs.
Resolution	Review KB 3000. ESXi upgrade might be required.

Table 154: ESXi APD handling check [106445]

Name	ESXi APD handling check
Description	Check if APD handling is as recommended.
Cause	APD handling is not configured to the recommended value.
Impact	If APD is not enabled, VMs could potentially end up in a critical state during a failure event.
Resolution	Refer KB 3328 for the steps to enable or disable APD handling.

Table 155: ESXi NFS heartbeat timeout check [106446]

Name	ESXi NFS heartbeat timeout check
Description	Check if NFS heartbeat timeout is set to recommended value.
Cause	NFS heartbeat parameter is set NOT set to 30 secs.
Impact	This paramter essentially decides the amount of time ESXI waits before aborting a heartbeat request. If set incorrectly, it may cause ESXI to percieve the datastore status incorrectly.
Resolution	Refer KB 3328 for steps to disable APD handling.

Table 156: ESXi AHCI Driver version check [106447]

Name	ESXi AHCI Driver version check



Description	Checks if the VMW_AHCI driver version is less than or equal to 1.0.0-34 in ESXi 6.5 or greater
Cause	The VMW_AHCI driver version 1.0.0-34 has many issues in ESXi 6.5 or greater. This check is for the driver version.
Impact	Esxi 6.5 with AHCI driver version less than or equal to 1.0.0-34 installed.
Resolution	Disable the AHCI driver in ESXi 6.5 or greater.

Table 157: ESXi storage access from Host [106449]

Name	ESXi storage access from Host
Description	Check if storage is accessible from the host
Cause	Storage is not properly configured on the host.
Impact	All storage I/O might go to a single node.
Resolution	Review KB 2263.

Table 158: Incorrect Kerberos Setup [106450] [A106450]

Name	Hyper-V Kerberos setup check
Description	Check if Kerberos is set up correctly
Alert message	Kerberos is not set up correctly.
Cause	Kerberos is not setup correctly.
Impact	Kerberos authentication is impacted.
Resolution	Check Kerberos setup.
Resolution	Review KB 3030 for additional details.

Table 159: Hyper-V VMQ Check [106451]

Name	Hyper-V VMQ Check
Description	Check if VMQ is disabled on host
Cause	VMQ is enabled on Hyper-V host.
Impact	Enabling VMQ on NICs with a faulty driver can cause VMs to fail during live migration.
Resolution	Nutanix recommends that you disable VMQ. To do so, find the VMQ status: Get-Netadaptervmq and for the component for which VMQ is enabled, disable VMQ. Disable-NetAdapterVmq -name <component>.

Table 160: Hyper-V Host pending reboots check [106452]

Name	Hyper-V Host pending reboots check



Description	Check that host has no pending reboots
Cause	Host has pending reboot.
Impact	Hyper-v Hosts pending reboots after Windows cluster aware or other updates/hyperv controller utilities fail to show any disks attached even though the disks are physically attached.
Resolution	Restart the Hyper-v Host.

Table 161: Orphaned VSS copies are detected. [106453] [A106453]

Name	Orphaned shadow copies
Description	Check if any orphan VSS copies are present.
Alert message	Orphaned VSS Copies Detected
Cause	Orphaned VSS copies are detected.
Impact	Performance degradation
Resolution	Please delete orphaned VSS copies. See KB 2798 for more information.

Table 162: Domain join check [106454]

Name	Domain join check
Description	Check if host is joined to a domain
Cause	Node(s) is/are not joined to a domain.
Impact	The host will not be manageable using Hyper-V Manager or SCVMM.
Resolution	Refer to KB 2231.
Resolution	Join the cluster and nodes to a domain.

Table 163: Hyper-V services check [106455]

Name	Hyper-V services check
Description	Check if essential services are running on the host.
Cause	NutanixHAService/DiskMonitoerService/NutanixHostAgent not running properly.
Impact	Cluster cannot do disk monitoring and storage traffic re-routing if the Controller VM is unreachable or down.
Resolution	Verify status of Nutanix services. Review KB2242 for detailed information.

Table 164: Window update check [106456]

Name	Window update check



Description	Check if automatic Windows updates are disabled on the host
Cause	Windows update mode is set to automatic.
Impact	Hyper-V hosts can reboot in an uncontrolled manner because of automatic windows updates.
Resolution	Set the update mode to Manual. Refer to KB article 2253

Table 165: MS Failover cluster check [106457]

Name	MS Failover cluster check
Description	Check if the Hyper-V failover cluster is properly configured
Cause	Hyper-V failover cluster is not configured with host disk drives.
Impact	It can cause outage as it will prevent disks from being available to CVMs after CVM and Host reboot.
Resolution	Refer to KB 2282.

Table 166: Required Windows features check [106458]

Name	Required Windows features check
Description	Check if the required Windows features are installed and the banned Windows features are not installed
Cause	Required windows roles are not enabled on the host.
Impact	Recommended features are not enabled.
Resolution	Review KB 2244.

Table 167: Host networking check [106459]

Name	Host networking check
Description	Check if networking is properly configured on the host
Cause	The Nutanix-specific virtual network switches are not correctly configured.
Impact	Nutanix services will not come up in the Controller VM.
Resolution	Create the missing vmswitches/network adapters.

Table 168: Hyper-V host power scheme [106460]

Name	Hyper-V host power scheme
Description	Check if high performance profile is enabled on the host
Cause	Power scheme on the host is not configured to High Performance.
Impact	Hosts can experience increased latency and performance might be degraded.



Resolution	Review KB 2217.

Table 169: Nutanix Utilities [106461]

Name	Nutanix Utilities
Description	Check if high performance profile is enabled on the host
Cause	Required Nutanix utilities are not found on the host.
Impact	Some Nutanix Services will fail to start.
Resolution	Review KB 2254.

Table 170: CVM configuration check [106462]

Name	CVM configuration check
Description	Check if the CVM is properly configured
Cause	A problem with the CVM configuration has been detected.
Impact	CVM services may not start after reboot.
Resolution	Review KB 2283 and 2219.

Table 171: HyperV storage access from Host [106463]

Name	HyperV storage access from Host
Description	Check if storage is accessible from the host
Cause	Storage is not properly configured on the host.
Impact	All storage I/O might go to a single node.
Resolution	Review KB 2263.

Table 172: WinRM service check [106464]

Name	WinRM service check
Description	Check if the winrm service on the host is listening on port 5985 by connecting to the host using WSMAN.
Cause	Either password is not updated in the cluster configuration or WinRM service not running in one or more Hyper-V hosts.
Impact	The CVM cannot connect to the host to manage it if wsman connection is failing or WinRM service not running in one or more Hyper-V hosts.
Resolution	Ensure that the WinRM service is running and listening on the port 5986.
Resolution	Update the correct password in the cluster configuration if wsman connection is failing.



Resolution	Review KB 2243 for more details.

Table 173: Crash dumps check [106465]

Name	Crash dumps check
Description	Check crash dumps on the host
Cause	Dedicated kernel crash dump is not enabled.
Cause	Kernel crash dump exists on the host.
Impact	Host kernel crash dump may not work correctly.
Resolution	Review KB 2249.

Table 174: High Availability Routes check [106466]

Name	High Availability Routes check
Description	Check if HA routes are setup correctly in the host in pre 4.0 releases.
Cause	The External switch may have to be recreated.
Impact	HA will be unavailable.
Resolution	You can recover the routes by running genesis restart on the Controller VM.

Table 175: UTC clock check [106467]

Name	UTC clock check
Description	Check if UTC Clock is enabled
Cause	UTC clocks on the node are not enabled or CVMs are not in sync.
Impact	User-defined cluster procedures that are expected to run synchronously may run at different times, which may lead to faulty outcomes.
Resolution	Review KB 2711.

Table 176: User VMs status [106468]

Name	User VMs status
Description	Check if the user virtual machines are properly configured
Cause	User VMs are in critical state.
Impact	User VMs in critical state can not be accessible.
Resolution	Review KB 2259.



Table 177: CVM Panics check [106469]

Name	CVM Panics check
Description	Check if any CVMs have panicked.
Cause	At least one CVM panicked recently.
Impact	Storage may have been unavailable while the CVM was restarting.
Resolution	Follow the instructions in KB 4379.

Table 178: Host Missing Critical Windows Updates [106470] [A106470]

Name	Critical Windows updates check
Description	Check if required windows updates are installed
Alert message	Host is missing critical Windows Updates.
Cause	Host is missing critical Windows Updates.
Impact	Host might experience problems regarding storage / live migration.
Resolution	Review KB 3011 for more information.

Table 179: HyperV VSS file limit check [106473]

Name	HyperV VSS file limit check
Description	Check containers file count within limit.
Cause	VSS can fail on containers with higher file count.
Impact	VSS may fail on the container.
Resolution	Reduce the number of VMs and associated files on the container below the file limit.

Table 180: Hyper-V Disk Pass-Through Check [106474]

Name	Hyper-V Disk Pass-Through Check
Description	Verify if the disks are pass-through to the Controller VM in Hyper-V 2016.
Cause	All the storage controllers are not pass-through to the Controller VM.
Impact	Misconfigured physical drives are not be accessible.
Resolution	Run Repair-CVM command to fix the Controller VM configuration.

Table 181: Active Directory DC(s) running on cluster [106476] [A106476]

Name	DC/DNS running on Nutanix Cluster Check
Description	Check if all active directory domain controllers and DNS servers are running on Nutanix cluster



Alert message	Active Directory DC(s) running on cluster which will prevent failover cluster from starting after Nutanix Cluster starts and stops
Cause	Active Directory Domain Controller(s) or DNS servers configured on the UVMs in the cluster
Impact	Failover cluster will fail to start after the Nutanix cluster starts and stops(due to power failure for example), as it depends on availability of domain controllers which will result in User VM downtime
Resolution	There should be at least one domain controller and DNS server running outside Nutanix cluster, that can be used by failover cluster in case locally running DC servers are failed

Table 182: Remote Site Time Sync [110001]

Name	Remote Site Time Sync
Description	Checks whether time drift between remote site is within limit.
Cause	Time drift between source and remote cluster is high.
Impact	Remote replication may fail to complete as expected.
Resolution	Use NTP to sync time

Table 183: Remote site connectivity not normal. [110002] [A110002]

Name	Remote Site Connectivity
Description	Check if connectivity to remote sites is normal.
Alert message	Connectivity to remote site <i>remote_name</i> is not normal.
Cause	Remote site may not have the local site configured as a remote.
Cause	Remote site might be stale (destroyed and re-created).
Cause	Cluster services at remote site might not be running.
Cause	Network connectivity to remote site might have been lost.
Cause	Firewall/VPN on source cluster, remote cluster, or both may prevent connection establishment.
Impact	Remote replication may fail to complete as expected.
Resolution	Configure the local site as a remote on the remote site.
Resolution	Remove and re-add the remote site on the local cluster.
Resolution	Ensure that cluster services at remote site are running.
Resolution	Check network connectivity to remote site.
Resolution	Check firewall/VPN on source cluster and remote cluster to allow them to connect to each other.

Table 184: Remote Site MTU settings [110005]

Name	Remote Site MTU settings
Description	Check MTU setting for connecting to remote cluster
Cause	MTU size mismatch on CVMs between remote sites.
Impact	Degraded replication performance may be observed if MTU size is not consistent between remote sites.
Resolution	Review KB 1949.

Table 185: Remote Site configuration [110006]

Name	Remote Site configuration
Description	Check remote site configuration on local and remote sites
Cause	Remote site is not properly configured.
Impact	Nutanix snapshot replication might fail to the remote site.
Resolution	Review KB 3335.

Table 186: Remote Site virtual external IP(VIP) [110007]

Name	Remote Site virtual external IP(VIP)
Description	Check if remote site configuration on remote sites has virtual ip of this cluster
Cause	Remote Cluster may not be using the virtual external IP address (VIP) of the local cluster for remote site connectivity.
Cause	Remote Cluster may not be using the DR network segmentation virtual IP address (VIP) of the local cluster for remote site connectivity.
Impact	CVM shutting down on a remote site results in replication problems.
Impact	Partial or incorrect DR network segmentation configuration.
Resolution	Add the virtual external IP address (VIP) of the local cluster to the remote site configuration of the remote cluster. Add the virtual external IP address (VIP) of the remote cluster to the remote site configuration of the local cluster.
Resolution	If the local cluster has DR network segmentation enabled, add the DR network segmentation virtual IP address (VIP) and the network segmented external subnet of the local cluster to the remote site configuration of the local cluster.

Table 187: AWS cloud instance not configured properly. [110008] [A110008]

Name	AWS Instance Check
Description	Verify that the AWS cloud instances have recommended configuration.
Alert message	AWS cloud instance at remote <code>remote_name</code> not configured properly.
Cause	AWS instance does not have at least two 190GB+ disks.



Impact	Curator scans may fail.
Resolution	Make sure that AWS instance has at least two 190GB+ disks.

Table 188: Old generation AWS instance configured. [110009] [A110009]

Name	AWS Instance Type Check
Description	Checks whether old generation AWS instance is configured for remote site.
Alert message	Old generation AWS instance is configured for remote site <code>remote_name</code>. Follow KB 3861 to change the instance type to m3.2xlarge and cloud CVM configuration for better RPO/RT0.
Cause	Present AWS instance has slow replication performance.
Impact	Present AWS instance has slow replication performance.
Resolution	Follow the KB 3861 to change the instance type to m3.2xlarge and cloud CVM configuration.

Table 189: AOS version of cloud remote site is less than source cluster. [110010] [A110010]

Name	Cloud Remote Version Check
Description	Check AOS version for cloud remote site.
Alert message	AOS version of cloud remote site <code>remote_name</code>: (<code>remote_nos_version</code>) is less than that of source cluster(<code>src_nos_version</code>).
Cause	AOS version of cloud remote site is lower than that of source cluster.
Impact	Performance improvements made in new AOS version will not be available for cloud CVM.
Resolution	Upgrade cloud remote site.

Table 190: Cloud cluster does not have all recommended gflags set. [110012] [A110012]

Name	Cloud Gflags Check
Description	Check if cloud remote sites have recommended gflags set.
Alert message	Cloud cluster does not have all recommended gflags set.
Cause	Cluster may have been upgraded to a newer version.
Impact	Some of the improvements made in new release will not be enabled.
Resolution	Set recommended gflags. The <code>cloud_connect_fix</code> script can be run.

Table 191: Egroup count on cloud disk is higher than the recommended threshold. [110013] [A110013]

Name	Cloud Egroup Count check



Description	Check if egroup count on cloud remote is within threshold.
Alert message	Egroup count on cloud disk is higher than the recommended threshold.
Cause	Large amount of data may have been stored on cloud appliance.
Impact	Slow performance can be seen.
Resolution	Spread the protection domains that you want to replicate among multiple cloud remote sites.

Table 192: Incorrect Cluster Information in the Remote Site {remote_name} [110014] [A110014]

Name	Incorrect Cluster Information in Remote Site
Description	Check if the cluster information in remote site is correct.
Alert message	Replications to this remote site may fail because the cluster information is incorrect in the remote site <code>remote_name</code>.
Cause	Incorrect IP addresses may have been found in remote site configuration.
Cause	Remote cluster may have been destroyed and recreated.
Impact	Replications to this remote site may fail.
Resolution	Remove incorrect IP addresses and add correct IP addresses in the remote site configuration.
Resolution	Delete the remote site and add it again. You need to configure the schedules again.
Resolution	If you cannot resolve the error, contact Nutanix support.

Table 193: Metadata disk usage is higher than the specified limit. [110015] [A110015]

Name	Cloud Metadata Disk Usage Check
Description	Check if the metadata disk usage on cloud Controller VM is within the specified limit.
Alert message	Metadata disk usage is higher than the specified limit: <code>Limit</code>.
Cause	Large amount of data may have been stored on cloud appliance.
Impact	Slow performance can be seen.
Resolution	Spread the protection domains that you want to replicate among multiple cloud remote sites.
Resolution	Follow the KB 4302 to add a new metadata disk to existing cloud Controller VM.

Table 194: Cloud disk usage is nearing the threshold. [110016] [A110016]

Name	Cloud Disk Usage Check
Description	Check if the cloud disk usage is within the threshold limit.



Alert message	Cloud disk usage: <code>usage_bytes</code> TB, is nearing the threshold: <code>threshold_bytes</code> TB. Cloud remote site will be turned to read-only once the threshold is crossed.
Cause	Large amount of data may have been stored on the cloud appliance.
Impact	Slow performance can be seen.
Resolution	Spread the protection domains that you want to replicate among multiple cloud remote sites.
Resolution	Reduce retention of snapshots on the cloud remote site.

Table 195: Azure cloud Controller VM has smaller disks. [110017] [A110017]

Name	Azure Disk Size Check
Description	Verify that the Azure cloud Controller VM has recommended configuration.
Alert message	Azure cloud Controller VM has smaller than the two 256GB+ recommended disks. Contact Nutanix support for further assistance.
Cause	Azure cloud Controller VM does not have minimum of two 256GB+ disks.
Impact	Slow performance can be seen.
Impact	Curator scans may fail.
Resolution	Contact Nutanix support so that the Azure cloud Controller VM has at least two 256GB+ disks.

Table 196: Remote site in same vCenter datacenter. [110018] [A110018]

Name	Remote Datacenter Check
Description	Check if any remote site is added under the same datacenter in vCenter.
Alert message	Hosts of current cluster and Async DR remote site '<code>remote_name</code>' are registered under the same datacenter in vCenter.
Cause	Source and remote site of an Async DR protection domain are under the same datacenter in the vCenter.
Impact	Migration to this remote site may fail.
Resolution	Register source and remote sites with different datacenter in vCenter.
Resolution	Source and remote site in same vCenter datacenter must have different vStore container names.
Resolution	If source and remote site are in same vCenter datacenter, make sure that source and remote container names are different.

Table 197: Checking if DR services are reachable or not. [110020] [A110020]

Name	DR Service Reachability Check
Description	Check if the DR services are reachable or not.



Alert message	message . Please Refer to kb_url for resolution.
Cause	Connectivity affected due to the VPN or the service is down.
Impact	DR workflows will not work.
Resolution	Restart the service.

Table 198: Checks if AOS version of the remote site configured for synchronous replication is on AOS version >= 5.12 [110021] [A110021]

Name	Remote AOS Version Check
Description	Checks if the remote site configured is on AOS version >= 5.12 to support Synchronous Replication
Alert message	For protection rule: <i>protection_rule_name</i>, target cluster(s): <i>remote_site_list</i> have version older than 5.12
Cause	Remote site is on AOS version < 5.12
Impact	Synchronous Replication will not be supported.
Resolution	Upgrade the cluster on the remote site to AOS version >= 5.12

Table 199: Latency to target Prism Element(s) is greater than maximum value allowed for synchronous replications. [110022] [A110022]

Name	Remote Site Latency Check
Description	Check if the the latency to target cluster is less than the maximum value allowed for synchronous replications.
Alert message	Latency to target Prism Element(s) <i>err_msg</i> is greater than <i>max_permitted_latency</i>
Cause	Target cluster is unreachable or latency to target cluster is greater than 5ms.
Impact	Sync Rep will be affected.
Resolution	Ensure that the target cluster is reachable or choose another taregt cluster.

Table 200: Active Backup Schedule Check [110200]

Name	Active Backup Schedule Check
Description	Checks if a backup schedule exists for protection domain protecting some entities.
Cause	No backup schedule exists for protection domain protecting some entities.
Cause	Backup schedules exist for protection domain not protecting any entity.
Cause	If schedule was intentionally deleted, an expired snapshot(s) could still consume storage space.
Impact	No backups will be made.
Impact	Any expired snapshots (if existing), would still consume storage space.



Resolution	Configure a backup schedule for protection domain protecting some entities.
Resolution	Delete backup schedule(s) for protection domain not protecting any entity. Delete the protection domain as well once the schedule is deleted.
Resolution	If the schedule(s) was intentionally deleted to reclaim storage space, delete this protection domain also to reclaim the space taken up by snapshots.

Table 201: Storage Container(s) is/are not mounted on all nodes. [110201] [A110201]

Name	Storage Container Mount Configuration
Description	Verify that the Storage Container(s) is/are mounted on all nodes.
Alert message	Storage container(s) <code>unmounted_containers</code> for protection domain <code>pd_name</code> is/are not mounted on all nodes.
Cause	Storage Container(s) is/are not mounted on all nodes.
Impact	Registration of VM(s) may fail during restoration from backup.
Resolution	Mount the Storage Container(s) on all nodes.

Table 202: Local Backup Configuration [110204]

Name	Local Backup Configuration
Description	Verify if recovery is possible from a local backup.
Cause	Backup schedule is not configured properly.
Cause	Backup is not happening as expected.
Impact	Recovery from local backup may not be possible.
Resolution	Verify backup schedule configuration.
Resolution	Verify if backups are happening without error.

Table 203: Cluster Certificate Expiring [110215] [A1115]

Name	SED CA Certificate
Description	Check if CA certificates are about to expire.
Alert message	Cluster CA certificate <code>certificate_name</code> expires on <code>expiration_date</code>.
Cause	Certificates have defined expiration dates.
Impact	If the cluster's certificate expires, it will be unable to verify trust of server certificates.
Resolution	Get a new signed certificate.

Table 204: Cloud remote check [110220]

Name	Cloud remote check
Description	Check backup and on-wire compression settings for cloud remote sites
Cause	
Impact	
Resolution	Contact Nutanix support.

Table 205: Cloud cluster does not have recommended configuration locally [110221] [A110221]

Name	Cloud local check
Description	Check that cloud cluster has recommended configuration locally.
Alert message	Cloud cluster does not have recommended configuration locally. Verify that cloud cluster has recommended configuration locally and then proceed.
Cause	AWS cluster does not have exactly two storage pools (backup-sp and DoNotUse-sp).
Cause	Storage Pool backup-sp does not have exactly one disk with CLOUD storage tier.
Cause	Some remote sites configured on cloud do not have onwire compression enabled.
Impact	
Resolution	Contact Nutanix support.

Table 206: Secondary PDs Not in Sync [110222] [A110222]

Name	Metro Availability Secondary PD sync check
Description	Check if secondary metro PD is in sync with primary
Alert message	Secondary Protection domains are no more in sync with primary Protection domains.
Cause	Link between primary and secondary is not healthy.
Impact	Secondary protection domain is no more in sync.
Resolution	Review KB 3032.

Table 207: Backup snapshots on metro secondary check [110223]

Name	Backup snapshots on metro secondary check
Description	Check if there are any backup snapshots on a metro secondary which can be used to reclaim space



Cause	A local snapshot of the Storage Container saved on the standby cluster in a metro availability configuration.
Impact	Snapshot exists on remote site and Free space on cluster is less than 15 percent.
Resolution	Upgrade the cluster to AOS 4.1.4 or later.

Table 208: Checkpoint snapshot on Metro configured Protection Domain [110224]

Name	Checkpoint snapshot on Metro configured Protection Domain
Description	Check that last auto checkpoint snapshot was taken in the last 4 hours
Cause	Automatic checkpoint snapshot might not be created or created with delay.
Impact	Time for data replication and resync will be longer when Metro Availability is re-enabled.
Resolution	Create a schedule manually for checkpoint snapshot.
Resolution	Review KB.

Table 209: ESX VM Virtual Hardware Version Compatible [110227]

Name	ESX VM Virtual Hardware Version Compatible
Description	Check Virtual hardware version compatibility of protected VMs
Cause	Virtual hardware version of the VM is not compatible with the maximum virtual hardware version supported by any of the nodes at the remote site.
Impact	VMs will not be registered on the remote site hosts upon PD failover.
Resolution	Upgrade the hypervisor version on the remote site to support virtual hardware version for the VM.

Table 210: Duplicate VM names [110228]

Name	Duplicate VM names
Description	Check for duplicate VM names
Cause	Duplicate named VMs are found.
Impact	Duplicate named VMs will not be protected by protection domains.
Resolution	Review KB 1907.

Table 211: Protection Domains File share [110229]

Name	Protection Domains File share
Description	Check for files/VMs protected by multiple protection domains
Cause	Some files/VMs are being protected by multiple protection domains.



Impact	Migrating such a PD would remove the VMs from the inventory.
Resolution	Review KB 2034.

Table 212: Vstore PDs with multiple CGs [110231]

Name	Vstore PDs with multiple CGs
Description	Check for multiple consistency groups in PDs
Cause	This is probably a result of SRA-SRM workflows creating multiple CGs in respective PDs.
Impact	Having more than one consistency group for a Storage Container level protection domain can result in issues when snapshot restore or DR failover actions are initiated.
Resolution	Contact Nutanix support.

Table 213: Snapshot file location check [110233]

Name	Snapshot file location check
Description	Check snapshot file location
Cause	VM protected by Nutanix Protection Domain has a trailing backslash in the snapshot file location.
Impact	Nutanix snapshots might be skipped for these VMs if they are protected by Nutanix Protection Domains.
Resolution	Review 2069.

Table 214: Latest Snapshot contains all the UVMs [110234]

Name	Latest Snapshot contains all the UVMs
Description	Check that all VMs are protected by snapshot
Cause	Snapshot doesn't include all the VMs protected by that protection domain.
Impact	Some VMs are not protected.
Resolution	Review KB 3336.

Table 215: Detected VMs with non local data [110237]

Name	Detected VMs with non local data
Description	Metro Availability VMs are accessing data from local cluster.
Cause	Detected VMs polling data from remote location.
Impact	Performance degradation
Resolution	Review KB 2093.



Table 216: Metro protection domain active on both sites [110239]

Name	Metro protection domain active on both sites
Description	Check that Metro protection domain active on both sites
Cause	Metro protection domain is active on both sites.
Impact	Synchronous replication will not happen when both sites are in Active State.
Resolution	Disable Metro protection from the primary site and then reenale.

Table 217: Files in a streched VMs should be in the same Storage Container [110241]

Name	Files in a streched VMs should be in the same Storage Container
Description	Check that VMs are consistently stretched
Cause	VM that is protected by Metro protection domain has some of its files in a different Storage Container from the Metro protected Storage Container.
Cause	In case of Hyper-V hosts, the VMs on stretch container may not be using Metro name.
Impact	VMs might not be properly recovered on the remote site.
Resolution	If the VM resides on a Metro enabled Storage Container, move all its associated files to the same Storage Container.
Resolution	Use the Metro name for all the VM components to avail the benefits of Metro features.

Table 218: Vstore VM Files Consistency Check [110253]

Name	Vstore VM Files Consistency Check
Description	Check that VMs in a Vstore are consistently protected
Cause	VM that is protected by Vstore protection domain has some of its files in a different Storage Container from the Vstore protected Storage Container.
Impact	VMs might not be properly recovered on the remote site.
Resolution	If the VM resides on a Vstore protection enabled Storage Container, move all its associated files to the same Storage Container.

Table 219: Internal Consistency Groups Check [110254]

Name	Internal Consistency Groups Check
Description	Check if a PD contains internal consistency groups
Cause	PDs may have some leftover internal consistency groups because metro availability enablement was aborted.
Impact	Snapshots of a PD in this state cannot be used to perform a restore operation.
Resolution	Review KB.



Table 220: Linked Clones in high frequency snapshot schedule [110256]

Name	Linked Clones in high frequency snapshot schedule
Description	Verify if the protection domain contains any linked clones that are configured with a high frequency snapshot schedule.
Cause	The protection domain might have a high frequency snapshot schedule configured with linked clones.
Impact	Snapshots of the protection domain in this state might lose data.
Resolution	Reduce the snapshot frequency or move the linked clones out of the protection domain.

Table 221: SSD Snapshot reserve space check [110257]

Name	SSD Snapshot reserve space check
Description	Check if the snapshot reserve space on SSD is filling up.
Cause	Too many snapshots are getting retained.
Cause	Very less SSD capacity is available.
Cause	Replication link is too slow.
Impact	High frequency snapshotting cannot be sustained for some/all Protection Domains
Resolution	Reduce the number of entities that are configured for high-frequency snapshots.
Resolution	Increase the SSD capacity in the cluster.
Resolution	System will automatically try to reclaim SSD space by falling back to a lower snapshot frequency.

Table 222: Aggressive break replication timeout. [110258] [A110258]

Name	Metro aggressive break replication timeout check
Description	Check if break replication timeout of metro protection domain is more than the recommended limit.
Alert message	Break replication timeout of metro protection domain is below the recommended minimum of 15 seconds.
Cause	Break replication timeout of a metro protection domain is below recommended limit.
Cause	Metro availability is disabled.
Impact	Performance degradation.
Resolution	Set break replication timeout of the affected protection domain to recommended value.
Resolution	Enable metro availability.

Table 223: Unplanned Failover and Failback can Cause Full Replication. [110259] [A110259]

Name	Nearsync unplanned failover check
Description	Unplanned failover and failback can cause full replication.
Alert message	Unplanned failover and failback can result in full replication of the protection domain <code>pd_name</code>. See KB 5695 for more information.
Cause	Unplanned failover and failback can cause full replication.
Impact	Activate of a protection domain and subsequent replication to the original site can result in a full replication.
Resolution	See KB 5695 for detailed information.

Table 224: Incomplete Metro setup. [110261] [A110261]

Name	Metro HyperV 2016 metro cluster name setup check
Description	Metro setup for HyperV 2016 is incomplete.
Alert message	Metro setup for HyperV 2016 is incomplete.
Cause	Metro cluster name is not setup.
Cause	Different Metro cluster name on two clusters.
Cause	AD or DNS entry is not created for Metro cluster name.
Cause	Password on two clusters may not match.
Impact	Unsupported Metro feature.
Resolution	Please use setup_hyperv.py script to setup Metro.
Resolution	Use a common Metro cluster name across clusters to setup Metro.
Resolution	Please use setup_hyperv.py to add Metro cluster name.
Resolution	Please provide correct password to setup_hyperv.py while setting up Metro.

Table 225: NearSync Replication of the Protection Domain {protection_domain_name} has not Progressed. [110264] [A110264]

Name	Nearsync Replication Stuck
Description	Nearsync replication of the protection domain has not progressed.
Alert message	NearSync replication of the protection domain <code>protection_domain_name</code> to the remote site <code>remote_name</code> for the snapshot <code>snapshot_id</code> has not made any progress in the last <code>time_string</code>.
Cause	Network outage or saturated network with the remote site.
Cause	Stargate service is not healthy on the source or remote site.
Cause	SSD snapshot reserve on the remote site does not have enough free disk space.
Impact	Protection domain may transition out of NearSync.
Impact	Snapshots of the protection domain will not be available on the remote site.



Resolution	Check the network connection with the remote site.
Resolution	Ensure that the remote site has enough disk space available in the SSD snapshot reserve.

Table 226: Latest snapshot missing entities check [110265]

Name	Latest snapshot missing entities check
Description	Check that all the VMs/VGs/Files in the protection domain at the time of snapshot creation are protected by the snapshot.
Cause	Snapshot doesn't include all the VMs/VGs/Files protected by that protection domain at the time of snapshot creation.
Impact	Some VMs/VGs/Files are not protected.
Resolution	Review KB 6258.

Table 227: Maximum Entities in a Consistency Group Check [110266]

Name	Maximum Entities in a Consistency Group Check
Description	Check if each CG in a PD does not contain more than 20 entities
Cause	Various
Impact	Snapshots of a PD in this state cannot be used to perform a restore operation.
Resolution	Review KB.

Table 228: Checking if ports of relevant services are open [110454] [A110454]

Name	Ports Open Check
Description	Checks if the ports for the relevant services required for Synchronous Replication are open
Alert message	<i>error_msg</i>
Cause	Relevant services cannot be reached
Impact	Synchronous Replication will not work.
Resolution	Check firewall rules on the CVM and in your environment and reach out to Nutanix Support for any service related issues

Table 229: E-mail alerts check [111000]

Name	E-mail alerts check
Description	Check email alerts
Cause	E-mail alerts disabled
Impact	Alerts will not be sent through email, asup cases might not be created.



Resolution	Review smtp configuration.
Resolution	Review KB 1586.

Table 230: E-mail alerts contacts configuration [111001]

Name	E-mail alerts contacts configuration
Description	Check email contacts
Cause	E-mail alerts do not have contacts configured.
Impact	Alerts are enabled, but no contacts are configured.
Resolution	Configure emails contacts where alerts should be sent.

Table 231: Cassandra nodes up [111002]

Name	Cassandra nodes up
Description	Check that all Cassandra nodes are running
Cause	
Impact	Meta data load is inconsistent.
Resolution	If no other Cassandra checks are failing, repair nodetool by following KB 1369. If other Cassandra checks are failing, contact Nutanix support.

Table 232: Cassandra tokens consistent [111003]

Name	Cassandra tokens consistent
Description	Check that all Cassandra nodes are consistent
Cause	
Impact	Meta data load is inconsistent.
Resolution	If no other Cassandra checks are failing, repair nodetool by following KB 1369. If other Cassandra checks are failing, contact Nutanix support.

Table 233: Cassandra metadata balanced across CVMs [111004]

Name	Cassandra metadata balanced across CVMs
Description	Check that Cassandra metadata balanced across CVMs
Cause	Metadata is not distributed evenly among nodes in the cluster.
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Correct imbalance by using the nCLI command 'cluster cassandra-token-range-skew-fixer-start'.



Table 234: Zookeeper fault tolerance check [111007]

Name	Zookeeper fault tolerance check
Description	Checks if Zookeeper services have the desired fault tolerance level
Cause	Zookeeper and/or Zookeeper monitor services are down on one or more nodes
Impact	The cluster may not be able to continue operating in the case of a node or block failure.
Resolution	Contact Nutanix support

Table 235: Time since last Curator scan is beyond threshold [111008] [A111008]

Name	Curator scan time elapsed check
Description	Check if time since last Curator Scan is beyond threshold.
Alert message	Time since last Curator scan is beyond threshold. Contact Nutanix support for assistance.
Cause	Background cluster maintenance tasks are not occurring.
Impact	Cluster operation may be affected.
Resolution	Contact Nutanix Support.

Table 236: Zookeeper Not Active on All CVMs [111009] [A111009]

Name	Zookeeper active on all CVMs
Description	Checks if number of active zookeepers is equal to the total number of zookeeper CVMs
Alert message	Zookeeper is not active on all CVMs.
Cause	Zookeeper service is not running on all CVMs.
Impact	Cluster resiliency is affected.
Resolution	Contact Nutanix support.

Table 237: Zookeeper nodes distributed in multi-block cluster [111010]

Name	Zookeeper nodes distributed in multi-block cluster
Description	Checks if zookeeper nodes are distributed in a multi block cluster
Cause	Zookeeper service has nodes running in the same block.
Impact	Cluster is not block aware.
Resolution	Review KB 1587.

Table 238: Zookeeper alias is incorrectly configured in the cluster. [111011] [A111011]

Name	Zookeeper Alias Check
Description	Check Zookeeper alias information configuration
Alert message	msg
Cause	ZK aliases in one or more sources do not match with the Zookeeper server config maintained by Zookeeper.
Impact	Nutanix utilities or Nutanix cluster services may not work as expected.
Resolution	Ensure the ZK aliases are consistent across all locations.

Table 239: GPU drivers installed [111015]

Name	GPU drivers installed
Description	Check if drivers are installed
Cause	GPU driver may be corrupt or missing or not functioning properly.
Impact	GPU card (NVIDIA) may not function properly.
Resolution	Confirm that GPU is installed and re-install or upgrade the driver.
Resolution	Review KB 2714.

Table 240: CVM passwordless SSH to host [111019] [A111019]

Name	Host passwordless SSH
Description	Check passwordless SSH into local hypervisor
Alert message	CVM <code>svm_ip</code> is not able to passwordless SSH to its local hypervisor
Cause	SSH disabled on hypervisor
Impact	Node will be removed from the cluster, leading to reduced capacity
Resolution	Enable SSH on hypervisor
Resolution	Please refer to KB 1501 for more information.

Table 241: Non default gflags check [111020]

Name	Non default gflags check
Description	Check for non-default settings and warn if values not consistent across CVMs
Cause	One or more settings have non-default values.
Cause	One or more non-default settings do not have the same value on all CVMs.
Impact	Non-default settings will be overwritten during upgrade.
Impact	Cluster services such as stargate may fail.
Resolution	Follow the instructions in KB 1530.



Table 242: Node is in degraded state [111021] [A1305]

Name	Degraded Node check
Description	Check if any node is in degraded state
Alert message	Node <code>svm_ip svm_id</code> has been marked as being in a degraded state
Cause	Unstable environment such as partially bad disks, network bandwidth reduction, network packet drops, soft lockups, hardware issues(such as unreliable DIMM with ECC errors).
Impact	Overall cluster and user VM performance may be adversely affected.
Resolution	Refer to KB4206.

Table 243: Name server configuration [111022]

Name	Name server configuration
Description	Check for DNS Server configuration
Cause	Name server is not configured on CVMs and hypervisor hosts or is not able to resolve queries.
Impact	Name resolution will not work.
Resolution	Verify if Name server is configured on CVMs and on hypervisor hosts and if it can resolves queries. Review KB 3005 for more details.

Table 244: HTTP proxy check [111023]

Name	HTTP proxy check
Description	Check if HTTP proxy is working
Cause	No proxy configured.
Cause	Unable to connect to proxy on the port.
Impact	
Resolution	Check the proxy server setting like network, port, user credentials, etc

Table 245: LDAP configuration [111024]

Name	LDAP configuration
Description	Check LDAP configuration.
Cause	LDAP not correctly configured in the cluster.
Impact	Directory users might not be able to log properly.
Resolution	Review KB 2997.



Table 246: Auto support check fails [111025] [A111025]

Name	Pulse configuration
Description	Check for auto support configuration
Alert message	Pulse is not configured. Connection Status: <code>auto_support_check_connection_status</code>, Error Message: <code>auto_support_check_message</code>
Cause	Pulse is not configured.
Impact	Pulse emails will not be sent.
Resolution	Verify if Pulse is configured, if emails can be sent via SMTP server or connections to nsc01.nutanix.net/nsc02.nutanix.net hosts are successful.
Resolution	Review KB 1585 for more details.

Table 247: Snapshots space utilization status [111026]

Name	Snapshots space utilization status
Description	Check that space occupied by snapshots is below the threshold
Cause	Snapshots disk space utilization exceeds the threshold.
Impact	Snapshots occupies most of the disk space on the cluster.
Resolution	Review KB 1853.

Table 248: Virtual IP Check [111027] [A111027]

Name	Virtual IP check
Description	Check if virtual IP is configured and reachable.
Alert message	<code>vm_type</code> Virtual IP is configured but unreachable.
Cause	Cluster virtual IP is not configured.
Cause	Cluster services are down or cluster is not started yet.
Impact	Nutanix features that use virtual IP address might be adversely affected.
Impact	
Resolution	Configure a valid virtual IP for the cluster and verify that all cluster services are up.

Table 249: Remote syslog server check [111028]

Name	Remote syslog server check
Description	Check remote syslog server connectivity
Cause	Remote syslog server is not reachable.
Impact	remote syslog will not be receiving log output from cluster.



Resolution	Check for network problem with rsyslog server or upstream network.

Table 250: Internal port configuration [111029]

Name	Internal port configuration
Description	Check if internal config port is disabled
Cause	Port 7777 is enabled on CVM.
Impact	This configuration is not recommended for security reasons.
Resolution	Disable port 7777 on the CVMs.

Table 251: Deduplication efficiency check [111030]

Name	Deduplication efficiency check
Description	Check that value of content_cache_dedup_ref_count is low
Cause	Deduplication is not working efficiently. Usually occurs when deduplication is enabled after Storage Container has been in use.
Impact	Space saved by deduplication is not optimal.
Resolution	Move guest VMs from one Storage Container to another to generate deduplication fingerprint, or use vdisk_manipulator tool to create fingerprint on a specific vdisk.

Table 252: Storage Pool SSD tier usage [111031]

Name	Storage Pool SSD tier usage
Description	Check that storage pool SSD tier usage under ILM threshold
Cause	Storage Pool SSD utilization consistely above 75%.
Impact	Performance degradation.
Resolution	Delete unnecessary data.
Resolution	Review knowledge article 2882 for more information.

Table 253: Cluster version check [111032]

Name	Cluster version check
Description	Check that cluster runs GA version of AOS
Cause	Non GA version of AOS detected.
Impact	Cluster is on Non GA version.
Resolution	Review KB 2720.



Table 254: Hardware configuration validation [111033]

Name	Hardware configuration validation
Description	Check that /etc/nutanix/factory_config.json is valid
Cause	factory_config.json does not contain valid configuration.
Impact	Service degradation.
Resolution	Review KB 3001.

Table 255: File permissions check [111034]

Name	File permissions check
Description	Check /tmp permissions
Cause	Proper permissions are not set on /tmp directory in the CVM.
Impact	Some Nutanix services will fail to read/write data into /tmp directory which results in service crashes.
Resolution	Set /tmp directory permissions to 1777. Run the command "sudo chmod 1777 /tmp"

Table 256: Advanced Vector Extensions (AVX) enabled [111035]

Name	Advanced Vector Extensions (AVX) enabled
Description	Checks AVX enabled on CPUs.
Cause	Intel's Advanced Vector Extensions (AVX) is disabled on the node.
Impact	If AVX is disabled, upgrade will fail.
Resolution	Enable AVX on the node. Review KB 8148.

Table 257: Hypervisor version check [111037]

Name	Hypervisor version check
Description	Check whether the same hypervisor version running on all cluster nodes.
Cause	Hypervisor version is not same across all hosts.
Impact	It is recommended to keep the hypervisor version same across all the hosts in the cluster.
Resolution	Keep hypervisor versions same across the hosts in the cluster.

Table 258: Zookeeper connections check [111038]

Name	Zookeeper connections check
Description	Check that Zookeeper is correctly configured for a cluster



Cause	
Impact	
Resolution	Review KB 1587.

Table 259: Check that cluster virtual IP address is part of cluster external subnet [111044] [A111044]

Name	Check that cluster virtual IP address is part of cluster external subnet
Description	Check that cluster virtual IP address is part of cluster external subnet
Alert message	Cluster virtual IP address <code>vip</code> is in a different subnet from the cluster external subnet <code>external_subnet</code>
Cause	Cluster virtual IP address is in a different subnet from the cluster external subnet
Impact	Data protection in between remote sites will be affected.
Resolution	Ensure that cluster virtual IP address is part of cluster external subnet

Table 260: Mellanox NIC Speed is not 10GbE and 40 GbE or both 10GbE and 40GbE NICs are installed on one node [111047] [A111047]

Name	Mellanox NIC Status check
Description	Checks if Mellanox NICs are down or if any Mellanox NIC has speed other than 10GbE or 40GbE. Checks if both 10GbE and 40GbE Mellanox NICs are installed on one node.
Alert message	Mellanox NIC on host <code>host_ip</code> has problem: <code>alert_msg</code>
Cause	The NIC is disconnected from the switch, or the switch port is not functioning correctly, or both 10GbE and 40GbE Mellanox NICs are installed on one node.
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Ensure that the NIC is connected to the switch and that the switch port is functioning properly. Ensure that only 10GbE or 40GbE Mellanox NIC is installed on one node.

Table 261: Prism Central using default password [111050] [A111050]

Name	PC using default password
Description	PC is using default password
Alert message	One or more Prism Central instances are still using the default credentials
Cause	SSH access credentials have not changed from the shipped default.
Impact	Insecure configuration.



Resolution	Ensure that Prism Central is using a custom password for improved security.

Table 262: Mellanox NIC not Installed or with Wrong Type on host machine [111051] [A111051]

Name	Mellanox NIC Mixed Family check
Description	Checks if Mellanox NICs belong to the same Mellanox NIC family.
Alert message	Mellanox NIC on host <i>host_ip</i> has problem: <i>alert_msg</i>
Cause	Mellanox NICs are not from the same Mellanox NIC family.
Impact	Cluster performance may be significantly degraded. Mixing Mellanox NICs from different families may cause the NICs to not function properly.
Resolution	Replace some Mellanox NICs to make them come from same Mellanox NIC family.

Table 263: Notifications dropped check [111052]

Name	Notifications dropped check
Description	Some alert notifications were dropped
Cause	One or more cluster services are not healthy.
Impact	Some of the alerts will not be raised. Cluster issues may stay unnoticed.
Resolution	Contact Nutanix support.

Table 264: Cluster SSD configuration compatibility check for high frequency snapshotting [111058]

Name	Cluster SSD configuration compatibility check for high frequency snapshotting
Description	Check if the SSD configuration of the cluster is compatible with high frequency snapshotting.
Cause	Cluster does not have an SSD configuration recommended for high frequency snapshotting. (All hybrid node SSDs must be greater than 1.2TB)
Impact	High frequency snapshotting may not be sustained for some or all of the protection domains.
Impact	High latency may be seen in the user VMs.
Resolution	Use a cluster with one of the recommended SSD configurations.
Resolution	Replace the smaller SSDs in the hybrid nodes with SSD size greater than 1.2TB.
Resolution	Use a lower snapshot frequency for entities in this cluster.

Table 265: Alert E-mail Failure [111066] [A111066]

Name	Alert E-mail Failure



Description	Failed to send alert emails
Alert message	Unable to send alert E-mail from Controller VM <code>service_vm_external_ip</code> due to <code>reason</code>.
Cause	Issue with the e-mail client
Impact	Alerts notification emails will not be sent.
Resolution	Contact Nutanix Support.

Table 266: Registered Prism Central check [111067]

Name	Registered Prism Central check
Description	Check if services on registered Prism Central IP addresses are up
Cause	Prism service is not running on the registered Prism Central VM
Impact	The cluster cannot be managed through registered Prism Central
Resolution	Start cluster services on Prism Central VM or re-register the cluster with a Prism Central

Table 267: The frequency of CPU {cpu_id} on host {host_ip} is extremely low [111068] [A111068]

Name	Host CPU Frequency Check
Description	Checks the frequencies of CPUs on a node.
Alert message	CPU <code>cpu_id</code> on host <code>host_ip</code> has extremely low frequency: <code>alert_msg</code>
Cause	Either a power supply or a fan has failed. A failed power supply reduces the power available to the CPU, and a failed fan leads to overheating of the CPU. Each condition causes the frequency of the CPU to drop below the frequency that is required for node operations.
Impact	When CPU frequency is extremely low, the node may fail and impact workload.
Resolution	Replace the failed power supply or fans. Refer to the Nutanix documentation for instructions.

Table 268: AOS Version EOL [111070] [A111070]

Name	AOS EOL check
Description	Check if the cluster nodes are running an AOS Version that will no longer be supported
Alert message	The AOS Version <code>#current_AOS_version</code> installed on this cluster is end of life (EOL) from <code>eol_date</code>. Upgrade AOS immediately.
Cause	End of Life (EOL) AOS Version is installed on the cluster nodes
Impact	Support may not be able to resolve issues since this AOS version will no longer be maintained



Resolution	Upgrade to a supported AOS version immediately. The AOS EOL Schedule can be found on the Nutanix support portal in KB2720: https://portal.nutanix.com/kb/2720 . To determine the AOS upgrade path, whether it be long-term supported (LTS) or short-term supported (STS) refer to https://portal.nutanix.com/#/page/upgradePaths . The Nutanix SW EOL policy can be found at https://www.nutanix.com/support-services/product-support/policies/

Table 269: Maximum supported HDD capacity per node for AOS version check [111073]

Name	Maximum supported HDD capacity per node for AOS version check
Description	Check to verify HDD capacity per node can be supported by AOS version.
Cause	HDD capacity per node exceeds maximum allowed threshold.
Impact	VMs might experience increased latencies.
Resolution	If the node's HDD capacity exceeds the maximum allowed threshold for AOS versions pre-5.11, Nutanix recommends upgrading the cluster to AOS 5.11 or later. Please refer to KB 7196 for additional details.

Table 270: Maximum supported HDD capacity per node for memory size check [111074]

Name	Maximum supported HDD capacity per node for memory size check
Description	Check to verify HDD capacity per node can be supported by available memory per node.
Cause	HDD capacity per node exceeds maximum allowed threshold.
Impact	VMs might experience increased latencies.
Resolution	If the node's HDD capacity exceeds the maximum threshold supported by available memory, a hardware configuration change may be required. Please refer to KB 7196 for additional details.

Table 271: SW Encryption Keys from {kms_name} are unavailable [111075] [A111075]

Name	SW Encryption Keys Unavailable
Description	Check to see if the KMS actually has the expected passwords for the container encryption keys.
Alert message	SW encryption keys for one or more containers are not available from external key manager <i>kms_name</i>.
Cause	If a 3rd party KMS gets restored from a backup or other incorrect management operation, it may have missing keys.
Impact	Cluster will have data unavailability if rebooted.
Resolution	Contact Nutanix and 3rd party key manager support.



Table 272: Cannot Remove Password Protected Disk(s) [130003] [A1106]

Name	Password Protected Disk Status
Description	Cannot remove password protected disk(s)
Alert message	Cannot remove password protected disks <code>disk_id_list</code> from the configuration as they cannot be cleaned.
Cause	Either the disks are offline, pulled out or the corresponding node is down.
Impact	The disks cannot be removed from the configuration.
Resolution	Ensure the disks are accessible.

Table 273: Automatic Addition Of Node To Metadata Ring Disabled [130004] [A1079]

Name	Node Auto Add Check
Description	Automatic Addition Of Node To Metadata Ring Disabled
Alert message	Automatic addition of node to metadata ring has been disabled on CVM <code>service_vm_id</code> with IP address <code>ip_address</code>.
Cause	Node found unstable again after automatically adding it back to metadata ring recently.
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	If the node was down for an extended period of time and is now running, add it back to the metadata store by going to host details. Otherwise, contact Nutanix support.

Table 274: Node Detached From Metadata Ring [130005] [A1055]

Name	Node Detached From Ring Check
Description	Node detached from metadata ring
Alert message	Node <code>ip_address</code> is now detached from metadata ring due to <code>reason</code>.
Cause	Either a metadata drive has failed, the node was down for an extended period of time, or an unexpected subsystem fault was encountered, so the node was removed from the metadata store.
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	If the metadata drive has failed, replace the metadata drive as soon as possible. Refer to the Nutanix documentation for instructions. If the node was down for an extended period of time and is now running, add it back to the metadata store by going to host details. Otherwise, contact Nutanix support.

Table 275: Metadata Dynamic Ring Change Operation Stuck [130006] [A1117]

Name	Metadata DynRingChangeOp Status
Description	Dynamic Ring Change operation not making progress
Alert message	operation for node ip_address is not progressing. Node performing the operation: ip_address_doing_ring_change
Cause	A node in the cluster is unhealthy.
Impact	Node addition and node removal operations will not proceed to completion.
Resolution	Ensure that all nodes in the cluster are healthy. If they are and the alert persists, contact Nutanix support.

Table 276: Metadata Dynamic Ring Change Operation Too Slow [130007] [A1116]

Name	Metadata DynRingChangeOp Slow Check
Description	Dynamic Ring Change operation too slow
Alert message	operation for node ip_address is progressing slowly. Total elapsed time: elapsed_time_mins min(s). Node performing the operation: ip_address_doing_ring_change
Cause	A node in the cluster is unhealthy, or the cluster contains an extraordinarily large amount of data.
Impact	Cluster performance may be degraded.
Resolution	Ensure that all nodes in the cluster are healthy. If they are and the alert persists, contact Nutanix support.

Table 277: Node Forwarding Metadata Requests [130008] [A1037]

Name	Node In Forwarding Mode Check
Description	Node Forwarding Metadata Requests
Alert message	Metadata service on node ip_address is running in forwarding mode due to reason .
Cause	Either a metadata drive has failed, node removal has been initiated, or an unexpected subsystem fault has been encountered.
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.

Table 278: Large Metadata Size Detected [130009] [A1119]

Name	Metadata Size



Description	Large metadata size detected
Alert message	Node <code>ip_address</code> contains a large amount of metadata.
Cause	A node in the cluster contains a large amount of metadata and has exceeded thresholds.
Impact	Cluster performance may be degraded.
Resolution	Ensure that all nodes in the cluster are healthy. If they are and the alert persists, contact Nutanix support.

Table 279: Node Marked To Be Auto Added To Metadata Ring [130010] [A1078]

Name	Node Marked To Be Auto Added Check
Description	Node Marked To Be Auto Added To Metadata Ring
Alert message	Node <code>service_vm_id</code> with IP address <code>ip_address</code> is marked to be added to metadata ring. <code>reason</code> .
Cause	Node was detached from metadata ring earlier. It is now stable and will be added back to the metadata store.
Impact	The metadata service will start to use the node.
Resolution	No action is necessary.

Table 280: Node Marked To Be Detached From Metadata Ring [130011] [A1054]

Name	Node Marked To Be Detached Check
Description	Node Marked To Be Detached From Metadata Ring
Alert message	Node <code>ip_address</code> is marked to be detached from metadata ring due to <code>reason</code> .
Cause	Either a metadata drive has failed, the node was down for an extended period of time, or an unexpected subsystem fault was encountered, so the node is marked to be removed from the metadata store.
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	If the metadata drive has failed, replace the metadata drive as soon as possible. Refer to the Nutanix documentation for instructions. If the node was down for an extended period of time and is now running, add it back to the metadata store by going to host details. Otherwise, contact Nutanix support.

Table 281: Metadata Ring Imbalance [130012] [A1072]

Name	Metadata Imbalance Check
Description	Cassandra Metadata Imbalance



Alert message	Metadata ring imbalance for <code>num_partitions</code> partitions. The partitions are <code>partition_list</code>
Cause	One or more nodes have a disproportionately larger token range size. This imbalance can cause performance bottlenecks on the node(s) affected.
Impact	Node performance may be significantly degraded.
Resolution	Execute cassandra ring skew fix operation using the NCLI command 'cluster cassandra-token-range-skew-fixer-start'.

Table 282: Curator Job Running Too Long [130015] [A1120]

Name	Curator Job Status
Description	Curator Job Running Too Long
Alert message	Curator job <code>name</code> with id <code>execution_id</code> has been running for a long time i.e. <code>elapsed_time_secs</code> seconds.
Cause	Various
Impact	Background cluster maintenance tasks might get affected in the future. The root cause should be addressed as soon as possible.
Resolution	Contact Nutanix support if this issue persists.

Table 283: Curator Scan Failure [130016] [A1081]

Name	Curator Scan Status
Description	Curator Scan Failure
Alert message	Curator <code>scan_type</code> scans have repeatedly failed to complete.
Cause	Various
Impact	Background cluster maintenance tasks are not occurring. While cluster operation will not be immediately affected, the root cause should be addressed as soon as possible.
Resolution	Contact Nutanix support.

Table 284: Duplicate IP Address Detected [130017] [A1096]

Name	IP Configuration
Description	Duplicate IP Address Detected
Alert message	Detected IP address conflict for cluster virtual IP address <code>duplicate_ip</code> .
Cause	The cluster virtual IP address may be configured on another host.
Impact	The cluster is not accessible through the virtual IP address. Other services that use the virtual IP address, such as remote replication or SCCM, may be affected.



Resolution	Either configure a different cluster virtual IP address or unconfigure the IP address from the other host.

Table 285: Duplicate Remote Cluster ID [130018] [A1038]

Name	Duplicate Remote Cluster ID Check
Description	Duplicate Remote Cluster ID
Alert message	Remote cluster '<i>remote_name</i>' is disabled because the name conflicts with remote cluster '<i>conflicting_remote_name</i>'.
Cause	Two remote sites with different names or different IP addresses have same cluster ID. This can happen in two cases: (a) A remote cluster is added twice under two different names (through different IP addresses) or (b) Two clusters have the same cluster ID.
Impact	Protected data is not replicated to the remote site.
Resolution	In case (a) remove the duplicate remote site. In case (b) verify that the both clusters have the same cluster ID and contact Nutanix support.

Table 286: IP Address Not Hosted [130023] [A1097]

Name	Virtual IP Configuration
Description	IP Address Not Hosted
Alert message	Unable to host virtual IP address <i>ip_address</i>.
Cause	The network configuration is incorrect.
Impact	The cluster is not accessible through the virtual IP address. Other services that use the virtual IP address, such as remote replication or SCCM, may be affected.
Resolution	Contact Nutanix support.

Table 287: SMTP Error [130060] [A1080]

Name	SMTP Configuration
Description	Failed to send email
Alert message	Failed to send email from Controller VM <i>ip_address</i> via <i>smtp_host</i> due to error: <i>error</i>.
Cause	The cluster cannot send Pulse and alert emails to Nutanix support with the current configuration.
Impact	Nutanix support cannot proactively monitor cluster health and offer assistance before issues become acute.
Resolution	Ensure an SMTP server is configured and reachable from the cluster. If no SMTP server is configured, ensure that the firewalls allow the cluster to send email messages through the Nutanix SMTP servers (nsc01.nutanix.net and nsc02.nutanix.net) on port 80.



Table 288: Remote site latency is high [130075] [A1168]

Name	RemoteSiteHighLatency
Description	Latency to a stretch remote site is high
Alert message	Latency to remote site ' <i>remote_name</i> ' has been more than ' <i>remote_site_high_latency_threshold_ms</i> ' ms , for ' <i>sustained_high_latency_secs</i> ' seconds
Cause	Various
Impact	Replication will fail
Resolution	Check if remote site is reachable.

Table 289: Possible degraded Node [130087] [A130087]

Name	Node Degradation Status
Description	Services on one of the nodes in the cluster are in a degraded state and possibly not making progress.
Alert message	Possible degraded node <i>ip_address</i> with Controller VM ID <i>service_vm_id</i> reported by component <i>component_name</i> .
Cause	Various
Impact	Cluster performance may be significantly degraded.
Resolution	Contact Nutanix support.

Table 290: CPS Deployment Evaluation Mode [130093] [A130093]

Name	CPS Deployment Evaluation Mode
Description	Checks if the CPS deployment is running in evaluation mode.
Alert message	CPS deployment of cluster with cluster id <i>cluster_id</i> is running in evaluation mode. Please apply a license by following the CPS deployment guide.
Cause	CPS is running in evaluation mode.
Impact	The deployment cannot be in evaluation and the user has to apply a license.
Resolution	Apply a license to the deployment.

Table 291: Stretch connectivity is lost [130119] [A130119]

Name	Stretch Connectivity Lost
Description	Stretch connectivity is lost
Alert message	Stretch connectivity from local site ' <i>local_stargate_handler</i> ' to remote site ' <i>remote_name</i> ' is lost due to ' <i>reason</i> '
Cause	Remote site(s) not reachable



Impact	Metro connection of the remote site is unavailable
Resolution	Check remote site(s) connectivity and reenable if necessary

Table 292: Alert raised on cloud remote site {remote_name}: {alert_message} [130135] [A130135]

Name	Cloud Remote Alert
Description	Alert generated on cloud remote site.
Alert message	Alert raised on cloud remote site <i>remote_name: alert_message</i>
Cause	Various
Impact	Various
Resolution	Resolve the issue stated in the alert. If you cannot resolve the issue, please contact Nutanix support.

Table 293: Snapshot reserve on SSD is full. [130139] [A130139]

Name	LWS Store Full
Description	LWS store has become full on the cluster.
Alert message	The SSD space reserved for snapshots is full.
Cause	Snapshot reserve on SSD is full.
Impact	Creation of the recovery points might be stopped.
Resolution	Reduce the retention time and/or the IO load.

Table 294: Two node cluster state changed [130151] [A130151]

Name	Two Node Cluster State Change
Description	Change in state of two node cluster
Alert message	Two node cluster state changed: <i>message</i>
Cause	There is a change in availability of a node (or network between the nodes).
Impact	If the cluster is now in stand-alone mode, it cannot handle the remaining nodes failure.
Resolution	Two node cluster will auto-heal when node (or network between the nodes) is available again.

Table 295: Two node cluster changed state to stand-alone mode. [130152] [A130152]

Name	Two Node Cluster changed state to stand-alone mode
Description	One of the nodes in a two node cluster is currently unavailable, the cluster is now operating in stand-alone mode.



Alert message	Two node cluster state change: <i>message</i>
Cause	There was an unplanned event that changed the availability of a node (or network between the nodes).
Impact	The cluster is not node fault tolerant until the cluster auto-heals.
Resolution	Two node cluster will auto-heal when node (or network between the nodes) is available again.

Table 296: Two node cluster state change to stand-alone mode [130153] [A130153]

Name	Two node cluster changed state to stand-alone mode.
Description	Change in state of two node cluster
Alert message	Two node cluster state change: <i>message</i>
Cause	There is a change in availability of a node (or network between the nodes). This event was planned.
Impact	The cluster is not node fault tolerant until the cluster auto-heals.
Resolution	Two node cluster will auto-heal when the planned event is completed successfully.

Table 297: State changed for two node cluster [130154] [A130154]

Name	Two node cluster changed state back to a two-node mode.
Description	Change in state of two node cluster
Alert message	Two node cluster state changed: <i>message</i>
Cause	A previously unavailable node or service is now available again.
Impact	No impact.
Resolution	No resolution necessary.

Table 298: Witness is unreachable from node. [130155] [A130155]

Name	Witness is unreachable for two node cluster from a node.
Description	Unable to reach the Witness VM over the network, the cluster may not be able to handle a node failure.
Alert message	The node cannot issue witness APIs to witness: <i>message</i>
Cause	Node is unable to issue APIs to witness VM.
Impact	The cluster may not be able to handle node failure until the witness is reachable again.
Resolution	Check the witness VM to verify it is up and verify the cluster nodes can communicate with the Witness VM over the network.

Table 299: Drive Removal Stuck [130156] [A130156]

Name	Drive Removal Stuck
Description	Drive Removal Stuck
Alert message	Removal of drive <code>disk_id</code> is stuck.
Cause	Various
Impact	Drive cannot be removed from the cluster.
Resolution	Contact Nutanix support.

Table 300: Unable to Retrieve the Availability Zone Endpoint [130173] [A130173]

Name	Unable to get Availability Zone Endpoint
Description	Unable to retrieve the endpoint information for the Availability Zone
Alert message	Unable to retrieve the endpoint information for the Availability Zone <code>availability_zone_physical_name</code>
Cause	Prism Central, Prism Element, or Xi - Availability Zone may be unreachable
Impact	VMs may not be fully protected
Resolution	Check if the Prism Central, Prism Element, or Xi - Availability Zone are reachable

Table 301: Unable to communicate with the Data Center Manager [130175] [A130175]

Name	Failure to communicate with DCM
Description	Unable to retrieve the endpoint information for the Availability Zone from the Data Center Manager
Alert message	Unable to retrieve the endpoint information for the Availability Zone <code>availability_zone_name</code> from the Data Center Manager
Cause	Data Center Manager may be unreachable or has provided an invalid response
Impact	VMs may not be fully protected
Resolution	Check if Data Center Manager is reachable from the Xi - Availability Zone

Table 302: Data At Rest Encryption key backup warning [130184] [A130184]

Name	Data At Rest Encryption key backup warning.
Description	New encryption keys not backed up.
Alert message	Data At Rest Encryption key backup warning: <code>message</code>
Cause	New encryption keys have been created since the last backup was downloaded.
Impact	If the encryption keys are lost, the data in the storage container cannot be accessed.



Resolution	Download the latest copy of the cluster's key backup.

Table 303: Local key manager master key rotation warning. [130185] [A130185]

Name	Local key manager master key rotation warning.
Description	Unable to rotate local key manager's master key.
Alert message	Local key manager master key rotation warning. Last failure reason: <i>failure_message</i>
Cause	One or more nodes in the cluster are unavailable.
Impact	If the cluster experiences any more node failures, the data on the cluster may be inaccessible.
Resolution	Ensure Mantle service is UP on all nodes. If the warning persists, contact Nutanix support.

Table 304: Remote Availability Zone Latency is High [130187] [A130187]

Name	Latency to Remote Availability Zone is High
Description	Latency to the remote Availability Zone is high
Alert message	Latency to the Remote Availability Zone '<i>availability_zone_physical_name</i>' exceeds the threshold limit.
Cause	This might occur due to multiple factors.
Impact	Replications will fail
Resolution	Check if the Remote Availability Zone is reachable over the network.

Table 305: Cluster Join to Domain Failure [130189] [A1154]

Name	ClusterJoinToDomainFailure
Description	Failed to join AOS cluster to domain
Alert message	Failed to join AOS cluster to domain <i>domain_name</i>: <i>reason</i>
Cause	Cluster could not be domain joined because of incorrect logon account credential or other causes
Impact	Workflows such as backup service (volume shadow copy service) and Kerberos for SMB will not work.
Resolution	Ensure that the the domain account credential specified is valid and that the virtual computer object representing the AOS cluster does not exist

Table 306: Background Encryption Stuck [130194] [A130194]

Name	Background Encryption Stuck



Description	Background Encryption Stuck
Alert message	Encryption of container <code>container_id</code> is not progressing. Contact Nutanix support for further assistance.
Cause	Oplog entries, extent groups, NFS entries or other components are not being encrypted.
Impact	Encryption is not progressing.
Resolution	Contact Nutanix support.

Table 307: Failed to reserve host memory for Atlas networking [130202] [A130202]

Name	Failed to reserve host memory for Atlas networking
Description	Failed to reserve host memory for Atlas networking.
Alert message	Failed to reserve host memory for Atlas networking
Cause	Failed to reserve host memory for Atlas networking.
Impact	If hypervisor memory is exhausted, random VMs may be powered off to reclaim memory.
Resolution	Power off VMs to reduce host memory consumption.

Table 308: Pulse cannot connect to REST server endpoint [140001] [A140001]

Name	REST endpoint connection status
Description	Check if Pulse can connect to REST server endpoint.
Alert message	Pulse cannot connect to REST server endpoint. Connection Status: <code>connection_status</code>, Pulse Enabled: <code>enable</code>, Error Message: <code>message</code>
Cause	Pulse cannot connect to REST server endpoint.
Impact	Data driven serviceability and customer support cannot be performed.
Resolution	Ensure that the REST server endpoint is reachable from Pulse.

Table 309: OVS Service Restart [150002] [A150002]

Name	Ovs Service Restarted
Description	OVS Vswitch has restarted on one or more AHV hosts. Rebuilt bridges and flow rules on these hosts.
Alert message	Node <code>uuid_node</code>: detected unexpected restart of ovs-vswitchd process on host <code>host_ip</code> from old_pid <code>old_pid</code> to new_pid <code>new_pid</code>. OVS service configuration was restored after restart. Guest VMs might have experienced temporary network disconnect.
Cause	ovs-vswitchd restarted
Impact	AHV networking
Resolution	Rebuilt bridges and flow rules on hosts



Table 310: File server unreachable check [160008] [A160008]

Name	File Server Network Reachable
Description	Checks that File server is reachable
Alert message	File server <i>file_server_name</i> is unreachable
Cause	File server Network is unreachable
Impact	File server operations will not work.
Resolution	Contact Nutanix support.

Table 311: File server down check [160009] [A160009]

Name	File Server Reachable
Description	Checks that File server is down
Alert message	File server <i>file_server_name</i> is down
Cause	All File server VMs are powered off
Impact	File server operations will not work.
Resolution	Contact Nutanix support

Table 312: File server multiple VMs on single node check [160010] [A160010]

Name	File Server VM Status
Description	Checks that multiple File server VMs are running on a single node
Alert message	File server <i>file_server_name</i> has multiple VMs on single node
Cause	File server VMs are running on a single node
Impact	Single node failure may bring down the File server
Resolution	Contact Nutanix support.

Table 313: File server services down check [160011] [A160011]

Name	File Server Status
Description	Checks that all File server services are running
Alert message	File server <i>file_server_name</i> services (<i>service_names</i>) are down
Cause	One or more services are not running on File server VMs.
Impact	File server operations will not work.
Resolution	Contact Nutanix support.



Table 314: File Server PD activates on multiple sites [160022] [A160022]

Name	File Server PD Active On Multiple Sites
Description	Checks that File Server PD Active On Multiple Sites
Alert message	File server <i>file_server_name</i> protection domain <i>protection_domain_name</i> activates on multiple site(s) <i>remote_site_names</i>
Cause	File Server PD activates on multiple sites
Impact	File server operations might not work
Resolution	Deactivate PD on other sites

Table 315: File server invalid snapshot [160027] [A160027]

Name	File Server Invalid Snapshot Warning
Description	Checks that File server is in healthy state to take snapshot
Alert message	File server <i>file_server_name</i> is not healthy and leads to invalid snapshot
Cause	File Server NVMs might deleted
Impact	File server operations will not work and creates the bad snapshots
Resolution	Cancel the snapshot schedule
Resolution	Avoid taking snapshot and migration
Resolution	Contact Nutanix support.

Table 316: File server entities not protected [160029] [A160029]

Name	File Server Entities Not Protected
Description	Checks all File server entities are protected
Alert message	Some of File server <i>file_server_name</i> entities are not protected
Cause	Some of the File Server entities are not protected
Impact	File server migration might fails and creates the bad snapshots
Resolution	Contact Nutanix support.

Table 317: Multiple File server versions are present in the cluster. [160030] [A160030]

Name	Multiple File Server Versions Check
Description	Checks whether multiple File server versions are running in the cluster
Alert message	File Servers are on different versions. Please upgrade all File Servers to the same version as soon as possible.
Cause	File server(s) were powered off or not reachable during the last fileservers upgrade.



Impact	FileServer upgrade may not upgrade some file servers.
Resolution	Upgrade individual File Server.

Table 318: File server upgrade task hung for too long [160031] [A160031]

Name	File Server Upgrade Task Stuck Check
Description	Check File server task is hung for a long time.
Alert message	File server <i>file_server_name</i> upgrade task hung for too long
Cause	Various. Check alert message for details.
Impact	FileServer upgrade task could fail.
Resolution	Contact Nutanix support

Table 319: File Server PD enabled on non-compatible Remote Site [160033] [A160033]

Name	Remote Site Not File Server Capable
Description	Checks File server remote sites are file server capable
Alert message	File server <i>file_server_name</i> protection domain <i>protection_domain_name</i> enabled on non-compatible Remote site <i>remote_site_name</i>
Cause	Remote site is not File server capable
Impact	File server migration/failover will fail
Resolution	Upgrade remote site cluster AOS to 5.0 and above version

Table 320: Remote Site AOS not compatible with File Server version [160056] [A160056]

Name	Remote Site AOS not compatible with File Server version
Description	Checks File server remote sites AOS are file server version compatible
Alert message	File server <i>file_server_name</i> protection domain <i>protection_domain_name</i> enabled on non-compatible Remote site <i>remote_site_name</i>: <i>reason</i>
Cause	Remote site AOS is not compatible with File server version
Impact	File server migration/failover will fail
Resolution	Upgrade remote site cluster AOS to 5.6 and above version

Table 321: File Analytics VM component failure. [160064] [A160064]

Name	File Analytics VM component failure
Description	One or more components of the File Analytics VM {avm_ip} are not functioning properly or have failed.



Alert message	One or more components of the File Analytics VM <code>avm_ip</code> are not functioning properly or have failed.
Cause	One or more File Analytics VM components have failed.
Impact	File Server analytics will not perform appropriately.
Resolution	Contact Nutanix Support.

Table 322: File Analytics VM high CPU usage [160065] [A160065]

Name	File Analytics VM high CPU usage.
Description	File Analytics VM {avm_ip} has a high CPU usage of {usage_percent}%
Alert message	File Analytics VM <code>avm_ip</code> has a high CPU usage of <code>usage_percent%</code>
Cause	File Analytics VM is performing too many CPU intensive operations.
Impact	File Analytics VM's performance will be impacted.
Resolution	Monitor which component has a high CPU usage and reduce the processing load.

Table 323: File Analytics VM high disk usage. [160066] [A160066]

Name	File Analytics VM high disk usage.
Description	File Analytics VM {avm_ip} has a high disk usage of {usage_percent}%
Alert message	File Analytics VM <code>avm_ip</code> has a high disk usage of <code>usage_percent%</code>
Cause	Extensive disk usage on the File Analytics VM
Impact	File Analytics VM's performance will be impacted.
Resolution	Add more capacity to the File Analytics VM or delete some data.

Table 324: File Analytics VM low memory available. [160067] [A160067]

Name	File Analytics VM low memory available.
Description	The memory available on File Analytics VM {avm_ip} is {memory_available} bytes which is low.
Alert message	The percent memory available on File Analytics VM <code>avm_ip</code> is <code>memory_available</code> which is low
Cause	Excessive memory usage on the File Analytics VM
Impact	File Analytics VM's performance will be impacted.
Resolution	Add more RAM



Table 325: File server version check [160070]

Name	File server version check
Description	Check that file server runs GA version of Nutanix Files
Cause	Non-GA version of file server detected.
Impact	File server is on non-GA version.
Resolution	Follow the instructions in KB article 6601.

Table 326: File server VM Hardware Clock Timezone not Supported [160071] [A160071]

Name	File server VM Hardware Clock Timezone not Supported
Description	Checks that file server VM hardware clock timezone is UTC
Alert message	File server <i>file_server_name</i> : File server VM <i>vm_name_list</i> hardware clock timezone are not UTC.
Cause	File server VM hardware clock timezone is set to unsupported value
Impact	Services on file server VM may fail due to time drift
Resolution	Contact Nutanix support.

Table 327: Nutanix Files Version EOL [160073] [A160073]

Name	File server version EOL check
Description	File server(s) running versions which have reached end-of-life: {fs_eol_msg}
Alert message	File server(s) running versions which have reached end-of-life. <i>fs_eol_msg</i>
Cause	End-of-life Nutanix Files version running on cluster.
Impact	Support may not be able to resolve issues since this Nutanix Files version is no longer maintained.
Resolution	Upgrade to a supported version of files. Refer to the Nutanix Files EOL policy for clarification. Follow the instructions in KB article 6601. If you have questions regarding which Nutanix Files release you need to upgrade to, whether it should be long-term supported (LTS) or short-term supported (STS), please refer to the upgrade paths on the support portal.

Table 328: File Server upgrade task stuck [160075] [A160075]

Name	File Server Upgrade Task Stuck
Description	File Server Upgrade Task Stuck.
Alert message	Upgrade of file server <i>file_server_name</i> is stuck due to <i>reason</i>.
Cause	Check alert message for details
Impact	File server has not been upgraded.



Resolution	Please resolve the issue and run afs infra.resume_fs_upgrade on cvm to resume Nutanix Files upgrade.

Table 329: Files cluster HA takeover failure [160076] [A160076]

Name	Files cluster HA Takeover Failure
Description	Files cluster HA takeover failed
Alert message	Files cluster <i>file_server_name</i> HA takeover has failed multiple times on FSVM <i>fsvm_internal_ip</i> due to <i>message</i>.
Cause	Files cluster HA takeover process failed for multiple attempts
Impact	Files cluster may be unavailable.
Resolution	Contact Nutanix Support

Table 330: File server container dedup check [160077] [A160077]

Name	File Server Container Dedup Enabled
Description	Nutanix Files container has dedup policy enabled which is not recommended
Alert message	File server(s) with dedup policy enabled on storage container: <i>fs_dedup_msg</i>
Cause	On-disk dedup policy is enabled on Nutanix Files container which is not recommended
Impact	Enabling dedup policy on file server container may impact storage performance.
Resolution	Please follow the instructions in KB article 6944 to disable deduplication on Nutanix Files container

Table 331: File server VM down check [160078] [A160078]

Name	File Server VM Status
Description	Checks if all File server VMs are reachable
Alert message	File server <i>file_server_name</i>'s VMs (<i>svm_names</i>) are down
Cause	One or more VMs are down on File server.
Impact	File server operations will not work.
Resolution	Contact Nutanix support.

Table 332: File server clone grace period [160079] [A160079]

Name	File Server Clone Grace Period Check
Description	Check for licensing of cloned file servers



Alert message	Cloned file server(s) in grace period: <code>fs_clone_msg</code>
Cause	One or more cloned file servers are within grace period
Impact	Cloned file servers will continue to operate and will be considered for licensing after end of grace period
Resolution	Ensure there is suitable file server licensing available by the end of grace period

Table 333: File Server Alert [160080] [A160080]

Name	File Server Alert
Description	Files server alert
Alert message	<code>file_server_alert_msg</code>
Cause	Check alert description
Impact	Check alert description
Resolution	Check alert description

Table 334: Cluster Connectivity Status [200001] [A200001]

Name	IDF DB to DB sync heartbeat check
Description	Checks if IDF replica is receiving requests from IDF master.
Alert message	IDF data from cluster <code>cluster_name</code> (uuid = <code>cluster_uuid</code>) is not up-to-date.
Cause	Cluster network connectivity or CVM services such as insights server, insights uploader, insights receiver, Aplos, or Prism gateway could be down.
Impact	IDF data that is replicated from master cluster is not up to date.
Resolution	Ensure that cluster network connectivity is up and all CVM services are up.

Table 335: Dynamic scheduling failure. [200201] [A200201]

Name	Acropolis Dynamic Scheduler Status
Description	One or more nodes have resource contention. This imbalance can cause performance bottlenecks on the node(s) affected.
Alert message	Dynamic scheduling failure. <code>reason</code>
Cause	Cluster may have insufficient CPU or Controller VM resources.
Cause	Cluster may have insufficient resources to satisfy VM group affinity policies or VM host affinity policies.
Impact	Cluster performance may be significantly degraded.
Resolution	Shut down unneeded VMs to free cluster resources.
Resolution	Expand the cluster to add resources.



Table 336: Policy not applicable to any Host. [200305] [A200305]

Name	No viable Hosts to place VMs in PE
Description	Policy as defined does not apply to any of the Hosts.
Alert message	Rule <code>rule_uuid</code> applies to no Host in cluster <code>cluster_uuid</code>
Cause	Addition of a new policy(s) or update to existing policy(s) or Attaching/ Detaching categories to Hosts.
Impact	VMs will not have any Hosts to run on.
Resolution	Make appropriate updates to affinity policies or Host categories in the affected PE.

Table 337: VM forcibly powered off. [200401] [A200401]

Name	VM forcibly powered off
Description	Acropolis has forcibly powered off a VM that should have been powered off
Alert message	VM named <code>vm_name</code> with id <code>vm_id</code> has been forcibly powered off. <code>reason</code>
Cause	VM does not have any attached disk
Impact	None expected. The VM should not have been running.
Resolution	No action required.

Table 338: Restart VMs before Performing Upgrade or Migrate Operation [200404] [A200404]

Name	Restart VMs before Performing Upgrade or Migrate Operation
Description	Restart all the VMs before you perform the upgrade or migrate operation.
Alert message	Restart all the VMs '<code>reboot_vms</code>' before you perform the upgrade or migration operation
Cause	A node with a lower feature set has been added.
Impact	You will not be able to migrate the existing VMs or upgrade the cluster.
Resolution	Ensure that all the VMs in the cluster are restarted.

Table 339: Report Generation Failure [200501] [A200501]

Name	Report Generation Failure
Description	Report Generation failed. Database may be down or cluster may have insufficient storage.
Alert message	Report generation failed for report <code>report_name</code> scheduled for <code>timestamp</code>. <code>reason</code>
Cause	Cluster may have insufficient storage.
Impact	Some of the reports can be missed.



Resolution	Free the space by deleting older reports that have been generated.
Resolution	Manually run the report instance scanner utility.

Table 340: Send Report Through E-mail Failure [200502] [A200502]

Name	Send Report Through E-mail Failure
Description	Sending of report through E-mail failed. SMTP might be down or report might have been deleted.
Alert message	Sending of report through E-mail failed for report <code>report_name</code>. <code>reason</code>
Cause	SMTP server might be unreachable or might not have been configured.
Cause	Generated report may have been deleted.
Impact	Some of the report emails can be missed.
Resolution	Check the SMTP configuration and ensure that its reachable.

Table 341: Report Quota Scan Failure [200503] [A200503]

Name	Report Quota Scan Failure
Description	Report Quota scan failed. Cassandra might be unreachable.
Alert message	Report quota scan failed for the report <code>report_name</code> at <code>timestamp</code>. <code>reason</code>
Cause	Cassandra may not be reachable for some time.
Impact	There might be some stale reports on the cluster that are taking unnecessary space on the cluster.
Resolution	Manually run the report instance scanner utility..

Table 342: Availability Zone Connection Failure [200801] [A200801]

Name	Availability Zone Connection Failure
Description	The remote availability zone {az_url} is unreachable.
Alert message	The remote availability zone <code>az_url</code> is unreachable.
Cause	Services on remote site might be down.
Cause	Connection might be missing on remote site.
Cause	Access tokens might have expired.
Impact	Disaster recovery operations might fail due to unreachability of remote site.
Resolution	Unpair from the remote availability zone, if already paired.
Resolution	Pair with the remote availability zone again.

Table 343: PE-PC Connection Failure [200802] [A200802]

Name	PE-PC Connection Failure
Description	The remote {remote_vm_type} {cluster_ip} is unreachable.
Alert message	The remote <code>remote_vm_type</code> <code>cluster_ip</code> is unreachable.
Cause	The name servers might be missing.
Cause	Services on remote site might be down.
Cause	Access tokens might have expired.
Impact	Management tasks for various entities, including VM, might fail due to unreachability of remote {remote_vm_type}.
Resolution	Ensure that the appropriate name servers are configured.
Resolution	Reregister PE with PC.

Table 344: VM Recovery Storage Container(s) is/are not mounted on all hosts. [300431] [A300431]

Name	VM Recovery Storage Container Mount Check
Description	Verify that the Storage Container(s) is/are mounted on all the hosts.
Alert message	Container(s) <code>container</code> part of recovery point of entities which is/are not mounted on all the hosts.
Cause	Storage Container(s) is/are not mounted on all the hosts.
Impact	Registration of ESX VM(s) will fail during data recovery workflows.
Resolution	Mount the Storage Container(s) on all the hosts.

Table 345: Entity Sync Failure [500101] [A500101]

Name	Entity Sync failed
Description	Entity Sync failed
Alert message	<code>message</code>
Cause	Various
Impact	Data Protection and Recovery might be impacted.
Resolution	Perform Forced Entity Sync after ensuring that all the remote entities can be overwritten. If the issue still persists, contact Nutanix support.

Table 346: Entity Sync Failure for the Protection Policy [500102] [A500102]

Name	Entity Sync failed for Protection Policy
Description	Entity Sync failed for the Protection Policy
Alert message	<code>message</code>



Cause	Various
Impact	Data Protection and Recovery might be impacted.
Resolution	Perform Forced Entity Sync after ensuring that all the remote entities can be overwritten. If the issue still persists, contact Nutanix support.

Table 347: Entity Sync Failure for the Recovery Plan [500103] [A500103]

Name	Entity Sync failed for Recovery Plan
Description	Entity Sync failed for the Recovery Plan
Alert message	message
Cause	Various
Impact	Data Protection and Recovery might be impacted.
Resolution	Perform Forced Entity Sync after ensuring that all the remote entities can be overwritten. If the issue still persists, contact Nutanix support.

Table 348: Entity Sync Failure for the Availability Zone [500104] [A500104]

Name	Entity Sync failed for Availability Zone
Description	Entity Sync failed for the Availability Zone
Alert message	message
Cause	Various
Impact	Data Protection and Recovery might be impacted.
Resolution	Ensure Availability Zone is reachable and do a 'Forced Entity Sync' to ensure that all the entities are in sync. If the issue still persists, contact Nutanix support.

Table 349: Translated Addresses Retrieval Failure in NGT [600102] [A600102]

Name	Unable to get translated addresses in NGT
Description	NGT is unable to retrieve the XAT address translations.
Alert message	Unable to retrieve translated addresses in NGT because the Prism Element IDF service may be unreachable.
Cause	Prism Element IDF may be unreachable.
Cause	Translation information is not available in IDF.
Impact	NGT features may not work on the User VMs.
Resolution	Check if the Prism Element IDF service is reachable.



Table 350: Security Planning is disabled [803001] [A803001]

Name	Security Planning is disabled
Description	Security Planning is disabled due to de-registration of Prism Element from Prism Central.
Alert message	Security Planning is disabled. Reason: details
Cause	The cluster that is hosting the Epoch data collector was de-registered from Prism Central.
Impact	The Security Planning feature has been automatically disabled.
Resolution	Enable Security Planning from the Prism Central UI to deploy a new instance of the Epoch data collector.

Table 351: Epoch SaaS service connectivity lost [803002] [A803002]

Name	Epoch SaaS connectivity is lost
Description	Epoch SaaS service can't be reached.
Alert message	Epoch SaaS service connectivity has been lost. Details: details.
Cause	The network connectivity between Prism Central and Epoch SaaS service is not reliable.
Cause	Epoch SaaS service may be experiencing issues.
Cause	The saved access key or info for Epoch SaaS service may need to be updated.
Impact	The Prism Central Security Planning feature will not function properly.
Resolution	Check the connectivity between Prism Central and Epoch SaaS service.
Resolution	Check that the entered Epoch credentials in Prism Central are valid and up-to-date.
Resolution	Check that the Epoch SaaS service is accessible and running normally by logging onto the service.

Table 352: ID Firewall lost connectivity to domain controller [803003] [A803003]

Name	ID Firewall is unable to connect to one or more domain controllers
Description	ID Firewall is unable to connect to one or more domain controllers used for scraping identities.
Alert message	Domain controller domain_controller on domain domain is not reachable and accepting LDAP or WMI connections. Details: details.
Cause	The domain controller may not be reachable via the network.
Cause	The domain controller may not be accepting LDAP/WMI connections due to permissions issues.
Cause	The domain controller may not be running or fully booted up.
Impact	ID Firewall will not function properly and Flow VDI security policies may not be enforced until access to the domain controller is restored.



Resolution	Check network connectivity from Prism Central to the domain controller.
Resolution	Check that the domain controller is fully up and accepting WMI and LDAP connections.
Resolution	Check that the service account registered for ID Firewall has permissions for both WMI and LDAP access, and that the service account is active and that the correct password is entered in Prism Central.

Table 353: ID Firewall unable to locate mapped Active Directory object [803004] [A803004]

Name	ID Firewall mapped object not found
Description	ID Firewall was unable to find a mapped object in Active Directory.
Alert message	The mapped <code>object_type</code> <code>object_dn</code> with GUID <code>object_guid</code> was not found in <code>domain</code>.
Cause	The mapped object in Active Directory which is referenced for ID Firewall could not be found because it was deleted from Active Directory.
Impact	Flow VDI security policies written around the mapped object will not function as expected.
Resolution	The mapping should be checked, and if the mapped object does not exist, or if it has been deleted and re-created, the mapping must also be deleted and possibly re-created. Security policies should be updated to use the new mapping or not reference the mapping at all if the object was permanently deleted.

Table 354: ID Firewall did not recover state after reconnecting to a temporarily unreachable domain controller [803005] [A803005]

Name	ID Firewall state recovery failed
Description	ID Firewall was unable to recover state after reconnecting to a temporarily unreachable domain controller.
Alert message	ID Firewall was unable to recover state from domain controller <code>domain_controller</code> on domain <code>domain</code>. Details: <code>details</code>.
Cause	Recovery failed because either too much time has passed since the domain controller was last reachable, or the event log on the domain controller has rolled over and the events necessary to recover are no longer present.
Impact	The state enforced by ID Firewall may be inconsistent, and therefore applied policies may not be properly enforced until action is taken.
Resolution	In order to ensure security policies are properly enforced, all active Nutanix VDI VM users should log out and log back into their VMs.

Table 355: Epoch Data Collector upgrade available [803006] [A803006]

Name	Epoch Data Collector needs upgrade
Description	Checks whether there is an upgrade available for the Epoch Data Collector.



Alert message	message
Cause	A new version of the Epoch Data Collector is available and not yet installed.
Impact	Some new functionality or bugfixes in the Epoch Data Collector may not be available.
Resolution	Go to Epoch Data Collector Enablement page to upgrade the Epoch Data Collector.

Controller VM

Table 356: CVM Boot Raid Degraded [1019] [A1107]

Name	Boot RAID Health
Description	Check the health of boot RAID volumes.
Alert message	Boot RAID volumes on CVM <code>service_vm_external_ip</code> are degraded.
Cause	RAID volume can be in an abnormal state due to power loss or temporary or permanent drive failure.
Impact	This might be due to a boot drive failure. The CVM might become inoperable if the issue is not resolved.
Resolution	Verify boot disk status from the web console Hardware page and replace the disk if required.

Table 357: Abnormal Host Boot RAID State [1020] [A1105]

Name	Host Boot RAID State
Description	Check the health of boot RAID volumes.
Alert message	Hypervisor RAID volume '<code>volume_label</code>' is in abnormal state '<code>volume_state</code>' on host <code>host_ip</code>.
Cause	RAID volume can be in an abnormal state due to power loss or temporary or permanent drive failure.
Impact	This might be due to a boot drive failure. The CVM might become inoperable if issue is not resolved.
Resolution	Use remote console to verify the RAID status to decide if replacement of drive is required.

Table 358: CVM or Prism Central VM RAM Usage High [3023] [A1056]

Name	CVM Memory Usage
Description	Check that CVM or Prism Central VM memory usage is not high.
Alert message	Main memory usage in Controller VM or Prism Central VM <code>ip_address</code> is high, <code>available_memory_kb</code> KB of memory is free.
Cause	The RAM usage on the Controller VM or Prism Central VM has been high.



Impact	Cluster performance may be significantly degraded.
Resolution	Check the memory utilisation of Prism Central VM or Controller VM. If abnormal behaviour is seen, please collect logs and contact Nutanix Support.
CVM Available Memory Threshold	Critical Threshold: 768000 KB

Table 359: The {vm_type} is not synchronizing time with any external servers. [3026] [A3026]

Name	CVM NTP Time Synchronized
Description	Checks to ensure that the Controller VM is synchronizing time with an NTP server.
Alert message	The {vm_type} is not synchronizing time with any external servers. <i>alert_msg</i>
Cause	External NTP servers are not configured or are not reachable
Impact	Workflows involving Kerberos may fail if the time difference between the Controller VM and the NTP server is greater than 5 minutes.
Resolution	Verify that the external NTP servers are configured and reachable.

Table 360: Kernel Memory Usage High [3027] [A1034]

Name	Kernel Memory Usage
Description	Checks if CVM's kernel memory usage is higher than the configured threshold
Alert message	Controller VM {ip_address}'s kernel memory usage is higher than configured threshold. Cluster performance will be significantly degraded
Cause	Various
Impact	Cluster performance will be significantly degraded. In the case of multiple nodes with the same condition, the cluster will be unable to service I/O requests.
Resolution	Contact Nutanix support.
Kernel Slab Cache Lost Size Maximum Threshold	Critical Threshold: 15 %

Table 361: CVM Rebooted [3028] [A1024]

Name	CVM Rebooted Check
Description	Check that Cvm is not rebooted recently
Alert message	{vm_type} {ip_address} has been rebooted on {reboot_timestamp_str}.



Cause	Cvm is rebooted.
Impact	During the time the Controller VM is down, cluster compute and storage capacity are reduced.
Resolution	Check cvm status.

Table 362: CVM Service(s) Restarting Frequently [3029] [A1032]

Name	CVM Services Status
Description	Check if services have restarted recently on the CVM.
Alert message	One or more cluster services have restarted atleast 10 times in 15 minutes on Controller VM <i>ip_address</i>.Following service(s) have restarted : <i>failed_components</i>.
Cause	Faulty behaviour in one or more service(s), leading to frequent crashes,in a short period of time.
Impact	Cluster performance may have significantly degraded. In case of multiple nodes with the same condition, the cluster may lose the ability to serve I/O requests.
Resolution	If this alert is generated frequently, please collect CVM logs and contact Nutanix support.

Table 363: CVM Python Service(s) Restarting Frequently [3037] [A3037]

Name	Python Services Fatal Check
Description	Check if python services have crashed/restarted recently in PCVM/CVM
Alert message	One or more cluster services have restarted within 15 minutes in the PCVM/CVM <i>ip_address</i>, the total restarts being ≥ 10. The following are the service(s) that have restarted : <i>failed_components</i>.
Cause	Faulty behaviour in one or more service(s) leading to frequent crashes in a short period of time.
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	If this alert occurs once or infrequently, no action is necessary. If this alert occurs frequently, contact Nutanix support.

Table 364: CVM CPU Utilization [3041]

Name	CVM CPU Utilization
Description	Checks whether CPU usage is high on CVM.
Cause	CPU usage is high.
Impact	High I/O latency may be experienced by some workloads.



Resolution	Reduce CPU-intensive processes.

Table 365: Multiple cpu_unblock processes running [3042] [A3042]

Name	cpu_unblock Process Check
Description	Check that there are no stale cpu_unblock processes running.
Alert message	Multiple cpu_unblock processes are running on <code>svm_ip</code>.
Cause	Zookeeper restarting frequently on cluster node.
Impact	Cluster performance may be significantly degraded.
Resolution	Kill all cpu_unblock processes and restart cluster services on node.

Table 366: Jumbo Frames Enabled for NIC {nic_name} on {service_vm_external_ip} [3063] [A1095]

Name	NIC MTU Configuration
Description	Check MTU of the CVM network interfaces.
Alert message	Controller VM <code>service_vm_external_ip</code> has NIC <code>nic_name</code> with MTU set to <code>mtu</code> instead of <code>desired_mtu</code>.
Cause	MTU is not set correctly on eth0/eth1
Impact	If MTU is not set correctly in the Controller VMs, all Nutanix services timeout causing the storage to be unavailable.
Resolution	Set MTU to the right value

Table 367: CVM Connectivity Failure [3064] [A1001]

Name	CVM Passwordless Connectivity
Description	Check that CVM has passwordless connection to each other.
Alert message	Controller VM <code>target_ip</code> is not reachable from Controller VM <code>source_ip</code> in the last 3 attempts.
Cause	CVM has no passwordless connection to each other
Impact	Cluster compute and storage capacity are reduced. Until data stored on this host is replicated to other hosts in the cluster, the cluster has one less copy of guest VM data.
Resolution	Check CVM network configuration

Table 368: MTU configuration across Controller VMs is not consistent. [3068] [A3068]

Name	NIC MTU Ping Check
Description	Check connectivity between Controller VMs by pinging with configured MTUs.



Alert message	Failed to ping other CVMs with MTU set to <code>mtu</code> on Controller VM <code>service_vm_external_ip</code> .
Cause	MTU is not consistent across Controller VM and host external interfaces, or upstream network.
Impact	An inconsistent MTU configuration can lead to network instability across Controller VMs and hosts, which in certain conditions can lead to storage availability issues.
Resolution	Set MTU consistently across Controller VM and host external interfaces. Ensure upstream physical network MTU supports the configuration.

Table 369: NIC MTU Uniformity Check [3069]

Name	NIC MTU Uniformity Check
Description	Check if the MTU setting is the same on the matching external interface across every CVM in the cluster.
Cause	MTU is not uniform across the CVMs' management interface
Impact	If MTU is not the same across Controller VM interfaces which are in the same layer 2 broadcast domain or VLAN, some packets could be dropped which could cause the storage to be unavailable.
Resolution	Set MTU on the matching external interface uniformly across the CVMs in the cluster

Table 370: CVM NIC Speed Low [6008] [A1058]

Name	TenGig Adaptor Status
Description	Checks whether CVM is uplinked to 10 GbE NIC.
Alert message	<code>alert_msg</code>
Cause	CVM not uplinked to 10 GbE NIC.
Cause	1 GbE NIC is part of the bond.
Cause	If hypervisor is ESXi, it may be in lockdown.
Impact	The Controller VM is not configured to use the 10 GbE NIC or is configured to share load with a slower NIC.
Resolution	Check network configuration.
Resolution	If hypervisor is ESXi, check for lockdown and unlock it.

Table 371: CVM Host Subnet Mismatch [6202] [A1048]

Name	CVM Subnet
Description	Checks that host and cvm share the same subnet.

Alert message	Controller VM <code>svm_ip</code> with network address <code>svm_subnet</code> is in a different network than the Hypervisor <code>hypervisor_ip</code> , which is in the network <code>hypervisor_subnet</code> .
Cause	The Controller VM and the hypervisor are not on the same subnet.
Impact	Controller VM high availability will not apply in the case of Controller VM failure, leading to guest VM unavailability.
Resolution	Reconfigure the cluster. Either move the Controller VMs to the same subnet as the hypervisor hosts or move the hypervisor hosts to the same subnet as the Controller VMs.

Table 372: Invalid SSD/HDD Drive Combination [6209] [A1164]

Name	Invalid Drive Configuration
Description	Checks whether combination of SSDs and HDDs on the node is a supported configuration.
Alert message	Invalid SSD/HDD combination seen on <code>ip_address</code> - <code>num_ssd</code> SSDs (<code>num_nvme</code> NVMeS) and <code>num_hdd</code> HDDs.
Cause	The combination of SSDs and HDDs on the node is not supported configuration.
Impact	Cluster performance will be significantly degraded.
Resolution	Add or remove SSDs or HDDs until their combination is a supported configuration.

Table 373: Node Avg Load - Non Critical [6511]

Name	Node Avg Load - Non Critical
Description	Checks whether the system load average is high.
Cause	I/O workload on the cluster is high.
Impact	High I/O latencies may be experienced by some workloads.
Resolution	Redistribute VMs to reduce load on the affected node.
Resolution	Tune the applications to reduce CPU demand.
Resolution	Reduce the number of VMs on the host.

Table 374: HDD IO Latency [6514]

Name	HDD IO Latency
Description	Verify whether the average time for I/O requests to the HDDs managed by the Controller VM is not high.
Cause	I/O workload on the cluster is high.
Impact	High I/O latencies may be experienced by some workloads.



Resolution	Redistribute VMs to reduce load on the affected Controller VM.

Table 375: CPU Average Load Critically High on Controller VM. [6516] [A6516]

Name	Node Avg Load - Critical
Description	Checks whether the system load average is critically high.
Alert message	Average CPU load on Controller VM <i>ip_address</i> is critically high (above critical average load threshold value <i>critical_high_value</i>)
Cause	Critically high load on the system.
Impact	Low performance may be experienced due to cluster-wide I/O latency.
Resolution	Contact Nutanix Support for assistance.

Table 376: Susceptible to FA61 check [20021]

Name	Susceptible to FA61 check
Description	Check if cluster is susceptible to FA61
Cause	Cluster is running version of AOS susceptible to FA61 and has EC enabled
Impact	Metadata Corruption
Resolution	Refer to FA61 guidance. Upgrade to version AOS-5.9.2.3 or higher for a critical fix which prevents data loss

Table 377: Affected by FA61 check [20022]

Name	Affected by FA61 check
Description	Checks if the cluster is affected by FA61
Cause	Signature of FA61 found in Curator logs or Curator returned non-zero value for InternalError on current or previously run susceptible AOS version
Impact	Metadata Corruption
Resolution	Refer to FA61 guidance. Upgrade to version AOS-5.9.2.3 or higher for a critical fix which prevents data loss

Table 378: Stargate's memory fragmentation check [20023]

Name	Stargate's memory fragmentation check
Description	Checks if stargate's memory fragmentation should be below certain threshold.
Cause	If memory fragmentation is high, this will result in lower cache hits and hence degrade performance.
Impact	Cluster performance may be significantly degraded.
Resolution	Contact Nutanix Support for diagnosis.



Table 379: Single active root partition check [101063]

Name	Single active root partition check
Description	Verifies that there only one active root partition in the CVM.
Cause	The storage partitioning on the CVM may have been modified manually after installation. A normal CVM should have one active root partition.
Impact	AOS upgrade will fail.
Resolution	Contact Nutanix Support.

Table 380: Root partition health check [101064]

Name	Root partition health check
Description	Checks the health of the RAID device backing the root partition.
Cause	A component in the node may have been replaced.
Impact	AOS upgrade will fail.
Resolution	Contact Nutanix Support.

Table 381: RAID health checks [101065]

Name	RAID health checks
Description	Checks the health of all the RAID devices in the CVM.
Cause	One or more storage devices on the node are failing.
Impact	AOS upgrade will fail.
Resolution	Contact Nutanix Support.

Table 382: Core dumps are enabled on this CVM or Prism Central VM. [101069] [A101069]

Name	Saltstate coreoff check
Description	Checks if saltstate is correctly set to default coreoff state to ensure services do not produce core dumps unnecessarily.
Alert message	Core dumps are enabled on the <code>vm_type</code> with IP <code>service_vm_external_ip</code>.
Cause	Core dumps are enabled for services running on this CVM or Prism Central VM.
Impact	Core dumps are useful for isolated troubleshooting but leaving enabled may unnecessarily fill the node /home system partition, and this can affect node stability and availability.
Resolution	Follow KB7586 to validate and fix coreoff saltstate. Resolution may involve graceful node reboot, so ensure the cluster is otherwise healthy and resiliency status is OK prior to this action.



Table 383: SED Keys Unavailable [101071] [A101071]

Name	SED Keys Unavailable
Description	Check to see if the KMS actually has the expected passwords for the SED disks.
Alert message	SED keys for one or more disks are not available from external key manager <i>kms_name</i>.
Cause	If a 3rd party KMS gets restored from a backup or other incorrect management operation, it may have missing keys.
Impact	Cluster will have data unavailability if rebooted.
Resolution	Contact Nutanix and 3rd party key manager support.

Table 384: The hypervisor is not synchronizing time with any external servers. [103090] [A103090]

Name	Hypervisor time synchronized
Description	Checks to ensure that the hypervisor is synchronizing time with an NTP server.
Alert message	The hypervisor is not synchronizing time with any external servers. <i>alert_msg</i>
Cause	External NTP servers are not configured or are not reachable.
Impact	Logs may have different timestamps in the hypervisor and the CVMs. Hypervisor may not work as expected.
Resolution	Verify if the NTP servers are configured and reachable from the hypervisor.

Table 385: Check for whitespace in network configuration. [103096]

Name	Check for whitespace in network configuration.
Description	Check if Network configuration files (ifcfg-ethX and /etc/hosts) are malformed, with whitespaces inserted.
Cause	Network configuration files are malformed, with whitespaces inserted.
Impact	Genesis may crash, causing services to restart.
Resolution	Follow the instructions in KB 5796.

Table 386: Check for empty network configuration files. [103097]

Name	Check for empty network configuration files.
Description	Check if there are empty network configuration files (ifcfg-ethX and /etc/hosts).
Cause	Network configuration files are malformed, with empty files created.
Impact	Genesis may crash, causing services to restart.



Resolution	Follow the instructions in KB 5796.

Table 387: Disk firmware needs upgrade [106059] [A106059]

Name	Disk Firmware Check
Description	Checks if any of the disks passthrough to CVM, requires firmware upgrade or not
Alert message	<code>alert_msg</code>
Cause	Current disk firmware is not recommended by Nutanix
Impact	Disk may not function properly, resulting in data unavailability.
Resolution	Upgrade Disk Firmware version. Perform inventory from Life Cycle Manager to check and update to the latest firmware version. Check KB 6937 for alternative update option.

Table 388: Controller VM Certificate Expiring [110214] [A1114]

Name	SED Node Certificate
Description	Check if node certificates are about to expire.
Alert message	Certificate for key management server <code>key_management_server_name</code> for Controller VM <code>service_vm_external_ip</code> expires on <code>expiration_date</code>.
Cause	Certificates have defined expiration dates.
Impact	If the node's certificate expires then it will be unable to authenticate with external servers.
Resolution	Get a new signed certificate.

Table 389: Node in Maintenance Mode [130001] [A1013]

Name	Node Maintenance Mode Status
Description	Node in Maintenance Mode.
Alert message	Controller VM <code>ip_address</code> is put in maintenance mode due to <code>reason</code>.
Cause	Node removal has been initiated.
Impact	The node will be removed from the cluster.
Resolution	No action is necessary.

Table 390: Kerberos Clock Skew Failure [130024] [A1083]

Name	Kerberos Clock Skew Status
Description	Kerberos Clock Skew Failure



Alert message	Kerberos authentication failure occurred on Controller VM <code>ip_address</code> due to clock skew between the Controller VM and the Domain Controller (reason: <code>reason</code>).
Cause	There is considerable clock skew between the Controller VM and the Domain Controller.
Impact	The Controller VM cannot manage the host, which may lead to node instability.
Resolution	Ensure that the time on the Controller VM is synchronized with the time of the host. It can be accomplished by configuring the NTP server properly.

Table 391: Metadata Volume Snapshot Persistent Failure [130026] [A1087]

Name	Metadata Volume Snapshot Persistent
Description	Metadata Volume Snapshot Persistent Failure
Alert message	The last <code>failure_count</code> metadata volume snapshots have failed. <code>reason</code>
Cause	The EBS snapshot service is unavailable, possibly because AWS could not be reached.
Impact	If the cloud instance fails, backups of the guest VM may be unavailable.
Resolution	Contact Nutanix support.

Table 392: Metadata Volume Snapshot Timeout [130027] [A1088]

Name	Metadata Volume Snapshot Status
Description	Metadata Volume Snapshot Timeout Failure
Alert message	Metadata volume snapshot could not be created in <code>duration_mins</code> minutes.
Cause	Metadata volume snapshot frequency is too high.
Impact	If the cloud instance fails, backups of the guest VM may be unavailable.
Resolution	Contact Nutanix support.

Table 393: Physical Disk Added To Slot [130034] [A1102]

Name	Physical Disk Add Check
Description	Physical Disk added to slot.
Alert message	Disk with serial <code>disk_serial</code> and model <code>disk_model</code> was added in drive bay <code>disk_location</code> on Controller VM <code>service_vm_external_ip</code>. Action Required: <code>action</code>. Visit Prism hardware page for further actions.
Cause	A drive was inserted into an empty slot.
Impact	The drive may not be available to the cluster until further configuration is completed.
Resolution	Ensure that the drive is added correctly from the Hardware page.



Table 394: Physical Disk Drive Has Failed [130035] [A1104]

Name	Physical Disk Status
Description	Physical drive has failed.
Alert message	Drive <code>disk_id</code> with serial <code>disk_serial</code> and model <code>disk_model</code> in drive bay <code>disk_location</code> located on node <code>node_serial_number</code> in position <code>node_position</code> on Controller VM <code>service_vm_external_ip</code> has failed.
Cause	The drive has failed.
Impact	Cluster storage capacity is reduced.
Resolution	Replace the failed drive. Please refer to KB 4158.

Table 395: Physical Disk Removed From Slot [130036] [A1103]

Name	Physical Disk Remove Check
Description	Physical disk removed from slot
Alert message	Disk with serial <code>disk_serial</code> and model <code>disk_model</code> was removed from drive bay <code>disk_location</code> on Controller VM <code>service_vm_external_ip</code>
Cause	A drive was physically disconnected from a slot.
Impact	Migration of data from the disk will start.
Resolution	Reseat the removed disk back in slot. Contact the Hardware Vendor if issue persists.

Table 396: Self Encrypting Drive Operation Failure [130051] [A1111]

Name	SED Operation Status
Description	Self encrypting drive operation failure
Alert message	Self Encrypted Drive operation <code>operation</code> has failed for disk <code>disk_id</code> with serial <code>disk_serial</code> on node <code>service_vm_external_ip</code> .
Cause	A self encrypting drive operation could not be performed.
Impact	The desired action could not be performed.
Resolution	Retest key management server configuration to ensure connectivity and certificates are valid then retry the command, or contact Nutanix support.

Table 397: Stargate Temporarily Down [130054] [A1030]

Name	Stargate Status
Description	Stargate Temporarily Down.
Alert message	Stargate on Controller VM <code>ip_address</code> is down for <code>downtime</code> seconds.
Cause	Various



Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Contact Nutanix support.

Table 398: Unsupported Configuration For Redundancy Factor 3 [130055] [A1092]

Name	FT2 Configuration
Description	Unsupported Configuration For Redundancy Factor 3.
Alert message	Controller VM <code>service_vm_id</code> with IP address <code>service_vm_external_ip</code> has <code>actual_ram_size_gb</code>GB RAM which does not meet the configuration requirement of <code>min_ram_required_gb</code>GB RAM to support 'Redundancy Factor 3' feature.
Cause	To support 'Redundancy Factor 3' feature all controller VMs in the cluster must meet the minimum requirements for RAM.
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Increase RAM on the Controller VM to meet the minimum requirement. Contact Nutanix support for assistance.

Table 399: Stargate Responsive [130069]

Name	Stargate Responsive
Description	Check if any Stargate process is still running after being killed.
Cause	Stargate got into an unresponsive state.
Impact	Nutanix datastore may become inaccessible.
Resolution	Restart the Stargate process.

Table 400: Disk Diagnostic Failure [130089] [A1139]

Name	Disk Diagnostic Status
Description	The drive diagnostic test has failed.
Alert message	Drive <code>disk_id</code> with model <code>disk_model</code> with serial <code>disk_serial</code> in drive bay <code>disk_location</code> on Controller VM <code>service_vm_external_ip</code> has failed diagnostic test. This drive is failing or has failed.
Cause	The drive is failing or has failed.
Impact	Cluster storage capacity has been reduced.
Resolution	Replace the failing drive. Refer to the Nutanix documentation for instructions.



Table 401: File server VM CPU Utilization [160060]

Name	File server VM CPU Utilization
Description	Checks whether CPU usage is high on file server VM.
Cause	CPU usage of one or all of file server VMs is high.
Impact	Some workloads may encounter high I/O latency.
Resolution	Ensure that any 3rd-party load balancer is configured according to the vendor's recommendations.

DR

Table 402: VM is protected in multiple PDs. [110217] [A110217]

Name	Protection Domains sharing VMs
Description	VMs are protected in multiple Protection Domains.
Alert message	VM <code>vm_name</code> is protected in multiple protection domains <code>pd_names</code>.
Cause	VMs are protected in multiple Protection Domains.
Impact	Recovery of VMs will fail.
Resolution	Verify that VMs are not protected in multiple Protection Domains.

Table 403: NGT installation required. [110242] [A110242]

Name	Cross Hypervisor NGT Installation Check
Description	Some VMs in protection domain replicating to cross-hypervisor remote site do not have NGT installed.
Alert message	Some VMs <code>vm_list</code> in cross-hypervisor protection domain <code>pd_name</code> do not have NGT installed.
Cause	Some VMs in protection domain replicating to cross-hypervisor remote site do not have NGT installed.
Impact	Recovery of VMs on the remote site might fail.
Resolution	Verify that NGT is installed on all the VMs of the protection domain that you want to recover on the remote site.

Table 404: Too many files in the Consistency Group. [110243] [A110243]

Name	Metro/Vstore Consistency Group File Count Check
Description	Checks if too many files are being protected by a single consistency group of any Metro/Vstore protection domain.
Alert message	The consistency group <code>cg_name</code> (part of protection domain <code>pd_name</code>) protects <code>num_files</code> files, which exceeds the currently set threshold of <code>threshold_num_files</code>.



Cause	Too many files are being protected by a single consistency group of a Metro/Vstore protection domain.
Impact	Performance on Metro-protected or Vstore-protected Storage Container may be degraded.
Resolution	Delete some VMs/files from the consistency group, or move to another Storage Container and protect the new Storage Container.
Resolution	Change alert threshold from the alert policies section.
Resolution	If needed, contact Nutanix Support for assistance.
Vstore consistency group file count threshold	Warning Threshold: 1200 files

Table 405: Too many files in the Protection Domain. [110244] [A110244]

Name	Metro/Vstore Protection Domain File Count Check
Description	Checks if too many files are being protected by any Metro/Vstore protection domain.
Alert message	The protection domain <code>pd_name</code> protects <code>num_files</code> files, which exceeds the currently set threshold of <code>threshold_num_files</code>.
Cause	Too many files are being protected by a Metro/Vstore protection domain.
Impact	Performance on Metro-protected or Vstore-protected Storage Container may be degraded.
Resolution	Delete some VMs/files from the protection domain, or move to another Storage Container and protect the new Storage Container.
Resolution	Change alert threshold from the alert policies section.
Resolution	If needed, contact Nutanix Support for assistance.
Vstore protection domain file count threshold	Warning Threshold: 3600 files

Table 406: Found old clones on cluster. [110245] [A110245]

Name	Clone Age Check
Description	Check for any clones restored from protection domain snapshots that are too old.
Alert message	Found <code>count</code> clones related to protection domain <code>pd_name</code> which are older than the currently configured threshold of <code>clone_age_threshold</code> days.
Cause	Clones were not removed automatically.
Impact	Storage performance may be degraded.



Resolution	Remove any clones restored from protection domain snapshots that are too old.
Age threshold for clones restored from protection domains.	Warning Threshold: 7 days

Table 407: Too many clones on cluster. [110246] [A110246]

Name	Clone Count Check
Description	Check if there are too many clones restored from protection domain snapshots.
Alert message	Found count clones related to protection domain pd_name which is more than the currently configured threshold of clone_count_threshold .
Cause	Clones were not removed automatically.
Impact	Storage performance may be degraded.
Resolution	Remove any clones restored from protection domain snapshots that are too old.
Count threshold for clones restored from protection domains.	Warning Threshold: 3 clones

Table 408: Protecting VMs that are using shared VHDX disks is unsupported. [110247] [A110247]

Name	Check VHDX Disks
Description	Check if VMs with shared VHDX disks are not part of any protection domain.
Alert message	Protecting VMs that are using shared VHDX disks is unsupported.
Cause	VMs with shared VHDX disks are protected, which is an unsupported configuration.
Impact	Snapshotting shared VHDX disks that are connected to protected VMs may fail.
Resolution	Unprotect VMs that use shared VHDX disks.

Table 409: Symlinks found on metro/vstore protected container. [110248] [A110248]

Name	Metro Vstore Symlinks Check
Description	Check for symlinks in metro/vstore protection domain.



Alert message	Symlinks found on metro/vstore protected container <code>ctr_name</code>, which might lead to cluster instability.
Cause	Symlinks present in metro/vstore protection domain.
Impact	Snapshot operations on the metro/vstore protected container may fail, leading to replication failures.
Resolution	Remove symlinks from metro/vstore protected container.

Table 410: Remote Stargate Version Check [110249]

Name	Remote Stargate Version Check
Description	Check for the AOS version running on Metro remote site.
Cause	Stargate running on Metro remote site is out of date.
Impact	Metro cluster may not be available due to crashing Stargate.
Resolution	Contact Nutanix support.

Table 411: Aged third-party backup snapshots present [110250] [A110250]

Name	Aged Third-party Backup Snapshot Check
Description	Check for aged third-party backup snapshots.
Alert message	<code>pd_name</code> has <code>num_snapshot</code> aged third-party backup snapshot(s) and they may unnecessarily consume storage space in the cluster.
Cause	Third-party backup snapshots are present in the cluster longer than the configured threshold.
Impact	Aged snapshots may unnecessarily consume storage space in the cluster.
Resolution	Contact Nutanix support.
Age threshold for third-party backup snapshots.	Warning Threshold: 7 days

Table 412: Replications are scheduled on entitites from the storage containers that have deduplication enabled. [110251] [A110251]

Name	Replication of deduped entities to non-compliant remote sites
Description	Replications are scheduled on entitites from the storage containers that have deduplication enabled
Alert message	Protection domain <code>pd_name</code> with entities from the storage containers <code>dedup_ctrs</code> that have deduplication enabled are being replicated to single-node backup or cloud remote sites <code>remote_names</code>.



Cause	Protection domain with entities from the storage containers that have deduplication enabled are being replicated to a single-node backup or cloud remote site.
Impact	Performance of replications and retrievals will be affected.
Resolution	Either disable deduplication on the storage containers, or use a multi-node AOS cluster as a replication target.

Table 413: Protection Domain activation may fail as conflicting files exist. [110255] [A110255]

Name	Protection Domain File Conflict Check
Description	Check if any existing file turns into a conflicting state after activating an inactive Protection Domain.
Alert message	Files <code>conflicting_files</code> conflict with activation of Protection Domain <code>pd_name</code>. All entities may not be recovered. Please remove these files before activating this Protection Domain.
Cause	This might occur if some entities were unprotected from the Protection Domain before making it inactive or some files were created outside the scope of the protection domain.
Impact	Some entities may not be recovered after you activate the Protection Domain.
Resolution	Remove all the conflicting files.

Table 414: Unplanned Failover and Failback can Cause Full Replication. [110259] [A110259]

Name	Nearsync unplanned failover check
Description	Unplanned failover and failback can cause full replication.
Alert message	Unplanned failover and failback can result in full replication of the protection domain <code>pd_name</code>. See KB 5695 for more information.
Cause	Unplanned failover and failback can cause full replication.
Impact	Activate of a protection domain and subsequent replication to the original site can result in a full replication.
Resolution	See KB 5695 for detailed information.

Table 415: Remote Replication is Lagging for Protection Domain Snapshot [110260] [A110259]

Name	PD Snapshot Lagging Remote Replication
Description	Remote replication of the protection domain snapshot is lagging.
Alert message	Snapshot '<code>snapshot_uuid</code>' of the protection domain '<code>protection_domain_name</code>' has replication to remote '<code>remote_name</code>' that is lagging by <code>lagging_minutes</code> minutes.
Cause	Replication of the protection domain snapshot to the remote cluster is too slow.



Impact	Recovery point objective for the protection domain will be affected.
Resolution	Ensure that the number of entities are within the supported limit.
Resolution	Ensure that replication bandwidth to remote cluster is unaffected.

Table 416: Old Entity-Centric Third-Party Backup Snapshots Present [110263] [A110263]

Name	Aged Entity-Centric Third-party Backup Snapshot Check
Description	Check for the old entity-centric third-party backup snapshots.
Alert message	VM <code>vm_name</code> has <code>num_snapshot</code> old third-party backup snapshot(s) and may unnecessarily consume storage space in the cluster.
Cause	Entity-centric third-party backup snapshots are present in the cluster longer than the configured threshold.
Impact	Old snapshots may unnecessarily consume storage space in the cluster.
Resolution	Contact Nutanix support.
Age threshold for third-party backup snapshots.	Warning Threshold: 7 days

Table 417: NearSync Replication of the Protection Domain {protection_domain_name} has not Progressed. [110264] [A110264]

Name	Nearsync Replication Stuck
Description	Nearsync replication of the protection domain has not progressed.
Alert message	NearSync replication of the protection domain <code>protection_domain_name</code> to the remote site <code>remote_name</code> for the snapshot <code>snapshot_id</code> has not made any progress in the last <code>time_string</code>.
Cause	Network outage or saturated network with the remote site.
Cause	Stargate service is not healthy on the source or remote site.
Cause	SSD snapshot reserve on the remote site does not have enough free disk space.
Impact	Protection domain may transition out of NearSync.
Impact	Snapshots of the protection domain will not be available on the remote site.
Resolution	Check the network connection with the remote site.
Resolution	Ensure that the remote site has enough disk space available in the SSD snapshot reserve.

Table 418: Protected VMs not CBR capable [110268] [A110268]

Name	Protected VMs CBR Incapable



Description	Some of the protected VMs are not capable of backup and recovery.
Alert message	Protected VM <code>vm_name</code> is not CBR capable. Reason: <code>cbr_incapable_reason</code>
Cause	Not all protected VMs have a valid configuration for backup and recovery
Impact	RPO will be affected.
Resolution	Review and update the VM configuration. If the configuration of the VM cannot be changed then unprotect the VMs.

Table 419: Cerebro stats collector failed. [111059] [A111059]

Name	Cerebro stats collector
Description	Collect Cerebro stats and publish to CFS.
Alert message	Cerebro stats collector failed.
Cause	Snapshot chain length has exceeded the threshold
Cause	Garbage collection is not able to keep up with the configured RPO
Resolution	Refer to KB 6200.

Table 420: Secure Tunnel To Remote Site Down [130000] [A1090]

Name	Remote Site Tunnel Status
Description	Secure Tunnel To Remote Site Down.
Alert message	Secure tunnel to remote site <code>remote_name</code> is down.
Cause	Incorrect remote site configuration or network connectivity issue.
Impact	Replication to remote site will fail.
Resolution	Check if the IP address specified in the remote site is reachable.

Table 421: Metro Availability Is Disabled [130002] [A1124]

Name	Metro Connectivity
Description	Metro availability is disabled.
Alert message	Metro availability for the protection domain '<code>protection_domain_name</code>' to the remote site '<code>remote_name</code>' is disabled because of '<code>reason</code>'.
Cause	Remote site unreachable.
Impact	Metro availability operation is disabled.
Resolution	Check if cluster service is healthy at remote cluster.

Table 422: Duplicate Remote Cluster ID [130018] [A1038]

Name	Duplicate Remote Cluster ID Check



Description	Duplicate Remote Cluster ID
Alert message	Remote cluster ' remote_name ' is disabled because the name conflicts with remote cluster ' conflicting_remote_name '.
Cause	Two remote sites with different names or different IP addresses have same cluster ID. This can happen in two cases: (a) A remote cluster is added twice under two different names (through different IP addresses) or (b) Two clusters have the same cluster ID.
Impact	Protected data is not replicated to the remote site.
Resolution	In case (a) remove the duplicate remote site. In case (b) verify that the both clusters have the same cluster ID and contact Nutanix support.

Table 423: Entities Restored But Unprotected [130019] [A1134]

Name	Restored Entities Protected
Description	Entities Restored But Unprotected
Alert message	Some entities became unprotected after restoring protection_domain_name . Count: num_entities . Entity names: entity_names . Reason: reason .
Cause	Some other protected entities conflict with the entities being recovered.
Impact	Some entities became unprotected after restoring the snapshot.
Resolution	Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.

Table 424: Entities Skipped During Restore [130020] [A1132]

Name	Entities Restored Check
Description	Entities Skipped During Restore
Alert message	Some entities skipped during restore of protection_domain_name . Count: num_entities . Entity names: entity_names . Reason: reason .
Cause	Existing files conflict with the files to be recovered.
Impact	Not all entities could be restored from the snapshot.
Resolution	Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.

Table 425: Invalid Consistency Group [130022] [A1136]

Name	Consistency Group Configuration
Description	Invalid Consistency Group
Alert message	Failed to create a snapshot of protection domain protection_domain_name as some files in consistency group consistency_group_name overlap other consistency groups.
Cause	VMs with common files are not in the same consistency group.



Impact	Snapshot will fail.
Resolution	Put VMs with common files in the same consistency group.

Table 426: Metadata Volume Snapshot Persistent Failure [130026] [A1087]

Name	Metadata Volume Snapshot Persistent
Description	Metadata Volume Snapshot Persistent Failure
Alert message	The last <code>failure_count</code> metadata volume snapshots have failed. <code>reason</code>
Cause	The EBS snapshot service is unavailable, possibly because AWS could not be reached.
Impact	If the cloud instance fails, backups of the guest VM may be unavailable.
Resolution	Contact Nutanix support.

Table 427: Metadata Volume Snapshot Timeout [130027] [A1088]

Name	Metadata Volume Snapshot Status
Description	Metadata Volume Snapshot Timeout Failure
Alert message	Metadata volume snapshot could not be created in <code>duration_mins</code> minutes.
Cause	Metadata volume snapshot frequency is too high.
Impact	If the cloud instance fails, backups of the guest VM may be unavailable.
Resolution	Contact Nutanix support.

Table 428: Metro Availability Configuration Failed [130028] [A1123]

Name	Metro Availability
Description	Metro availability start failed
Alert message	Failed to <code>operation</code> for the protection domain ' <code>protection_domain_name</code> ' to the remote site ' <code>remote_name</code> '. Reason: ' <code>reason</code> '.
Cause	Various
Impact	Metro availability operation could not be started.
Resolution	Contact Nutanix support if this issue persists.

Table 429: Stale NFS Mount [130029] [A1135]

Name	Metro Vstore Mount Status
Description	Stale NFS mount
Alert message	Unable to process NFS requests for vstore id <code>vstore_id</code> due to <code>reason</code> .
Cause	Stale NFS file handle.



Impact	Unable to process NFS requests.
Resolution	Unmount the datastore from all hosts and remount again.

Table 430: Protected VM Is Not Nutanix Backup And Recovery Compliant [130037] [A1109]

Name	Protected VM CBR Capability
Description	Protected VM is not Nutanix backup and recovery compliant.
Alert message	message
Cause	Protected VM may have been modified to include files that are not on a Nutanix Storage Container.
Impact	Any data associated with the VM may not be backed up or replicated to a remote site.
Resolution	Remove or detach files from the VM that are not on Nutanix Storage Containers and make the VM Nutanix backup and recovery compliant.

Table 431: Protected VM Not Found [130038] [A1010]

Name	Protected VM Not Found
Description	Protected VM not Found
Alert message	Unable to locate VM '<i>vm_name</i>' (ID <i>vm_id</i>)' protected by protection domain '<i>protection_domain_name</i>'.
Cause	The protected VM cannot be found and may have been deleted.
Impact	Any data associated with the VM may not be backed up or replicated to a remote site.
Resolution	Remove the VM from the protection domain.

Table 432: Protected VM Renamed During Clone [130039] [A1118]

Name	Clone VM Renamed Check
Description	Protected VM renamed during clone
Alert message	Protected VM with name '<i>vm_name</i>' in protection domain '<i>protection_domain_name</i>' will be registered as '<i>new_vm_name</i>' during DR clone operation to avoid duplicate VM names.
Cause	A VM with the same name of the VM to be cloned, already exists.
Impact	The cloned VM will be registered under the new name.
Resolution	No action is necessary.



Table 433: Protection Domain Activation [130040] [A1043]

Name	PD Active
Description	Protection Domain Activation
Alert message	Unable to make protection domain ' <i>protection_domain_name</i> ' active on remote site ' <i>remote_name</i> ' due to ' <i>reason</i> '.
Cause	Various
Impact	Protected VMs could not be started during failover to a remote site.
Resolution	Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.

Table 434: Protection Domain Activation or Migration Failure [130041] [A1060]

Name	PD Change Mode Status
Description	Protection Domain Activation or Migration Failure
Alert message	Protection domain <i>protection_domain_name</i> activate/deactivate failed with the error : <i>reason</i>
Cause	Protection domain cannot be activated or migrated.
Cause	Protection domain with same name might be active on remote site.
Cause	Remote sites might not be configured correctly, please check the remote sites on source and target clusters.
Cause	Protection Domain has one or more missing VMs/VGs.
Impact	Protected VMs could not be started during failover to a remote site.
Resolution	Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.

Table 435: Failed To Change State Of One Or More VMs [130042] [A1094]

Name	PD VM Action Status
Description	Failed to change state of one or more VMs.
Alert message	Failed to <i>action</i> one or more VMs in protection domain <i>protection_domain_name</i> .
Cause	The hypervisor may have insufficient resources or the configuration of the VM restored from the snapshot may be invalid.
Impact	If VMs failed to stop, no snapshot can be taken for the protection domain. If VMs failed to start, the VMs will not be powered on for use.
Resolution	Check error reported by the hypervisor management software. If you cannot resolve the error, contact Nutanix support.



Table 436: Registration Of One Or More VMs Failed [130043] [A1093]

Name	PD VM Registration Status
Description	Registration of one or more VMs failed
Alert message	Failed to register one or more VMs of the protection domain <code>protection_domain_name</code>.
Cause	VMs in the snapshot have invalid VM configuration or the local hypervisor reported an error.
Impact	VMs restored from the snapshot will not be unavailable.
Resolution	Check the existence of the VM in the snapshot and check any errors reported in the hypervisor management software. If you cannot resolve the error, contact Nutanix support.

Table 437: Protection Domain Full Replication Performed [130044] [A1108]

Name	PD Full Replication Status
Description	Full Replication was done.
Alert message	The replication for snapshot with id <code>snap_id</code> of protection domain <code>protection_domain_name</code> to remote site <code>remote_name</code> may be a full replication.
Cause	A reference snapshot could not be found.
Impact	The replication may take longer.
Resolution	No action is necessary.

Table 438: Protection Domain Replication Expired [130045] [A1003]

Name	PD Replication Expiry Status
Description	Protection Domain Replication Expired
Alert message	Protection domain <code>protection_domain_name</code> replication to the remote site <code>remote_name</code> has expired before it is started.
Cause	Replication is taking too long to complete before the snapshots expire.
Impact	The snapshot that was scheduled to be replicated could not be replicated. Data that is expected to be protected may not be protected.
Resolution	Review replication schedules taking into account bandwidth and overall load on systems. Confirm retention time on replicated snapshots.

Table 439: Protection Domain Replication Failure [130046] [A1015]

Name	PD Replication Status
Description	Protection Domain Replication Failure



Alert message	Protection domain <code>protection_domain_name</code> replication to remote site <code>remote_name</code> failed. <code>reason</code> .
Cause	Various
Impact	Snapshots of protected VMs were not replicated to the remote site.
Resolution	Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.

Table 440: Skipped Replication Of Snapshot For Protection Domain [130047] [A1113]

Name	PD Replication Skipped Status
Description	Skipped replication of the snapshot.
Alert message	Replication was skipped for protection domain <code>protection_domain_name</code> of snapshot <code>snapshot_name</code> to remote site <code>remote_name</code> . <code>reason</code>
Cause	Snapshot replication was skipped because a newer snapshot is available.
Impact	Snapshot is not present on the remote site.
Resolution	If this issue persists, consider reducing the replication frequency.

Table 441: Failed to Receive Snapshot for the Protection Domain [130048] [A1127]

Name	Protection Domain Receive Snapshot Failure
Description	Protection Domain Receive Snapshot Failure.
Alert message	Failed to receive snapshot for the protection domain ' <code>protection_domain_name</code> ' from the remote site <code>remote_name</code> . <code>reason</code>
Cause	Various
Impact	The snapshot receive operation could not be completed.
Resolution	Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.

Table 442: Protection Domain Snapshot Failure [130049] [A1064]

Name	PD Snapshot Status
Description	Protection Domain Snapshot Failure.
Alert message	Protection domain <code>protection_domain_name</code> snapshot ' <code>snapshot_id</code> ' failed. <code>reason</code> .
Cause	Protection domain cannot be snapshotted.
Cause	Metro protection domain has more entities than supported.
Impact	A requested snapshot of guest VMs and files in the protection domain did not succeed.
Resolution	Make sure all VMs and files are available.



Resolution	In case of Metro protection domain, make sure number of entities are within the supported limit.

Table 443: Remote Site Is Unhealthy [130050] [A1125]

Name	Remote Site Health
Description	Remote site is unhealthy.
Alert message	Remote ' remote_name ' is unhealthy due to reason ' reason '.
Cause	Various
Impact	Replication to remote site can fail.
Resolution	Check if the IP address specified in the remote site is reachable.

Table 444: Snapshot Partially Crash Consistent [130052] [A1110]

Name	Snapshot Crash Consistent
Description	Snapshot partially crash consistent.
Alert message	Failed to create an application-consistent snapshot for one or more VMs in snapshot snapshot_name of protection domain protection_domain_name . A crash-consistent snapshot has been created for these VMs instead.
Cause	VSS or hypervisor error
Impact	Recovery may take longer and involve manual steps.
Resolution	No action is necessary if this issue is intermittent. If this error persists, contact Nutanix support.

Table 445: VM Action Error [130057] [A1033]

Name	VM Action Status
Description	VM Action Error
Alert message	Failed to action VM ' vm_name (ID vm_id)' due to reason
Cause	A VM could not be restored because of a hypervisor error, or could not be deleted because it is still in use.
Impact	The requested VM action (restore or delete) could not be completed.
Resolution	Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.

Table 446: VM Virtual Hardware Version Not Compatible [130058] [A1133]

Name	VM Virtual Hardware Version Compatible
Description	VM Virtual Hardware Version not Compatible



Alert message	Virtual hardware version of the VM <code>vm_name</code> is not supported by any of the nodes on the remote site <code>remote_name</code> . Protection domain: <code>protection_domain_name</code> . Reason: <code>reason</code> .
Cause	Virtual hardware version of the VM is not compatible with the maximum virtual hardware version supported by any of the nodes at the remote site.
Impact	VM may not register properly on the remote site on restore/clone.
Resolution	Upgrade the hypervisor version on the remote site to support virtual hardware version for the VM.

Table 447: Remote Site {remote_name} Network Mapping Missing [130062] [A1156]

Name	Remote Site Network Configuration
Description	No Network Mapping Specified.
Alert message	No network mapping specified for remote site <code>remote_name</code>.
Cause	No network mapping specified when remote site is created.
Impact	When the VM is restored or cloned on the remote site, networking configuration may not be recovered.
Resolution	Specify network mapping for the remote site.

Table 448: Remote Site {remote_name} Network Mapping Invalid [130063] [A1157]

Name	Remote Site Network Mapping Configuration
Description	Invalid Network Mapping Specified.
Alert message	Invalid network mapping specified for remote site <code>remote_name</code>: <code>reason</code>.
Cause	Either the source or destination network configuration is not present in network mapping specified in the remote site, or the networks specified in the network mapping do not exist.
Impact	When the VM is restored or cloned on the remote site, networking configuration may not be recovered.
Resolution	Create network mapping with the networks present on the source and the destination cluster for the remote site.

Table 449: Protection Domain Snapshot Operation Skipped [130066] [A1147]

Name	PD Snapshot Skipped Status
Description	Skipped snapshot operation.
Alert message	Replication was skipped for protection domain <code>protection_domain_name</code> of snapshot <code>snapshot_name</code> to remote site <code>remote_name</code>. <code>reason</code>
Cause	Snapshot operation was skipped.
Impact	Snapshot is not present on both the local and remote sites.



Resolution	Wait for the completion of the upgrade process before initiating snapshot schedules.

Table 450: Protected Volume Group not Found [130070] [A1162]

Name	Protected Volume Group Not Found
Description	Protected Volume Group not Found
Alert message	Unable to locate volume group with name ' <i>vg_name</i> ' and internal ID ' <i>vg_uuid</i> ' protected by protection domain ' <i>protection_domain_name</i> '.
Cause	The protected volume group cannot be found and may have been deleted.
Impact	Any data associated with the volume group may not be backed up or replicated to a remote site.
Resolution	Remove the volume group from the protection domain.

Table 451: Volume Group Action Error [130071] [A1163]

Name	Volume Group Action Status
Description	Volume Group Action Error
Alert message	Failed to <i>action</i> volume group with name ' <i>vg_name</i> ' and internal ID ' <i>vg_uuid</i> ' because <i>reason</i> in protection domain ' <i>protection_domain_name</i> '
Cause	A volume group could not be restored or could not be deleted because it is still in use.
Impact	The requested volume group action (restore or delete) could not be completed.
Resolution	Detach the volume group from VMs and external initiators before recovery.
Resolution	Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.

Table 452: VSS Snapshot Failed [130073] [A130073]

Name	VSS Snapshot Status
Description	VSS snapshot failed.
Alert message	VSS snapshot failed for the VM(s) <i>vm_names</i> protected by the <i>protection_domain_name</i> in the snapshot <i>snapshot_id</i> because <i>reason</i> .
Cause	Guest is not able to quiesce VM due to internal error.
Impact	Crash consistent snapshot is taken instead of application consistent snapshot.
Resolution	Look at the logs in the guest VM. If the VM failed to quiesce, reduce the load on the VM and try again.



Table 453: Nutanix Guest Tools Not Installed [130074] [A130074]

Name	NGT Configuration
Description	Nutanix Guest Tools not installed.
Alert message	Vss is enabled but Nutanix Guest Tools are not installed on the guest VM(s) <code>vm_names</code> protected by <code>protection_domain_name</code>.
Cause	VSS is enabled but Nutanix Guest Tools are not installed on the VM(s).
Impact	Crash consistent snapshot is taken instead of application consistent snapshot.
Resolution	Install Nutanix Guest Tools on the guest VM(s).

Table 454: Protection Domain might have symlinks. [130077] [A130077]

Name	Snapshot Symlink Check
Description	Protection domain snapshot has symlinks.
Alert message	Protection Domain <code>protection_domain_name</code> might have symlinks that have been skipped in the snapshot.
Cause	The protection domain might have symlinks that have been skipped in the snapshot.
Impact	The symlinks in the protection domain have been skipped in the snapshot.
Resolution	Delete the symlinks or protect the entities individually.

Table 455: Entity Restore Aborted [130078] [A130078]

Name	EntityRestoreAbort
Description	Entity restore aborted.
Alert message	Restoring VMs failed for snapshot <code>snapshot_id</code> protected by <code>protection_domain_name</code> because <code>reason</code>.
Cause	Unable to restore VMs due to internal error.
Impact	Snapshot expiry is removed to prevent data loss, and you should remove it manually.
Resolution	Restart the restore process. Snapshot expiry is removed to allow subsequent restore operations. Snapshot will need to be deleted manually to free up the space if required. If this error persists, contact Nutanix support.

Table 456: Skipped Replication Of The Snapshot [130079] [A130079]

Name	Remote Site Snapshot Replication Status
Description	Skipped replication of the snapshot.
Alert message	Replication was skipped for protection domain <code>protection_domain_name</code> to remote site <code>remote_name</code>. <code>reason</code>



Cause	Snapshot replication was skipped because the remote is not fully upgraded.
Impact	Snapshot is not present on the remote site.
Resolution	Wait for the completion of the upgrade process on the remote before initiating snapshot schedules.
Resolution	Manually replicate the snapshot from Prism UI after the remote site upgrade process completes or wait for next scheduled snapshot to replicate data.

Table 457: Nutanix Guest Tools Agent is not reachable on the VM [130081] [A130081]

Name	VSS VM Reachable
Description	VM not reachable.
Alert message	VM(s) <i>vm_names</i> Guest Agent Service is not reachable, protected by <i>protection_domain_name</i>.
Cause	The communication link to the VMs Nutanix Guest Agent service seems to be down
Impact	Crash consistent snapshot is taken instead of application consistent snapshot.
Resolution	Please check and restart Nutanix Guest Agent service inside the VM

Table 458: Nutanix Guest Tools is not supported on remote site [130082] [A130082]

Name	Remote Site NGT Support
Description	Remote site does not support NGT.
Alert message	Remote site <i>remote_site</i> does not support Nutanix Guest Tools. <i>reason</i>.
Cause	Remote cluster is not upgraded
Impact	Nutanix Guest Tools information for the VM will be lost.
Resolution	Upgrade the remote cluster.

Table 459: Protected Vms Not Found [130083] [A130083]

Name	Protected VMs Not Found
Description	Protected Vms Not Found
Alert message	Unable to locate VM(s) <i>vm_names</i> protected by protection domain '<i>protection_domain_name</i>'.
Cause	Protected VM(s) cannot be found and may have been deleted.
Impact	Any data associated with the VMs may not be backed up or replicated to a remote site.
Impact	Next transition into NearSync will be delayed if the protection domain is configured with high frequency schedule.



Resolution	Verify that the VM(s) are registered on the host. If the VM(s) were unregistered or deleted - Remove VM(s) from the protection domain. If the CVM on the host the VM(s) are registered on was shutdown or rebooting, retry the snapshot operation.

Table 460: Protected Volume Groups Not Found [130084] [A130084]

Name	Protected Volume Groups Not Found
Description	Protected Volume Groups Not Found
Alert message	Unable to locate volume group(s) <i>vg_names</i> protected by protection domain '<i>protection_domain_name</i>'.
Cause	Protected volume group(s) cannot be found and may have been deleted.
Impact	Any data associated with the volume group(s) may not be backed up or replicated to a remote site.
Impact	Next transition into NearSync will be delayed if the protection domain is configured with high frequency schedule.
Resolution	Remove volume group(s) from the protection domain.

Table 461: VSS Software or (pre_freeze/post_thaw) Scripts Not Installed [130085] [A130085]

Name	VSS Scripts Not Installed
Description	VSS software or pre_freeze/post_thaw Scripts Not Installed
Alert message	VSS is enabled but VSS software or pre_freeze/post_thaw scripts are not installed on the guest VM(s) <i>vm_names</i> protected by <i>protection_domain_name</i>.
Cause	VSS is enabled but VSS software or pre_freeze/post_thaw scripts are not installed on the VM(s).
Impact	Crash consistent snapshot is taken instead of application consistent snapshot.
Resolution	Install VSS software or (pre_freeze/post_thaw) scripts on the guest VM(s).

Table 462: VStore Snapshot Status [130086] [A130086]

Name	VStore Snapshot Status
Description	VStore Snapshot Status
Alert message	Snapshot status for vstore <i>vstore_name</i>: <i>reason</i>.
Cause	Various
Impact	The requested VStore snapshot action could not be completed.
Resolution	Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.



Table 463: Failed to snapshot entities [130088] [A130088]

Name	Failed To Snapshot Entities
Description	Failed to snapshot entities
Alert message	Failed to snapshot the entities <code>entity_names</code> in the protection domain <code>protection_domain_name</code>, snapshot <code>snapshot_id</code>.
Cause	Conflicting hypervisor tasks might be running on the entities.
Impact	The entities are not captured in the snapshot.
Resolution	Retry the operation after the hypervisor tasks are complete.

Table 464: Related Entity Protection Status [130090] [A130090]

Name	Related Entity Protection Status
Description	Protection status of a related entity.
Alert message	<code>error_message</code>.
Cause	Related entity is not protected in the same protection domain.
Impact	Related VM/Volume Group will not be snapshotted and recovered.
Resolution	Protect the related entity in the same consistency group.

Table 465: Remote Site Operation Mode ReadOnly [130092] [A1189]

Name	Remote Site Operation Mode ReadOnly
Description	Operation Mode of Remote Site changed to kReadOnly.
Alert message	Remote site <code>remote_name</code> moved to read-only state, outgoing replications to this remote will fail.
Cause	SSD failure on remote site in case of single node or two-node cluster.
Cause	Cloud virtual appliance ran out of memory in case of cloud connect.
Cause	Data stored on cloud appliance exceeded 20TB.
Impact	All outgoing replications to remote site will be aborted.
Resolution	Replace the failed SSD on the remote site.
Resolution	Delete some remote snapshots to reduce memory or data usage in case of cloud connect and wait for some time.
Resolution	Increase the cloud CVM size. Refer to the Nutanix documentation for instructions.

Table 466: VM Renamed On Conversion [130094] [A130094]

Name	VM Renamed On Conversion
Description	VM Renamed On Conversion



Alert message	VM ' <i>vm_name</i> ' (ID <i>vm_id</i>)' will be registered as ' <i>new_vm_name</i> ' during VM conversion to avoid duplicate VM names.
Cause	Directory with same name as VM name exists on Storage Container.
Impact	The converted VM will be registered under the new name.
Resolution	No action is necessary.

Table 467: Failed to reconfigure Nutanix Guest Tools during VMs recovery for protection domain [130095] [A130095]

Name	Failed To Recover NGT Information
Description	Failed to enable Nutanix Guest Tools during VMs recovery for protection domain.
Alert message	Failed to reconfigure Nutanix Guest Tools for VMs in protection domain <i>protection_domain_name</i> due to <i>reason</i>.
Cause	Nutanix Guest Tools service might be down.
Impact	VM
Resolution	Enable and mount Nutanix Guest Tools, and restart Nutanix Guest Agent service within the VM, for VMs in protection domain. If you cannot resolve the error, contact Nutanix support.

Table 468: Failed to reconfigure Nutanix Guest Tools for the recovered VM [130096] [A130096]

Name	Failed To Recover NGT Information for VM
Description	Failed to reconfigure Nutanix Guest Tools for a VM in protection domain.
Alert message	Failed to reconfigure Nutanix Guest Tools for VM <i>vm_name</i> in protection domain <i>protection_domain_name</i>.
Cause	Virtual IP address of the cluster might not have been configured.
Cause	Error in generating certificate for VM.
Cause	Guest VM information could not be retrieved.
Impact	VM
Resolution	Configure the virtual IP address of the cluster (if not configured).
Resolution	Enable and Mount Nutanix Guest Tools on the failed VM, and restart the Nutanix Guest Agent service within the VM.
Resolution	Resolve the issue. If you cannot resolve the error, contact Nutanix support.

Table 469: Failed to reconfigure Nutanix Guest Tools for a VM in protection domain [130097] [A130097]

Name	Failed To Mount NGT ISO On Recovery of VM



Description	Failed to mount ISO image as part of Nutanix Guest Tools reconfiguration for a VM in protection domain
Alert message	Failed to reconfigure Nutanix Guest Tools for VM <code>vm_name</code> in protection domain <code>protection_domain_name</code> due to ISO mount failure.
Cause	Virtual IP address of the cluster might not have been configured.
Cause	Controller VM Nutanix Guest Tools version is lower than guest VM Nutanix Guest Tools version.
Cause	Guest VM information could not be retrieved.
Cause	NGT ISO could not be mounted.
Impact	VM
Resolution	Configure the virtual IP address of the cluster (if not configured).
Resolution	Mount Nutanix Guest Tools ISO on the failed VM, and restart Nutanix Guest Agent service within the VM.
Resolution	Make sure there is at least one free CD-ROM slot on the Guest VM. If you cannot resolve the error, contact Nutanix support.

Table 470: External iSCSI Attachments Not Snapshotted [130099] [A130099]

Name	External iSCSI Attachments Not Snapshotted
Description	External iSCSI Attachments Not Snapshotted.
Alert message	<code>error_message</code>. Volume Group external iSCSI attachments cannot be snapshotted for these VMs.
Cause	VMs have duplicate IQNs.
Cause	VMs have configured iSCSI target IP addresses that do not belong to the local Nutanix cluster.
Cause	IQN could not be resolved to VM ID.
Impact	Volume group external iSCSI attachments cannot be snapshotted and restored for these VMs.
Resolution	Make sure all VMs in the protection domain have different IQNs.
Resolution	Remove iSCSI target IP addresses that do not belong to the the local Nutanix cluster from the VMs.
Resolution	Install Nutanix Guest Tools on all the VMs using volume group external iSCSI attachments.

Table 471: Volume Group Attachments Not Restored [130101] [A130101]

Name	Volume Group Attachments Not Restored
Description	Volume Group Attachments Not Restored.
Alert message	<code>error_message</code> for the protection domain '<code>protection_domain_name</code>'.
Cause	Acropolis failed to attach volume groups to the VMs.

Impact	VMs will not be able to access volume group iSCSI disks.
Resolution	Manually attach volume groups to the VMs.

Table 472: Self service restore operation failed [130102] [A130102]

Name	Self service restore operation Failed
Description	Self service restore operation failed
Alert message	Unable to ' <i>operation</i> ' disk from VM ' <i>vm_name</i> ' (ID <i>vm_id</i>)' due to ' <i>reason</i> '.
Cause	Snapshot disk could not be recovered.
Cause	Recovered disk could not be attached to the VM.
Cause	Self-service restore disk could not be detached from VM.
Impact	In the case of attachment failure, the user will not be able to restore files from snapshot.
Impact	In the case of detachment failure, the guest VM will have an extraneous disk attached.
Resolution	Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.

Table 473: Witness is not configured [130114] [A130114]

Name	Witness Not Configured
Description	Witness is not configured
Alert message	Failed to <i>operation</i> for the protection domain ' <i>protection_domain_name</i> ' to the remote site ' <i>remote_name</i> '. Reason: Witness is not configured
Cause	Witness is not configured
Impact	Metro availability operation failed
Resolution	Configure Witness

Table 474: Witness is not reachable [130115] [A130115]

Name	Witness Not Reachable
Description	Witness is not reachable
Alert message	Failed to <i>operation</i> since <i>unreachable_since</i> for the protection domain ' <i>protection_domain_name</i> ' to the remote site ' <i>remote_name</i> '. Reason: Witness is not reachable or the Witness credentials are incorrect
Cause	Witness or the network is down
Cause	Witness credentials are incorrect
Impact	Metro availability failures will not be automatically handled
Resolution	Check if Witness is up on the network



Resolution	Check Witness credentials

Table 475: Metro Availability Is Promoted [130116] [A130116]

Name	Automatic Promote Metro Availability
Description	Metro availability is promoted.
Alert message	Metro availability for the protection domain '<i>protection_domain_name</i>' to the remote site '<i>remote_name</i>' is promoted because of '<i>reason</i>'.
Cause	Remote site unreachable.
Impact	Metro availability is promoted on standby cluster.
Resolution	Check if cluster service is healthy at remote cluster.

Table 476: Failed to update Metro Availability failure handling [130117] [A130117]

Name	Updating Metro Failure Handling Failed
Description	Error in updating failure-handling on Metro Availability protection domain.
Alert message	Error in updating Metro Availability failure handling for the protection domain '<i>protection_domain_name</i>' to the remote site '<i>remote_name</i>'. Reason: '<i>reason</i>'.
Cause	Various
Impact	Metro Availability failure-handling could not be updated.
Resolution	Contact Nutanix support if this issue persists.

Table 477: Failed to update Metro Availability failure handling on the remote site [130118] [A130118]

Name	Updating Metro Failure Handling Remote Failed
Description	Error in updating failure handling on the remote Metro Availability protection domain
Alert message	Error in updating Metro Availability failure handling on the remote site '<i>remote_name</i>' for the protection domain '<i>protection_domain_name</i>'. Reason: '<i>reason</i>'.
Cause	Various
Impact	Failure handling on this Metro Availability protection domain might be affected
Resolution	Perform a local-only update of failure handling on the remote site by using ncli

Table 478: Authentication failed in Witness [130123] [A130123]

Name	Authentication Failed in Witness



Description	Authentication failed in Witness
Alert message	Failed to <i>operation</i> for the protection domain ' <i>protection_domain_name</i> ' to the remote site ' <i>remote_name</i> '. Reason: Authentication failure in Witness
Cause	Witness configured with incorrect credentials
Impact	Metro availability failures will not be automatically handled
Resolution	Configure Witness with correct credentials

Table 479: Metro Availability Is Promoted [130125] [A130125]

Name	Manual Promote Metro Availability
Description	Metro availability is promoted.
Alert message	Metro availability for the protection domain ' <i>protection_domain_name</i> ' to the remote site ' <i>remote_name</i> ' is promoted.
Cause	An administrator promoted metro availability.
Impact	Metro availability is promoted on standby node.
Resolution	Reenable to get the benefit of Metro protection domain.

Table 480: Metro Availability Is Disabled [130126] [A130126]

Name	Manual Break Metro Availability
Description	Metro availability is disabled.
Alert message	Metro availability for the protection domain ' <i>protection_domain_name</i> ' to the remote site ' <i>remote_name</i> ' is disabled.
Cause	An administrator disabled metro availability.
Impact	Metro availability operation is disabled.
Resolution	Reenable to get the benefit of Metro protection domain.

Table 481: VMs in the standby site of a Metro Availability protection domain are running at suboptimal performance. [130129] [A130129]

Name	Metro Protection domain VMs running at Sub-optimal performance
Description	VMs in the standby site of a Metro Availability protection domain are running at suboptimal performance.
Alert message	Warning: <i>num_of_vms</i> VMs are running on the standby site ' <i>remote_name</i> ' of protection domain ' <i>protection_domain_name</i> '. However, all the I/O operations are being forwarded to this site. This can significantly affect the performance of VMs. It is recommended to move VMs to this site or update the DRS rules. The affected VMs list is: <i>vm_names</i>



Cause	VMs are hosted at the standby site of the Metro Availability protection domain.
Impact	Experience of decreased performance
Resolution	Host the VMs at the active site of the Metro Availability protection domain.

Table 482: Protection domain contains more than the specified VMs [130130] [A130130]

Name	Protection Domain VM Count Check
Description	Protection domain contains multiple VMs.
Alert message	It is not recommended to protect more than one VM in a protection domain <code>protection_domain_name</code> if you are replicating to the backup only remote site <code>remote_name</code>.
Cause	Replicating a protection domain with multiple VMs to a backup only remote site.
Impact	Selective retrieval of any VM will take more time as the snapshot will capture all the VMs in the protection domain.
Impact	Replications may be skipped because of a large snapshot size.
Resolution	Protect each VM in a separate protection domain.

Table 483: Snapshot contains entities from the storage container that have deduplication enabled. [130131] [A130131]

Name	Replication Of Deduped Entity
Description	Snapshot contains entities from the storage container that have deduplication enabled.
Alert message	Do not replicate protection domain <code>protection_domain_name</code> comprising entities from the storage container that have deduplication enabled to a <code>remote_site_type</code> remote site <code>remote_name</code>.
Cause	Replicating a snapshot with entities from the storage container that have deduplication enabled to a single node backup or cloud remote site.
Impact	Performance of replications and retrievals will be affected.
Resolution	Either disable dedup on the Storage Containers, or use a multi-node NOS cluster as a replication target.

Table 484: Cloud Appliance Deployment Failed [130133] [A130133]

Name	Cloud Appliance Deployment Failed
Description	Cloud Appliance Deployment Failed
Alert message	Failed to configure cloud appliance for remote site <code>remote_site_name</code>. <code>reason</code>.



Cause	Resource quota for buckets/instances/storage accounts exhausted for cloud account.
Cause	AWS credentials do not have permissions to perform S3/SG/EC2 related operations.
Cause	Cloud CVM deployed was not able to reach the Internet.
Impact	Cloud appliance could not be configured as remote site.
Resolution	Make sure that account has enough quota to create new buckets/instances.
Resolution	Make sure that the subnet in which cloud appliance is being deployed allows Internet connectivity.
Resolution	Make sure that AWS IAM user has permissions to perform S3/SG/EC2 operations.
Resolution	Resolve the issue. If you cannot resolve the error, contact Nutanix support.

Table 485: vStore is being replicated to backup only remote site. [130134] [A130134]

Name	Vstore Replication To Backup Only Remote
Description	vStore is being replicated to backup only remote site.
Alert message	vStore protection domain <code>protection_domain_name</code> is being replicated to backup only remote site <code>remote_name</code>.
Cause	vStore is being replicated to backup only remote site.
Impact	Low RTO can be seen as selective retrieve of VMs is not supported and snapshot of whole protection domain, having all the VMs in it, has to be retrieved.
Resolution	One VM per protection domain is the recommended configuration to achieve higher RTO.

Table 486: Replication of protection domain {protection_domain_name} has not progressed. [130137] [A130137]

Name	Protection Domain Replication Stuck
Description	Replication of protection domain has not progressed.
Alert message	Replication of protection domain <code>protection_domain_name</code> to remote site <code>remote_name</code> for the snapshot <code>snapshot_id</code> has not made any progress in last one hour.
Cause	Network outage with remote.
Cause	Remote cluster does not have enough free disk space available.
Cause	Single-node cluster might be in read-only mode, due to one of the metadata SSDs failure.
Impact	Snapshots of protection domain will not be available on the remote site.
Resolution	Check network connection with remote.
Resolution	Make sure that the remote cluster has enough disk space available.



Resolution	Replace the failed SSD.

Table 487: Snapshot reserve on SSD is full. [130139] [A130139]

Name	LWS Store Full
Description	LWS store has become full on the cluster.
Alert message	The SSD space reserved for snapshots is full.
Cause	Snapshot reserve on SSD is full.
Impact	Creation of the recovery points might be stopped.
Resolution	Reduce the retention time and/or the IO load.

Table 488: Failed to allocate snapshot reserve [130140] [A130140]

Name	LWS store allocation too long
Description	Failed to allocate snapshot reserve.
Alert message	Failed to allocate snapshot reserve in <i>duration_mins</i> minutes.
Cause	Snapshot reservation timed out.
Impact	Protection domain will not be able to meet less than 15 minutes RPO.
Resolution	Contact Nutanix support

Table 489: Failed to allocate snapshot reserve on the remote site [130141] [A130141]

Name	LWS store allocation in remote too long
Description	Failed to allocate snapshot reserve on the remote site.
Alert message	Failed to allocate snapshot reserve on the remote site '<i>remote_name</i>' in <i>duration_mins</i> minutes.
Cause	Snapshot reserve allocation on the remote site timed out.
Impact	Failed to meet less than 15 minute RPO for the remote site.
Resolution	Contact Nutanix support

Table 490: Protection Domain transitioning to lower frequency snapshotting. [130142] [A130142]

Name	Protection Domain transitioning to lower frequency snapshot schedule
Description	Protection Domain transitioning to lower frequency snapshot schedule.
Alert message	Protection Domain '<i>protection_domain_name</i>' has degraded to '<i>rpo_string</i>' RPO. <i>reason</i>.
Cause	High RPO cannot be sustained due to system constraints.



Cause	New entity has been protected in the protection domain which needs to be seeded to the remote.
Impact	All ongoing high frequency replications will be aborted and RPO for data protection will be affected.
Resolution	Protection Domain will automatically attempt to transition to a higher RPO after system constraints are resolved.
Resolution	Protection Domain will automatically attempt to transition to a higher RPO after seeding for newly added entities completes.

Table 491: Protection Domain falling back to lower frequency replications to remote. [130143] [A130143]

Name	Protection Domain fallback to lower frequency replications to remote
Description	Protection Domain fallback to lower frequency replications to remote.
Alert message	Protection Domain '<i>protection_domain_name</i>' has degraded to '<i>rpo_string</i>' RPO to remote '<i>remote_name</i>'. <i>reason</i>.
Cause	High frequency replications to remote site cannot be sustained due to system constraints.
Impact	Recovery Points available on the remote for Data Protection will be affected
Resolution	Protection Domain will automatically attempt to transition to higher frequency replications when resources are available.

Table 492: Protection Domain transitioning to higher frequency snapshotting. [130144] [A130144]

Name	Protection Domain transitioning to higher frequency snapshot schedule
Description	Protection Domain transitioning to higher frequency snapshot schedule.
Alert message	Protection Domain '<i>protection_domain_name</i>' has upgraded to <i>rpo_string</i> RPO.
Cause	Higher RPO can be sustained as the system is stable.
Cause	High frequency snapshot schedule has been enabled.
Impact	RPO for Data Protection will be affected.
Resolution	Reduce snapshot schedule frequency or delete the schedule if high frequency snapshots are not needed.

Table 493: Snapshot queued for replications to remote site [130145] [A130145]

Name	Protection Domain Snapshots Queued for Replication to Remote Site
Description	Protection domain { <i>protection_domain_name</i> } has one or more snapshots queued for replication to remote site { <i>remote_name</i> }



Alert message	Protection domain <code>protection_domain_name</code> has one or more snapshots queued for replication to remote site <code>remote_name</code>
Cause	Various
Impact	One or more snapshots cannot be replicated to the remote site to meet the configured recovery point objective
Resolution	Maintain enough cross-cluster network bandwidth and throttle primary IO

Table 494: The Protection Domain {`protection_domain_name`} will maintain {`rpo_string`} RPO until seeding process completes to the remote_site {`remote_name`} [130146] [A130146]

Name	Protection Domain Seeding Process Pending to Remote Site
Description	Protection Domain { <code>protection_domain_name</code> } will maintain { <code>rpo_string</code> } RPO until seeding process completes to the remote_site { <code>remote_name</code> }
Alert message	The Protection Domain <code>protection_domain_name</code> will maintain <code>rpo_string</code> RPO until seeding process completes to the remote_site <code>remote_name</code> . The seeding process will be completed when snapshot reserve has been allocated on the local and remote site
Cause	Various
Impact	The Protection Domain { <code>protection_domain_name</code> } will maintain { <code>rpo_string</code> } RPO until seeding process completes to the remote_site { <code>remote_name</code> }
Resolution	Maintain enough network bandwidth and throttle primary IO to help the seeding transfer to succeed. Protection domain will wait for the seeding transfer and automatically enable configured RPO afterwards

Table 495: VM registration failed [130147] [A130147]

Name	VM Registration Failure
Description	VM Registration Failure
Alert message	Registration of virtual machine, <code>vm_name</code> , failed because <code>reason</code>
Cause	This might happen for several reasons. Check the alert message for more details.
Impact	VMs cannot be registered on the cluster.
Resolution	Verify the existence of the VM in the snapshot and check whether any errors are reported in the hypervisor management software. Check the VM version and the maximum VM version supported by the host. If you cannot resolve the error, contact Nutanix support

Table 496: Snapshot creation delayed [130148] [A130148]

Name	Protection Domain Snapshots Delayed
Description	Protection domain { <code>protection_domain_name</code> } has creation of one or more snapshots delayed



Alert message	Protection domain <code>protection_domain_name</code> has creation of one or more snapshots delayed
Cause	Various
Impact	One or more snapshots cannot be created and hence replicated to remote sites to meet the configured recovery point objective
Resolution	Throttle other administrative operations in the cluster

Table 497: VM Recovery Point Creation Failed. [130157] [A130157]

Name	VM Recovery Point Failed
Description	VM Recovery Point creation failed.
Alert message	Failed to capture the Recovery Point for VM '<code>vm_name</code>'.
Cause	Nutanix Data Protection service error.
Impact	VM Recovery Point will not be available to recover the VM.
Resolution	Retry the VM Recovery Point operation.

Table 498: VM Protection Failed. [130158] [A130158]

Name	VM Protection Failed
Description	VM protection failed.
Alert message	Failed to protect VM: '<code>vm_name</code>', because of '<code>reason</code>'.
Cause	Nutanix Data Protection service error.
Impact	VM cannot be recovered in case of disaster.
Resolution	VM protection will get retrIGGERED. In case if this situation persists, please contact nutanix support.

Table 499: VM Recovery Point Replication Failed. [130159] [A130159]

Name	VM Recovery Point Replication Failed
Description	VM recovery point replication failed.
Alert message	Failed to replicate the VM recovery point '<code>recovery_point_uuid</code>' of the VM with ID '<code>vm_id</code>', due to '<code>reason</code>'.
Cause	Nutanix Data Protection service error.
Impact	VM recovery point is not available to recover the VM on the remote cluster.
Resolution	Retry the VM recovery point replication operation.



Table 500: Protection Rule Conflict Occurred. [130161] [A130161]

Name	Protection Rule Conflicts
Description	Unable to protect VM due to conflicting protection rules.
Alert message	Failed to apply the protection rule on the VM '<i>vm_name</i>', because of '<i>reason</i>'.
Cause	One or more protection rules protecting the same entity.
Impact	VM cannot be recovered if disaster occurs.
Resolution	Resolve all the conflicts between the protection rules or modify the VM to ensure that it is protected under one protection rule.

Table 501: Recovery Point Replication Skipped. [130162] [A130162]

Name	Recovery Point Replication Skipped
Description	Recovery Point Replication Skipped.
Alert message	Skipped replication of Recovery point created at '<i>recovery_point_create_time</i>' UTC of the VM: '<i>vm_name</i>' because of '<i>reason</i>'.
Cause	Latest recovery point is in the queue for the replication.
Impact	One of the recovery point replication will be skipped.
Resolution	Trigger the replication of the snapshot manually.

Table 502: Failed to find the Target cluster for Replication [130163] [A130163]

Name	Fetch Replication Target Failure
Description	Cannot find a target cluster for replication on the Remote Availability Zone.
Alert message	Cannot find the replication target on the Availability Zone '<i>availability_zone_physical_name</i>'.
Cause	There might be multiple reasons for this issue.
Impact	Replication of recovery points to this remote Availability Zone will be affected.
Resolution	Retry the process. If this error persists, contact Nutanix support.

Table 503: VM Replication Failure [130164] [A130164]

Name	VM Replication Failure
Description	VM Replication Failure
Alert message	Replication of VM Recovery Point created at '<i>recovery_point_create_time</i>' UTC for VM '<i>vm_name</i>' to the Recovery Location '<i>availability_zone_physical_name</i>' failed.
Cause	Various
Impact	VM Recovery Point was not replicated to the Recovery Location.



Resolution	Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.

Table 504: Application-Consistent Recovery Point Failed. [130165] [A130165]

Name	VSS Recovery Point Status
Description	Application-consistent Recovery Point failed.
Alert message	Application-Consistent Recovery Point created at <code>recovery_point_create_time</code> UTC for the VM <code>vm_name</code> failed because <code>reason</code>.
Cause	Failed to quiesce the VM due to some internal error.
Impact	Crash-consistent Recovery Point is taken instead of application-consistent Recovery Point.
Resolution	Look at the logs in the guest VM. If the VM failed to quiesce, reduce the load on the VM and try again.

Table 505: VM Replication Expired. [130166] [A130166]

Name	VM Replication Expiry Status
Description	VM Replication Expired
Alert message	Replication of Recovery Point created at <code>recovery_point_create_time</code> UTC for VM <code>vm_name</code> to the Recovery Location <code>az_name</code> has expired before it could be started.
Cause	Replication is taking too long to complete before the Recovery Points expire.
Impact	The Recovery Point that was scheduled to be replicated could not be replicated. Data that is expected to be protected may not be protected.
Resolution	Review replication schedules taking into account bandwidth and overall load on systems. Confirm retention time on replicated Recovery Points.

Table 506: Replication Time Exceeded the RPO [130167] [A130167]

Name	Replication Time Exceeds Rpo
Description	Replication Time Exceeded the RPO Limit.
Alert message	Replication time of the snapshot created at '<code>recovery_point_create_time</code>' UTC for VM <code>vm_name</code> exceeded the RPO limits.
Cause	Latest replication took more time than the RPO limit.
Impact	Replications for the next snapshots will be delayed.
Resolution	Please check your network.

Table 507: NGT on VM is not Reachable [130168] [A130168]

Name	NGT is unreachable
Description	NGT on the VM is not reachable.
Alert message	Communication link to Nutanix Guest Tools on the VM <code>vm_name</code> is not working.
Cause	Communication link to Nutanix Guest Tools on the VM is not working.
Impact	Crash consistent snapshot will be taken instead of application consistent snapshot and static address IP preservation will fail.
Resolution	Restart the Nutanix Guest Tools service in the VM.

Table 508: VM Replication has not progressed. [130169] [A130169]

Name	VM Recovery Point Replication Stuck
Description	Replication of the VM has not progressed.
Alert message	Replication of the Recovery Point created at <code>recovery_point_create_time</code> UTC for VM <code>vm_name</code> to Recovery Location <code>az_name</code> has not made any progress in last one hour.
Cause	Various.
Impact	VM Recovery Point will not be available at the Recovery Location.
Resolution	Various.

Table 509: Nutanix Guest Tools Not Installed on the VM [130170] [A130170]

Name	VM NGT Configuration
Description	Nutanix Guest Tools not installed.
Alert message	Nutanix VSS is enabled but Nutanix Guest Tools is not installed on the VM <code>vm_name</code>.
Cause	Nutanix VSS is enabled but Nutanix Guest Tools is not installed on the VM.
Impact	Application Consistent Recovery Points will fail. Instead Crash consistent recovery point will be taken.
Resolution	Install Nutanix Guest Tools on the guest VM.

Table 510: Nutanix VSS provider or (pre_freeze/post_thaw) Scripts Not Installed. [130171] [A130171]

Name	VSS Scripts Not Installed On VM
Description	Nutanix VSS provider or pre_freeze/post_thaw Scripts Not Installed.
Alert message	VSS is enabled, but Nutanix VSS provider or pre_freeze/post_thaw scripts are not installed on the VM <code>vm_name</code>.



Cause	VSS is enabled, but Nutanix VSS provider or pre_freeze/post_thaw scripts are not installed on the VM.
Impact	Crash-consistent Recovery Point will be taken instead of application-consistent Recovery Point.
Resolution	Install Nutanix VSS provider or (pre_freeze/post_thaw) scripts on the guest VM.

Table 511: Partial Recovery Point. [130172] [A130172]

Name	Partial Recovery Point
Description	Partial Recovery Point.
Alert message	Recovery Point for VM <code>vm_name</code> failed to capture associated policies and categories because <code>reason</code>.
Cause	Recovery Point for VM failed to capture associated policies and categories.
Impact	VM will not be restored with associated categories and policies
Resolution	Check the connection between PC and the PE, then retry the Recovery Point operation.

Table 512: Recovery Point Replication Failed [130174] [A130174]

Name	Recovery Point Replication Failed
Description	Recovery Point replication failed.
Alert message	Failed to replicate recovery point created at: '<code>recovery_point_create_time</code>' UTC of the VM: '<code>vm_name</code>' to the recovery location: <code>availability_zone_physical_name</code>.
Cause	Various.
Impact	Recovery Point will be unavailable for recovery at the recovery location.
Resolution	Retry the Recovery Point replication operation.

Table 513: VM registered without any Network [130176] [A130176]

Name	VM Registered Without Network
Description	Failed to attach vNIC for the recovered VM.
Alert message	Failed to attach the vNIC for recovered VM: <code>vm_name</code>.
Cause	Network Mapping not found or there might be some internal error.
Impact	VM will come up without any network.
Resolution	Specify network mapping for the VM.



Table 514: Availability Zone Validation Failed [130177] [A130177]

Name	Availability Zone Validation Failed
Description	Availability Zone referenced in the Protection Policy is not Valid.
Alert message	Protection Policy: '<i>protection_rule_name</i>' is processed but following issue is observed - '<i>reason</i>'.
Cause	Invalid Availability Zone configured in the Protection Policy.
Impact	Protection Policy will not be processed.
Resolution	Update the Availability Zone in Protection Policy.

Table 515: Entity Unprotection Failed [130178] [A130178]

Name	Entity Unprotection Failed
Description	Unable to unprotect the VM because of some internal error.
Alert message	Failed to unprotect the VM: '<i>vm_name</i>', because of '<i>reason</i>'.
Cause	Data Protection service may be down.
Impact	VM will remain protected.
Resolution	If this issue persists for a long time, contact Nutanix support for further assistance.

Table 516: Data Protection tasks are not progressing. [130179] [A130179]

Name	Cerebro tasks Stuck
Description	Cerebro Tasks got stuck.
Alert message	Multiple tasks are not making progress on PE:<i>cluster_name</i> because of long running replication tasks.
Cause	Replication Tasks are running for long.
Impact	Tasks can stuck forever.
Resolution	Abort the long running tasks.

Table 517: Recovery Location is in Read Only mode [130180] [A130180]

Name	Entity Centric Remote Site Operation Mode ReadOnly
Description	Operation Mode of the Recovery Location changed to Read Only mode.
Alert message	Availability zone: <i>availability_zone_physical_name</i> is in read-only state. Replications to this site will fail.
Cause	Various
Impact	Replications to this Recovery Location will be aborted.
Resolution	Various



Table 518: Snapshot has Entities from the Storage Container that have Deduplication Enabled [130181] [A130181]

Name	Replication Of Deduped Entity
Description	Snapshot contains entities from the storage container that have deduplication enabled.
Alert message	Do not replicate the recovery point comprising entities from the storage container that have deduplication enabled to the availability zone: <code>availability_zone_physical_name</code>.
Cause	A snapshot with entities from the storage container that have deduplication enabled was replicated to a single node backup or cloud recovery location.
Impact	Performance of replications and retrievals will be affected.
Resolution	Either disable the deduplication on the storage containers or use a multi-node AOS cluster as a replication target.

Table 519: Nutanix Guest Tools not upgraded [130182] [A130182]

Name	NGT needs upgrade
Description	Nutanix Guest Tools not upgraded.
Alert message	Nutanix Guest Tools are not upgraded to version <code>version</code> or higher on the guest VM(s) <code>vm_names</code> protected by <code>protection_domain_name</code>.
Cause	VSS is enabled but Nutanix Guest Tools are not upgraded on the VM(s).
Impact	Crash-consistent snapshot will be taken instead of application consistent snapshot.
Resolution	Install newer version of Nutanix Guest Tools on the guest VM(s).

Table 520: Protected VM(s) May Not Be Recoverable [130183] [A130183]

Name	Protected VMs May Not Be Recoverable
Description	Protected VM(s) May Not Be Recoverable
Alert message	VM(s) <code>vm_names</code> protected by protection domain '<code>protection_domain_name</code>' are running from VMWare snapshots. Recovery of the VMs from the PD snapshot '<code>snapshot_id</code>' may not succeed. Manual intervention is required to bootup the VMs.
Cause	Protected VM(s) are running from VMWare snapshots and may not be recoverable.
Impact	Recovery of the VMs from the PD snapshot may not succeed
Resolution	Ensure that the VMs are not running from VMWare snapshot at the time of scheduled Protection Domain snapshot. Refer Nutanix KB: 7023 for further details.



Table 521: VM Protection Failed [130186] [A130186]

Name	VM Protection Failed
Description	Failed to protect the VM by using the defined Protection Policy.
Alert message	Failed to apply protection policy '<i>protection_rule_name</i>' on the VM '<i>vm_name</i>' because '<i>reason</i>'
Cause	Various
Impact	VM will not be protected.
Resolution	Fix the issue as displayed in the alert message.

Table 522: VM Recovery May Fail [130188] [A130188]

Name	VM Recovery May Fail
Description	VM may not start on the recovery location.
Alert message	VM <i>vm_name</i> may not start on the recovery location '<i>availability_zone_name</i>'
Cause	There might be various reasons for this issue.
Impact	VM will not start on the recovery location.
Resolution	Resolve the stated reason for the issue.

Table 523: Virtual IP Address not Configured on the Cluster. [130191] [A130191]

Name	Virtual IP not set on Cluster
Description	Virtual IP address is not configured on the cluster.
Alert message	Virtual IP address is not configured on the cluster <i>cluster_name</i>. Protection and replication of VMs will be affected.
Cause	Virtual IP address has not been configured on the cluster.
Impact	Protection and replication of VMs will be affected.
Resolution	Configure the virtual IP address on the cluster.

Table 524: Full Replication Started for the VM [130193] [A130193]

Name	VM Recovery Point Full Replication
Description	A full replication was triggered for the VM.
Alert message	Replication for the VM <i>vm_name</i> recovery point with the create time <i>recovery_point_create_time</i> to recovery AZ <i>az_name</i> will be a full replication.
Cause	A reference recovery point could not be found.
Impact	Replication will take longer to finish.



Resolution	No action is necessary.

Table 525: Protected VM Not Recoverable [130195] [A130195]

Name	VM Not Recoverable
Description	Protected VM Not Recoverable
Alert message	VM <code>vm_name</code> is running from the VMware snapshots. Recovery of the VM from the Recovery Point created at <code>recovery_point_create_time</code> UTC may not succeed. Manual intervention is required to bootup the VM.
Cause	Protected VM is running from the VMware snapshot and may not be recoverable.
Impact	Recovery of the VM from the Recovery Point may not succeed.
Resolution	Ensure that the VM is not running from the VMware snapshot at the time of scheduled Recovery Point.

Table 526: VM Virtual Hardware Version Incompatible [130196] [A130196]

Name	VM Virtual Hardware Version Incompatible
Description	VM Virtual Hardware Version Incompatible
Alert message	Virtual hardware version of the VM <code>vm_name</code> is not supported on the availability zone <code>availability_zone_name</code>.
Cause	Virtual hardware version of the VM is not compatible with the maximum virtual hardware version supported by any of the nodes at the remote Availability Zone.
Impact	VM may not register properly on the remote Availability Zone on restore/clone.
Resolution	Upgrade the hypervisor version on the remote Availability Zone to support the virtual hardware version of the VM.

Table 527: VM Restoration Failed. [130197] [A130197]

Name	VM Restoration Failed
Description	VM Restoration Failed.
Alert message	Failed to restore VM: '<code>vm_name</code>' due to <code>reason</code>.
Cause	Various.
Impact	VM will not be restored.
Resolution	Retry the VM Restoration operation.



Table 528: VM Migration Failed [130199] [A130199]

Name	VM Migration Failed
Description	VM Migration Failed.
Alert message	Failed to migrate VM: '<i>vm_name</i>' due to internal error.
Cause	Internal Error.
Impact	VM will not be migrated.
Resolution	Retry the VM Migration operation. If this situation persists, contact Nutanix support.

Table 529: Recovery Location is not in Good State [130203] [A130203]

Name	Recovery Location Not In Good State
Description	Recovery location is not in good state.
Alert message	Availability zone <i>availability_zone_physical_name</i> is not in good state. Replications to this site will fail.
Cause	There might be various reasons for this issue.
Impact	Replications to this recovery location will fail.
Resolution	Resolve the stated reason for the failure.

Table 530: Entity is being Transitioned to a Higher Frequency Snapshot Schedule. [130204] [A130204]

Name	Entity transitioning to higher frequency snapshot schedule
Description	Entity is being transitioned to a higher frequency snapshot schedule.
Alert message	Entity '<i>vm_name</i>' has been upgraded to <i>rpo_string</i> RPO.
Cause	High frequency snapshot schedule has been enabled and Higher RPO can be sustained as the system is stable.
Impact	RPO for Data Protection will be affected.
Resolution	If high-frequency snapshots are not required, reduce the snapshot schedule frequency in the Protection Rule or remove the VM from the Protection Rule with high frequency schedule.

Table 531: Entity is being Transitioned to a Lower Frequency Snapshot Schedule. [130205] [A130205]

Name	Entity transitioning to lower frequency snapshot schedule
Description	Entity is being transitioned to a lower frequency snapshot schedule.
Alert message	Entity '<i>vm_name</i>' has degraded to '<i>rpo_string</i>' RPO. <i>reason</i>.
Cause	High RPO cannot be sustained due to system constraints.



Impact	All ongoing high frequency replications will be aborted and RPO for data protection will be affected.
Resolution	Entity will automatically attempt to transition to a higher RPO after system constraints are resolved.

Table 532: Remote Subnet is not Configured Appropriately [130206] [A130206]

Name	Remote Subnet Mismatch With Provided Value
Description	Remote subnet is not configured appropriately.
Alert message	Remote <code>remote_name</code> is configured with the subnet <code>user_specified_subnet</code>, which is not matching with the subnet <code>remote_returned_subnet</code> of the remote cluster.
Cause	Subnet specified in the remote site is not matching with the subnet of the remote cluster.
Impact	Improper firewall and routing settings.
Resolution	Reconfigure the remote subnet parameter.

Table 533: Hosting of Virtual IP of the Network Segmentation DR Service Failed. [130207] [A130207]

Name	Failed Hosting Network Segmented VIP
Description	Hosting of virtual IP of the DR service network segmentation failed.
Alert message	Virtual IP <code>virtual_ip</code> of the DR service network segmentation could not be hosted on the network interface <code>network_interface</code>.
Cause	Network configuration might have been changed or has not been configured properly.
Impact	Replication of the protected entities will be affected.
Resolution	Reconfigure the DR service network configuration.

Table 534: Snapshot Replication to Remote Site is Lagging. [130210] [A130210]

Name	Recovery Point Lagging Remote Replication
Description	Replication of the snapshot to the remote cluster is lagging.
Alert message	Replication of the recovery point for the VM '<code>vm_name</code>' to the remote site '<code>remote_az</code>' is lagging by '<code>Lagging_minutes</code>' minutes.
Cause	There may be multiple reasons for this issue.
Impact	Recovery point objective for the protection domain will be affected.
Resolution	Retry the process. If this error persists, contact Nutanix support.



Table 535: Recovery Plan Execution Failure [300401] [A300401]

Name	Recovery Plan Execution Failure
Description	Execution of the Recovery Plan failed.
Alert message	<code>operation_type</code> failed for the Recovery Plan <code>recovery_plan_name</code>. One or more VMs in this Recovery Plan might have failed to start.
Cause	One or more VMs were not recovered as part of the Recovery Plan execution.
Impact	One or more VMs or applications in this Recovery Plan might not have started completely.
Resolution	Review the Recovery Plan execution report for detailed step by step information on the recovery plan execution and failures.

Table 536: Validation Warnings found during Recovery Plan Execution [300402] [A300402]

Name	Recovery Plan Execution Failure due to Validation Warnings
Description	Execution of the Recovery Plan failed.
Alert message	<code>operation_type</code> for Recovery plan <code>recovery_plan_name</code> resulted in validation warnings. You can choose to execute the Recovery Plan anyway or you can fix the issues.
Cause	The Recovery Plan validation resulted in warnings.
Impact	The Recovery Plan execution could not be started due to warnings in the validation report.
Resolution	Review the Recovery Plan validation report for detailed information.

Table 537: Recovery Plan Validation Failed With Errors [300408] [A300408]

Name	Validation Failed For Recovery Plan with Errors
Description	Recovery Plan validation failed with errors.
Alert message	Validation failed with errors for the Recovery Plan <code>recovery_plan_name</code>.
Cause	There might be various reasons for this issue.
Impact	VM recovery might fail.
Resolution	Review the Recovery Plan execution report for detailed step by step information on the recovery plan execution and failures.

Table 538: Recovery Plan Validation Failed with Warnings [300409] [A300409]

Name	Validation Failed For Recovery Plan with Warnings
Description	Recovery Plan validation failed with warnings.
Alert message	Validation failed with warnings for the Recovery Plan <code>recovery_plan_name</code>.



Cause	There might be various reasons for this issue.
Impact	VM recovery might fail.
Resolution	Review the Recovery Plan execution report for detailed step by step information on the recovery plan execution and failures.

Table 539: Failover or Failback operations are not possible [300411] [A300411]

Name	Delete the failed over VMs from either the Primary or the Recovery Location
Description	Delete the failed over VMs on either the primary or the recovery location.
Alert message	Unplanned Failover for the Recovery Plan <code>recovery_plan_name</code> has completed successfully. Failover or Failback operations won't be possible because the VMs might be present on both the primary and target locations.
Cause	Unplanned failover for the Recovery Plan completed successfully.
Impact	Failover and failback operations are not possible without cleanup.
Resolution	Validate the VMs that are part of this Recovery Plan and delete them from the site where they are no longer required.

Table 540: Floating IPs Allocation Failed [300420] [A300420]

Name	Floating IPs Allocation Failure.
Description	Floating IPs allocation failed.
Alert message	Failed to allocate <code>num_failed_floating_ips</code> floating IPs, required for the dynamic allocation, for <code>operation_type</code> on the Recovery Plan <code>recovery_plan_name</code>.
Cause	The system may not have enough floating IP addresses.
Impact	VMs will come up without floating IPs.
Resolution	Try to manually assign pre-allocated floating IPs for the vNICs.

Table 541: Floating IPs Deallocation Failed [300421] [A300421]

Name	Floating IPs Deallocation Failure.
Description	Floating IPs deallocation failed.
Alert message	Failed to deallocate <code>num_failed_floating_ips</code> floating IPs, which were allocated dynamically but are not associated with any vNIC, for <code>operation_type</code> on the Recovery Plan <code>recovery_plan_name</code>.
Cause	The floating IPs might be in use.
Impact	Floating IPs will remain reserved. Reserved floating IPs are billable.
Resolution	If you are not using the floating IPs, deallocate them manually.



Table 542: Recovery Plan Execution Failure due to Validation Errors [300422] [A300422]

Name	Recovery Plan Execution Failure due to Validation Errors
Description	Execution of the Recovery Plan failed.
Alert message	The operation_type for the recovery plan recovery_plan_name could not be started due to errors in validation.
Cause	The Recovery Plan validation failed.
Impact	The Recovery Plan execution could not be started.
Resolution	Review the Recovery Plan validation report for detailed information.

Table 543: Floating IPs Deallocation Failed after Failback from Xi. [300427] [A300427]

Name	Unassociated Floating IPs Deallocation Failure after Failback from Xi.
Description	Failed to deallocate unassociated floating IPs after failback from Xi that were earlier dynamically created.
Alert message	Failed to deallocate unassociated floating IPs failed_floating_ips after failback from Xi that were earlier dynamically created, for operation_type on the Recovery Plan recovery_plan_name .
Cause	The floating IPs might be in use.
Impact	Floating IPs will continue to remain reserved. Reserved floating IPs are billable.
Resolution	If you are not using the floating IPs, deallocate them manually.

Entity Sync

Table 544: Entity Sync Failure [500101] [A500101]

Name	Entity Sync failed
Description	Entity Sync failed
Alert message	message
Cause	Various
Impact	Data Protection and Recovery might be impacted.
Resolution	Perform Forced Entity Sync after ensuring that all the remote entities can be overwritten. If the issue still persists, contact Nutanix support.

Guest VM

Table 545: VM IO Latency [2001]

Name	VM IO Latency
Description	Checks that the Average I/O Latency is not high.



Cause	Cluster capacity may not be enough causing the average I/O latency to be high.
Impact	Application performance on the VM may be degraded.
Resolution	Add capacity to cluster.

Table 546: VM Swap Rate [3020]

Name	VM Swap Rate
Description	Checks whether swap rate for VM is not high.
Cause	Swap rate is high. VMs may not have enough memory allocated.
Impact	VM performance may be degraded.
Resolution	Increase memory, reduce memory intensive VMs, or reduce the total number of VMs.

Table 547: VM Memory Usage High [3021]

Name	VM Memory Usage High
Description	Checks whether memory usage is high on VMs.
Cause	Memory usage is high.
Impact	Application performance on the VM may be degraded.
Resolution	Configure additional memory for the VM. Tune the applications to reduce memory demand.

Table 548: Datastore VM Count High [3030] [A1170]

Name	Datastore VM Count Check
Description	Checks for high VM count on datastores
Alert message	The number of VMs on Datastore: <code>datastore</code> is <code>vm_count</code>
Cause	High number of VMs on at least one datastore
Impact	HA service might be affected as VMware HA protects only 2048 VMs per datastore.
Resolution	Move some VMs to another datastore
VMs per Datastore Threshold	Warning Threshold: 1945 VMs per datastore Critical Threshold: 2048 VMs per datastore

Table 549: VM CPU Utilization [3040]

Name	VM CPU Utilization



Description	Checks whether CPU usage is high on VMs.
Cause	CPU usage is high.
Impact	Application performance on the VM may be degraded.
Resolution	Reduce CPU-intensive processes.

Table 550: VM Transmit Packet Drop Check [3060]

Name	VM Transmit Packet Drop Check
Description	Checks whether transmitted packet drop rate is high.
Cause	Transmitted packet drop rate is high.
Impact	Application performance on the VM may be degraded.
Resolution	Check network configuration/connections.

Table 551: VM Receive Packet Drop Check [3061]

Name	VM Receive Packet Drop Check
Description	Checks whether received packet drop rate is high.
Cause	Received packet drop rate is high.
Impact	Application performance on the VM may be degraded.
Resolution	Check network configuration/connections.

Table 552: VM has non-ASCII name. [106472] [A106472]

Name	Non-ASCII VM Names
Description	Check if any VMs have non-ASCII names.
Alert message	VM <i>vm_name</i> has non-ASCII name, which may cause snapshotting problems.
Cause	Some VMs have non-ASCII names.
Impact	Snapshot operations may fail.
Resolution	Update the VM names to consist only of ASCII characters.

Table 553: Node Failure [130030] [A1137]

Name	Node Status
Description	Node Failure
Alert message	Host <i>hypervisor_address</i> appears to have failed. High Availability is restarting VMs on <i>failover_host_info</i>.
Cause	Host is not accessible.



Impact	VMs will restart on another host.
Resolution	Check for network connectivity or a hardware failure.

Table 554: Protected VM Not Found [130038] [A1010]

Name	Protected VM Not Found
Description	Protected VM not Found
Alert message	Unable to locate VM '<i>vm_name</i>' (ID <i>vm_id</i>)' protected by protection domain '<i>protection_domain_name</i>'.
Cause	The protected VM cannot be found and may have been deleted.
Impact	Any data associated with the VM may not be backed up or replicated to a remote site.
Resolution	Remove the VM from the protection domain.

Table 555: VM Action Error [130057] [A1033]

Name	VM Action Status
Description	VM Action Error
Alert message	Failed to <i>action</i> VM '<i>vm_name</i>' (ID <i>vm_id</i>)' due to <i>reason</i>
Cause	A VM could not be restored because of a hypervisor error, or could not be deleted because it is still in use.
Impact	The requested VM action (restore or delete) could not be completed.
Resolution	Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.

Table 556: Failure To Restart VMs For HA Event [130064] [A1145]

Name	VMs Restart Status
Description	Failure to restart VMs for HA event
Alert message	Failed to restart one or more VMs that were running on failed host <i>hypervisor_address</i>.
Cause	Not enough memory or CPU resources within the cluster.
Impact	Some VMs will remain powered off.
Resolution	Shut down unneeded VMs to free cluster resources; expand the cluster to add resources; or wait until the failed host recovers. Once resources are available, power on the VMs.



Table 557: HA Host Evacuation Failure [130065] [A1146]

Name	Host Evacuation Status
Description	Failure to evacuate host while entering maintenance mode or reserving host for HA.
Alert message	Failed to evacuate some VMs from host <code>hypervisor_address</code> while <code>reason</code>.
Cause	Not enough memory/CPU resources within the cluster.
Impact	The host will not enter requested state.
Resolution	Shut down unneeded VMs to free cluster resources or expand the cluster to add resources.

Table 558: HA Healing Failure [130067] [A1155]

Name	VM HA Healing Status
Description	HA healing failure
Alert message	Could not restore VM High Availability.
Cause	Not enough memory/CPU resources within the cluster.
Impact	VMs are no longer protected against host failure.
Resolution	Shut down unneeded VMs to free cluster resources, expand the cluster to add resources, or disable protection against host failures. VMs cannot be powered on or migrated until the system has sufficient resources to protect against the configured number of host failures.

Table 559: Vm High Availability Failure [130068] [A130068]

Name	VM HA Status
Description	High Availability Failure
Alert message	VMs are no longer protected against host failure. Reason: <code>reason</code>
Cause	Not enough memory/CPU resources within the cluster.
Impact	VMs are no longer protected against host failure.
Resolution	Shut down unneeded VMs to free cluster resources, expand the cluster to add resources, or disable protection against host failures. VMs cannot be powered on or migrated until the system has sufficient resources to protect against the configured number of host failures.

Table 560: iSCSI Configuration Failed [130100] [A130100]

Name	iSCSI Configuration Failed
Description	iSCSI Configuration Failed
Alert message	Failed to re-configure iSCSI settings on the recovered VM '<code>vm_name</code>'. <code>reason</code>.



Cause	Nutanix Guest Tools failed to execute some iSCSI commands on the guest VM.
Impact	iSCSI disks may become unavailable on the guest VM.
Resolution	If IQN and iSCSI target IP addresses of the VM have been updated by Nutanix Guest Agent, Discover and connect to new targets after rebooting the VM.
Resolution	Manually configure iSCSI settings on the guest VM.
Resolution	Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.

Table 561: Nutanix Guest Tools mount failed. [130103] [A130103]

Name	NGT Mount Failure
Description	Nutanix Guest Tools mount failed.
Alert message	Nutanix Guest Tools mount failed. <i>reason.</i>
Cause	Failed to attach ISO image to the VM, please check the logs for exact cause of failure.
Impact	Nutanix Guest Tools is not mounted on VM.
Resolution	Check the logs for error information or contact Nutanix support.

Table 562: NGT version of the VM is incompatible with the NGT version of the cluster. [130104] [A130104]

Name	NGT Version Incompatible
Description	NGT version of the VM is incompatible with the NGT version of the cluster.
Alert message	<i>message.</i>
Cause	VM is restored from the snapshot.
Impact	Nutanix Guest Tools features will not work.
Resolution	Uninstall NGT, Mount the NGT ISO, and then re-install NGT on the VM.

Table 563: VSS Snapshot is not supported for the VM. [130105] [A130105]

Name	VSS Snapshot Not Supported
Description	VSS snapshot is not supported for the VM.
Alert message	VSS snapshot is not supported for the VM 'vm_name', because <i>reason.</i>
Cause	VM has unsupported configuration.
Impact	Hypervisor based application consistent snapshot is taken on ESX. Crash consistent snapshot is taken for other hypervisors.
Resolution	Please look at the alert message and fix the invalid configuration.



Table 564: VSS Snapshot Aborted. [130106] [A130106]

Name	VSS Snapshot Aborted
Description	VSS snapshot is aborted by the Guest VM.
Alert message	VSS snapshot is aborted for the VM '<i>vm_name</i>' (ID <i>vm_id</i>)', because <i>reason</i>.
Cause	Guest VM aborted VSS snapshot operation.
Impact	Crash consistent snapshot is taken.
Resolution	Please look at the alert message and fix the issue.

Table 565: Recovered VM Disk Configuration Update Failed [130107] [A130107]

Name	Disk Configuration Update Failed
Description	Disk Configuration Update Failed
Alert message	Failed to make some disks online on the recovered VM '<i>vm_name</i>'. <i>reason</i>.
Cause	Nutanix Guest Tools failed to automatically bring the disks online.
Impact	Disks may become offline on the recovered VM.
Resolution	Manually bring the disks online on the recovered VM.
Resolution	Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.

Table 566: Execution of the PostThaw Script Failed [130113] [A130113]

Name	PostThaw Script Execution Failed
Description	Execution of the PostThaw Script Failed
Alert message	Failed to execute the post_thaw script during the creation of the application consistent snapshot with uuid '<i>snapshot_uuid</i>' for the VM '<i>vm_name</i>'. Error : <i>error_message</i>.
Cause	Guest VM failed to execute the post_thaw script during the creation of the application consistent snapshot.
Impact	Some of the applications that got stopped by the pre_freeze script may not start after application consistent snapshot is created for the VM.
Resolution	Manually run the post_thaw script on the guest VM.
Resolution	Fix the cause of the script failure to avoid any further execution failures.
Resolution	Resolve the specified reason for the failure. If you still cannot resolve the error, contact the Nutanix support.

Table 567: Failed to take the application-consistent snapshot for the VM. [130127] [A130127]

Name	Application Consistent Snapshot Skipped
Description	Failed to take the application-consistent snapshot for the VM.



Alert message	Failed to take the application-consistent snapshot for the VM '<i>vm_name</i>', because <i>reason</i>.
Cause	VMware tools is not installed on the VM.
Cause	VMware tools is not running on the VM.
Cause	Removal of temporary hypervisor-based snapshot of the VM failed.
Impact	Crash consistent snapshot is taken for the VM.
Resolution	Install the VMware tools on the VM.
Resolution	Start the VMware tools on the VM.
Resolution	Manually remove temporary hypervisor-based snapshot of the VM.
Resolution	Resolve the issue. If you cannot resolve the error, contact Nutanix support.

Table 568: VM Registration Warning [130132] [A130132]

Name	VM Registration Warning
Description	VM Registration caused warning
Alert message	VM with name <i>vm_name</i> has been recovered successfully, however the following issue was observed - <i>reason</i>
Cause	Processor features may not be compatible
Impact	VM registered successfully during disaster recovery but with warnings.
Resolution	No action is necessary.

Table 569: Guest Power Operation through NGT Failed [130149] [A130149]

Name	VM Guest Power Op Failed
Description	Guest Power Operation Failed
Alert message	Failed to perform operation '<i>operation</i>' on VM '<i>vm_name</i>'. <i>reason</i>.
Cause	Power operation failed in guest operating system.
Impact	Desired power operation may not be completed.
Resolution	Manually shutdown the VM by logging into the VM and running the appropriate shutdown command.
Resolution	Resolve the issue for the failure. If you cannot resolve the issue, contact Nutanix support.

Table 570: Removal of the temporary hypervisor snapshot failed for the VM. [130150] [A130150]

Name	Temporary Hypervisor Snapshot Cleanup Failed
Description	Removal of the temporary hypervisor snapshot that got created while taking the host-based application-consistent snapshot has failed for the VM.



Alert message	Removal of the temporary hypervisor snapshot that got created while taking the Nutanix host-based application-consistent snapshot has failed for the VM <code>vm_name</code> .
Cause	VM might be experiencing high I/O when host-based application-consistent snapshot cleanup is attempted.
Impact	Future application-consistent snapshots might fail.
Resolution	Wait until the merging of the snapshots is complete and then verify that there are no temporary Nutanix host-based snapshots before taking an application-consistent snapshot for the VM.

Table 571: Nutanix Guest Tools failed to initiate VM reboot. [130190] [A130190]

Name	NGT VM Reboot Failure
Description	Nutanix Guest Tools failed to initiate reboot on the VM.
Alert message	Nutanix Guest Tools failed to initiate VM reboot after NGT install/upgrade. <code>reason</code>.
Cause	Failed to initiate reboot on the VM, please check the logs for exact cause of failure.
Impact	VM is not rebooted after NGT install/upgrade.
Resolution	Check the logs for error information or contact Nutanix support.

Table 572: Conflicting NGT policies [130192] [A130192]

Name	NGT policy conflicts
Description	There are conflicting NGT policies for the VM.
Alert message	Conflicting NGT policies found for the VM.
Cause	More than one NGT policy is applicable for this VM. Please resolve the conflicts.
Impact	VM is not rebooted after NGT install/upgrade.
Resolution	Check the logs for error information or contact Nutanix support.

Table 573: Node Maintenance Mode Failure [130198] [A130198]

Name	Node Removal Failure
Description	VMs failed to evacuate off of the node.
Alert message	Failed to evacuate VMs off of node <code>node_uuid</code>: <code>error</code>
Cause	Other nodes in this cluster may not have enough resources available for VMs to migrate.
Impact	VMs can not migrate to other nodes in the cluster.
Resolution	Free the resources by switching off some VMs.



Table 574: NGT CD-ROM not Unmounted on the VM. [130334] [A130334]

Name	NGT CD-ROM Not Unmounted
Description	NGT CD-ROM not unmounted on the VM.
Alert message	NGT CD-ROM not unmounted on the VM 'vm_name (ID vm_id)'
Cause	NGT CD-ROM not unmounted on the VM because NGT version of the VM is incompatible with the NGT version of the cluster.
Impact	NGT CD-ROM is not unmounted on VM.
Resolution	Unmount the NGT ISO on the VM.

Table 575: File Analytics VM component failure. [160081] [A160081]

Name	File Analytics VM component failure
Description	One or more components of the File Analytics VM {avm_ip} are not functioning properly or have failed.
Alert message	One or more components of the File Analytics VM avm_ip are not functioning properly or have failed.
Cause	One or more File Analytics VM components have failed.
Impact	File Server analytics will not perform appropriately.
Resolution	Contact Nutanix Support.

Table 576: File Analytics VM high CPU usage [160082] [A160082]

Name	File Analytics VM high CPU usage.
Description	The File Analytics VM {avm_ip} has a high CPU usage of {usage_percent}%
Alert message	The File Analytics VM avm_ip has a high CPU usage of usage_percent%
Cause	The File Analytics VM is performing too many CPU intensive operations.
Impact	The File Analytics VM's performance will be impacted.
Resolution	Monitor which component has a high CPU usage and reduce the processing load.

Table 577: File Analytics VM high disk usage. [160083] [A160083]

Name	File Analytics VM high disk usage.
Description	The File Analytics VM {avm_ip} has a high disk usage of {usage_percent}%
Alert message	The File Analytics VM avm_ip has a high disk usage of usage_percent%
Cause	Extensive disk usage on the File Analytics VM
Impact	The File Analytics VM's performance will be impacted.
Resolution	Add more capacity to the File Analytics VM or delete some data.



Table 578: File Analytics VM low memory available. [160084] [A160084]

Name	File Analytics VM low memory available.
Description	The memory available on File Analytics VM {avm_ip} is {memory_available} bytes which is low.
Alert message	The percent memory available on File Analytics VM <code>avm_ip</code> is <code>memory_available</code> which is low
Cause	Excessive memory usage on the File Analytics VM
Impact	The File Analytics VM's performance will be impacted.
Resolution	Add more RAM

Table 579: vNUMA VM Pinning Failure. [200402] [A200402]

Name	vNUMA VM Pinning Failure
Description	Cluster was unable to provide requested vNUMA pinning for VM.
Alert message	VM '<code>vm_name</code>' with uuid <code>vm_id</code> is running with a non-compliant vNUMA pinning following <code>operation</code>.
Cause	Insufficient memory and/or CPU resources available.
Impact	VM performance may degrade.
Resolution	Migrate VM to a host which has adequate CPU and memory resources.

Table 580: NGT Update Available [600101] [A600101]

Name	NGT update available
Description	Checks whether VMs have latest NGT version installed.
Alert message	NGT on <code>vm_name</code> should be upgraded. It has improvements to allow simultaneous upgrades of NGT on multiple VMs when the next NGT upgrade is available.
Cause	NGT version on VM is not latest
Impact	This NGT update contains bug fixes and improvements to enable simultaneous upgrade of NGT on multiple VMs when the next NGT upgrade is available.
Resolution	Upgrade NGT version on VM to latest

Hardware

Table 581: Storage Device Health Bad [1005] [A1069]

Name	SmartCtl Health
Description	Smartctl health bad.



Alert message	Smartctl reports attribute <code>attribute</code> having value <code>value</code> for device <code>device</code> on CVM <code>service_vm_external_ip</code>.
Cause	Smartctl health bad.
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Check or replace disk.

Table 582: SATA DOM has failed. [1025] [A1165]

Name	SATADOM Status
Description	Checks that host SATA DOM is functioning.
Alert message	SATA DOM on host <code>ip_address</code> has failed.
Cause	SATA DOM has lost power connection.
Cause	SATA DOM has failed.
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Check the SATA DOM physical connection.
Resolution	Replace the SATA DOM as soon as possible. Refer to the Nutanix documentation for instructions.

Table 583: SATA DOM not reachable. [1030] [A1177]

Name	SATADOM Connection Status
Description	Checks that host SATA DOM is reachable.
Alert message	SATA DOM on host <code>ip_address</code> cannot be reached. The SATA DOM disconnection is determined by using sata dom information data: <code>data</code>
Cause	SATA DOM has lost power connection.
Cause	SATA DOM is not installed.
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Check if SATA DOM is installed.
Resolution	Check the SATA DOM physical connection.
Resolution	Replace the SATA DOM as soon as possible. Refer to the Nutanix documentation for instructions.



Table 584: SATA DOM has worn out. [1031] [A1180]

Name	SATADOM Wearout Status
Description	Checks the wearout of SATA DOM via SMART data.
Alert message	SATA DOM on host <code>ip_address</code> has PE cycles above 4500 or PE cycles above 3000 and daily PE cycles above 15.
Cause	SATA DOM has been in use for a long time.
Impact	Cluster performance may be significantly degraded or or the cluster may become unusable. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Replace SATA DOM as soon as possible. Refer to the Nutanix documentation for instructions.

Table 585: SATA DOM needs firmware version upgrade to S170119. [1033] [A1183]

Name	SATA DOM 3ME Date and Firmware Status
Description	Checks the firmware version of SATA DOM.
Alert message	SATA DOM on host <code>ip_address</code> has old firmware version: <code>alert_msg</code>
Cause	SATA DOM firmware version is not the latest version.
Impact	Impacted by Shift read retry issue. The problem can cause node to fail to boot.
Resolution	Update the firmware to the latest version. Perform inventory from Life Cycle Manager to check and update to the latest firmware version. Check KB 6937 for alternative update option.

Table 586: SATA DOM has Guest VM. [1034] [A1184]

Name	SATA DOM Guest VM Check
Description	Checks that no guest VM is installed on SATA DOM.
Alert message	SATA DOM on host <code>ip_address</code> contains a guest VM, which will accelerate the degradation of the SATA DOM.
Cause	A guest VM is installed on SATA DOM.
Cause	A VM is incorrectly configured.
Impact	Degradation of the SATA DOM will be accelerated, leading to unavailability of the node.
Resolution	Remove guest VMs from SATA DOM.
Resolution	Reconfigure guest VMs on the host machine.

Table 587: SATADOM-SL 3IE3 has high wear. [1035] [A1185]

Name	SATADOM-SL 3IE3 Wearout Status



Description	Checks the wearout of SATADOM-SL 3IE3 via SMART data.
Alert message	SATADOM-SL 3IE3 on host <code>ip_address</code> has device life smaller than 5%
Cause	SATADOM-SL 3IE3 Device Life too short (<5).
Impact	Cluster performance may be significantly degraded or or the cluster may become unusable. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Replace SATADOM-SL 3IE3 as soon as possible. Refer to the Nutanix documentation for instructions.

Table 588: {model} firmware version is not the latest firmware version. [1055] [A1192]

Name	SATA DOM Firmware Status
Description	Checks the firmware version of SATA DOM.
Alert message	<code>model</code> on host <code>ip_address</code> has firmware version <code>fwv</code>, need to upgrade to firmware version <code>Lfwv</code>.
Cause	SATA DOM firmware version is not the latest version.
Impact	The node may fail to start.
Resolution	Update the firmware to the latest version.

Table 589: Samsung PM1633 drive has worn out. [1056] [A1193]

Name	Samsung PM1633 Wearout Status
Description	Checks the status of Samsung PM1633 drive via SMART data.
Alert message	Samsung PM1633 drive <code>dev_name</code> on host <code>ip_address</code> has the following problems: <code>err_msg</code>
Cause	Samsung PM1633 drive revision is wrong or Samsung PM1633 drive has been in use for a long time and close to the maximum write endurance.
Impact	Cluster performance may be significantly degraded or or the cluster may become unusable. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Replace Samsung PM1633 drive as soon as possible. Refer to the Nutanix documentation for instructions.

Table 590: Toshiba PM3 drive has worn out. [1057] [A1194]

Name	Toshiba PM3 Status
Description	Checks the status of Toshiba PM3 drive via SMART data.
Alert message	Toshiba PM3 drive <code>dev_name</code> on host <code>ip_address</code> has the following problems: <code>err_msg</code>
Cause	Toshiba PM3 drive revision is wrong or Toshiba PM3 drive has been in use for a long time and close to the maximum write endurance.



Impact	Cluster performance may be significantly degraded or or the cluster may become unusable. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Replace Toshiba PM3 drive as soon as possible. Refer to the Nutanix documentation for instructions.

Table 591: NVMe Drive has errors. [1059] [A1196]

Name	NVMe Status Check
Description	Check that host NVMe drive is functioning properly.
Alert message	NVMe Drive on host <code>ip_address</code> has errors: <code>alert_msg</code>
Cause	NVMe Drive is damaged or worn out.
Impact	Cluster performance may be degraded. The integrity of data from NVMe drive may be harmed.
Resolution	Replace the problematic NVMe drive as soon as possible.

Table 592: Toshiba PM4 drive has worn out. [1060] [A1197]

Name	Toshiba PM4 Status
Description	Check the status of Toshiba PM4 drive.
Alert message	Toshiba PM4 drive <code>dev_name</code> on host <code>ip_address</code> has the following problems: <code>err_msg</code>
Cause	Toshiba PM4 drive revision is unsupported or Toshiba PM4 drive has been in use for a long time and close to the maximum write endurance.
Impact	Cluster performance may be significantly degraded or or the cluster may become unusable. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Replace Toshiba PM4 drive as soon as possible. Refer to the Nutanix documentation for instructions.

Table 593: Hypervisor Boot Drive Wear High [1061] [A1198]

Name	Hyve Boot Disk Status
Description	Check Hyve boot disk status
Alert message	Hypervisor boot drive on the Controller VM <code>ip_address</code> has problem: <code>alert_msg</code>
Cause	The drive is close to the maximum write endurance.
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.



Resolution	Replace the drive as soon as possible. Refer to the Nutanix documentation for instructions.

Table 594: More than one type of Toshiba PM4 drives installed on the node. [1062] [A1199]

Name	Toshiba PM4 Config
Description	Check the Configuration of Toshiba PM4 drives.
Alert message	Toshiba PM4 drive <code>dev_name</code> on host <code>ip_address</code> has the following problems: <code>err_msg</code>
Cause	More than one type of Toshiba PM4 drives are installed on the node.
Impact	Toshiba PM4 drives may not function properly.
Resolution	Keep only one type of Toshiba PM4 drives.

Table 595: Toshiba PM4 drives not compatible with AOS or Foundation version on the node. [1063] [A1200]

Name	Toshiba PM4 Version Compatibility
Description	Check the compatibility of Toshiba PM4 drives with AOS and Foundation versions.
Alert message	Toshiba PM4 drive <code>dev_name</code> on host <code>ip_address</code> has the following problems: <code>err_msg</code>
Cause	Toshiba PM4 drives are not compatible with the current AOS or foundation version.
Impact	Toshiba PM4 drives may not function properly.
Resolution	Upgrade AOS or Foundation version.

Table 596: Firmware version of Toshiba PM4 drives is old. [1064] [A1201]

Name	Toshiba PM4 FW Version
Description	Check the firmware of Toshiba PM4 drives.
Alert message	Toshiba PM4 drive <code>dev_name</code> on host <code>ip_address</code> has the following problems: <code>err_msg</code>
Cause	Firmware version of Toshiba PM4 drives is not the latest version.
Impact	Toshiba PM4 drives may not function properly.
Resolution	Upgrade Toshiba PM4 firmware version.

Table 597: Samsung PM1633 drives not compatible with AOS or Foundation version on the node. [1065] [A1202]

Name	Samsung PM1633 Version Compatibility



Description	Check the compatibility of Samsung PM1633 drives with AOS and Foundation versions.
Alert message	Samsung PM1633 drive <code>dev_name</code> on host <code>ip_address</code> has the following problems: <code>err_msg</code>
Cause	Samsung PM1633 drives are not compatible with the current AOS or foundation version.
Impact	Samsung PM1633 drives may not function properly.
Resolution	Upgrade AOS or Foundation version.

Table 598: Firmware version of Samsung PM1633 drives is old. [1066] [A1203]

Name	Samsung PM1633 FW Version
Description	Check the firmware of Samsung PM1633 drives.
Alert message	Samsung PM1633 drive <code>dev_name</code> on host <code>ip_address</code> has the following problems: <code>err_msg</code>
Cause	Firmware version of Samsung PM1633 drives is not the latest version.
Impact	Samsung PM1633 drives may not function properly.
Resolution	Upgrade Samsung PM1633 firmware version. Perform inventory from Life Cycle Manager to check and update to the latest firmware version. Check KB 6937 for alternative update option.

Table 599: NIC Flaps [3066] [A1067]

Name	NIC Flapping Check
Description	Check that all nics have no flapping.
Alert message	NIC <code>nic_name</code> in host <code>host_ip</code> has flapped <code>nic_flap_count</code> times in one hour.
Cause	A physical networking component is failing.
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Check all physical networking components (including switch,cables, and NIC) and replace the failed component.
NIC Flap Threshold	Warning Threshold: 5 flaps

Table 600: NIC Error Rate Check [6011]

Name	NIC Error Rate Check
Description	Checks whether the NICs on hypervisor are experiencing errors.
Cause	NIC may be failing.



Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Check NIC health and if required, replace the NIC as soon as possible. Refer to the Nutanix documentation for instructions.

Table 601: NIC RX_CRC Error Rate High [6013] [A6013]

Name	NIC CRC Error Rate Check
Description	Checks whether the NICs on hypervisor are experiencing rx_crc errors.
Alert message	NIC <code>nic_name</code> in host <code>host_ip</code> has encountered many <code>error_type</code>.
Cause	NIC or its peripherals may be failing.
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Replace the NIC, cable, SFP+ adapter, or switch port. Refer to the Nutanix documentation for instructions.
RX CRC Error Percentage Threshold	Critical Threshold: 1 %

Table 602: NIC RX_Missed Error Rate High [6014] [A6014]

Name	NIC Missed Error Rate Check
Description	Checks whether the NICs on hypervisor are experiencing rx_missed errors.
Alert message	NIC <code>nic_name</code> in host <code>host_ip</code> has encountered many <code>error_type</code>.
Cause	RX buffer size or transient condition.
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	RX buffer size might be too small to absorb current inbound traffic. This can be a transient condition. If the alert reappears frequently, contact Nutanix support for instructions on how to increase RX buffer.
RX Missed Error Percentage Threshold	Warning Threshold: 1 %

Table 603: SSD IO Latency [6515]

Name	SSD IO Latency
Description	Checks whether the system load average is high.



Cause	There may be too many VMs on one host.
Impact	High I/O latencies may be experienced by some workloads.
Resolution	Redistribute VMs to reduce load on the affected Controller VM.

Table 604: Fan Speed Low [15006] [A1045]

Name	Fan Speed Low Check
Description	Check that fan speed is not low
Alert message	<code>fan_id</code> has stopped or its speed is low on Controller VM <code>ip_address</code>.
Cause	A fan has failed.
Impact	The device may fail or be permanently damaged.
Resolution	Replace the fan as soon as possible. Refer to the Nutanix documentation for instructions.
Lower threshold for fan speed in RPM	Critical Threshold: 500 rpm

Table 605: Fan Speed High [15007] [A1020]

Name	Fan Speed High Check
Description	Check that fan speed is not too high
Alert message	Speed of <code>fan_id</code> exceeded <code>fan_rpm</code> RPM on Controller VM <code>ip_address</code>.
Cause	The device is overheating to the point of imminent failure.
Impact	The device may fail or be permanently damaged.
Resolution	Ensure that the fans in the block are functioning properly and that the environment is cool enough.
Upper threshold for fan speed in RPM	Critical Threshold: 25000 rpm

Table 606: RAM Fault [15008] [A1052]

Name	Ram Fault Status
Description	Check if current available memory has gone below the installed size.
Alert message	DIMM fault detected on host <code>ip_address</code>. The node is running with <code>current_memory_gb</code> GB whereas <code>installed_memory_gb</code> GB was installed. <code>dimmsg</code>
Cause	A DIMM has failed, or DIMMs have been removed from the node.
Impact	Memory capacity on the node is less than expected.



Resolution	Determine if a failed DIMM needs to be replaced. Refer to the Nutanix documentation for instructions.

Table 607: CPU Temperature High [15010] [A1049]

Name	CPU Temperature
Description	Check that cpu temperature is not too high.
Alert message	Temperature of <code>cpu_id</code> exceeded <code>temperatureC</code> on Controller VM <code>ip_address</code>. <code>fan_speed</code>
Cause	The device is overheating to the point of imminent failure.
Impact	The device may fail or be permanently damaged.
Resolution	Ensure that the fans in the block are functioning properly and that the environment is cool enough.

Table 608: RAM Temperature High [15011] [A1022]

Name	DIMM Temperature
Description	Check that DIMM temperature is not high.
Alert message	Temperature of <code>dimm_id</code> exceeded <code>temperatureC</code> on Controller VM <code>ip_address</code>
Cause	DIMM temperature is high.
Impact	The device may fail or be permanently damaged.
Resolution	Check ram temperature.

Table 609: System Temperature High [15012] [A1012]

Name	System Temperature
Description	Check that system temperature is not high
Alert message	System temperature exceeded <code>temperatureC</code> on Controller VM <code>ip_address</code>
Cause	The device is overheating to the point of imminent failure.
Impact	The device may fail or be permanently damaged.
Resolution	Ensure that the fans in the block are functioning properly and that the environment is cool enough.

Table 610: GPU Temperature High [15014] [A1070]

Name	GPU Temperature
Description	Check that gpu temperature is not high
Alert message	Temperature of <code>gpu_id</code> exceeded <code>temperatureC</code> on host <code>ip_address</code>



Cause	The device is overheating to the point of imminent failure.
Impact	The device may fail or be permanently damaged.
Resolution	Ensure that the fans in the block are functioning properly and that the environment is cool enough.

Table 611: Hardware Clock Failure [15015] [A1059]

Name	Hardware Clock Status
Description	Check if Hardware clock has failed
Alert message	Hardware clock in host <code>host_ip</code> has failed.
Cause	The RTC clock on the host has failed or the RTC battery has died.
Impact	The host may fail after restart and remain unavailable until it is replaced.
Resolution	Replace the node. Follow the instructions mentioned in KB 2462.

Table 612: IPMI Error [15016] [A1050]

Name	IPMI SDR Status
Description	Checks whether IPMI SDR has failures.
Alert message	Controller VM <code>ip_address</code> is unable to fetch IPMI SDR repository.
Cause	The IPMI interface is down.
Cause	There is a network connectivity issue.
Cause	The hypervisor password does not match what is configured in the cluster.
Cause	The cluster did not correctly fail back after a Controller VM was temporarily down.
Impact	Nutanix support staff are able to access the cluster to assist with any issue.
Resolution	Ensure that the IPMI interface is functioning.
Resolution	Ensure that physical networking, VLANs, and virtual switches are configured correctly.
Resolution	Ensure that the hypervisor password configured in the cluster matches the actual hypervisor password.
Resolution	If the alert persists, contact Nutanix support.

Table 613: GPU Fault [15017] [A1071]

Name	GPU Status
Description	Check that GPU have not faulted.
Alert message	<code>gpu_model</code> GPU fault detected on host <code>host_ip</code>. The node is running with <code>current_gpu_count</code> GPU(s) whereas <code>installed_gpu_count</code> GPU(s) were installed.



Cause	The device is failing.
Impact	Graphics acceleration is not available to guest VMs.
Resolution	Replace the device. Refer to the Nutanix documentation for instructions.

Table 614: Power Supply Status Down/Unretrievable [15018] [A1046]

Name	Power Supply Status
Description	Check that Power supply has no errors.
Alert message	<code>per_psu_msg</code>
Cause	1. The power supply has failed or status not retrievable because BMC can not communicate with PSU properly2. BMC version may not be up to date and/or able to communicate with PSU for accurate status
Impact	The affected block does not have power supply redundancy.
Resolution	1. Restore the power supply as soon as possible.Refer to the Nutanix documentation for instructions.2. If PSU info unavailable, upgrade to the latest BMC firmware, and apply factory default after BMC upgrade. If it does not solve the issue, Refer to Nutanix KB 7386

Table 615: High number of correctable ECC errors found. [15019] [A1187]

Name	IPMI SEL CECC Check
Description	Check for number of CECC errors in the IPMI SEL.
Alert message	<code>per_dimm_msg</code>
Cause	A memory module in the node is failing.
Impact	The node may stop running, reducing cluster capacity.
Resolution	Replace the failing memory module as soon as possible. Refer to the Nutanix documentation for instructions.

Table 616: CPU Voltage Abnormal [15024] [A15024]

Name	CPU Voltage
Description	Check that CPU voltage is within normal range.
Alert message	Voltage of <code>cpu_id</code> not within normal range on Controller VM <code>ip_address</code>
Cause	Server board component problem may cause the CPU voltage to be too high or too low.
Impact	The device may fail or be permanently damaged.
Resolution	Replace the node. Refer to the Nutanix documentation for instructions.



Table 617: RAM Voltage Abnormal [15025] [A15025]

Name	DIMM Voltage
Description	Check that DIMM voltage is within normal range.
Alert message	Voltage of <code>dim_id</code> is not within normal range on Controller VM <code>ip_address</code>
Cause	Memory problem and server board problem may cause DIMM voltage to be too high or too low.
Impact	The device may fail or be permanently damaged.
Resolution	Replace the memory. If the problem persists, replace the node. Refer to the Nutanix documentation for instructions

Table 618: CPU-VRM Temperature High [15026] [A15026]

Name	CPU-VRM Temperature
Description	Check that CPU-VRM temperature is not too high.
Alert message	Temperature of CPU-VRM <code>cpu_vrm_id</code> exceeded <code>temperatureC</code> on Controller VM <code>ip_address</code>
Cause	Server board component problem may cause the CPU-VRM to be overheating.
Impact	The device may fail or be permanently damaged.
Resolution	Replace the node. Refer to the Nutanix documentation for instructions.

Table 619: RAM-VRM Temperature High [15027] [A15027]

Name	DIMM-VRM Temperature
Description	Check that DIMM-VRM temperature is not high.
Alert message	Temperature of DIMM-VRM <code>dim_vrm_id</code> exceeded <code>temperatureC</code> on Controller VM <code>ip_address</code>
Cause	Server board component problem may cause DIMM-VRM temperature to be high.
Impact	The device may fail or be permanently damaged.
Resolution	Replace the node. Refer to the Nutanix documentation for instructions.

Table 620: CPU Temperature Reading Error [15028] [A15028]

Name	CPU Temperature Fetch
Description	Check that CPU temperature can be fetched.
Alert message	Failed to fetch temperature of <code>cpu_id</code> on Controller VM <code>ip_address</code>
Cause	CPU temperature sensor may be unstable or BMC firmware may be reporting temperature inaccurately.



Impact	CPU temperature cannot be monitored.
Resolution	Check BMC SEL log. If CPU temperature in SEL log is 0, check CPU temperature sensor. If CPU temperature in SEL is not 0, ensure that the BMC firmware version is correct.

Table 621: CPU Temperature Low [15031] [A15031]

Name	CPU Temperature Low
Description	Check that cpu temperature is not too low.
Alert message	Temperature of <code>cpu_id</code> is below <code>temperatureC</code> on Controller VM <code>ip_address</code>
Cause	BMC or the sensor on CPU may have problem
Impact	The device may fail or be permanently damaged.
Resolution	Do BMC factory default. If the problem still persists, replace the CPU.

Table 622: RAM Temperature Low [15032] [A15032]

Name	DIMM Temperature Low
Description	Check that DIMM temperature is not low.
Alert message	Temperature of <code>dimm_id</code> is below <code>temperatureC</code> on Controller VM <code>ip_address</code>
Cause	BMC or the sensor on DIMM may have problem
Impact	The device may fail or be permanently damaged.
Resolution	Do BMC factory default. If the problem still persists, replace the chassis.

Table 623: System Temperature Low [15033] [A15033]

Name	System Temperature Low
Description	Check that system temperature is not low
Alert message	System temperature is below <code>temperatureC</code> on Controller VM <code>ip_address</code>
Cause	BMC or the sensor on chassis may have problem
Impact	The device may fail or be permanently damaged.
Resolution	Do BMC factory default. If the problem still persists, replace the chassis.

Table 624: System Temperature Reading Error [15034] [A15034]

Name	System Temperature Fetch
Description	Check that System temperature can be fetched.
Alert message	Failed to fetch system temperature on Controller VM <code>ip_address</code>



Cause	System temperature sensor may be unstable or BMC firmware may be reporting temperature inaccurately.
Impact	System temperature cannot be monitored.
Resolution	Check BMC SEL log. If System temperature in SEL log is 0, check System temperature sensor. If System temperature in SEL is not 0, ensure that the BMC firmware version is correct.

Table 625: [15039] [A15039]

Name	IPMI SEL UECC Check
Description	Check for number of UECC errors for last one day in the IPMI SEL.
Alert message	DIMM UECC errors found for <code>dimm_id</code> on node <code>host_ip</code>
Cause	A memory module in the node is failing.
Impact	The node may stop running, reducing cluster capacity.
Resolution	Replace the failing memory module as soon as possible. Refer to the Nutanix documentation for instructions.

Table 626: Hypervisor boot drive RAID is in an unhealthy state. [101058] [A101058]

Name	Hypervisor boot drive RAID status
Description	Checks the status of Dell Marvell BOSS hypervisor boot drive.
Alert message	Hypervisor boot drive RAID on <code>ip_address</code> is in <code>raid_state</code> state.
Cause	Hypervisor boot drive RAID is not functioning properly.
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Contact Support to help restore the hypervisor boot drive RAID.

Table 627: SATA DOM needs firmware upgrade. [101060] [A101060]

Name	SATA DOM 3IE3 Firmware Status
Description	Checks the firmware version of SATA DOM SL 3IE3.
Alert message	SATA DOM on host <code>ip_address</code> has old firmware version: <code>alert_msg</code>
Cause	SATA DOM SL 3IE3 firmware version is not the latest version.
Impact	Impacted by Shift read retry issue. The problem can cause node to fail to boot.
Resolution	Update the firmware to the latest version. Perform inventory from Life Cycle Manager to check and update to the latest firmware version. Check KB 6937 for alternative update option.



Table 628: Free block count of SATA DOM has gone below threshold. [101062] [A101062]

Name	SATA DOM 3IE3 Free Block Count Check
Description	Checks the free block count of SATA DOM SL 3IE3.
Alert message	<i>alert_msg</i>
Cause	SATA DOM SL 3IE3 free block count is less than 64.
Impact	As FBC approaches zero the host boot disk performance may be impacted.
Resolution	Ensure SATADOM firmware is up to date via LCM or KB6937. Refer to KB4395 for using TRIM script to proactively reclaim FBC.

Table 629: Toshiba PM5 drive has worn out. [101066] [A101066]

Name	Toshiba PM5 Status
Description	Check the status of Toshiba PM5 drive.
Alert message	Toshiba PM5 drive <i>dev_name</i> on host <i>ip_address</i> has the following problems: <i>err_msg</i>
Cause	Toshiba PM5 drive has been in use for a long time and close to the maximum write endurance.
Impact	Cluster performance may be significantly degraded or or the cluster may become unusable. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Replace Toshiba PM5 drive as soon as possible. Refer to the Nutanix documentation for instructions.

Table 630: More than one type of Toshiba PM5 drives installed on the node. [101067] [A101067]

Name	Toshiba PM5 Config
Description	Check the Configuration of Toshiba PM5 drives.
Alert message	Toshiba PM5 drive <i>dev_name</i> on host <i>ip_address</i> has the following problems: <i>err_msg</i>
Cause	More than one type of Toshiba PM5 drives are installed on the node.
Impact	Toshiba PM5 drives may not function properly.
Resolution	Keep only one type of Toshiba PM5 drives.

Table 631: SM863 drive has worn out. [106017] [A106017]

Name	Samsung SM863 SSD status check
Description	Check the status of SM863 and SM863a SSD drive
Alert message	<i>dev_type</i> drive <i>dev_name</i> on host <i>ip_address</i> has the following problems: <i>err_msg</i>



Cause	SM863a SSD drive revision is unsupported or SM863 and SM863a has been in use for a long time and close to maximum write endurance.
Impact	Cluster performance may be significantly degraded or or the cluster may become unusable. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Replace SM863 or SM863a SSD drive as soon as possible. Refer to the Nutanix documentation for instructions.

Table 632: SAS Connectivity Not Normal [106019] [A106019]

Name	SAS Connectivity
Description	Check SAS connectivity
Alert message	SAS cables missing or not properly connected on host.
Cause	SAS cables are not properly connected.
Impact	Cluster performance may be significantly degraded.
Resolution	Ensure that SAS cables are properly connected.

Table 633: CPUs of different types or models are in the same chassis. [106026] [A1190]

Name	CPU type on chassis check
Description	Checks whether CPUs within a chassis are of the same type.
Alert message	Chassis <code>chassis_id</code> has CPU configuration problem: <code>alert_msg</code>. This problem may affect the performance and stability of the nodes in the chassis. Chassis cpu info: <code>chassis_cpu_info</code>
Cause	A node of a different type was added to the chassis.
Impact	Having nodes of different CPU types in the same chassis is not supported.
Resolution	Remove the node that has a different CPU type from the chassis. Before moving a node to a chassis, make sure its CPU type is the same as that of other nodes in the chassis.

Table 634: Micron5100 drive has worn out. [106029] [A106029]

Name	Micron5100 SSD status check
Description	Check the status of Micron5100 SSD drive
Alert message	Micron5100 drive <code>dev_name</code> on host <code>ip_address</code> has the following problems: <code>err_msg</code>
Cause	Micron5100 SSD drive revision is unsupported or is close to maximum write endurance.
Impact	Cluster performance may be significantly degraded or or the cluster may become unusable. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.



Resolution	Replace Micron5100 SSD drive as soon as possible. Refer to the Nutanix documentation for instructions.

Table 635: SM863a drives not compatible with AOS or Foundation version on the node.
[106030] [A106030]

Name	Samsung SM863a version compatibility check
Description	Check the compatibility of SM863a drives with AOS and Foundation versions.
Alert message	<code>dev_type</code> drive <code>dev_name</code> on host <code>ip_address</code> has the following problems: <code>err_msg</code>
Cause	SM863a drives are not compatible with the current AOS or foundation version.
Impact	SM863a drives may not function properly.
Resolution	Upgrade AOS or Foundation version.

Table 636: Firmware version of SM863 or SM863a drives is old. [106031] [A106031]

Name	Samsung SM863/SM863a FW version check
Description	Check the firmware of SM863 or SM863a drives.
Alert message	<code>dev_type</code> drive <code>dev_name</code> on host <code>ip_address</code> has the following problems: <code>err_msg</code>
Cause	Firmware version of SM863 or SM863a drives is not the latest one.
Impact	SM863 or SM863a drives may not function properly.
Resolution	Upgrade SM863 or SM863a firmware version. Perform inventory from Life Cycle Manager to check and update to the latest firmware version. Check KB 6937 for alternative update option.

Table 637: PM863a drives not compatible with AOS or Foundation version on the node.
[106033] [A106033]

Name	Samsung PM863a version compatibility check
Description	Check the compatibility of PM863a drives with AOS and Foundation versions.
Alert message	<code>PM863a</code> drive <code>dev_name</code> on host <code>ip_address</code> has the following problems: <code>err_msg</code>
Cause	PM863a drives are not compatible with the current AOS or foundation version.
Impact	Expand cluster and foundation workflows will not work.
Resolution	Upgrade AOS or Foundation version.

Table 638: Firmware version of PM863a drives is old. [106034] [A106034]

Name	Samsung PM863a FW version check



Description	Check firmware of PM863a drives.
Alert message	PM863a drive <code>dev_name</code> of model <code>dev_model</code> on host <code>ip_address</code> has the following problems: <code>err_msg</code>
Cause	Firmware version of PM863a drives is not the latest one.
Impact	PM863a drives may not function properly.
Resolution	Upgrade PM863a firmware version. Perform inventory from Life Cycle Manager to check and update to the latest firmware version. Check KB 6937 for alternative update option.

Table 639: PM863a drive has worn out. [106035] [A106035]

Name	Samsung PM863a status check
Description	Check the status of PM863a SSD
Alert message	PM863a drive <code>dev_name</code> of model <code>dev_model</code>, serial id <code>serial_id</code>, location <code>location</code> on host <code>ip_address</code> has the following problems: <code>err_msg</code>
Cause	PM863a is approaching maximum write endurance.
Impact	Cluster performance may be significantly degraded or the cluster may become unusable. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Replace PM863a SSD as soon as possible. Refer to the Nutanix documentation for instructions.

Table 640: SATADOM-ML 3SE has high wear. [106037] [A106037]

Name	SATADOM-ML 3SE Wearout Status
Description	Checks the wearout of SATADOM-ML 3SE via SMART data.
Alert message	SATADOM-ML 3SE on host <code>ip_address</code> has the following problem: <code>alert_msg</code>
Cause	SATADOM-ML 3SE Remaining Device Life too short.
Impact	Cluster performance may be significantly degraded or or the cluster may become unusable. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Replace SATADOM-ML 3SE as soon as possible. Refer to the Nutanix documentation for instructions.

Table 641: M.2 Micron5100 host boot drive has worn out. [106039] [A106039]

Name	M.2 Micron5100 host boot drive status check
Description	Check the status of M.2 Micron5100 host boot drive
Alert message	M.2 Micron5100 host boot drive <code>dev_name</code>, serial id <code>serial_id</code> on host <code>ip_address</code> has the following problems: <code>err_msg</code>
Cause	M.2 Micron5100 host boot drive is close to maximum write endurance.



Impact	Cluster performance may be significantly degraded or or the cluster may become unusable. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Replace M.2 Micron5100 host boot drive as soon as possible. Refer to the Nutanix documentation for instructions.

Table 642: M.2 Intel S3520 host boot drive has worn out. [106040] [A106040]

Name	M.2 Intel S3520 host boot drive status check
Description	Check the status of M.2 Intel S3520 host boot drive
Alert message	M.2 Intel S3520 host boot drive <code>dev_name</code> on host <code>ip_address</code> has the following problems: <code>err_msg</code>
Cause	M.2 Intel S3520 host boot drive is close to maximum write endurance.
Impact	Cluster performance may be significantly degraded or or the cluster may become unusable. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Replace M.2 Intel S3520 host boot drive as soon as possible. Refer to the Nutanix documentation for instructions.

Table 643: Host boot disk serial number has changed [106042] [A106042]

Name	M.2 Boot Disk change check
Description	Check the change of host boot disk
Alert message	Host boot disk on host <code>ip_address</code> has the following problems: <code>err_msg</code>
Cause	Active host boot disk failed and needs to switch to spare host boot disk to boot the host.
Impact	Host may be experiencing boot problems
Resolution	Replace the active host boot disk in future. Refer to the Nutanix documentation for instructions.

Table 644: Intel S4600 drives not compatible with AOS or Foundation version on the node. [106043] [A106043]

Name	Intel S4600 version compatibility check
Description	Check the compatibility of Intel S4600 drives with AOS and Foundation versions.
Alert message	Intel S4600 drive <code>dev_name</code> on host <code>ip_address</code> has the following problems: <code>err_msg</code>
Cause	Intel S4600 drives are not compatible with the current AOS or foundation version.
Impact	Expand cluster and foundation workflows will not work.



Resolution	Upgrade AOS or Foundation version.

Table 645: Firmware version of Intel S4600 drives is old. [106044] [A106044]

Name	Intel S4600 FW version check
Description	Check firmware of Intel S4600 drives.
Alert message	Intel S4600 drive <code>dev_name</code> of model <code>dev_model</code> on host <code>ip_address</code> has the following problems: <code>err_msg</code>
Cause	Firmware version of Intel S4600 drives is not the latest one.
Impact	Intel S4600 drives may not function properly.
Resolution	Upgrade Intel S4600 firmware version.

Table 646: Intel S4600 drive has worn out. [106045] [A106045]

Name	Intel S4600 status check
Description	Check the status of Intel S4600 SSD
Alert message	Intel S4600 drive <code>dev_name</code> of model <code>dev_model</code> on host <code>ip_address</code> has the following problems: <code>err_msg</code>
Cause	Intel S4600 is approaching maximum write endurance.
Impact	Cluster performance may be significantly degraded or the cluster may become unusable. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Replace Intel S4600 SSD as soon as possible. Refer to the Nutanix documentation for instructions.

Table 647: Firmware version of Intel S4600 drives is old or invalid. [106049] [A106049]

Name	Intel S4600 FW version check on Lenovo
Description	Check firmware of Intel S4600 drives on Lenovo platform.
Alert message	Intel S4600 drive <code>dev</code> of <code>model</code> on Lenovo host <code>ip_address</code> has the following problems: <code>err_msg</code>
Cause	Firmware version of Intel S4600 drives on Lenovo platform is an old one that has bugs.
Impact	Intel S4600 drives may not function properly.
Resolution	Upgrade Intel S4600 firmware version to at least version LU41 on Lenovo platform.

Table 648: PM883 drive has worn out. [106051] [A106051]

Name	Samsung PM883 status check



Description	Check the status of PM883 SSD
Alert message	PM883 drive <code>dev_name</code> of model <code>dev_model</code>, serial id <code>serial_id</code>, location <code>location</code> on host <code>ip_address</code> has the following problems: <code>err_msg</code>
Cause	PM883 is approaching maximum write endurance.
Impact	Cluster performance may be significantly degraded or the cluster may become unusable. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Replace PM883 SSD as soon as possible. Refer to the Nutanix documentation for instructions.

Table 649: SM883 drive has worn out. [106052] [A106052]

Name	Samsung SM883 SSD status check
Description	Check the status of SM883 SSD drive
Alert message	<code>dev_type</code> drive <code>dev_name</code> on host <code>ip_address</code> has the following problems: <code>err_msg</code>
Cause	SM883 has been in use for a long time and close to maximum write endurance.
Impact	Cluster performance may be significantly degraded or or the cluster may become unusable. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Replace SM883 SSD drive as soon as possible. Refer to the Nutanix documentation for instructions.

Table 650: M10 GPU Configuration Wrong on the Node [111040] [A111040]

Name	M10 GPU check
Description	Check number of M10 GPUs on a node
Alert message	M10 GPU configuration has problem on node <code>ip_address</code>: <code>err_msg</code>
Cause	The number of M10 GPU cards on a node is not 1 or 2 or 3, or the number of M10 GPUs on a card is not 4.
Impact	If the number of M10 GPU cards is not 1 or 2 or 3, the node cannot reboot. If the number of M10 GPU on a card is not 4, the VMs that the missing M10 GPU is assigned to may not run.
Resolution	Add or remove M10 GPU cards so that either 1 or 2 or 3 are installed on the node. Replace M10 GPU card that does not have 4 M10 GPUs.

Table 651: M60 GPU Configuration Wrong on the Node [111041] [A111041]

Name	M60 GPU check
Description	Check number of M60 GPUs on a node
Alert message	M60 GPU configuration has problem on node <code>ip_address</code>: <code>err_msg</code>



Cause	The number of M60 GPU cards on a node is not 1 or 2 or 3, or the number of M60 GPUs on a card is not 2.
Impact	If the number of M60 GPU cards is not 1 or 2 or 3, the node cannot reboot. If the number of M60 GPU on a card is not 2, the VMs that the missing M60 GPU is assigned to may not run.
Resolution	Add or remove M60 GPU cards so that either 1 or 2 or 3 are installed on the node. Replace M60 GPU card that does not have 2 M60 GPUs.

Table 652: More than one type of GPUs Installed on the same Node [111049] [A111049]

Name	Mixed GPU check
Description	Check if more than one type of GPUs installed on the same node
Alert message	More than one type of GPUs are installed on node <code>ip_address: alert_msg</code>
Cause	More than one type of GPUs are installed on the same node
Impact	VMs running on the node may experience performance degradation or become temporarily unavailable
Resolution	Keep only one type of GPUs on the node

Table 653: P40 GPU Configuration Wrong on the Node [111057] [A111057]

Name	P40 GPU check
Description	Check number of P40 GPUs on a node
Alert message	P40 GPU configuration has problem on node <code>ip_address: err_msg</code>
Cause	The number of P40 GPUs on a node is not 1 or 2 or 3.
Impact	If the number of P40 GPUs is not 1 or 2 or 3, the node cannot reboot.
Resolution	Add or remove P40 GPU so that either 1 or 2 or 3 are installed on the node.

Table 654: P40 GPU BMC Version is old on the Node [111060] [A111060]

Name	P40 GPU BMC check
Description	Check BMC version of P40 GPUs on a node
Alert message	P40 GPU BMC Version has problem on node <code>ip_address: err_msg</code>
Cause	P40 BMC on a node is not up to the required version.
Impact	Thermal functions for P40 would not work if BMC version is before 3.56.
Resolution	Upgrade P40 BMC to the latest version, at least BMC with version 3.56.

Table 655: V100 GPU Configuration Wrong on the Node [111065] [A111065]

Name	V100 GPU check



Description	Check number of V100 GPUs on a node
Alert message	V100 GPU configuration has problem on node <code>ip_address</code>: <code>err_msg</code>
Cause	The number of V100 GPUs on a node is not 1 or 2 or 3.
Impact	If the number of V100 GPUs is not 1 or 2 or 3, the node cannot reboot.
Resolution	Add or remove V100 GPU so that either 1 or 2 or 3 are installed on the node.

Table 656: P4 GPU Configuration Wrong on the Node [111069] [A111069]

Name	P4 GPU check
Description	Check number of P4 GPUs on a node
Alert message	P4 GPU configuration has problem on node <code>ip_address</code>: <code>err_msg</code>
Cause	The number of P4 GPUs on a node is not 1 to 6.
Impact	If the number of P4 GPUs is not 1 to 6, the node cannot reboot.
Resolution	Add or remove P4 GPUs so that 1 to 6 P4 GPUs are installed on the node.

Table 657: Node Failure [130030] [A1137]

Name	Node Status
Description	Node Failure
Alert message	Host <code>hypervisor_address</code> appears to have failed. High Availability is restarting VMs on <code>failover_host_info</code>.
Cause	Host is not accessible.
Impact	VMs will restart on another host.
Resolution	Check for network connectivity or a hardware failure.

Table 658: Non Self Encryption Drive Disk Inserted [130031] [A1122]

Name	Non SED Disk Inserted Check
Description	Non Self Encrypting Drive disk inserted
Alert message	Non encrypted disk with serial <code>disk_serial</code> was added in drive bay <code>disk_location</code> on Controller VM <code>service_vm_external_ip</code>. It is not usable because the rest of the cluster is protected using encrypted drives.
Cause	A cluster with self-encrypting drives has a non-self encrypting drive installed.
Impact	The new disk was not mounted because the cluster cannot be fully secured without all drives supporting encryption.
Resolution	Ensure all installed drives are of self-encrypting type. If more help is needed call Nutanix support.



Table 659: Disk Failed/Marked offline [130033] [A1044]

Name	Disk Offline Status
Description	Disk Bad
Alert message	Disk <code>disk_position</code> with id <code>disk_id</code> on node <code>node_position</code> of block <code>block_position</code> has been marked offline due to IO errors. Serial number of the disk is <code>disk_serial</code> in host <code>host_ip</code> of block <code>block_serial</code> .
Cause	A drive on the node has failed and was marked offline.
Impact	Cluster storage capacity will be reduced until the drive is replaced.
Resolution	Replace the failed drive. Refer to the Nutanix documentation for instructions.

Table 660: Physical Disk Drive Has Failed [130035] [A1104]

Name	Physical Disk Status
Description	Physical drive has failed.
Alert message	Drive <code>disk_id</code> with serial <code>disk_serial</code> and model <code>disk_model</code> in drive bay <code>disk_location</code> located on node <code>node_serial_number</code> in position <code>node_position</code> on Controller VM <code>service_vm_external_ip</code> has failed.
Cause	The drive has failed.
Impact	Cluster storage capacity is reduced.
Resolution	Replace the failed drive. Please refer to KB 4158.

Table 661: Physical Disk Removed From Slot [130036] [A1103]

Name	Physical Disk Remove Check
Description	Physical disk removed from slot
Alert message	Disk with serial <code>disk_serial</code> and model <code>disk_model</code> was removed from drive bay <code>disk_location</code> on Controller VM <code>service_vm_external_ip</code>
Cause	A drive was physically disconnected from a slot.
Impact	Migration of data from the disk will start.
Resolution	Reseat the removed disk back in slot. Contact the Hardware Vendor if issue persists.

Table 662: Disk Diagnostic Failure [130089] [A1139]

Name	Disk Diagnostic Status
Description	The drive diagnostic test has failed.
Alert message	Drive <code>disk_id</code> with model <code>disk_model</code> with serial <code>disk_serial</code> in drive bay <code>disk_location</code> on Controller VM <code>service_vm_external_ip</code> has failed diagnostic test. This drive is failing or has failed.



Cause	The drive is failing or has failed.
Impact	Cluster storage capacity has been reduced.
Resolution	Replace the failing drive. Refer to the Nutanix documentation for instructions.

Table 663: Node Maintenance Mode Failure [130198] [A130198]

Name	Node Removal Failure
Description	VMs failed to evacuate off of the node.
Alert message	Failed to evacuate VMs off of node <code>node_uuid</code>: error
Cause	Other nodes in this cluster may not have enough resources available for VMs to migrate.
Impact	VMs can not migrate to other nodes in the cluster.
Resolution	Free the resources by switching off some VMs.

License

Table 664: Files License Invalid [160086] [A160086]

Name	File Server License Usage Check
Description	Check for Files license capacity violation
Alert message	Cluster does not have enough Files license capacity. <code>fs_license_msg</code>
Cause	The Files license capacity for the cluster is non-compliant.
Impact	Cluster does not have enough Files license capacity.
Resolution	Apply for a new license.
Resolution	Increase Files license capacity.

Node

Table 665: SATA DOM has failed. [1025] [A1165]

Name	SATADOM Status
Description	Checks that host SATA DOM is functioning.
Alert message	SATA DOM on host <code>ip_address</code> has failed.
Cause	SATA DOM has lost power connection.
Cause	SATA DOM has failed.
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Check the SATA DOM physical connection.



Resolution	Replace the SATA DOM as soon as possible. Refer to the Nutanix documentation for instructions.

Table 666: SATA DOM not reachable. [1030] [A1177]

Name	SATADOM Connection Status
Description	Checks that host SATA DOM is reachable.
Alert message	SATA DOM on host <code>ip_address</code> cannot be reached. The SATA DOM disconnection is determined by using sata dom information data: <code>data</code>
Cause	SATA DOM has lost power connection.
Cause	SATA DOM is not installed.
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Check if SATA DOM is installed.
Resolution	Check the SATA DOM physical connection.
Resolution	Replace the SATA DOM as soon as possible. Refer to the Nutanix documentation for instructions.

Table 667: SATA DOM has worn out. [1031] [A1180]

Name	SATADOM Wearout Status
Description	Checks the wearout of SATA DOM via SMART data.
Alert message	SATA DOM on host <code>ip_address</code> has PE cycles above 4500 or PE cycles above 3000 and daily PE cycles above 15.
Cause	SATA DOM has been in use for a long time.
Impact	Cluster performance may be significantly degraded or or the cluster may become unusable. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Replace SATA DOM as soon as possible. Refer to the Nutanix documentation for instructions.

Table 668: SATA DOM needs firmware version upgrade to S170119. [1033] [A1183]

Name	SATA DOM 3ME Date and Firmware Status
Description	Checks the firmware version of SATA DOM.
Alert message	SATA DOM on host <code>ip_address</code> has old firmware version: <code>alert_msg</code>
Cause	SATA DOM firmware version is not the latest version.
Impact	Impacted by Shift read retry issue. The problem can cause node to fail to boot.



Resolution	Update the firmware to the latest version. Perform inventory from Life Cycle Manager to check and update to the latest firmware version. Check KB 6937 for alternative update option.

Table 669: SATA DOM has Guest VM. [1034] [A1184]

Name	SATA DOM Guest VM Check
Description	Checks that no guest VM is installed on SATA DOM.
Alert message	SATA DOM on host <code>ip_address</code> contains a guest VM, which will accelerate the degradation of the SATA DOM.
Cause	A guest VM is installed on SATA DOM.
Cause	A VM is incorrectly configured.
Impact	Degradation of the SATA DOM will be accelerated, leading to unavailability of the node.
Resolution	Remove guest VMs from SATA DOM.
Resolution	Reconfigure guest VMs on the host machine.

Table 670: SATADOM-SL 3IE3 has high wear. [1035] [A1185]

Name	SATADOM-SL 3IE3 Wearout Status
Description	Checks the wearout of SATADOM-SL 3IE3 via SMART data.
Alert message	SATADOM-SL 3IE3 on host <code>ip_address</code> has device life smaller than 5%
Cause	SATADOM-SL 3IE3 Device Life too short (<5).
Impact	Cluster performance may be significantly degraded or or the cluster may become unusable. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Replace SATADOM-SL 3IE3 as soon as possible. Refer to the Nutanix documentation for instructions.

Table 671: {model} firmware version is not the latest firmware version. [1055] [A1192]

Name	SATA DOM Firmware Status
Description	Checks the firmware version of SATA DOM.
Alert message	<code>model</code> on host <code>ip_address</code> has firmware version <code>fww</code>, need to upgrade to firmware version <code>Lfww</code>.
Cause	SATA DOM firmware version is not the latest version.
Impact	The node may fail to start.
Resolution	Update the firmware to the latest version.



Table 672: Samsung PM1633 drive has worn out. [1056] [A1193]

Name	Samsung PM1633 Wearout Status
Description	Checks the status of Samsung PM1633 drive via SMART data.
Alert message	Samsung PM1633 drive <code>dev_name</code> on host <code>ip_address</code> has the following problems: <code>err_msg</code>
Cause	Samsung PM1633 drive revision is wrong or Samsung PM1633 drive has been in use for a long time and close to the maximum write endurance.
Impact	Cluster performance may be significantly degraded or or the cluster may become unusable. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Replace Samsung PM1633 drive as soon as possible. Refer to the Nutanix documentation for instructions.

Table 673: Toshiba PM3 drive has worn out. [1057] [A1194]

Name	Toshiba PM3 Status
Description	Checks the status of Toshiba PM3 drive via SMART data.
Alert message	Toshiba PM3 drive <code>dev_name</code> on host <code>ip_address</code> has the following problems: <code>err_msg</code>
Cause	Toshiba PM3 drive revision is wrong or Toshiba PM3 drive has been in use for a long time and close to the maximum write endurance.
Impact	Cluster performance may be significantly degraded or or the cluster may become unusable. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Replace Toshiba PM3 drive as soon as possible. Refer to the Nutanix documentation for instructions.

Table 674: NVMe Drive has errors. [1059] [A1196]

Name	NVMe Status Check
Description	Check that host NVMe drive is functioning properly.
Alert message	NVMe Drive on host <code>ip_address</code> has errors: <code>alert_msg</code>
Cause	NVMe Drive is damaged or worn out.
Impact	Cluster performance may be degraded. The integrity of data from NVMe drive may be harmed.
Resolution	Replace the problematic NVMe drive as soon as possible.

Table 675: Toshiba PM4 drive has worn out. [1060] [A1197]

Name	Toshiba PM4 Status



Description	Check the status of Toshiba PM4 drive.
Alert message	Toshiba PM4 drive <code>dev_name</code> on host <code>ip_address</code> has the following problems: <code>err_msg</code>
Cause	Toshiba PM4 drive revision is unsupported or Toshiba PM4 drive has been in use for a long time and close to the maximum write endurance.
Impact	Cluster performance may be significantly degraded or or the cluster may become unusable. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Replace Toshiba PM4 drive as soon as possible. Refer to the Nutanix documentation for instructions.

Table 676: More than one type of Toshiba PM4 drives installed on the node. [1062] [A1199]

Name	Toshiba PM4 Config
Description	Check the Configuration of Toshiba PM4 drives.
Alert message	Toshiba PM4 drive <code>dev_name</code> on host <code>ip_address</code> has the following problems: <code>err_msg</code>
Cause	More than one type of Toshiba PM4 drives are installed on the node.
Impact	Toshiba PM4 drives may not function properly.
Resolution	Keep only one type of Toshiba PM4 drives.

Table 677: Toshiba PM4 drives not compatible with AOS or Foundation version on the node. [1063] [A1200]

Name	Toshiba PM4 Version Compatibility
Description	Check the compatibility of Toshiba PM4 drives with AOS and Foundation versions.
Alert message	Toshiba PM4 drive <code>dev_name</code> on host <code>ip_address</code> has the following problems: <code>err_msg</code>
Cause	Toshiba PM4 drives are not compatible with the current AOS or foundation version.
Impact	Toshiba PM4 drives may not function properly.
Resolution	Upgrade AOS or Foundation version.

Table 678: Firmware version of Toshiba PM4 drives is old. [1064] [A1201]

Name	Toshiba PM4 FW Version
Description	Check the firmware of Toshiba PM4 drives.
Alert message	Toshiba PM4 drive <code>dev_name</code> on host <code>ip_address</code> has the following problems: <code>err_msg</code>
Cause	Firmware version of Toshiba PM4 drives is not the latest version.



Impact	Toshiba PM4 drives may not function properly.
Resolution	Upgrade Toshiba PM4 firmware version.

Table 679: Samsung PM1633 drives not compatible with AOS or Foundation version on the node. [1065] [A1202]

Name	Samsung PM1633 Version Compatibility
Description	Check the compatibility of Samsung PM1633 drives with AOS and Foundation versions.
Alert message	Samsung PM1633 drive <code>dev_name</code> on host <code>ip_address</code> has the following problems: <code>err_msg</code>
Cause	Samsung PM1633 drives are not compatible with the current AOS or foundation version.
Impact	Samsung PM1633 drives may not function properly.
Resolution	Upgrade AOS or Foundation version.

Table 680: Firmware version of Samsung PM1633 drives is old. [1066] [A1203]

Name	Samsung PM1633 FW Version
Description	Check the firmware of Samsung PM1633 drives.
Alert message	Samsung PM1633 drive <code>dev_name</code> on host <code>ip_address</code> has the following problems: <code>err_msg</code>
Cause	Firmware version of Samsung PM1633 drives is not the latest version.
Impact	Samsung PM1633 drives may not function properly.
Resolution	Upgrade Samsung PM1633 firmware version. Perform inventory from Life Cycle Manager to check and update to the latest firmware version. Check KB 6937 for alternative update option.

Table 681: Latency between CVMs is high [6005] [A6005]

Name	Inter CVM Latency
Description	Checks whether latency among CVMs is high.
Alert message	Latency between <code>vm_types</code> is higher than <code>ping_threshold</code> ms.
Cause	Network configuration/connectivity may not be proper.
Impact	If problem persists I/O performance will be degraded.
Resolution	Check network connectivity/configuration.



Table 682: Wsman Connectivity lost [6210] [A1186]

Name	WSMan Connectivity Check
Description	Checks if CVM can connect to the host using WSMan.
Alert message	Wsman Remote Shell Connectivity lost between CVM and the host.
Cause	Password is not updated in the cluster configuration.
Impact	The CVM cannot connect to the host to manage it.
Resolution	Update the correct password in the cluster configuration.

Table 683: CVM memory reservation is incorrectly configured. [6211] [A6211]

Name	CVM Memory Pinned Check
Description	Verify CVM memory reservation and pinning.
Alert message	CVM memory reservation is incorrectly configured on CVM with IP <code>service_vm_external_ip</code>
Cause	CVM memory reservation is configured incorrectly.
Impact	This can potentially result in cluster instability.
Resolution	Please contact Nutanix Support to verify memory resolution.

Table 684: Host CPU Utilization [6409]

Name	Host CPU Utilization
Description	Checks wheather cpu contention is high on host.
Cause	High CPU utilization on host.
Impact	I/O performance may be degraded.
Resolution	Reduce CPU intensive VMs or reduce the total number of VMs.

Table 685: Host Swap Rate [6413]

Name	Host Swap Rate
Description	Checks whether swap rate is not high for host.
Cause	Swap rate is high. Host may not have enough memory.
Impact	VM performance is degraded.
Resolution	Increase memory on host, reduce memory intensive VMs, or reduce the total number of VMs.



Table 686: Metadata disk(s) not mounted on CVM [101055] [A101055]

Name	Metadata Disk Mounted Check
Description	Check that all metadata disks are mounted.
Alert message	Metadata disk(s) on Controller VM <code>svm_ip_address</code> are not mounted: <code>unmounted_disks</code>
Cause	Metadata disk(s) not mounted on CVM.
Impact	Cluster performance may be significantly degraded. If this condition continues, there is a potential for data corruption and/or loss.
Resolution	Contact Nutanix Support to mount the metadata disk.

Table 687: Hypervisor boot drive RAID is in an unhealthy state. [101058] [A101058]

Name	Hypervisor boot drive RAID status
Description	Checks the status of Dell Marvell BOSS hypervisor boot drive.
Alert message	Hypervisor boot drive RAID on <code>ip_address</code> is in <code>raid_state</code> state.
Cause	Hypervisor boot drive RAID is not functioning properly.
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Contact Support to help restore the hypervisor boot drive RAID.

Table 688: SATA DOM needs firmware upgrade. [101060] [A101060]

Name	SATA DOM 3IE3 Firmware Status
Description	Checks the firmware version of SATA DOM SL 3IE3.
Alert message	SATA DOM on host <code>ip_address</code> has old firmware version: <code>alert_msg</code>
Cause	SATA DOM SL 3IE3 firmware version is not the latest version.
Impact	Impacted by Shift read retry issue. The problem can cause node to fail to boot.
Resolution	Update the firmware to the latest version. Perform inventory from Life Cycle Manager to check and update to the latest firmware version. Check KB 6937 for alternative update option.

Table 689: Free block count of SATA DOM has gone below threshold. [101062] [A101062]

Name	SATA DOM 3IE3 Free Block Count Check
Description	Checks the free block count of SATA DOM SL 3IE3.
Alert message	<code>alert_msg</code>
Cause	SATA DOM SL 3IE3 free block count is less than 64.
Impact	As FBC approaches zero the host boot disk performance may be impacted.



Resolution	Ensure SATADOM firmware is up to date via LCM or KB6937. Refer to KB4395 for using TRIM script to proactively reclaim FBC.

Table 690: Toshiba PM5 drive has worn out. [101066] [A101066]

Name	Toshiba PM5 Status
Description	Check the status of Toshiba PM5 drive.
Alert message	Toshiba PM5 drive <code>dev_name</code> on host <code>ip_address</code> has the following problems: <code>err_msg</code>
Cause	Toshiba PM5 drive has been in use for a long time and close to the maximum write endurance.
Impact	Cluster performance may be significantly degraded or or the cluster may become unusable. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Replace Toshiba PM5 drive as soon as possible. Refer to the Nutanix documentation for instructions.

Table 691: More than one type of Toshiba PM5 drives installed on the node. [101067] [A101067]

Name	Toshiba PM5 Config
Description	Check the Configuration of Toshiba PM5 drives.
Alert message	Toshiba PM5 drive <code>dev_name</code> on host <code>ip_address</code> has the following problems: <code>err_msg</code>
Cause	More than one type of Toshiba PM5 drives are installed on the node.
Impact	Toshiba PM5 drives may not function properly.
Resolution	Keep only one type of Toshiba PM5 drives.

Table 692: CVM dvPortGroup Binding Check [103098]

Name	CVM dvPortGroup Binding Check
Description	Check for CVM dVportgroup binding
Cause	One (or more) Nutanix cluster Controller VM NIC is connected to a dVportgroup with ephemeral binding.
Impact	CVM may lose network connectivity on reboot
Resolution	Change the Dvportgroup binding to static

Table 693: Unsupported SFP is installed on host machine [103099] [A103099]

Name	Unsupported SFP message Check
Description	Check if an unsupported SFP is plugged in



Alert message	Unsupported SFP is installed on host machine <code>host_ip</code>: <code>alert_msg</code>
Cause	An unsupported SFP is plugged in.
Impact	NICs may not be recognized and network communication is affected.
Resolution	Remove the unsupported SFP

Table 694: SM863 drive has worn out. [106017] [A106017]

Name	Samsung SM863 SSD status check
Description	Check the status of SM863 and SM863a SSD drive
Alert message	<code>dev_type</code> drive <code>dev_name</code> on host <code>ip_address</code> has the following problems: <code>err_msg</code>
Cause	SM863a SSD drive revision is unsupported or SM863 and SM863a has been in use for a long time and close to maximum write endurance.
Impact	Cluster performance may be significantly degraded or or the cluster may become unusable. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Replace SM863 or SM863a SSD drive as soon as possible. Refer to the Nutanix documentation for instructions.

Table 695: CPUs of different types or models are in the same chassis. [106026] [A1190]

Name	CPU type on chassis check
Description	Checks whether CPUs within a chassis are of the same type.
Alert message	Chassis <code>chassis_id</code> has CPU configuration problem: <code>alert_msg</code>. This problem may affect the performance and stability of the nodes in the chassis. Chassis cpu info: <code>chassis_cpu_info</code>
Cause	A node of a different type was added to the chassis.
Impact	Having nodes of different CPU types in the same chassis is not supported.
Resolution	Remove the node that has a different CPU type from the chassis. Before moving a node to a chassis, make sure its CPU type is the same as that of other nodes in the chassis.

Table 696: Micron5100 drive has worn out. [106029] [A106029]

Name	Micron5100 SSD status check
Description	Check the status of Micron5100 SSD drive
Alert message	Micron5100 drive <code>dev_name</code> on host <code>ip_address</code> has the following problems: <code>err_msg</code>
Cause	Micron5100 SSD drive revision is unsupported or is close to maximum write endurance.



Impact	Cluster performance may be significantly degraded or or the cluster may become unusable. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Replace Micron5100 SSD drive as soon as possible. Refer to the Nutanix documentation for instructions.

Table 697: SM863a drives not compatible with AOS or Foundation version on the node.
[106030] [A106030]

Name	Samsung SM863a version compatibility check
Description	Check the compatibility of SM863a drives with AOS and Foundation versions.
Alert message	<code>dev_type</code> drive <code>dev_name</code> on host <code>ip_address</code> has the following problems: <code>err_msg</code>
Cause	SM863a drives are not compatible with the current AOS or foundation version.
Impact	SM863a drives may not function properly.
Resolution	Upgrade AOS or Foundation version.

Table 698: Firmware version of SM863 or SM863a drives is old. [106031] [A106031]

Name	Samsung SM863/SM863a FW version check
Description	Check the firmware of SM863 or SM863a drives.
Alert message	<code>dev_type</code> drive <code>dev_name</code> on host <code>ip_address</code> has the following problems: <code>err_msg</code>
Cause	Firmware version of SM863 or SM863a drives is not the latest one.
Impact	SM863 or SM863a drives may not function properly.
Resolution	Upgrade SM863 or SM863a firmware version. Perform inventory from Life Cycle Manager to check and update to the latest firmware version. Check KB 6937 for alternative update option.

Table 699: PM863a drives not compatible with AOS or Foundation version on the node.
[106033] [A106033]

Name	Samsung PM863a version compatibility check
Description	Check the compatibility of PM863a drives with AOS and Foundation versions.
Alert message	PM863a drive <code>dev_name</code> on host <code>ip_address</code> has the following problems: <code>err_msg</code>
Cause	PM863a drives are not compatible with the current AOS or foundation version.
Impact	Expand cluster and foundation workflows will not work.
Resolution	Upgrade AOS or Foundation version.



Table 700: Firmware version of PM863a drives is old. [106034] [A106034]

Name	Samsung PM863a FW version check
Description	Check firmware of PM863a drives.
Alert message	PM863a drive <code>dev_name</code> of model <code>dev_model</code> on host <code>ip_address</code> has the following problems: <code>err_msg</code>
Cause	Firmware version of PM863a drives is not the latest one.
Impact	PM863a drives may not function properly.
Resolution	Upgrade PM863a firmware version. Perform inventory from Life Cycle Manager to check and update to the latest firmware version. Check KB 6937 for alternative update option.

Table 701: PM863a drive has worn out. [106035] [A106035]

Name	Samsung PM863a status check
Description	Check the status of PM863a SSD
Alert message	PM863a drive <code>dev_name</code> of model <code>dev_model</code> , serial id <code>serial_id</code> , location <code>location</code> on host <code>ip_address</code> has the following problems: <code>err_msg</code>
Cause	PM863a is approaching maximum write endurance.
Impact	Cluster performance may be significantly degraded or the cluster may become unusable. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Replace PM863a SSD as soon as possible. Refer to the Nutanix documentation for instructions.

Table 702: SATADOM-ML 3SE has high wear. [106037] [A106037]

Name	SATADOM-ML 3SE Wearout Status
Description	Checks the wearout of SATADOM-ML 3SE via SMART data.
Alert message	SATADOM-ML 3SE on host <code>ip_address</code> has the following problem: <code>alert_msg</code>
Cause	SATADOM-ML 3SE Remaining Device Life too short.
Impact	Cluster performance may be significantly degraded or or the cluster may become unusable. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Replace SATADOM-ML 3SE as soon as possible. Refer to the Nutanix documentation for instructions.

Table 703: M.2 Micron5100 host boot drive has worn out. [106039] [A106039]

Name	M.2 Micron5100 host boot drive status check
Description	Check the status of M.2 Micron5100 host boot drive



Alert message	M.2 Micron5100 host boot drive <code>dev_name</code>, serial id <code>serial_id</code> on host <code>ip_address</code> has the following problems: <code>err_msg</code>
Cause	M.2 Micron5100 host boot drive is close to maximum write endurance.
Impact	Cluster performance may be significantly degraded or or the cluster may become unusable. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Replace M.2 Micron5100 host boot drive as soon as possible. Refer to the Nutanix documentation for instructions.

Table 704: M.2 Intel S3520 host boot drive has worn out. [106040] [A106040]

Name	M.2 Intel S3520 host boot drive status check
Description	Check the status of M.2 Intel S3520 host boot drive
Alert message	M.2 Intel S3520 host boot drive <code>dev_name</code> on host <code>ip_address</code> has the following problems: <code>err_msg</code>
Cause	M.2 Intel S3520 host boot drive is close to maximum write endurance.
Impact	Cluster performance may be significantly degraded or or the cluster may become unusable. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Replace M.2 Intel S3520 host boot drive as soon as possible. Refer to the Nutanix documentation for instructions.

Table 705: Host boot disk serial number has changed [106042] [A106042]

Name	M.2 Boot Disk change check
Description	Check the change of host boot disk
Alert message	Host boot disk on host <code>ip_address</code> has the following problems: <code>err_msg</code>
Cause	Active host boot disk failed and needs to switch to spare host boot disk to boot the host.
Impact	Host may be experiencing boot problems
Resolution	Replace the active host boot disk in future. Refer to the Nutanix documentation for instructions.

Table 706: Intel S4600 drives not compatible with AOS or Foundation version on the node. [106043] [A106043]

Name	Intel S4600 version compatibility check
Description	Check the compatibility of Intel S4600 drives with AOS and Foundation versions.
Alert message	Intel S4600 drive <code>dev_name</code> on host <code>ip_address</code> has the following problems: <code>err_msg</code>



Cause	Intel S4600 drives are not compatible with the current AOS or foundation version.
Impact	Expand cluster and foundation workflows will not work.
Resolution	Upgrade AOS or Foundation version.

Table 707: Firmware version of Intel S4600 drives is old. [106044] [A106044]

Name	Intel S4600 FW version check
Description	Check firmware of Intel S4600 drives.
Alert message	Intel S4600 drive <code>dev_name</code> of model <code>dev_model</code> on host <code>ip_address</code> has the following problems: <code>err_msg</code>
Cause	Firmware version of Intel S4600 drives is not the latest one.
Impact	Intel S4600 drives may not function properly.
Resolution	Upgrade Intel S4600 firmware version.

Table 708: Intel S4600 drive has worn out. [106045] [A106045]

Name	Intel S4600 status check
Description	Check the status of Intel S4600 SSD
Alert message	Intel S4600 drive <code>dev_name</code> of model <code>dev_model</code> on host <code>ip_address</code> has the following problems: <code>err_msg</code>
Cause	Intel S4600 is approaching maximum write endurance.
Impact	Cluster performance may be significantly degraded or the cluster may become unusable. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Replace Intel S4600 SSD as soon as possible. Refer to the Nutanix documentation for instructions.

Table 709: Firmware version of Intel S4600 drives is old or invalid. [106049] [A106049]

Name	Intel S4600 FW version check on Lenovo
Description	Check firmware of Intel S4600 drives on Lenovo platform.
Alert message	Intel S4600 drive <code>dev</code> of <code>model</code> on Lenovo host <code>ip_address</code> has the following problems: <code>err_msg</code>
Cause	Firmware version of Intel S4600 drives on Lenovo platform is an old one that has bugs.
Impact	Intel S4600 drives may not function properly.
Resolution	Upgrade Intel S4600 firmware version to at least version LU41 on Lenovo platform.



Table 710: PM883 drive has worn out. [106051] [A106051]

Name	Samsung PM883 status check
Description	Check the status of PM883 SSD
Alert message	PM883 drive <i>dev_name</i> of model <i>dev_model</i> , serial id <i>serial_id</i> , location <i>location</i> on host <i>ip_address</i> has the following problems: <i>err_msg</i>
Cause	PM883 is approaching maximum write endurance.
Impact	Cluster performance may be significantly degraded or the cluster may become unusable. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Replace PM883 SSD as soon as possible. Refer to the Nutanix documentation for instructions.

Table 711: SM883 drive has worn out. [106052] [A106052]

Name	Samsung SM883 SSD status check
Description	Check the status of SM883 SSD drive
Alert message	<i>dev_type</i> drive <i>dev_name</i> on host <i>ip_address</i> has the following problems: <i>err_msg</i>
Cause	SM883 has been in use for a long time and close to maximum write endurance.
Impact	Cluster performance may be significantly degraded or or the cluster may become unusable. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Replace SM883 SSD drive as soon as possible. Refer to the Nutanix documentation for instructions.

Table 712: Boot drive is in degraded state [106053] [A106053]

Name	Host Hardware RAID status Check
Description	Check status of hardware raid and alert user if raid is degeraded
Alert message	<i>alert_msg</i>
Cause	Boot drive mentioned in alert message may be offline and require attention
Impact	Node will operate without host boot drive redundancy but will be at risk of host unavailability if further hardware raid issues are seen
Resolution	Boot drive mentioned in alert message requires attention anf triage the root cause of it being unavailable

Table 713: Temperature of M2 disk is out of range [106054] [A106054]

Name	Host Hardware RAID M2 Temperature Check
Description	Checks temperature of each drive and alert if it is not within 0C and 75C



Alert message	<code>alert_msg</code>
Cause	Ambient temperature in the affected host environment may be too high for effective cooling, or airflow through host may be suboptimal
Impact	Continued operation above temperature thresholds may damage the host boot drives and lead to unplanned downtime
Resolution	Review the physical environment of the affected host and ensure air-conditioning systems are functioning and correctly calibrated and confirm the absence of any airflow obstructions around vents and the host itself

Table 714: M2 disk returned UECC errors [106055] [A106055]

Name	Host Hardware RAID M2 UECC Check
Description	Checks if disk has returned UECC errors
Alert message	<code>alert_msg</code>
Cause	Host boot disk has UECC errors
Impact	Continued operation with UECC errors, may lead to unplanned downtime
Resolution	Review the UECC errors of each host boot disk and triage accordingly

Table 715: M2 disk has worn out [106056] [A106056]

Name	Host Hardware RAID M2 wearout Check
Description	Checks wearout level of the disk.
Alert message	<code>alert_msg</code>
Cause	Disk is too old, or disk has gone through too many write cycles.
Impact	Slow Performance.
Resolution	Disk needs to be replaced.

Table 716: Firmware of raid M2 disk needs to be upgraded. [106057] [A106057]

Name	Host Hardware RAID M2 Firmware Check
Description	Checks if hardware raid M2 disk firmware needs to be upgraded
Alert message	<code>alert_msg</code>
Cause	Hardware raid drive is running outdated firmware and requires update
Impact	Host boot drive may not have up-to-date performance and stability features of more recent firmware
Resolution	Review KB6627 for instructions on upgrading the host boot drive firmware

Table 717: RAID card BIOS or Firmware or boot loader needs to be updated. [106058] [A106058]

Name	Host Hardware RAID Card Firmware Check
Description	Checks if RAID card BIOS, Firmware or boot loader needs to be updated.
Alert message	<i>alert_msg</i>
Cause	Hardware raid controller is running outdated BIOS or Firmware or Bootloader version. Please refer to the alert message
Impact	Hardware RAID controller may not have up-to-date performance and stability features of more recent firmware.
Resolution	Review KB6627 for instructions on upgrading BIOS or Firmware or Bootloader of hardware RAID controller

Table 718: SMART parameters of disk are out of range [106061] [A106061]

Name	SATA Disk Status Check
Description	Disk status check for SATA drives
Alert message	<i>alert_msg</i>
Cause	One or more S.M.A.R.T. parameters of the disk may be out of range
Impact	Cluster performance may be significantly degraded or node may become unavailable. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests. Available storage space is reduced If the drive gets marked offline.
Resolution	Based on which S.M.A.R.T. parameters is out of range, refer to the Nutanix documentation for instructions.

Table 719: SMART parameters of disk are out of range [106062] [A106062]

Name	SAS Disk Status Check
Description	Disk status check for SAS drives
Alert message	<i>alert_msg</i>
Cause	One or more S.M.A.R.T. parameters of the disk may be out of range
Impact	Cluster performance may be significantly degraded or node may become unavailable. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests. Available storage space is reduced If the drive gets marked offline.
Resolution	Based on which S.M.A.R.T. parameters is out of range, refer to the Nutanix documentation for instructions.

Table 720: SMART parameters of disk(host) are out of range [106063] [A106063]

Name	Host Boot Disk Status Check



Description	Disk status check for host boot disks
Alert message	<code>per_disk_msg</code>
Cause	One or more S.M.A.R.T. parameters of the disk may be out of range
Impact	Cluster performance may be significantly degraded or node may become unavailable. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests. Available storage space is reduced If the drive gets marked offline.
Resolution	Based on which S.M.A.R.T. parameters is out of range, refer to the Nutanix documentation for instructions.

Table 721: Disk firmware needs upgrade [106064] [A106064]

Name	Host Boot Disk Firmware Check
Description	Checks if any of the host boot disks, requires firmware upgrade or not
Alert message	<code>per_disk_msg</code>
Cause	Current host boot disk firmware is not recommended by Nutanix
Impact	Host boot Disk may not function properly, resulting in data unavailability.
Resolution	Upgrade host boot disk firmware version. Perform inventory from Life Cycle Manager to check and update to the latest firmware version. Check KB 6937 for alternative update option.

Table 722: BMC BIOS version check [111039]

Name	BMC BIOS version check
Description	Check BMC/Bios version of all nodes
Cause	Current BIOS/BMC version is not recommended for this node due to known issue
Impact	Current BIOS/BMC version may impact node behaviour, skew sensor data, or lack enhanced features and security fixes.
Resolution	Upgrade BMC and BIOS version. Perform inventory from Life Cycle Manager to check and update to the latest BMC/BIOS version. Check KB 2896 and 2905 for alternative update option.

Table 723: Mellanox NIC Driver version check [111045]

Name	Mellanox NIC Driver version check
Description	Check if Mellanox NIC driver at a recommended version
Cause	Mellanox NIC driver is not at a recommended version
Impact	Current driver version may not support all functionality. In some cases using a non-recommended driver version may lead to network and service instability.



Resolution	Check LCM for an updated Mellanox NIC driver or otherwise refer to KB4289 for further detail

Table 724: AHV Crash File Check [111053]

Name	AHV Crash File Check
Description	Check if /var/crash is empty.
Cause	Found kernel crash file(s) on AHV host(s).
Impact	Kernel issue may affect hypervisor functionalities.
Resolution	Notify Nutanix support to investigate the kernel issues.

Table 725: AHV Read-Only Filesystem Check [111055]

Name	AHV Read-Only Filesystem Check
Description	Check if any filesystem is read-only.
Cause	Found read-only filesystem(s) on AHV host(s).
Impact	The host may become unable to service I/O requests.
Resolution	Follow the instructions in the associated KB article to make the filesystems readable and writeable.

Table 726: AHV CVM Startup Dependency Check [111056]

Name	AHV CVM Startup Dependency Check
Description	Check if AHV CVM is configured to auto start.
Cause	Missing auto start file(s) on AHV host(s).
Impact	AHV CVM might not be started properly.
Resolution	Notify Nutanix support to investigate the issue.

Table 727: Detected incompatible AHV version [111061] [A111061]

Name	AHV Version Check
Description	Check if compatible AHV version is installed on host.
Alert message	Upgrade the version of AHV on the host to <code>desired_ver</code> or <code>later</code>.
Cause	The installed AHV version is not compatible with the current AOS version.
Impact	AHV is not running with the full set of features.
Resolution	Upgrade the version of AHV on the host to a version which is compatible with the current AOS version.



Table 728: AHV Time Zone Check [111071]

Name	AHV Time Zone Check
Description	Check whether the host's time zone is UTC, and warn that upgrading will change it.
Cause	The AHV host's time zone was changed/set to something other than UTC.
Impact	The host time zone will not affect the hypervisor appliance, but the time zone will be changed to UTC upon upgrade.
Resolution	Upgrading the host will resolve this notification, as will changing the host's time zone to UTC.

Table 729: Certificate Creation Error [130013] [A1112]

Name	Certificate Creation Status
Description	Certificate signing request creation failure
Alert message	Failed to create a certificate signing request on node <code>service_vm_external_ip</code> in the cluster.
Cause	Either the OpenSSL library had a failure or a non-supported subject field, subject value, or subject alternative name was provided.
Impact	Without a CSR then data at rest encryption cannot be enabled.
Resolution	Ensure that the certificate signing request fields are valid and retry the command, or contact Nutanix support if the OpenSSL configuration file is missing or corrupted.

Table 730: Agent VM Restoration Failure [130091] [A130091]

Name	Agent VM Restoration
Description	Failure to Restore Agent VM
Alert message	Failed to restore agent VM with UUID <code>vm_uuid</code>.
Cause	Not enough memory/CPU resources on this node.
Impact	VMs on this host may not function correctly because of the missing agent VM.
Resolution	Shut down unneeded VMs to free cluster resources or expand the cluster to add resources.

Table 731: VM group Snapshot and Current Mismatch [130110] [A130110]

Name	VM group Snapshot and Current Mismatch
Description	Check that snapshot to restore matches current VM group
Alert message	VM group <code>vm_group_uuid</code> Snapshot and Current Mismatch: <code>message</code>
Cause	VM group changed after snapshot was created.



Impact	VMs might be scheduled to run on same host instead of different host.
Resolution	Recreate VM snapshot to ensure up-to-date VM group information.

Table 732: Metro takeover - old primary site is hosting VM(s) [130111] [A130111]

Name	Metro Old Primary Site Hosting VMs
Description	Metro takeover - old primary site is hosting VMs
Alert message	Old primary site '<i>remote_name</i>' for protection domain '<i>protection_domain_name</i>' is hosting VM(s) while secondary is taking over.
Cause	The old primary site is hosting VMs while the secondary is taking over the metro protection domain.
Impact	VM performance might be degraded.
Resolution	Unregister the VMs hosted on the old primary and register them on the secondary site.

Table 733: Protection domain is in decoupled state [130112] [A130112]

Name	Protection Domain Decoupled Status
Description	Protection domain is in decoupled state
Alert message	Protection domain '<i>protection_domain_name</i>' protecting Storage Container '<i>container_name</i>' is in decoupled state.<i>impact</i>
Cause	Standby site has been promoted to primary
Impact	Metro availability for the protection domain on the primary is unavailable
Resolution	Reenable metro availability from the remote site

Table 734: Host Network Uplink Configuration Failed. [130160] [A130160]

Name	ClusterVswitch Configuration Update Failed
Description	Failed to apply host uplink configuration on the host(s).
Alert message	Host network uplink configuration update on Host(s):'<i>node_uuid</i>' failed due to the reason: <i>cause</i>
Cause	Host uplink configuration failed to apply on the host(s).
Impact	Host OVS configuration may not in sync with cluster vswitch configuration.
Resolution	Use manage_ovs to configure the host individually after correcting errors.

Table 735: Failed to configure host for Atlas networking [130201] [A130201]

Name	Failed to configure host for Atlas networking



Description	Failed to configure host for Atlas networking.
Alert message	Failed to configure host <code>hypervisor_address</code> for Atlas networking: <code>message</code>
Cause	Failed to configure host for Atlas networking.
Impact	VMs on Atlas networks cannot run on this host.
Resolution	Check the hypervisor version, and upgrade if necessary.

Table 736: Conntrack Stats Collector Check [200604]

Name	Conntrack Stats Collector Check
Description	Checks if Flow records are being transferred to Prism Central from hypervisor
Cause	Conntrack stats collector on node hypervisor is offline or unconfigured
Cause	Conntrack stats collector is blocked by firewall
Cause	Kafka is offline on Prism Central
Impact	Network Flow traffic for cluster node will not be visible in Prism Central
Resolution	Ensure Kafka is running on Prism Central.
Resolution	Restart Conntrack stats collector.
Resolution	Ensure that the firewall allows connection between registered Prism Central and AHV Hypervisor of the cluster.

Table 737: Flow policy-hit logging is not functional [200605] [A200605]

Name	Conntrack Mode Check
Description	Checks whether AHV is logging network traffic hitting Flow security policies.
Alert message	Network traffic hitting Flow security policies is not being logged.
Cause	Memory requirements may not be met for policy-hit logging to be enabled.
Cause	Prism Central, Prism Element & AHV versions may be incompatible.
Impact	Network traffic hitting security policies will not be logged when Flow is enabled on the cluster.
Resolution	Increase memory available on AHV and upgrade to the latest AHV version.
Resolution	Ensure AOS version is at least 5.11 on Prism Central and on Prism Element clusters managed by it.

Table 738: Openflow Table Getting Full [200607] [A200607]

Name	Openflow Table Full Check
Description	Checks that number of OVS rules is not reaching maximum number of rules
Alert message	Number of OVS rules is <code>num_flow_table_rules</code> for bridge <code>bridge_name</code> on AHV host <code>host_ip</code>. This is more than <code>usage_msg%</code> of maximum allowed number of rules.



Cause	Openflow table is getting full
Impact	When number of OVS rules will reach maximum, some VMs may lose network connectivity.
Resolution	Consider turning off unnecessary user VMs or remove unused vNICs from running user VMs on AHV host. For higher number of allowed OVS rules, ensure AOS version at least 5.10.5 and above on cluster.

Network

Table 739: CVM Connectivity Failure [3064] [A1001]

Name	CVM Passwordless Connectivity
Description	Check that CVM has passwordless connection to each other.
Alert message	Controller VM <code>target_ip</code> is not reachable from Controller VM <code>source_ip</code> in the last 3 attempts.
Cause	CVM has no passwordless connection to each other
Impact	Cluster compute and storage capacity are reduced. Until data stored on this host is replicated to other hosts in the cluster, the cluster has one less copy of guest VM data.
Resolution	Check CVM network configuration

Table 740: Host IP Not Reachable [3065] [A1006]

Name	Host IP Pingable
Description	Check that all host ips are pingable from local SVM.
Alert message	Hypervisor <code>target_ip</code> is not reachable from Controller VM <code>source_ip</code> in the last 6 attempts.
Cause	The hypervisor host is down or there is a network connectivity issue.
Impact	Cluster compute and storage capacity are reduced. Until data stored on this host is replicated to other hosts in the cluster, the cluster has one less copy of guest VM data.
Resolution	Ensure that the hypervisor host is running and that physical networking, VLANs, and virtual switches are configured correctly.

Table 741: NIC Link Down [3067] [A1082]

Name	NIC Link Status
Description	Checks whether any nic is down.
Alert message	Link on NIC <code>nic_name</code> of host <code>host_ip</code> is down. NIC description: <code>nic_description</code>
Cause	The NIC is disconnected from the switch, or the switch port is failing.



Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Ensure that the NIC is connected to the switch and that the switch port is functioning properly.

Table 742: Latency between CVMs is high [6005] [A6005]

Name	Inter CVM Latency
Description	Checks whether latency among CVMs is high.
Alert message	Latency between <code>vm_types</code> is higher than <code>ping_threshold</code> ms.
Cause	Network configuration/connectivity may not be proper.
Impact	If problem persists I/O performance will be degraded.
Resolution	Check network connectivity/configuration.

Table 743: CVM NIC Speed Low [6008] [A1058]

Name	TenGig Adaptor Status
Description	Checks whether CVM is uplinked to 10 GbE NIC.
Alert message	<code>alert_msg</code>
Cause	CVM not uplinked to 10 GbE NIC.
Cause	1 GbE NIC is part of the bond.
Cause	If hypervisor is ESXi, it may be in lockdown.
Impact	The Controller VM is not configured to use the 10 GbE NIC or is configured to share load with a slower NIC.
Resolution	Check network configuration.
Resolution	If hypervisor is ESXi, check for lockdown and unlock it.

Table 744: CVM Host Subnet Mismatch [6202] [A1048]

Name	CVM Subnet
Description	Checks that host and cvm share the same subnet.
Alert message	Controller VM <code>svm_ip</code> with network address <code>svm_subnet</code> is in a different network than the Hypervisor <code>hypervisor_ip</code>, which is in the network <code>hypervisor_subnet</code>.
Cause	The Controller VM and the hypervisor are not on the same subnet.
Impact	Controller VM high availability will not apply in the case of Controller VM failure, leading to guest VM unavailability.

Resolution	Reconfigure the cluster. Either move the Controller VMs to the same subnet as the hypervisor hosts or move the hypervisor hosts to the same subnet as the Controller VMs.

Table 745: Transmit Packet Drop Check [6404]

Name	Transmit Packet Drop Check
Description	Checks whether the host transmitted packets drop rate is high.
Cause	Transmitted packet drop rate is high.
Impact	Cluster performance may be degraded.
Resolution	Check network hardware/configuration.

Table 746: Receive Packet Drop Check [6405]

Name	Receive Packet Drop Check
Description	Checks whether the received packets drop rate is high.
Cause	Received packet drop rate is high.
Impact	Cluster performance may be degraded.
Resolution	Check network hardware/configuration.

Table 747: Unable to connect to vCenter. [6449] [A6449]

Name	vCenter Connection Check
Description	Check if vCenter connection is established.
Alert message	vCenter connection is not established. <code>cause_msg</code>, Check KB 3815 for detailed information.
Cause	vCenter server is unreachable on port 80.
Impact	Certain VM operations might fail.
Resolution	Confirm that vCenter server is reachable on port 80.

Table 748: CVM NIC Link Down [103094] [A103094]

Name	CVM NIC Link Status
Description	Checks whether any CVM NIC is down.
Alert message	Link on NIC <code>nic_name</code> of cvm <code>svm_ip</code> is down.
Cause	The rdma0 NIC is disconnected from the switch, or the switch port is failing.



Impact	Cluster performance may be significantly degraded. When rdma0 port is down, the CVM will switch to TCP mode via host vSwitch, and the RDMA mode of entire cluster is turned off. Critical large storage io service will suffer in non-RDMA mode.
Resolution	Ensure that the rdma0 NIC is connected to the switch and that the switch port is functioning properly.

Table 749: Network switch DCBX or NIC PFC is not configured correctly or enabled for RDMA NICs on CVM [103095] [A103095]

Name	RDMA network switch and NIC PFC config check
Description	Checks if network switch DCBX and NIC PFC are correctly configured and enabled.
Alert message	Network switch DCBX or NIC PFC is not configured correctly or enabled for RDMA NICs on CVM <code>cvm_ip: alert_msg</code>
Cause	Network switch DCBX is disabled or corresponding NIC PFC bit is not enabled.
Impact	RDMA traffic between CVMs is not enabled, which degrades cluster performance.
Resolution	Enable DCBX on network switch and configure PFC bit on RDMA NICs.

Table 750: Inconsistent Bridge/vSwitch configuration [103101] [A103101]

Name	NS Config Consistency Check
Description	Checks if the host bridge/vSwitch configuration matches NS config in zookeeper.
Alert message	Bridge/vSwitch <code>switch_name</code> on host <code>host_ip</code> has different configuration comparing to other hosts or Zeus configuration which can result in service degradation.
Cause	In case the customer modifies configuration during cluster lifetime, but does not restart genesis, this will result in inconsistent configuration.
Impact	Potentially will break configured services (ABS/DR).
Resolution	Please follow steps from KB 8018 to resolve this issue.

Table 751: Key Management Server Unavailable [110213] [A1099]

Name	Key Management Server Availability
Description	Check Key Management Server is available.
Alert message	Cannot connect to key management server <code>key_management_server_name</code> from Controller VM <code>service_vm_external_ip</code>.
Cause	The node cannot connect to a key management server either because of hardware failures, certificate validation failure, or network issues.



Impact	Secure clusters may not be able to unlock their drives on boot or change passwords.
Resolution	Ensure the key management server and nodes in a cluster have network connectivity to each other and all certificates are valid.

Table 752: Mellanox NIC Speed is not 10GbE and 40 GbE or both 10GbE and 40GbE NICs are installed on one node [111047] [A111047]

Name	Mellanox NIC Status check
Description	Checks if Mellanox NICs are down or if any Mellanox NIC has speed other than 10GbE or 40GbE. Checks if both 10GbE and 40GbE Mellanox NICs are installed on one node.
Alert message	Mellanox NIC on host <code>host_ip</code> has problem: <code>alert_msg</code>
Cause	The NIC is disconnected from the switch, or the switch port is not functioning correctly, or both 10GbE and 40GbE Mellanox NICs are installed on one node.
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Ensure that the NIC is connected to the switch and that the switch port is functioning properly. Ensure that only 10GbE or 40GbE Mellanox NIC is installed on one node.

Table 753: Mellanox NIC not Installed or with Wrong Type on host machine [111051] [A111051]

Name	Mellanox NIC Mixed Family check
Description	Checks if Mellanox NICs belong to the same Mellanox NIC family.
Alert message	Mellanox NIC on host <code>host_ip</code> has problem: <code>alert_msg</code>
Cause	Mellanox NICs are not from the same Mellanox NIC family.
Impact	Cluster performance may be significantly degraded. Mixing Mellanox NICs from different families may cause the NICs to not function properly.
Resolution	Replace some Mellanox NICs to make them come from same Mellanox NIC family.

Table 754: OVA Upload Interrupted [111076] [A111076]

Name	OVA Upload Interrupted
Description	OVA upload interrupted
Alert message	Upload interrupted for OVA <code>ova_name</code>
Cause	Network might have been disconnected.
Cause	Prism UI may have been closed.
Cause	OVA file might not be accessible.



Impact	OVA will not be available for use.
Resolution	Validate stable network connection.
Resolution	Ensure Prism UI is open and logged in.
Resolution	Ensure OVA file is accessible.

Table 755: Failed to configure host for Atlas networking [130201] [A130201]

Name	Failed to configure host for Atlas networking
Description	Failed to configure host for Atlas networking.
Alert message	Failed to configure host <i>hypervisor_address</i> for Atlas networking: <i>message</i>
Cause	Failed to configure host for Atlas networking.
Impact	VMs on Atlas networks cannot run on this host.
Resolution	Check the hypervisor version, and upgrade if necessary.

Table 756: Failed to reserve host memory for Atlas networking [130202] [A130202]

Name	Failed to reserve host memory for Atlas networking
Description	Failed to reserve host memory for Atlas networking.
Alert message	Failed to reserve host memory for Atlas networking
Cause	Failed to reserve host memory for Atlas networking.
Impact	If hypervisor memory is exhausted, random VMs may be powered off to reclaim memory.
Resolution	Power off VMs to reduce host memory consumption.

Table 757: Flow Rule Failed. [200601] [A200601]

Name	Flow rule failed
Description	Programming a Flow rule failed.
Alert message	Flow policy operation failed for <i>rule_name</i>. <i>reason</i>
Cause	Microseg service is unreachable or the PC is unavailable or the PE cluster is unavailable or Remote Connection from PC to PE failed or AHV host is unavailable or connection from PE to host failed. Check description.
Impact	VMs will not be protected by the Flow rule.
Resolution	Check PC microseg service, PE acropolis service, PC to PE connection, PE to AHV connection.



Table 758: Flow Control Plane Failed. [200602] [A200602]

Name	Flow control plane failed
Description	Flow failure event.
Alert message	Flow operation failed on <code>remote_uuid</code>. <code>reason</code>
Cause	Microseg service is unreachable or the PE cluster is unavailable or Remote Connection from PC to PE failed or AHV host is unavailable or connection from PE to host failed. Check description.
Impact	Flow control plane failure, rule cannot be programmed.
Resolution	Check PC microseg service, Check PE acropolis service, AHV microseg module.

Table 759: Flow Mode Change Failed [200606] [A200606]

Name	Flow policy hit mode change failed
Description	Flow control plane failure event
Alert message	Flow mode change failed on <code>remote_uuid</code>. <code>reason</code>
Cause	Flow mode configuration could not be programmed from PC to PE or PE to AHV. In Default mode network traffic hitting security policies won't be logged on AHV.
Impact	Flow running in default mode. Network traffic hitting security policies won't be logged by AHV.
Resolution	Ensure AHV host is reachable from Prism Central. Ensure all the services in Prism Central and Prism Element are in a healthy state and AHV host has more than 4 GB memory available. If the issue is still not resolved, collect <code>cvm_logs</code> and contact Nutanix Support

Table 760: Maximum VPN BGP route limit reached [801001] [A801001]

Name	Max VPN eBGP route limit reached
Description	On-prem VPN gateway is advertising more than the maximum number of routes accepted by the Xi VPN gateway over eBGP.
Alert message	A total of <code>num_routes</code> routes received from the on-prem VPN gateway. Only <code>max_num_routes</code> routes will be processed.
Cause	On-prem VPN gateway is advertising more than the maximum number of routes accepted by the Xi VPN gateway over eBGP.
Impact	VPN connectivity between Xi and the on-prem datacenter might be impacted.
Resolution	Consider aggregating more specific prefixes into larger prefixes to reduce the number of routes advertised.



Table 761: VPN IPSEC tunnel between on-prem and Xi datacenter is down [801002] [A801002]

Name	VPN IPSEC Tunnel down
Description	The IPSEC tunnel between the on-prem VPN gateway and Xi VPN gateway is down.
Alert message	VPN IPSEC tunnel between on-prem and Xi datacenter is down. Detailed error text: <code>error_detail</code>.
Cause	The IPSEC tunnel between the on-prem VPN gateway and Xi VPN gateway is down.
Impact	VPN connectivity between Xi and the on-prem datacenter will be impacted.
Resolution	Follow the on-prem VPN device vendor's troubleshooting steps. If you suspect that the problem is not with the on-prem VPN device, please contact Nutanix support.

Table 762: eBGP session between on-prem and Xi datacenter is down [801003] [A801003]

Name	eBGP session between VPN gateway peers down
Description	The eBGP session between the on-prem VPN gateway and Xi VPN gateway is down.
Alert message	eBGP session between on-prem and Xi datacenter is down. Detailed error text: <code>error_detail</code>.
Cause	The eBGP session between the on-prem VPN gateway and Xi VPN gateway is down.
Impact	VPN connectivity between Xi and the on-prem datacenter will be impacted because routes cannot be exchanged.
Resolution	Follow the on-prem VPN device vendor's troubleshooting steps. If you suspect that the problem is not with the on-prem VPN device, please contact Nutanix support.

Table 763: Epoch SaaS service connectivity lost [803002] [A803002]

Name	Epoch SaaS connectivity is lost
Description	Epoch SaaS service can't be reached.
Alert message	Epoch SaaS service connectivity has been lost. Details: <code>details</code>.
Cause	The network connectivity between Prism Central and Epoch SaaS service is not reliable.
Cause	Epoch SaaS service may be experiencing issues.
Cause	The saved access key or info for Epoch SaaS service may need to be updated.
Impact	The Prism Central Security Planning feature will not function properly.
Resolution	Check the connectivity between Prism Central and Epoch SaaS service.
Resolution	Check that the entered Epoch credentials in Prism Central are valid and up-to-date.



Resolution	Check that the Epoch SaaS service is accessible and running normally by logging onto the service.

Table 764: ID Firewall lost connectivity to domain controller [803003] [A803003]

Name	ID Firewall is unable to connect to one or more domain controllers
Description	ID Firewall is unable to connect to one or more domain controllers used for scraping identities.
Alert message	Domain controller <code>domain_controller</code> on domain <code>domain</code> is not reachable and accepting LDAP or WMI connections. Details: <code>details</code>.
Cause	The domain controller may not be reachable via the network.
Cause	The domain controller may not be accepting LDAP/WMI connections due to permissions issues.
Cause	The domain controller may not be running or fully booted up.
Impact	ID Firewall will not function properly and Flow VDI security policies may not be enforced until access to the domain controller is restored.
Resolution	Check network connectivity from Prism Central to the domain controller.
Resolution	Check that the domain controller is fully up and accepting WMI and LDAP connections.
Resolution	Check that the service account registered for ID Firewall has permissions for both WMI and LDAP access, and that the service account is active and that the correct password is entered in Prism Central.

Table 765: ID Firewall did not recover state after reconnecting to a temporarily unreachable domain controller [803005] [A803005]

Name	ID Firewall state recovery failed
Description	ID Firewall was unable to recover state after reconnecting to a temporarily unreachable domain controller.
Alert message	ID Firewall was unable to recover state from domain controller <code>domain_controller</code> on domain <code>domain</code>. Details: <code>details</code>.
Cause	Recovery failed because either too much time has passed since the domain controller was last reachable, or the event log on the domain controller has rolled over and the events necessary to recover are no longer present.
Impact	The state enforced by ID Firewall may be inconsistent, and therefore applied policies may not be properly enforced until action is taken.
Resolution	In order to ensure security policies are properly enforced, all active Nutanix VDI VM users should log out and log back into their VMs.

Nutanix Files



Table 766: File Server Space Usage High [160000] [A160000]

Name	File Server Space Usage
Description	File Server Space Usage High
Alert message	File Server space usage for <i>file_server_name</i> is at <i>usage_pct</i>%.
Cause	File Server Storage utilization is high.
Impact	If storage usage keeps growing, the shares will be marked as read-only.
Resolution	Expand File Server storage size, or ask users to delete unused files.

Table 767: File Server Space Usage Critical. [160001] [A160001]

Name	File Server Space Usage Critical
Description	File Server Space Usage Critical
Alert message	File Server space usage for <i>file_server_name</i> is at <i>usage_pct</i>%.
Cause	File Server storage utilization has reached a critical value.
Impact	The shares will be marked as read-only until there is more free space on the File Server.
Resolution	Expand the File Server storage size, or ask users to delete unused files.

Table 768: File Server Unreachable. [160002] [A160002]

Name	File Server Unavailable Check
Description	File server is unreachable.
Alert message	File server <i>file_server_name</i> is unreachable.
Cause	FSVMs are not reachable.
Impact	No operations can be executed against the File Server.
Resolution	Contact Nutanix support.

Table 769: File Server storage is not available. [160003] [A160003]

Name	File Server Storage Status
Description	File Server storage is not available.
Alert message	Storage for File Server <i>file_server_name</i> is unavailable.
Cause	File Server storage is unavailable due to network connectivity issues between FSVM and CVM.
Impact	Read and write operations on the File Server will fail.
Resolution	Contact Nutanix support.



Table 770: File Server scale-out failed. [160004] [A160004]

Name	File Server Scale-out Status
Description	File Server scale-out failed.
Alert message	File Server <i>file_server_name</i> scale-out failed because of <i>reason</i>
Cause	See details in the alert message.
Impact	Unable to scale-out the file server.
Resolution	Contact Nutanix support.

Table 771: File Server could not join the AD Domain [160005] [A160005]

Name	File Server Join Domain Status
Description	File Server could not join the AD Domain
Alert message	File Server <i>file_server_name</i> could not join the Domain because <i>reason</i>
Cause	See details in the alert message. Error could be due to one of the following reasons:AD is not reachable, or Domain credentials are invalid, or Entity with the specified File Server name already exists in the domain
Impact	No operations can be performed on the File Server as it has not joined the domain
Resolution	The resolution will depend on the cause, and could be either to check that the Domain is actually reachable or to provide the correct administrator credentials or to provide a different name for the File Server to avoid conflicts in the name.

Table 772: Node failed to join domain. [160006] [A160006]

Name	File Server Node Join Domain Status
Description	The node could not join the domain.
Alert message	The node could not join the domain for file server <i>file_server_name</i> as <i>reason</i>
Cause	AD is not reachable or Domain administrator credentials are invalid or a file server with the same name already exists in AD.
Impact	The node could not join the domain.
Resolution	Provide the correct domain administrator credentials or a different name for the file server.

Table 773: File Server Time Difference High [160007] [A160007]

Name	FSVM Time Drift Status
Description	The time drift between the FSVMs is beyond the acceptable range



Alert message	For <i>file_server_name</i> , time drift between the FSVMs (<i>lower_time_ip</i> and <i>higher_time_ip</i>) is more than the acceptable value of <i>time_difference_limit_secs</i>
Cause	NTP is not configured correctly or the NTP service is not functioning.
Impact	The Nutanix Files cluster may become inaccessible
Resolution	Check that the NTP service is running and is reachable from the FSVMs

Table 774: File Server storage cleanup failed [160012] [A160012]

Name	File Server Storage Cleanup Failure
Description	Failed to clean up storage for the File Server
Alert message	Storage <i>message</i> for the File Server <i>file_server_name</i> is not cleaned up
Cause	Acropolis service may be down on the cluster
Impact	File server storage will not be released
Resolution	Delete the volume groups associated with the File Server using 'acli'

Table 775: File server cannot connect with AD server [160013] [A160013]

Name	File Server AD Connectivity Failure
Description	File server cannot connect with AD server with configured information
Alert message	File server <i>file_server_name</i> cannot connect with AD server
Cause	The machine account credentials may have been changed, or there is a network connectivity issue for the AD server.
Impact	File server is in non-operational state. Client connectivity with file-server is disabled.
Resolution	Check that the AD server is reachable.
Resolution	Try leaving and joining the domain again.

Table 776: File Server performance optimization recommended [160015] [A160015]

Name	File Server Performance Optimization Recommended
Description	File server has a recommendation to optimize performance by using scale-up, scale-out or rebalance.
Alert message	A recommendation is available to optimize the performance on one or more nodes of file server <i>file_server_name</i>.
Cause	File server has one or more nodes under extensive load.
Impact	File server performance may be impacted and new client connections may be refused.
Resolution	Run 'Performance optimization' for the specified file server.



Table 777: User Quota Assignment Failed [160016] [A160016]

Name	File Server Quota allocation failed for user
Description	Failed to assign the specified quota to the user
Alert message	Failed to apply quota for user: <i>user_name</i> on share: <i>share_name</i> of File Server <i>file_server_name</i>
Cause	The user has used more space on the File Server than the specified quota value.
Impact	User has no quota limit
Resolution	Notify user to reduce their usage below the quota value, or increase the quota allocation for the user.

Table 778: Share utilization reached configured limit [160017] [A160017]

Name	Share Utilization Reached Configured Limit
Description	Share is no longer writeable.
Alert message	Utilization on <i>share_name</i> on file server <i>file_server_name</i> has reached the configured limit
Cause	Share utilization reached its configured limit
Impact	Writes to the share will fail until one of the actions is taken.
Resolution	Ask users to free up space or increase storage allocation to the share.

Table 779: File Server failed to get updated CVM IP address. [160018] [A160018]

Name	File Server CVM IP update failed
Description	File server cvm ip update failed.
Alert message	The file server <i>file_server_name</i> could not get updated CVM IP address.
Cause	Failed to contact file server .
Impact	File server is not reachable
Resolution	Contact Nutanix support.

Table 780: Appropriate Site Not Found in Active Directory [160019] [A160019]

Name	File Server Site Not Found
Description	Unable to determine an appropriate site in Active Directory
Alert message	File server <i>file_server_name</i>'s client network is not mapped to a site on the Active Directory.
Cause	File server client network is not part of any site in Active Directory.
Impact	File server may take a long time to join the domain, depending on which domain controller is selected.



Resolution	Add the client network of file server to a local site in Active Directory for optimal performance.

Table 781: File server DNS Updates Pending [160020] [A160020]

Name	File Server DNS Updates Pending
Description	DNS updates are pending after a file server operation.
Alert message	File server <i>file_server_name</i>'s DNS entries need to be programmed on nameservers. <i>message</i>
Cause	Failed to add or remove DNS entries
Impact	The file server cannot be used by clients by its name, but can be used with the IP addresses.
Resolution	Correct the required DNS entries under 'DNS' on the file server page.

Table 782: File Server activation failed [160021] [A160021]

Name	File Server Activation Failed
Description	File Server Activation Failed
Alert message	File server <i>file_server_name</i> activation failed due to <i>reason</i>
Cause	Check alert message for details
Impact	File Server is not usable.
Resolution	Try to activate file server again, and if the failure persists, then contact Nutanix support.

Table 783: Failed to set VM-to-VM anti-affinity rule [160023] [A160023]

Name	Failed To Set VM-to-VM Anti Affinity Rule
Description	Failed to set VM-to-VM anti affinity rule
Alert message	Failed to set VM-to-VM anti-affinity rule for <i>file_server_name</i>. <i>reason_and_resolution_msg</i>
Cause	Contact Nutanix support.
Impact	More than one FSVM may be deployed on a single physical host, which will cause data unavailability in case that host fails.
Resolution	Various. Check alert message for details.

Table 784: File server home share creation failed [160024] [A160024]

Name	File Server Home Share Creation Failed
Description	Failed to create home share during file server creation.



Alert message	Failed to create Home share for <code>file_server_name</code> because <code>message</code>
Cause	Check alert message for details
Impact	The default home share is not created.
Resolution	Check alert message to understand cause of failure and take action accordingly.

Table 785: Discovery of iSCSI targets failed. [160025] [A160025]

Name	File Server Iscsi Discovery Failure
Description	Failed to discover iSCSI targets on the CVM during the discovery process.
Alert message	Discovery of iSCSI targets failed for file server <code>file_server_name</code>
Cause	The 'external_data_services_ip' or the CVM IP addresses are not reachable from the FSVM.
Impact	Share operations and File server HA will fail.
Resolution	Check the configured IP addresses and ensure connectivity from the FSVM's Storage Network to the external_data_services_ip and the CVM IPs.

Table 786: File Server upgrade failed [160026] [A160026]

Name	File Server Upgrade Failed
Description	File Server Upgrade Failed.
Alert message	Upgrade of file server <code>file_server_name</code> failed due to <code>reason</code>.
Cause	Check alert message for details
Impact	File server has not been upgraded.
Resolution	Try again. If the failure persists, contact Nutanix support

Table 787: Incompatible File Server activated [160028] [A160028]

Name	Incompatible File Server Activation
Description	Activating an incompatible File Server
Alert message	Incompatible File server <code>file_server_name</code> is activated
Cause	File server is activated which is incompatible with current AOS version
Impact	Next AOS upgrade will fail.
Resolution	Upgrade file server

Table 788: File Server in heterogeneous state. [160032] [A160032]

Name	File Server In Heterogeneous State



Description	File server in heterogeneous state. Nodes do not match in their CPU or memory configuration.
Alert message	File server <code>file_server_name</code> is in heterogeneous state, as File Server VM configuration update operation failed due to <code>reason</code>
Cause	Various. Check alert message for details.
Impact	Performance of the file server is less than optimal.
Resolution	Get the largest configured CPU and memory values among VMs of reported File Server in VM page. Free up resources on hypervisor hosts and retry File Server VM configuration update with these values in File Server page.
Resolution	Contact Nutanix support

Table 789: Failed to correct File Server data and meta data inconsistencies [160034] [A160034]

Name	Failed To Run File Server Metadata Fixer Successfully
Description	Failed to Run File Server Metadata Fixer tool successfully
Alert message	Failed to run File Server <code>file_server_name</code> Metadata fixer task successfully. <code>reason_and_resolution_msg</code>
Cause	Metada data fixer task might timedout
Impact	Some of the top level directory might not shown on file server
Resolution	Rerun the Metadata fixer scli/afs command on File server. Usage: 'scli meta-data-check' or 'afs share.fix_tld_metadata'

Table 790: File server share deletion failed [160035] [A160035]

Name	File Server Share Deletion Failed
Description	Failed to delete share.
Alert message	Failed to delete share <code>share_name</code> on File Server <code>file_server_name</code>. <code>message</code>
Cause	Multiple causes possible. Most likely, a failure in destroying file systems from FSVMs.
Impact	Contents of share may not have been fully deleted after delete request.
Resolution	Please retry the operation. If operation still fails, contact Nutanix Support.

Table 791: File server compatibility check skipped [160036] [A160036]

Name	Skipped File Server Compatibility Check
Description	File server compatibility check skipped
Alert message	File server compatibility check skipped
Cause	File server compatibility check was skipped during AOS upgrade
Impact	Some management functions of the file server may not be available.



Resolution	Upgrade file servers that reside on the cluster.

Table 792: Failed to add one or more file server administrator users or groups [160038] [A160038]

Name	Failed to add one or more file server admin users or groups
Description	Failed to add one or more users or groups as file server administrators
Alert message	File server <i>file_server_name</i>: <i>failure_msg</i> Update file server administrators on the file server page
Cause	One or more users or groups could not be resolved on Active Directory.
Impact	File server administrator users or groups may not function as administrators
Resolution	Update file server administrators on the file server page

Table 793: Maximum connections limit about to reach on a file server VM [160039] [A160039]

Name	Maximum connections limit reached on a file server VM
Description	Maximum connections limit is about to reach on a file server VM
Alert message	File server VM <i>fsvm_name</i> on file server <i>file_server_name</i> has reached <i>conn_percentage</i>connections limit.
Cause	File server has one or more nodes under extensive load.
Impact	File server performance may be impacted and new client connections may be refused.
Resolution	Run 'Performance optimization' for the specified file server.

Table 794: File Server Clone failed [160040] [A160040]

Name	File Server Clone Failed
Description	File Server Clone Failed
Alert message	File server <i>file_server_name</i> clone failed due to <i>reason</i>
Cause	Check alert message for details
Impact	Clone File Server not available.
Resolution	Check alert message, Retry clone operation after rectifying, and if the failure persists, then contact Nutanix support.

Table 795: File Server rename failed [160041] [A160041]

Name	File Server Rename Failed
Description	File Server Rename Failed



Alert message	File server rename from <i>file_server_name</i> to <i>file_server_new_name</i> failed due to <i>reason</i>
Cause	Check alert message for details
Impact	Rename File Server not available.
Resolution	Check alert message, Retry rename operation after rectifying, and if the failure persists, then contact Nutanix support.

Table 796: File Server AntiVirus - ICAP Server Down [160042] [A160042]

Name	File Server AntiVirus - ICAP Server Down
Description	ICAP server is not responding to scan requests.
Alert message	ICAP server <i>scan_server_host</i> is not responding from file server <i>file_server_name</i> .
Cause	Failed to reach ICAP server from FSVM.
Impact	File server antivirus scanning is degraded.
Resolution	Check ICAP server and connectivity between FSVMs to ICAP server.

Table 797: File Server AntiVirus - All ICAP Servers Down [160043] [A160043]

Name	File Server AntiVirus - All ICAP Servers Down
Description	All configured ICAP servers are not responding to scan requests.
Alert message	All configured ICAP servers are not responding from file server <i>file_server_name</i> .
Cause	Failed to reach all configured ICAP server from FSVM.
Impact	File server antivirus scanning is not possible. Access is decided by scan policy.
Resolution	Check ICAP server and connectivity between FSVMs to ICAP server.

Table 798: File Server network change failed [160044] [A160044]

Name	File Server Network Change Failed
Description	File Server Network Change Failed
Alert message	File server <i>file_server_name</i> network change failed due to <i>reason</i>
Cause	Check alert message for details
Impact	File Server is not usable.
Resolution	Try to change file server network again, and if the failure persists, then contact Nutanix support.

Table 799: File Server AntiVirus - Scan Queue Full on FSVM [160045] [A160045]

Name	File Server AntiVirus - Scan Queue Full on FSVM
Description	Antivirus Scan Queue is Full on FSVM.
Alert message	Antivirus Scan Queue is Full on FSVM <i>nvm_name</i> of file server <i>file_server_name</i>.
Cause	Icap Daemon is not able to keep up with in coming requests.
Impact	File server antivirus scanning is slow, clients may observe high latency.
Resolution	Check ICAP server status. Check number of ICAP servers configured. Check workload on FSVM.

Table 800: File Server AntiVirus - Scan Queue Piling Up on FSVM [160046] [A160046]

Name	File Server AntiVirus - Scan Queue Piling Up on FSVM
Description	Antivirus Scan Queue is Piling Up on FSVM.
Alert message	Antivirus Scan Queue is Piling Up on FSVM <i>nvm_name</i> of file server <i>file_server_name</i>.
Cause	Incoming scan request rate is higher than the scan processing rate.
Impact	File server antivirus scanning is slow, clients may observe high latency.
Resolution	Check ICAP server status. Check number of ICAP servers configured. Check workload on FSVM.

Table 801: File Server AntiVirus - Excessive Quarantined / Unquarantined Files [160047] [A160047]

Name	File Server AntiVirus - Excessive Quarantined / Unquarantined Files
Description	Excessive Quarantined / Unquarantined Files on File Server.
Alert message	File Server <i>file_server_name</i> has a lot of quarantined / unquarantined files.
Cause	Number of quarantined / unquarantined files on file server is approaching limit.
Impact	File server will stop scanning the files once the limit is reached. Default action specified in the scan policy will be applied.
Resolution	Delete unwanted quarantined / unquarantined files, check icap server for false positives or check system for virus activity.

Table 802: File Server AntiVirus - Quarantined / Unquarantined Files Limit Reached [160048] [A160048]

Name	File Server AntiVirus - Quarantined / Unquarantined Files Limit Reached
Description	Quarantined / Unquarantined Files Limit is Reached on File Server.



Alert message	File Server <i>file_server_name</i> has reached quarantined / unquarantined files limit.
Cause	Number of quarantined / unquarantined files on file server reached limit.
Impact	File server stopped scanning files. Default action in the scan policy is applied.
Resolution	Delete unwanted quarantined / unquarantined files, check icap server for false positives or check system for virus activity.

Table 803: File Server time is out of sync with the Active Directory [160049] [A160049]

Name	File Server time is out of sync with the Active Directory
Description	File server time is out of sync with the Active Directory domain controllers.
Alert message	Time drift between the file server VMs and the Active Directory is at <i>time_drift_secs</i> seconds on file server <i>file_server_name</i>. If the drift becomes more than 300 seconds, the clients may get disconnected from the fileserver.
Cause	NTP is not configured correctly or the NTP service is not functioning.
Impact	The Nutanix Files cluster may become inaccessible.
Resolution	Check that the NTP service is running and the NTP servers are reachable from the file server VMs.

Table 804: File Server User Management Configuration Failed [160050] [A160050]

Name	File Server User Management Configuration Failed
Description	Given user management options could not be configured for the file server.
Alert message	The values provided for File server <i>file_server_name</i>'s user management options are incorrect.
Cause	Failed to configure user management with given options
Impact	The file server may have restricted or insecure access from various clients.
Resolution	Please check and try again with correct entries for various fields in the user management tab.

Table 805: /home partition usage on a file server VM higher than threshold [160051] [A160051]

Name	/home partition usage on a file server VM higher than threshold
Description	Space consumption under /home partition on file server VM is higher than the set threshold
Alert message	/home partition usage on File server VM <i>fsvm_name</i> on file server <i>file_server_name</i> at <i>usage_pct</i> (greater than threshold <i>threshold_pct</i>).
Cause	File server has one or more nodes under extensive load.
Impact	Deleting a share from a File server fails.



Resolution	Remove unwanted data or files from under /home for the specified file server.

Table 806: File server DNS records cannot be refreshed [160052] [A160052]

Name	File server DNS records refresh failed
Description	File server DNS records periodic refresh failed
Alert message	File server <i>file_server_name</i> DNS records cannot be updated
Cause	DNS entries may not be created with Nutanix Files computer account due to insufficient permissions.
Impact	If scavenging is active on the DNS servers, then Nutanix Files DNS records may get removed. This can cause clients to not connect with Nutanix Files using its DNS hostname.
Resolution	Create static DNS entries for the fileserver manually using the 'DNS' button or update required permissions for the Nutanix Files computer account.

Table 807: File Server Share Backup diff path translation failed [160053] [A160053]

Name	File Server Share Backup diff path translation failed
Description	Backup diff found inodes whose path translation failed
Alert message	Backup diff found inodes whose path translation failed <i>file_server_name</i>, error information: <i>failure_msg</i>
Cause	File server share may have hardlinks.
Impact	Some files may not have been backed up.
Resolution	Contact Nutanix Support

Table 808: File Server Partner Server Connectivity Down [160054] [A160054]

Name	File Server Partner Server Connectivity Down
Description	Partner server is not responding to file notifications.
Alert message	Partner server <i>partner_server_host</i> is not responding from file server <i>file_server_name</i>.
Cause	Failed to reach partner server from File Server VM.
Impact	File server stopped notifying file operation events.
Resolution	Ensure that the partner server is functioning and that there is connectivity between File Server VMs and the partner server.

Table 809: File Server DNS resolver IP connectivity Failure [160055] [A160055]

Name	File Server DNS resolver IP connectivity Failure



Description	File server cannot resolve its domain name using configured DNS resolver IP addresses.
Alert message	File server <code>file_server_name</code> cannot resolve its domain name using configured DNS resolver IP addresses. <code>message</code>
Cause	None of the configured DNS resolver IP addresses are reachable.
Cause	None of the configured DNS resolver IP addresses are resolving the file server domain name.
Impact	File server services may be unavailable and new clients may not be authenticated.
Resolution	Ensure that the servers specified by DNS resolver IP addresses are functional.
Resolution	Change the file server configuration to properly functioning DNS resolver IP addresses.

Table 810: File Server PD action to incompatible Remote Site AOS [160057] [A160057]

Name	File Server PD action warning for non compatible Remote Site AOS
Description	Remote Site AOS version is not compatible with File Server version.
Alert message	File server <code>file_server_name</code> protection domain <code>protection_domain_name</code> enabled on non-compatible Remote site <code>remote_site_name</code>
Cause	Remote Site AOS version is not compatible with File Server version.
Impact	Migrate-to or Activate-on Remote Site would fail.
Resolution	Upgrade Remote Site AOS to 5.6 or above.

Table 811: File Server NTP servers connectivity failure [160058] [A160058]

Name	File Server NTP servers connectivity failure
Description	File server cannot connect to NTP server.
Alert message	File Server <code>file_server_name</code> cannot reach NTP servers specified. <code>reason</code>
Cause	None of the configured NTP servers are reachable.
Impact	File server services may become inaccessible.
Resolution	Ensure that the NTP servers are functional and are reachable from the file server VMs.
Resolution	Update NTP servers using file server 'update'.

Table 812: File server services got interrupted [160059] [A160059]

Name	File server services got interrupted
Description	One of the fileserver services crashed.
Alert message	File server <code>file_server_name</code> services got interrupted



Cause	One of the fileserver services crashed.
Impact	Some fileserver service would be unavailable to NAS clients momentarily.
Resolution	Look at /home/nutanix/minerva_ha.log for details on crashed services.If the service continues to be down, consider restarting minerva_ha service.

Table 813: Share usage reaching to configured limit [160061] [A160061]

Name	Share Space Usage Reaching Configured Limit
Description	Share will be no longer writeable.
Alert message	Share <code>share_name</code> on file server <code>file_server_name</code> is approaching capacity
Cause	Share usage reaching its configured limit
Impact	Writes to the share will fail if share gets full.
Resolution	Increase storage allocation to the share or direct users to free up space by removing unneeded files.

Table 814: Pulse cannot connect to REST server endpoint on File Server [160062] [A160062]

Name	File Server Pulse REST endpoint connection status
Description	Check if Pulse can connect to REST server endpoint on File Server.
Alert message	Pulse cannot connect to REST server endpoint on File Server <code>file_server_name</code>. Connection Status: <code>connection_status</code>, Pulse Enabled: <code>enabled</code>, Error Message: <code>message</code>
Cause	Pulse cannot connect to REST server endpoint on File Server.
Impact	Data driven serviceability and customer support cannot be performed for affected File Server.
Resolution	Ensure that the REST server endpoint is reachable from Pulse on File Server.

Table 815: Storage pool usage reaching its limit [160063] [A160063]

Name	Storage Pool Usage Is Reaching Its Limit
Description	Shares on storage pool will be no longer writeable.
Alert message	<code>share_names</code> shares on file server <code>file_server_name</code> are about to run out of space. <code>command</code>
Cause	Storage pool usage is reaching its limit
Impact	Writes to the shares may fail if their storage pool gets full.
Resolution	Clean up space by removing unnecessary files in the listed shares to ensure there is no interruption creating new data.



Table 816: Duplicate IP address detected for a File Server VM. [160068] [A160068]

Name	File Server Duplicate IP Address Detected
Description	Check if any file server VMs have IP address conflict.
Alert message	Duplicate IP address detected for a file server VMs for <code>file_server_name</code>. Error Message: <code>message</code>
Cause	Another host has started using a file server VM IP address.
Impact	Some file server operations may fail due to invalid network configuration.
Resolution	Update IP address on host which has an IP address conflict with file server VM.

Table 817: File Server Unique Fsid failure [160069] [A160069]

Name	File Server Unique Fsid Failure
Description	Assigning Unique Fsid failed
Alert message	Found unique fsid failure <code>file_server_name</code>, error information: <code>failure_msg</code>
Cause	Duplicate Unique Fsid present
Impact	NFS handles may not be unique.
Resolution	Contact Nutanix Support

Table 818: File Server VM Time Drift from NTP Servers [160072] [A160072]

Name	File Server VM Time Drift from NTP Servers
Description	Difference between FSVM clock and NTP servers is too large
Alert message	File Server VM <code>fsvm_name</code> clock is different from NTP servers for <code>time_diff_secs</code> seconds. Acceptable time difference <code>time_diff_limit_secs</code> seconds.
Cause	NTP servers are not publishing time correctly.
Cause	File server VM time is not syncing with NTP servers.
Impact	File server cluster may become inaccessible
Resolution	Verify NTP servers are functioning

Table 819: File Server Service in Crash Loop [160074] [A160074]

Name	File server service in crash loop
Description	File server service in crash loop
Alert message	File server service <code>crashed_service_name</code> has crashed <code>crash_count</code> times in <code>crash_interval_secs</code> seconds.
Cause	File server service has crashed multiple times in short interval



Impact	Some file server features may not be functioning
Resolution	Contact Nutanix Support

Other

Table 820: CVM renamed incorrectly. [3035] [A3035]

Name	CVM name check
Description	Check that CVM has been named properly.
Alert message	CVM renamed incorrectly.
Cause	CVM was renamed after installation
Impact	Cluster upgrade will get stuck.
Resolution	Rename CVM back to its original name. CVM name is nonconfigurable by the end user.

Table 821: Orphan VM Snapshot Check [3200]

Name	Orphan VM Snapshot Check
Description	Checks for VM snapshots whose associated VMs are removed.
Cause	The VM associated with a snapshot has been removed.
Impact	Cluster storage resources are consumed by unneeded data.
Resolution	Delete the snapshots based on the VM that has been removed.

Table 822: Network adapter setting check [103068]

Name	Network adapter setting check
Description	Check Network adapter setting
Cause	Network adapter is not set to not connect on power on
Impact	CVM may not boot up after the host reboot.
Resolution	Check if "ethernet0.startConnected = "true"" is present in the CVM's .vmx file.

Table 823: PYNFS dependency check [106431]

Name	PYNFS dependency check
Description	Check PYNFS dependency
Cause	PYNFS is in use and is not present.
Impact	CVM may not boot up after the host reboot.
Resolution	Validate PYNFS configuration.



Table 824: localcli check [106432]

Name	localcli check
Description	Check that esxcli command not used in rc local script
Cause	esxcli command is used in rc local script.
Impact	CVM may not boot up after the host reboot.
Resolution	Check if 'esxcli' is used instead of 'localcli' in /etc/rc.local.d/local.sh.

Table 825: vim command check [106433]

Name	vim command check
Description	Check that vim-cmd vmshvc/power.on command is present in local script
Cause	vim-cmd vmshvc/power.on\ command not present in local script.
Impact	CVM may not boot up after the host reboot.
Resolution	Check if 'vim-cmd vmshvc/power.on' entry is present in 'local.sh'.

Table 826: Autobackup check [106436]

Name	Autobackup check
Description	Check that /sbin/auto-backup.sh has run successfully
Cause	/sbin/auto-backup.sh has not run successfully.
Impact	CVM may not boot up after the host reboot.
Resolution	Make sure '/bootbank/state.tgz' has a newer timestamp.

Table 827: EOF check [106437]

Name	EOF check
Description	Check that there is no line "EOF" at end of rc local script
Cause	"EOF" statement at end of rc local.
Impact	CVM may not boot up after the host reboot.
Resolution	Check that 'local.sh' does not have 'EOF'

Table 828: RC local script exit statement present [106438]

Name	RC local script exit statement present
Description	Check that top level exit statement is not present in RC local script
Cause	Top level exit statement present in script rc local preventing script lines from being run.



Impact	CVM may not boot up after the host reboot.
Resolution	Check if 'local.sh' has an 'exit' statement. Generate INFO if the exit statement is NOT within the if..fi statement

Table 829: .dvsData directory in local datastore [106439]

Name	.dvsData directory in local datastore
Description	Check that .dvsData directory is present on pynfs mounted local datastore
Cause	.dvsData directory is not persistent yet.
Impact	CVM may not boot up after the host reboot.
Resolution	Check if .dvsData directory exists in the local datastore.

Table 830: Svmboot mount check [106440]

Name	Svmboot mount check
Description	Checks if CVM ISO is on mounted local datastore
Cause	ServiceVM_Centos.iso is not on mounted local datastore.
Impact	CVM may not start after the host reboot.
Resolution	Check if ServiceVM_Centos.iso exists in the local datastore.

Table 831: ESXi vSphere Cluster config check [106448]

Name	ESXi vSphere Cluster config check
Description	Check if the nodes in a Nutanix cluster are configured correctly on a vSphere Cluster
Cause	One or more ESXi hosts which are part of the Nutanix cluster may be disconnected from the registered vCenter Server, not part of inventory, or registered across multiple vSphere Clusters
Impact	Prism UVM management operations, 1-click hypervisor upgrade, DRS, HA and other vSphere automation may be hampered.
Resolution	Review KB 5369

Table 832: XenServer OVS Bridge Check [106515]

Name	XenServer OVS Bridge Check
Description	Check if XenServer Hosts are using an OpenVSwitch bridge
Cause	XenServer host is configured to use a Linux bridge
Impact	Using the Linux bridge as the network backend is not supported for enterprise network management of XenServer.



Resolution	Use the OpenVSwitch Bridge

Table 833: XenServer log directory size check [106516]

Name	XenServer log directory size check
Description	Check if XenServer host's log directory /var/log is full
Cause	New partition scheme in XenServer 7.x assigns 4 GB for /var/log which can get full quickly
Impact	User VMs running on the cluster become unresponsive
Resolution	Delete some log files from /var/log on host

Table 834: Protection domain contains volume group(s) using client IP address to whitelist a client. [110262] [A110262]

Name	Protected Volume Group with IP address in Whitelist Check
Description	Check if any volume group with IP address in whitelist is protected in protection domain
Alert message	Volume groups <code>vg_names</code> in protection domain <code>pd_name</code> use client IP address to whitelist a client.
Cause	Protection domain contains volume group(s) using client IP address to whitelist a client.
Impact	The attachment between VM and VG is not recoverable from snapshot
Resolution	Update the client properties to include the client IQN instead of the IP address to configure the client list for the volume group(s).

Table 835: Protected VM Name Too Long or Contains Restricted Characters. [110267] [A110267]

Name	VM name does not meet requirements.
Description	Check if the protected VM name is either too long, or contains restricted characters.
Alert message	Restore of the protected VM <code>vm_name</code> will fail due to VM name containing restricted characters or VM name exceeding the character limit.
Cause	VM name is either too long (longer than 80 characters) or contains restricted characters. Restricted characters include [,], ?, *, ", ', ,.
Impact	The VM(s) will not be restored during failovers.
Resolution	Rename the VM(s).

Table 836: Detected snapshots on cluster with high density nodes [110450] [A110450]

Name	High density nodes snapshots check



Description	Checks if snapshots are present on a cluster with high density nodes
Alert message	Protection domain snapshots are detected on cluster with high density nodes
Cause	Detected snapshots on the cluster with high density nodes
Impact	Snapshots are not currently supported on clusters with high density nodes. They may impact cluster operations
Resolution	Delete snapshots on the cluster with high density nodes

Table 837: Data protection is configured on cluster with high density nodes [110451] [A110451]

Name	High density nodes DR configuration check
Description	Checks if disaster recovery is configured on a cluster with high density nodes
Alert message	Disaster recovery entities like protection domains, remote sites are detected on cluster with high density nodes
Cause	Remote sites or protection domains are detected on cluster with high density nodes
Impact	Data protection is currently not supported on clusters with high density nodes, this may impact cluster operations
Resolution	Delete disaster recovery entities like remote sites, protection domains on cluster with high density nodes

Table 838: NGT CA Setup Check [111046]

Name	NGT CA Setup Check
Description	Checks that the NGT CA setup is same on all nodes.
Cause	/home/ngt/ca.tar does not have the same checksum on all CVMs.
Impact	NGT operations may not function properly.
Resolution	Contact Nutanix support.

Table 839: RPO script validation on storage heavy cluster [111048]

Name	RPO script validation on storage heavy cluster
Description	Check that RPO validation script is installed and running on storage heavy cluster.
Cause	RPO validation scripts are not installed.
Impact	RPO may not met.
Resolution	Install RPO validation script on storage heavy cluster.
Resolution	Contact Nutanix support.



Table 840: CVM active upgrade check [111062]

Name	CVM active upgrade check
Description	Check if there are any active upgrades
Cause	An upgrade was requested for this entity
Impact	Maintenance activities should wait until upgrade is finished.
Resolution	Manually upgrade AOS. Contact Nutanix support for assistance.

Table 841: AOS upgrades are disabled on {cvm_ip} [111063] [A111063]

Name	AOS upgrades are disabled check
Description	Check if AOS upgrades are disabled on CVM
Alert message	AOS upgrades are disabled on {cvm_ip}
Cause	AOS upgrade has been paused manually.
Impact	Upgrades might fail because of stale entries
Resolution	Manually upgrade AOS. Contact Nutanix support for assistance.

Table 842: Firmware upgrades are disabled on {cvm_ip} [111064] [A111064]

Name	Firmware upgrades are disabled on this CVM
Description	Check if firmware upgrades are disabled on the CVM
Alert message	Firmware upgrades are disabled on {cvm_ip}
Cause	Firmware upgrade has been paused manually.
Resolution	Manually upgrade the firmware. Contact Nutanix support for assistance.

Table 843: Metro Availability Prechecks Failed [130124] [A130124]

Name	Metro Availability Prechecks Failed
Description	Metro availability prechecks failed
Alert message	Prechecks for {operation} failed for the protection domain '{protection_domain_name}' to the remote site '{remote_name}'. Reason: '{reason}'.
Cause	Various
Impact	Metro availability operation could not be started.
Resolution	Contact Nutanix support if this condition persists



Table 844: Recovery Point Objective Cannot Be Met [130138] [A130138]

Name	Recovery Point Objective Cannot Be Met
Description	Recovery Point Objective Cannot Be Met
Alert message	Recovery point objective cannot be met because 'reason'
Cause	Various
Impact	Recovery plan could be affected.
Resolution	Contact Nutanix support if this condition persists

Table 845: Free Xi Account Expired [200701] [A200701]

Name	Free Xi Account Expired
Description	Free period using Xi expired. If this issue persists, you may experience an interruption in service.
Alert message	Free period using Xi expired. If this issue persists, you may experience an interruption in service.
Cause	The free period is limited by time.
Resolution	Select a payment plan and method in the billing section under the top right gear menu.

Table 846: Xi Subscription Expired [200702] [A200702]

Name	Xi Subscription Expired
Description	Xi subscription period expired. If this issue persists, you may experience an interruption in service.
Alert message	Xi subscription period expired. If this issue persists, you may experience an interruption in service.
Cause	The subscription period is limited by time.
Resolution	Select a new payment plan in the billing section under the top right gear menu.

Table 847: Xi Payment Missed [200703] [A200703]

Name	Xi Payment Missed
Description	Unable to process payment. If this issue persists, you may experience an interruption in service.
Alert message	Unable to process payment. If this issue persists, you may experience an interruption in service.
Cause	This is likely due to an expired or canceled credit card.
Resolution	Update payment method in the billing section under the top right gear menu.



Table 848: Xi Payment Missed [200704] [A200704]

Name	Xi Payment Missed
Description	Unable to process payment. If this issue persists, you may experience an interruption in service.
Alert message	Xi tenant <code>tenant_uuid</code> invoice payment for <code>plan_name</code> plan failed at <code>timestamp_usecs</code>
Cause	This is likely due to a change in address or a forgotten or missplaced invoice.
Resolution	Ensure invoices are sent to the correct contact information provided in the billing section under the top right gear menu. Confirm with your team that the invoice wasn't misplaced or forgotten. Contact support if needed.

Table 849: External Client Authentication [300400] [A300400]

Name	Admin User API Authentication Check
Description	Checks if any external client is authenticated to the REST API with the admin user.
Alert message	External Client Authentication: IP <code>client_ip_address</code> is authenticating to the REST API with the admin user.
Cause	An external client is configured to authenticate to the REST API with the admin user.
Impact	Admin user will be disabled in a forthcoming release. REST API calls from the admin user will stop working.
Resolution	Authenticate to the REST API with a different user that has administrative privileges.

Table 850: Nucalm Internal Service has Stopped Working [400101] [A400101]

Name	Nucalm Service Down
Description	Nucalm internal service is down
Alert message	Discovered that the nucalm internal service is not working: <code>'service_names'</code>
Cause	Nucalm internal service may have stopped working
Impact	You will not be able to perform App Management related operations.
Resolution	Please refer to KB-7080

Table 851: Epsilon Internal Service has Stopped Working [400102] [A400102]

Name	Epsilon Service Down
Description	Epsilon internal service is down



Alert message	Discovered that the epsilon internal service is not working: 'service_names'
Cause	Epsilon internal service may have stopped working
Impact	You will not be able to perform App Management related operations.
Resolution	Please refer to KB-7080

Table 852: Entity count exceed discovered [400103] [A400103]

Name	Entity Count Limit Reached
Description	Entity count exceed discovered
Alert message	Discovered 'count' entities of type 'entity_type', limit is 'limit'
Cause	Entities in this Nutanix cluster has reached recommended upper limit.
Impact	Performiace of App Management operation impacted.
Resolution	Clean up some un-used entities.
Resolution	Resolve the issue for the failure. If you cannot resolve the issue, contact Nutanix support.

Table 853: Applications Archive ready for download. [400104] [A400104]

Name	Applications Archive Ready
Description	Archive ready for download.
Alert message	Application logs from 'start_time' to 'end_time' are now available for download and archival. Please go to Calm > Settings to download them.
Cause	Archival process archived the runlogs.
Impact	You will not be able to download it after a week.
Resolution	Make sure you download the archive.

Table 854: Calm showback is unable to reach beam service [400107] [A400107]

Name	Beam is unreachable
Description	Beam service is unreachable. Calm showback calculations will be affected
Alert message	Beam service is unreachable. Calm showback calculations will be affected
Cause	Prism central may not have internet connectivity or beam service is unavailable.
Impact	Impact on Calm Showback data.
Resolution	Review KB-6664 for further troubleshooting.



Table 855: AHV Prism Element attached [400108] [A400108]

Name	AHV Prism Element attached
Description	A new AHV Prism Element is attached.
Alert message	AHV Prism Element attached name: <i>pe_name</i>
Cause	A new AHV Prism Element is attached.
Impact	New Calm provider setting is added without showback pricelist.
Impact	Prism Element unusable in Calm if not associated with Calm project.
Resolution	Associate the Prism Element in Calm project.
Resolution	Define the pricelist for added Calm provider setting.

Table 856: AHV Prism Element detached [400109] [A400109]

Name	AHV Prism Element detached
Description	One of the existing AHV Prism Element is detached.
Alert message	AHV Prism Element detached name: <i>pe_name</i>
Cause	One of the existing AHV Prism Element is detached.
Impact	Actions other than soft delete cannot be run on Calm applications that are provisioned on this Prism Element.
Impact	Calm provider settings corresponding to this Prism Element is deleted.
Resolution	In case the Prism Element is permanently detached run soft delete on the Calm applications that are provisioned on this Prism Element.

Table 857: Calm Version Mismatch [400110] [A400110]

Name	Calm Version Mismatch
Description	Calm Version Mismatch
Alert message	All nodes are not running on same Calm version
Cause	Failed LCM upgrade
Impact	Calm will not work as expected
Resolution	Refer to KB 7528 for resolution

Table 858: Epsilon Version Mismatch [400111] [A400111]

Name	Epsilon Version Mismatch
Description	Epsilon Version Mismatch
Alert message	All nodes are not running on same Epsilon version
Cause	Failed LCM upgrade



Impact	Epsilon will not work as expected
Resolution	Refer to KB 7528 for resolution

Table 859: Compute-only cluster sizing: hyperconverged node count check [405000] [A405000]

Name	Compute-only cluster sizing: hyperconverged node count check
Description	Check that cluster with compute-only nodes has at least 4 hyperconverged nodes.
Alert message	Cluster with compute-only nodes should have min of 4 hyperconverged nodes
Cause	Compute-only node was added to a cluster with fewer than 4 hyperconverged nodes.
Impact	Can lead to cluster not delivering sufficient storage IO.
Resolution	Remove all compute-only nodes from cluster or add additional hyperconverged nodes.

Table 860: Compute-only cluster sizing: compute-only to hyperconverged node ratio check [405001] [A405001]

Name	Compute-only cluster sizing: compute-only:hyperconverged node ratio check
Description	Ratio of compute-only to hyperconverged nodes to not exceed 1:2
Alert message	Cluster with compute-only nodes should not have the ratio of compute-only to hyperconverged nodes exceed 1:2
Cause	High ratio of compute-only to hyperconverged nodes.
Impact	Can lead to cluster not delivering sufficient storage IO.
Resolution	Remove all compute-only nodes from cluster or add additional hyperconverged nodes.

Table 861: Compute-only cluster sizing: CVM size check [405002] [A405002]

Name	Compute-only cluster sizing: CVM size check
Description	Check that cluster with compute-only: Total CVM vCPUs must be greater than total cores on all compute-only nodes.
Alert message	Cluster with compute-only nodes should have total CVM vCPUs greater than total cores on all compute-only nodes
Cause	Compute-only node was added to a cluster with small CVMs.
Impact	Can lead to cluster not delivering sufficient storage IO.
Resolution	Remove all compute-only nodes from cluster or add additional hyperconverged nodes or resize CVMs



Table 862: Insufficient host bandwidth [405003] [A405003]

Name	Compute-only minimum bandwidth check
Description	Checks whether the total hyperconverged network bandwidth is twice the compute-only bandwidth.
Alert message	For a cluster with compute-only nodes the total amount of hyperconverged host bandwidth must be double the amount of compute-only host bandwidth.
Cause	Cluster is configured with too many compute-only nodes.
Cause	Hyperconverged/compute-only host NIC is misconfigured.
Impact	Cluster will not be able to deliver sufficient storage I/O.
Resolution	Remove compute-only nodes from the cluster.
Resolution	Add hyperconverged nodes to the cluster.
Resolution	Reconfigure hyperconverged host NICs to increase the amount of bandwidth.

Storage

Table 863: Data Disk Space Usage High [1003] [A1005]

Name	Data Disk Usage
Description	Checks if any individual disk (SSD or HDD) storage space utilization is sustained for long above a high threshold.
Alert message	Disk space usage for <code>mount_path</code> on entity <code>ip_address</code> is <code>disk_usage%</code> which exceeds the threshold of <code>disk_usage_threshold%</code>. <code>action_str</code>
Cause	Individual disk storage is being highly utilized, possibly due to overall storage pool usage pressure.
Cause	Sustained high intensity IO.
Cause	Issues with disk/node balancing.
Impact	In isolation a spike in individual disk usage is not an immediate cause for concern but if sustained or seen across multiple disks it can indicate an underlying issue with data balancing and may lead to significant cluster performance degradation or issues with storage maintenance.
Resolution	Ensure storage pool usage is below critical threshold.
Resolution	Ensure Curator service is running and periodic scans are completing.
Resolution	Consider any high intensity write I/O operations that may be in progress, such as inbound data migrations.
Disk Usage Threshold	Warning Threshold: 90 % Critical Threshold: 95 %



Table 864: Disk Inode Usage High [1004] [A1027]

Name	Inode Usage
Description	Checks if current inode usage is high.
Alert message	Inode usage for one or more disks on Controller VM <code>svm_ip_address</code> has exceeded <code>inode_usage_threshold</code>%.
Cause	Inode usage is high.
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Reduce disk usage or replace disk.
Inode Usage Threshold	Warning Threshold: 75 % Critical Threshold: 90 %

Table 865: Storage Device Health Bad [1005] [A1069]

Name	SmartCtl Health
Description	Smartctl health bad.
Alert message	Smartctl reports attribute <code>attribute</code> having value <code>value</code> for device <code>device</code> on CVM <code>service_vm_external_ip</code>.
Cause	Smartctl health bad.
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Check or replace disk.

Table 866: Disk Unused [1009] [A1063]

Name	Disk Configuration
Description	Checks that all disks in the cluster are assigned to a storage pool
Alert message	Drive with disk id <code>disk_id</code> on Controller VM <code>service_vm_external_ip</code> is not part of any storage pool.
Cause	The cluster was not configured correctly during the installation or a drive replacement was not completed correctly.
Impact	The cluster was not configured correctly during the installation or a drive replacement was not completed correctly
Resolution	Add the unused disk to the storage pool.



Table 867: Hypervisor Disk usage is above the recommended threshold. [1015] [A1161]

Name	Host Disk Usage
Description	Checks that host disk usage is not above the recommended threshold%.
Alert message	Disk ' <i>dev_name</i> ' space usage (<i>usage%</i>) on host <i>ip_address</i> has exceeded <i>threshold%</i> .
Cause	Too much data is stored on the host disk.
Impact	Cluster performance will be significantly degraded. In the case of multiple nodes with the same condition,the cluster may be unable to service I/O requests.
Resolution	Delete unneeded data from the host disk.
Host Disk Usage Threshold	Warning Threshold: 90 % Critical Threshold: 95 %

Table 868: Storage Pool Space Usage Exceeded the Configured Threshold [1016] [A1128]

Name	StoragePool Space Usage
Description	Checks the storage pool space usage to determine if it's above the threshold
Alert message	Storage Pool space usage for <i>storage_pool_name</i> is at <i>usage_pct%</i>. <i>alert_msg</i>
Cause	Space usage in the storage pool has exceeded the configured threshold.
Impact	Storage pool will run out of space, and the cluster will be unable to service I/O requests.
Resolution	Add more storage to the cluster or delete data to free up storage space.
Storage Pool Usage Threshold	Warning Threshold: 75 % Critical Threshold: 90 %

Table 869: Storage Container Space Usage Exceeded [1017] [A1130]

Name	Storage Container Space Usage
Description	Check high space usage on Storage Containers
Alert message	Storage Container space usage for <i>container_name</i> is at <i>usage_pct%</i>.
Cause	Excessive space usage in the Storage Container.
Impact	Storage Container will run out of space, and the cluster may become unable to service I/O requests.
Resolution	Delete some data to free storage space from the Storage Conatiner.
Storage Container Usage Threshold	Warning Threshold: 75 % Critical Threshold: 90 %



Table 870: System Root Partition Space Usage High [1021] [A1036]

Name	System Root Partition Usage
Description	Checks if CVM system root partition usage is within threshold to ensure uninterrupted operations. Does not refer to cluster shared storage.
Alert message	Disk space usage for <code>mount_path</code> on <code>entity ip_address</code> has exceeded <code>threshold%</code>.
Cause	Increased CVM system root partition usage due to excessive logging or incomplete maintenance operation.
Cause	Non-standard use of / (root) partition and sub-directories on the affected CVM local filesystem.
Impact	If CVM system root partition is highly utilised then certain maintenance operations, such as upgrades, may be impacted. If CVM system root partition is 100% utilised then services may stop and impact cluster storage availability.
Resolution	Remove unneeded/temporary/non-standard introduced files from affected CVM root partition.
Resolution	Refer to KB 1523 for further details.
CVM System Root Partition Usage Threshold	Warning Threshold: 80 % Critical Threshold: 90 %

Table 871: Flash Mode Configuration has exceeded the threshold limit [1022] [A1158]

Name	Flash Mode Configuration
Description	Check the total configured size of flash-mode-enabled vDisks with respect to the threshold limit.
Alert message	Flash-mode-enabled vDisks are configured to use more flash tier space than the allotted limit.
Cause	Flash-mode-enabled vDisks are configured to use more flash tier space than the allotted limit.
Impact	Flash mode policies may not be enforced properly.
Resolution	Reduce the number of flash-mode-enabled vDisks or increase the capacity of the flash tier.

Table 872: Flash Mode Usage Limit Exceeded [1023] [A1159]

Name	Flash Mode Usage
Description	Check the amount of usage by flash-mode-enabled vDisks with respect to the threshold limit.
Alert message	Flash Mode usage has exceeded the threshold limit.
Cause	Flash-mode-enabled vDisks are using more flash tier space than the allotted limit.



Impact	System may down-migrate flash mode data.
Resolution	Reduce the number of flash-mode-enabled vDisks or increase the capacity of the flash tier.

Table 873: Flash-mode-enabled VM Power Status [1024] [A1160]

Name	Flash Mode Enabled VM Power Status
Description	Check if any flash-mode-enabled VMs are Powered Off.
Alert message	One or more Powered Off flash-mode-enabled VMs detected.
Cause	One or more flash-mode-enabled VMs are Powered Off.
Impact	Flash tier space might be getting wasted.
Resolution	Disable Flash Mode on the Powered Off VMs, or power them on.

Table 874: High Garbage due to erasure coding [1026] [A1174]

Name	Erasure Code Garbage
Description	Checks if erasure coding garbage is below a safe threshold.
Alert message	Garbage due to erasure coding on <code>container_name</code> is currently <code>usage</code> which is above the safety threshold.
Cause	Suboptimal performance of erasure coding is generating garbage that is consuming additional storage space.
Impact	Erasure coding on Nutanix clusters may experience suboptimal storage savings.
Resolution	Review KB 2912 for corrective actions. Contact Nutanix support to discuss possible fixes before enabling or disabling erasure coding.

Table 875: Erasure coding on Acropolis base software versions less than 4.5.2 may perform suboptimally. [1027] [A1175]

Name	Erasure Code Configuration
Description	Checks erasure coding is enabled on cluster with version < 4.5.2
Alert message	Erasure Coding is enabled on Nutanix clusters with Acropolis base software version less than 4.5.2.
Cause	Erasure coding was enabled on the cluster
Impact	Erasure coding on Nutanix clusters with version < 4.5.2 may experience suboptimal storage savings.
Resolution	Upgrade to Acropolis base software 4.5.2 or higher. See KB 2912 for details.



Table 876: Invalid erasure code delay parameter. [1028] [A1176]

Name	Erasure-Code-Delay Configuration
Description	Validate the EC param value for erasure-code-delay
Alert message	The erasure coding parameter erasure-code-delay for Storage Container <code>container_name</code> <code>msg</code>
Cause	Invalid setting for erasure-code-delay
Impact	Current Erasure coding parameter value can lead to poor performance and suboptimal storage savings.
Resolution	Please fix the erasure-code-delay value. See KB 2912 for details.

Table 877: Storage Pool Flash Mode Configuration [1029] [A1173]

Name	Storage Pool Flash Mode Configuration
Description	Flash Mode is not supported when multiple storage pools are in use.
Alert message	Flash Mode configuration with Multiple Storage Pools detected.
Cause	Multiple Storage Pools and Flash Mode are in use together.
Impact	This configuration is not supported.
Resolution	Disable Flash Mode for all VMs and Volume Groups, or consolidate all Storage Pools into one.

Table 878: Hypervisor Boot Drive Wear High [1061] [A1198]

Name	Hyve Boot Disk Status
Description	Check Hyve boot disk status
Alert message	Hypervisor boot drive on the Controller VM <code>ip_address</code> has problem: <code>alert_msg</code>
Cause	The drive is close to the maximum write endurance.
Impact	Cluster performance may be significantly degraded. In the case of multiple nodes with the same condition, the cluster may become unable to service I/O requests.
Resolution	Replace the drive as soon as possible. Refer to the Nutanix documentation for instructions.

Table 879: VSS Snapshot of Container Share Failed [20020] [A20020]

Name	VSS snapshot of container share failed
Description	A VSS initiated snapshot of the container share failed due to too many files in the container.



Alert message	A backup application requires a container id <code>container_id</code> to be backed up, however the VSS snapshot operation on it failed as there are <code>files_present</code> files on the container. Only a container with a maximum of <code>files_max</code> files can be snapshotted through VSS.
Cause	Too many files in the container.
Impact	Third party backup applications using VSS to backup VMs from a container share will fail.
Resolution	Reduce the number of files in each container to less than or equal to 1000.

Table 880: Erasure coding pending check [101054]

Name	Erasure coding pending check
Description	Check if EC undo is pending on any of the Storage Containers
Cause	EC egroups are still pending to be converted to non-EC egroups after undo EC on the container.
Impact	Node/Disk remove will be delayed until the EC egroups are undone to non-EC egroups are completed.
Resolution	Contact Nutanix Support

Table 881: Volume Group Space Usage Exceeded [101056] [A101056]

Name	Volume Group Space Usage
Description	Check high space usage on volume groups
Alert message	Volume Group space usage for <code>volume_group_name</code> is at <code>usage_pct</code>%.
Cause	Excessive space usage in the volume group.
Impact	Volume Group will run out of space, and the cluster may become unable to service I/O requests to that volume group.
Resolution	Add more capacity to the volume group or delete some data to free storage space. The Guest OS and/or the application using the volume group should be configured to send the list of freed blocks to AOS using SCSI UNMAP command. This can be done as part of in-band or out of band TRIM operation(s). Refer to 'Thin Provisioning' section in the Nutanix Volumes - Nutanix Best Practices guide.
Volume Group Usage Threshold	Critical Threshold: 90 %

Table 882: System Non-Root Partitions Space Usage High [101061] [A1031]

Name	System Non-Root Partition Usage
Description	Checks if CVM system non-root partition usage is within threshold to ensure uninterrupted operations. Does not refer to cluster shared storage.



Alert message	Disk space usage for <code>mount_path</code> on <code>entity ip_address</code> has exceeded <code>threshold%</code> .
Cause	Increased CVM system non-root partition usage due to excessive logging, storage of large log bundles or unused 1-click binaries.
Cause	Non-standard use of <code>/home</code> directories on the affected CVM.
Impact	If CVM system non-root partitions are highly utilised then certain maintenance operations, such as upgrades, may be impacted. If CVM system non-root partitions are 100% utilised then services may stop and impact cluster storage availability.
Resolution	Remove unneeded 1-click bundles via Prism, log collector bundles and temporary/non-standard introduced large files from affected CVM partition.
Resolution	If unsure, please engage Nutanix Support for assistance.
Resolution	Please refer to KB 1540 for reference.
CVM System Non-Root Partition Usage Threshold	Warning Threshold: 75 % Critical Threshold: 90 %

Table 883: `/scratch` location Space Usage is High [101068] [A101068]

Name	Scratch space usage
Description	Checks that scratch space usage is not high.
Alert message	<code>/scratch space usage (usage%) on host ip_address</code> has exceeded <code>threshold%</code> .
Cause	Too much data is stored in scratch space.
Impact	If scratch space is not available, temporary data is stored on a RAM disk with limited space, causing issues during upgrade
Resolution	Delete unneeded data or add nodes to the cluster.
Scratch space Usage Threshold	Warning Threshold: 90 % Critical Threshold: 95 %

Table 884: SAS Connectivity Not Normal [106019] [A106019]

Name	SAS Connectivity
Description	Check SAS connectivity
Alert message	SAS cables missing or not properly connected on host.
Cause	SAS cables are not properly connected.
Impact	Cluster performance may be significantly degraded.
Resolution	Ensure that SAS cables are properly connected.



Table 885: Storage Container Replication Factor Low [110210] [A1074]

Name	Storage Container RF Status
Description	Check that the Storage Container region replication factor is not low
Alert message	Replication factor of <code>container_region</code> on Storage Container <code>container_name</code> is set to 1.
Cause	Replication factor on the Storage Container was changed to 1.
Impact	Only one copy of guest VM data exists on the cluster. If any node fails, the data stored on that node is unavailable. If any drive fails, the data stored on that drive is unrecoverable.
Resolution	Change the replication factor on the Storage Container with the nCLI command "container edit name='ctr_name' rf='2'" (or higher if appropriate for your environment).

Table 886: Compression Disabled [130014] [A1121]

Name	Compression Status
Description	Compression is Disabled.
Alert message	Disabling compression for future writes. Calculated metadata usage of <code>metadata_usage</code> bytes exceeds safety limit of <code>metadata_limit</code> bytes on Controller VM <code>service_vm_id</code>. The metadata disk size is <code>metadata_disk_size</code>.
Cause	Compression metadata usage has exceeded the safety limits. Compression metadata usage can increase due to a variety of reasons such as having too many snapshots.
Impact	Advantages of compression are not available.
Resolution	Once the metadata size is brought under the safety limit (for example, by removing unneeded snapshots, or by expanding the cluster size), contact Nutanix support to manually enable compression.

Table 887: Fingerprinting Disabled [130021] [A1068]

Name	Finger Printing Status
Description	Fingerprinting Disabled
Alert message	Disabling fingerprinting (deduplication) for future writes. Calculated metadata usage of <code>metadata_usage</code> bytes exceeds safety limit of <code>metadata_limit</code> bytes on Controller VM <code>service_vm_id</code>. The metadata disk size is <code>metadata_disk_size</code>.
Cause	The metadata disk is used for metadata and data. Once the metadata portion on the disk exceeds the safety threshold we disable fingerprinting (deduplication) of future writes. Metadata sizes can increase due to a variety of reasons such as having too many snapshots, or simply because too much data is fingerprinted.
Impact	The advantages of deduplication are not available.



Resolution	Once the metadata size is brought under the safety limit (for example, by removing unneeded snapshots, or by expanding the cluster size), contact Nutanix support to manually enable fingerprinting.

Table 888: Metadata Usage High [130025] [A1101]

Name	Metadata Usage
Description	Metadata Usage High.
Alert message	Metadata usage on Controller VM <code>service_vm_external_ip</code> has exceeded <code>critical_pct%</code>.
Cause	The cluster either has too many snapshots or too much data is being fingerprinted.
Impact	Cluster performance may be significantly degraded.
Resolution	Reduce metadata size by removing unneeded snapshots, creating snapshots less frequently, creating snapshots of fewer VMs, or expanding the cluster size.

Table 889: Non Self Encryption Drive Disk Inserted [130031] [A1122]

Name	Non SED Disk Inserted Check
Description	Non Self Encrypting Drive disk inserted
Alert message	Non encrypted disk with serial <code>disk_serial</code> was added in drive bay <code>disk_location</code> on Controller VM <code>service_vm_external_ip</code>. It is not usable because the rest of the cluster is protected using encrypted drives.
Cause	A cluster with self-encrypting drives has a non-self encrypting drive installed.
Impact	The new disk was not mounted because the cluster cannot be fully secured without all drives supporting encryption.
Resolution	Ensure all installed drives are of self-encrypting type. If more help is needed call Nutanix support.

Table 890: On-Disk Dedup Disabled [130032] [A1089]

Name	On-Disk Dedup Status
Description	On-Disk Dedup Disabled
Alert message	Disabling further on-disk deduplication (fingerprinting). Controller VM with <code>ram_size</code>GB RAM and <code>ssd_size</code>GB SSD does not meet the minimum requirements of <code>required_ram_size</code>GB RAM and <code>required_ssd_size</code>GB SSD to support on-disk deduplication feature.
Cause	To support on-disk deduplication (fingerprinting) feature, the Controller VM must meet the minimum requirements for RAM and SSD. When this safety threshold is not satisfied, on-disk deduplication is disabled.
Impact	The advantages of on-disk deduplication are not available.



Resolution	Once the Controller VMs are upgraded to meet the minimum requirements on RAM and SSD, on-disk deduplication can be enabled. Contact Nutanix support for assistance.

Table 891: Disk Failed/Marked offline [130033] [A1044]

Name	Disk Offline Status
Description	Disk Bad
Alert message	Disk disk_position with id disk_id on node node_position of block block_position has been marked offline due to IO errors. Serial number of the disk is disk_serial in host host_ip of block block_serial .
Cause	A drive on the node has failed and was marked offline.
Impact	Cluster storage capacity will be reduced until the drive is replaced.
Resolution	Replace the failed drive. Refer to the Nutanix documentation for instructions.

Table 892: Physical Disk Drive Has Failed [130035] [A1104]

Name	Physical Disk Status
Description	Physical drive has failed.
Alert message	Drive disk_id with serial disk_serial and model disk_model in drive bay disk_location located on node node_serial_number in position node_position on Controller VM service_vm_external_ip has failed.
Cause	The drive has failed.
Impact	Cluster storage capacity is reduced.
Resolution	Replace the failed drive. Please refer to KB 4158.

Table 893: Physical Disk Removed From Slot [130036] [A1103]

Name	Physical Disk Remove Check
Description	Physical disk removed from slot
Alert message	Disk with serial disk_serial and model disk_model was removed from drive bay disk_location on Controller VM service_vm_external_ip
Cause	A drive was physically disconnected from a slot.
Impact	Migration of data from the disk will start.
Resolution	Reseat the removed disk back in slot. Contact the Hardware Vendor if issue persists.

Table 894: Space Reservation Violated [130053] [A1021]

Name	Space Reservation Status



Description	Space Reservation Violated
Alert message	Space reservation configured on vdisk <code>vdisk_name</code> belonging to Storage Container id <code>container_id</code> could not be honored due to insufficient disk space resulting from a possible disk or node failure.
Cause	A drive or a node has failed.
Impact	The space reservations on the cluster can no longer be met. Writes by guest VMs may fail if expected storage space is not available.
Resolution	Change space reservations to total less than 90% of the available storage, and replace the drive or node as soon as possible. Refer to the Nutanix documentation for instructions.

Table 895: vDisk Block Map Usage High Critical [130056] [A1061]

Name	vDisk Block Map Usage
Description	vDisk Block Map Usage High Critical.
Alert message	Too many snapshots have been allocated in the system. This may cause perceivable performance degradation.
Cause	Too many vdisks or snapshots are present in the system.
Impact	Cluster performance may be significantly degraded.
Resolution	Remove unneeded snapshots and vdisks. If using remote replication, try to lower the frequency of taking snapshots. If you cannot resolve the error, contact Nutanix support.

Table 896: Protected Volume Group not Found [130070] [A1162]

Name	Protected Volume Group Not Found
Description	Protected Volume Group not Found
Alert message	Unable to locate volume group with name '<code>vg_name</code>' and internal ID '<code>vg_uuid</code>' protected by protection domain '<code>protection_domain_name</code>'.
Cause	The protected volume group cannot be found and may have been deleted.
Impact	Any data associated with the volume group may not be backed up or replicated to a remote site.
Resolution	Remove the volume group from the protection domain.

Table 897: Volume Group Action Error [130071] [A1163]

Name	Volume Group Action Status
Description	Volume Group Action Error
Alert message	Failed to <code>action</code> volume group with name '<code>vg_name</code>' and internal ID '<code>vg_uuid</code>' because <code>reason</code> in protection domain '<code>protection_domain_name</code>'



Cause	A volume group could not be restored or could not be deleted because it is still in use.
Impact	The requested volume group action (restore or delete) could not be completed.
Resolution	Detach the volume group from VMs and external initiators before recovery.
Resolution	Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.

Table 898: vDisk Block Map Usage High Warning [130072] [A130072]

Name	vDisk Block Map Usage Warning
Description	vDisk Block Map Usage High.
Alert message	Too many snapshots have been allocated in the system. This may cause perceivable performance degradation.
Cause	Too many vdisks or snapshots are present in the system.
Impact	Cluster performance may be significantly degraded.
Resolution	Remove unneeded snapshots and vdisks. If using remote replication, try to lower the frequency of taking snapshots. If you cannot resolve the error, contact Nutanix support.

Table 899: Unable to remount datastore [130076] [A130076]

Name	Datastore Remount Status
Description	Unable to remount datastore
Alert message	Unable to operation Datastore '<i>datastore_name</i>' host_info.
Cause	Various
Impact	VM HA and DRS functionality might be impaired.
Resolution	Please refer to KB2898 at the Nutanix Portal.

Table 900: Remounted datastore [130080] [A130080]

Name	Datastore Remount Success
Description	Remounted datastore
Alert message	Remounted Datastore '<i>datastore_name</i>' as metro availability is being enabled.
Cause	Various
Impact	VM HA and DRS functionality might be impaired.
Resolution	Please refer to KB2898 at the Nutanix Portal.



Table 901: Protected Vms Not Found [130083] [A130083]

Name	Protected VMs Not Found
Description	Protected Vms Not Found
Alert message	Unable to locate VM(s) <i>vm_names</i> protected by protection domain '<i>protection_domain_name</i>'.
Cause	Protected VM(s) cannot be found and may have been deleted.
Impact	Any data associated with the VMs may not be backed up or replicated to a remote site.
Impact	Next transition into NearSync will be delayed if the protection domain is configured with high frequency schedule.
Resolution	Verify that the VM(s) are registered on the host. If the VM(s) were unregistered or deleted - Remove VM(s) from the protection domain. If the CVM on the host the VM(s) are registered on was shutdown or rebooting, retry the snapshot operation.

Table 902: Protected Volume Groups Not Found [130084] [A130084]

Name	Protected Volume Groups Not Found
Description	Protected Volume Groups Not Found
Alert message	Unable to locate volume group(s) <i>vg_names</i> protected by protection domain '<i>protection_domain_name</i>'.
Cause	Protected volume group(s) cannot be found and may have been deleted.
Impact	Any data associated with the volume group(s) may not be backed up or replicated to a remote site.
Impact	Next transition into NearSync will be delayed if the protection domain is configured with high frequency schedule.
Resolution	Remove volume group(s) from the protection domain.

Table 903: VStore Snapshot Status [130086] [A130086]

Name	VStore Snapshot Status
Description	VStore Snapshot Status
Alert message	Snapshot status for vstore <i>vstore_name</i>: <i>reason</i>.
Cause	Various
Impact	The requested VStore snapshot action could not be completed.
Resolution	Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.



Table 904: Disk Diagnostic Failure [130089] [A1139]

Name	Disk Diagnostic Status
Description	The drive diagnostic test has failed.
Alert message	Drive <code>disk_id</code> with model <code>disk_model</code> with serial <code>disk_serial</code> in drive bay <code>disk_location</code> on Controller VM <code>service_vm_external_ip</code> has failed diagnostic test. This drive is failing or has failed.
Cause	The drive is failing or has failed.
Impact	Cluster storage capacity has been reduced.
Resolution	Replace the failing drive. Refer to the Nutanix documentation for instructions.

Table 905: Related Entity Protection Status [130090] [A130090]

Name	Related Entity Protection Status
Description	Protection status of a related entity.
Alert message	<code>error_message</code> .
Cause	Related entity is not protected in the same protection domain.
Impact	Related VM/Volume Group will not be snapshotted and recovered.
Resolution	Protect the related entity in the same consistency group.

Table 906: iSCSI Configuration Failed [130100] [A130100]

Name	iSCSI Configuration Failed
Description	iSCSI Configuration Failed
Alert message	Failed to re-configure iSCSI settings on the recovered VM ' <code>vm_name</code> '. <code>reason</code> .
Cause	Nutanix Guest Tools failed to execute some iSCSI commands on the guest VM.
Impact	iSCSI disks may become unavailable on the guest VM.
Resolution	If IQN and iSCSI target IP addresses of the VM have been updated by Nutanix Guest Agent, Discover and connect to new targets after rebooting the VM.
Resolution	Manually configure iSCSI settings on the guest VM.
Resolution	Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.

Table 907: Recovered VM Disk Configuration Update Failed [130107] [A130107]

Name	Disk Configuration Update Failed
Description	Disk Configuration Update Failed
Alert message	Failed to make some disks online on the recovered VM ' <code>vm_name</code> '. <code>reason</code> .



Cause	Nutanix Guest Tools failed to automatically bring the disks online.
Impact	Disks may become offline on the recovered VM.
Resolution	Manually bring the disks online on the recovered VM.
Resolution	Resolve the stated reason for the failure. If you cannot resolve the error, contact Nutanix support.

Table 908: System-Defined Flash Mode Usage Limit Exceeded [130120] [A130120]

Name	System Defined Flash Mode Usage Limit
Description	Check that usage for flash-mode-enabled vDisks is within system limits.
Alert message	System has down-migrated the data of flash-mode-enabled vDisks.
Cause	Too many vDisks are assigned to the flash tier, or the vDisks assigned to the flash tier are too large.
Impact	The advantages of flash mode may not be available.
Resolution	Reduce the number of flash-mode-enabled vDisks or increase the capacity of the flash tier.

Table 909: NFS metadata size overshoot [130122] [A130122]

Name	NFS Metadata Size Overshoot
Description	NFS Metadata Usage High
Alert message	NFS metadata usage is too high. Calculated metadata usage of <code>metadata_usage</code> bytes exceeds safety limit of <code>metadata_limit</code> bytes on Controller VM <code>service_vm_id</code>.
Cause	Either NFS datastore is used as a file store which is not supported, or too many files are present in the datastore.
Impact	Cluster performance and stability may be significantly degraded.
Resolution	If NFS datastore is used as a file store, remove the corresponding files and wait for up to 24 hours to allow compaction.

Table 910: Guest Power Operation through NGT Failed [130149] [A130149]

Name	VM Guest Power Op Failed
Description	Guest Power Operation Failed
Alert message	Failed to perform operation '<code>operation</code>' on VM '<code>vm_name</code>'. <code>reason</code>.
Cause	Power operation failed in guest operating system.
Impact	Desired power operation may not be completed.
Resolution	Manually shutdown the VM by logging into the VM and running the appropriate shutdown command.



Resolution	Resolve the issue for the failure. If you cannot resolve the issue, contact Nutanix support.

Table 911: Protected VM(s) May Not Be Recoverable [130183] [A130183]

Name	Protected VMs May Not Be Recoverable
Description	Protected VM(s) May Not Be Recoverable
Alert message	VM(s) <code>vm_names</code> protected by protection domain ' <code>protection_domain_name</code> ' are running from VMWare snapshots. Recovery of the VMs from the PD snapshot ' <code>snapshot_id</code> ' may not succeed. Manual intervention is required to bootup the VMs.
Cause	Protected VM(s) are running from VMWare snapshots and may not be recoverable.
Impact	Recovery of the VMs from the PD snapshot may not succeed
Resolution	Ensure that the VMs are not running from VMWare snapshot at the time of scheduled Protection Domain snapshot. Refer Nutanix KB: 7023 for further details.

Table 912: External Repository Access Failure [130208] [A130208]

Name	External Repository Access Failure
Description	Failed to access an NFS v3 external repository.
Alert message	Failed to access external repository <code>external_repository_name</code> . Failure reason: <code>message</code> .
Cause	Network connectivity issues between the AOS CVM subnet and the external repository. This can be either intermittent or permanent.
Impact	Mounting an NFS v3 external repository fails
Resolution	Make sure that the CVM subnet has access to the NFS server.

Table 913: IO failures to a data source in an external repository [130209] [A130209]

Name	IO failures to data sources in external repositories
Description	IO to an external data source within an external repository failed.
Alert message	IO to the external data source with path <code>relative_uri</code> in the external repository <code>external_repository_name</code> failed. Failure reason: <code>message</code> .
Cause	IO to the external data source failed with IO error from the NFS server.
Impact	IO to an data source within an external repository fails
Resolution	Make sure the NFS server is in a good state.

System Indicator



Table 914: VSS Snapshot is not supported for some VMs. [130200] [A130200]

Name	VSS Snapshot Not Supported
Description	VSS snapshot is not supported for some VMs.
Alert message	VSS snapshot is not supported for the following VMs protected by Protection Domain '<i>protection_domain_name</i>' VMs list: '<i>vm_names</i>'
Cause	Some VMs in the Protected Domain have unsupported configuration.
Impact	Hypervisor based application consistent snapshot is taken on ESX. Crash consistent snapshot is taken for other hypervisors.
Resolution	For more details, please review the alerts generated at VM page and fix the unsupported configuration.

Table 915: Tomcat is restarting frequently. [700101] [A700101]

Name	Tomcat Frequent Restart
Description	Tomcat is restarting frequently.
Alert message	Tomcat is restarting frequently on <i>cvm_ip</i>. Reason: <i>message</i>.
Cause	Tomcat is restarting frequently, please check prism, prism_monitor and catalina logs for exact cause of failure.
Impact	Performance and availability impact for Prism
Resolution	Check the logs for error information or contact Nutanix support.

VM

Table 916: Non Compliance with Host Affinity policies. [200403] [A200403]

Name	VM not in compliance with the defined affinity policies.
Description	VM not in compliance with the defined affinity policies.
Alert message	VM <i>vm_uuid</i> in cluster <i>cluster</i> not in compliance with policies <i>policy_list_str</i>
Cause	A new policy was added, an existing policy was updated, or a category was attached to or detached from the VM or host.
Impact	VM might not run on hosts it is expected to run on.
Resolution	Make appropriate updates to affinity policies OR VM or Host categories.



TASK STATUS

The web console displays detailed information about all tasks that have been performed on the cluster.

Task Page Navigation


- To view the Task dashboard, select **Home > Tasks**.
- An icon also appears in the main menu when one or more tasks are active (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 drop-down list of active tasks; clicking the **View All Tasks** button at the bottom of that list displays a details screen with information about all tasks for this cluster.

Note: The drop-down list of active tasks may include a **Clean Up** button (top right). Clicking this button removes from the list any tasks that are no longer running. However, this applies to the current session only. The full active list (including the non-running tasks) appears when you open a new Prism web console session.

- When multiple tasks are active, you can filter the list by entering a name in the filter by field.
- You can also filter the list by clicking the **Filters** button and selecting the desired filter options

Each task appears in the list for a minimum of one hour after completion, but how long that task remains in the list depends on several factors. In general, the maximum duration is two weeks. However, tasks are rotated off the list as new tasks arrive, so a task might disappear from the list much sooner when activity is high. In some cases a task appears for longer than two weeks because the last task for each component is retained in the listing.

View Task Status



Operation Message	Entity	Percent	Status	Create Time	Duration
Image update	Image	100%	Succeeded	03/14/19, 8:03:37 AM	4 minutes 59 seconds
Image create	Image	100%	Succeeded	03/14/19, 8:03:35 AM	Under 1 second
Image update	Image	100%	Succeeded	03/14/19, 5:57:08 AM	3 minutes 20 seconds
Image create	Image	100%	Succeeded	03/14/19, 5:57:05 AM	Under 1 second
Image delete	Image	100%	Succeeded	03/14/19, 5:51:00 AM	Under 1 second
Image delete	Image	100%	Succeeded	03/14/19, 5:51:00 AM	Under 1 second
Image delete	Image	100%	Succeeded	03/14/19, 5:51:00 AM	Under 1 second
Image delete	Image	100%	Succeeded	03/14/19, 5:50:59 AM	Under 1 second
Image update	Image	100%	Succeeded	03/13/19, 3:14:57 AM	2 minutes 46 seconds
Image create	Image	100%	Succeeded	03/13/19, 3:14:52 AM	Under 1 second
Image delete	Image	100%	Succeeded	03/13/19, 3:13:18 AM	Under 1 second
Image create	Image	100%	Succeeded	03/13/19, 2:54:50 AM	Under 1 second
Image update	Image	100%	Succeeded	03/12/19, 1:19:12 AM	11 minutes 37 seconds
Image create	Image	100%	Succeeded	03/12/19, 1:19:07 AM	Under 1 second
Image delete	Image	100%	Succeeded	03/12/19, 1:11:00 AM	Under 1 second
Image create	Image	100%	Succeeded	03/11/19, 8:41:21 AM	Under 1 second
Update cluster intentful	Cluster	100%	Succeeded	03/11/19, 6:41:38 AM	1 second
Create remote connection intentful	Remote connection	100%	Succeeded	03/11/19, 6:23:33 AM	2 seconds
cluster registration	Details	100%	Succeeded	03/11/19, 6:23:27 AM	10 seconds
Host restore VM locality	Host Details	100%	Succeeded	03/11/19, 5:16:58 AM	6 seconds

Figure 221: Task Dashboard

The following table describes the fields in the tasks list.

Table 917: Tasks List Fields

Parameter	Description	Values
Operation Message	Specifies which type of operation the task is performing.	upgrade <other>
Entity	Display the entity on which task has been performed. If the link appears on the entity, click it to display the details.	Entity description
Percent	Indicates the current percentage complete for the task.	0%-100%
Status	Indicates the task status, which can be pending, running, completed, or failed.	pending, running, completed, failed
Create Time	Displays when the task began.	seconds, minutes, hours
Duration	Displays how long the task took to complete.	seconds, minutes, hours

SYSTEM MANAGEMENT

The web console allow you to configure various system settings.

- You can specify one or more name servers (see [Configuring Name Servers](#) on page 595)
- If Acropolis is enabled, you can configure one or more network connections (see [Network Configuration for VM Interfaces](#) on page 138).
- You can create a whitelist of IP addresses that are allowed access (see [Configuring a Filesystem Whitelist](#) on page 594).
- You can specify one or more NTP servers for setting the system clock (see [Configuring NTP Servers](#) on page 597).
- You can configure one or more network switches for statistics collection (see [Configuring Network Switch Information](#) on page 144).
- You can specify an SMTP mail server (see [Configuring an SMTP Server](#) on page 598).
- You can configure SNMP (see [Configuring SNMP](#) on page 599).
- You can configure a login banner page (see [Configuring a Banner Page](#) on page 615).

Configuring a Filesystem Whitelist

About this task

A whitelist is a set of addresses that are allowed access to the cluster. Whitelists are used to allow appropriate traffic when unauthorized access from other sources is denied.

CAUTION: Using a Nutanix storage container as a general-purpose NFS or SMB share is not supported. Because the Nutanix solution is VM-centric, the preferred mechanism is to deploy a VM that provides file share services.

To add (or delete) an address to (from) the filesystem whitelist, do the following:

Procedure

1. Click the gear icon in the main menu and then select Filesystem Whitelists in the **Settings** page.
The Filesystem Whitelists dialog box appears.




Figure 222: Filesystem Whitelists Window

2. To add an address to the whitelist, do the following in the indicated fields:
 - a. **IP address:** Enter the IP address.
 - b. **Netmask:** Enter the netmask value.
 - c. Click the **+Add** button.The entry is added to the **Whitelist Entry** list. An NFS whitelist is created when the hypervisor is ESXi or AHV; a CIFS whitelist is created when the hypervisor is Hyper-V.
3. To delete an entry from the whitelist, click the X icon for that entry in the **Whitelist Entry** list.
A window prompt appears to verify the action; click the **OK** button. The entry is removed from the list.

Configuring Name Servers

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. To add (or delete) a name server, do the following:

Procedure

1. Click the gear icon in the main menu and then select **Name Servers** in the **Settings** page. The Name Servers dialog box appears.




Figure 223: Name Servers Window

2. To add a name server, enter the server IP address in the **Server IP** field and then click the **Add** button to the right of that field. The server is added to the **IP Address** list (below the **Server** field).

Note: Changes in name server configuration may take up to 5 minutes to take effect. Functions that rely on DNS may not work properly during this time.

3. To delete a name server entry, click the X icon for that server in the **IP Address** list. A window prompt appears to verify the action; click the **OK** button. The server is removed from the list.

Cluster Time Synchronization

Network Time Protocol (NTP) is a protocol for clock synchronization between computers. The hosts and CVMs in a Nutanix cluster must be configured to synchronize their system clocks with a list of stable NTP servers. Accurate timestamps are important for troubleshooting interactions with third-party software products such as Veeam or CommVault, which might require time synchronization between the hypervisor and the Controller VM to determine which files to back up. Accurate time synchronization between Nutanix clusters paired in Disaster Recovery (DR) configurations also ensures that snapshots do not expire too quickly or too late.

Graphs in the Prism interface rely on CVM time, and incorrect time skews graphs, especially in relation to other monitoring platforms such as vCenter, which rely on other clock sources.

Recommendations for Time Synchronization

Adhere to the following guidelines when configuring time synchronization on a Nutanix cluster:

- Where possible, synchronize Nutanix clusters with internal NTP sources to ensure stability from both a network and a security vulnerability perspective. When you cannot avoid using an external NTP source, prefer a time source maintained by the government.

- Make sure to specify at least five stable time sources that have a high degree of accuracy and that can be reached over a reliable network connection. Generally, the lower the stratum of an NTP source, the higher its accuracy.

Note that three is the minimum to identify one time source as a false ticker but provides no redundancy, four is the minimum for redundancy, and five is the recommended minimum for a good configuration.

If you want to use off-site NTP servers, see <http://support.ntp.org/bin/view/Support/SelectingOffsiteNTPServers> for various recommendations.

- Do not use rate-limited NTP servers.
- Synchronizing a Nutanix cluster with a Windows time source is known to cause issues over a period of time, so Nutanix recommends that you not synchronize a cluster's time with Windows NTP sources. Use reliable non-Windows time sources instead. In an Active Directory domain, the best practice (a design that both works around and improves upon having to include domain controllers in the list of NTP sources) is to bypass the domain controllers and to synchronize the Nutanix hosts and CVMs directly with the NTP sources with which the domain controllers synchronize their time. Specify a common list of at least five reliable non-Windows NTP sources for both the domain controllers and the Nutanix cluster.

Bypassing the domain controller as a time source is not an option for Hyper-V clusters owing to Kerberos requirements. When being joined to a domain, Nutanix clusters running Hyper-V detect local domain controllers and add them to all CVMs as NTP sources. For Hyper-V clusters, supplement the list of detected domain controllers with as many reliable non-Windows NTP sources as are required to meet the recommendation of a minimum of five NTP time sources. For example, if the Nutanix cluster adds two domain controllers as time sources, specify at least three reliable non-Windows NTP sources in the NTP server list. Specifying additional non-Windows NTP sources is necessary even if the domain controllers synchronize their time with a time source that is considered to be reliable.

- Specify public NTP servers by using their FQDN to help mitigate issues caused by IP address changes.

Configuring NTP Servers

About this task

To add (or delete) an NTP server entry, do the following:

Procedure

1. Click the gear icon in the main menu and then select **NTP Servers** in the **Settings** page. The NTP Servers dialog box appears.




Figure 224: NTP Servers Dialog Box

2. To add an NTP server entry, enter the server IP address or fully qualified host name in the **NTP Server** field and then click the **Add** button to the right of that field.
The name or address is added to the **HOST NAME OR IP ADDRESS** list (below the **NTP Server** field).
3. To delete an NTP server entry, click the cross icon for that server in the **Servers** list.
A window prompt appears to verify the action; click the **OK** button. The server is removed from the list.

Configuring an SMTP Server

About this task

Simple Mail Transport Protocol (SMTP) is an Internet standard protocol for electronic mail transmission across Internet Protocol (IP) networks, and Nutanix systems use SMTP to send alert emails and to exchange emails with Nutanix technical support.

If you have an external firewall deployed in your organization, see "" () to allow traffic from the cluster to Nutanix Support servers.

Note: Ensure that the **nos-alerts** and **nos-asup** recipients are added to the accepted domain of your SMTP server.

To configure an SMTP server entry, do the following:

Procedure

1. Click the gear icon in the main menu and then select **SMTP Server** in the **Settings** page.
The SMTP Server Settings dialog box appears.




Figure 225: SMTP Server Settings Window

2. 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).
- c. **Security Mode:** Enter the desired security mode from the pull-down list.
The options are NONE (unencrypted), STARTTLS (use TTL encryption), and SSL (use SSL encryption).
- d. **User:** Enter a user name.
The **User** and **Password** fields apply only when a secure option (STARTTLS or SSL) is selected. The user name might need to include the domain depending on the authentication process.
- e. **Password:** Enter the user password.
- a. **From Email Address** (optional): Enter an e-mail address that appears as the sender address.
By default, alert (see [Configuring Alert Emails](#) on page 335) and cluster status information (see [Configuring Pulse](#) on page 642) 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.

3. When all the fields are correct, click the **Save** button.

Configuring SNMP

About this task

The Simple Network Management Protocol (SNMP) is an application layer protocol that facilitates the exchange of management information between network devices. Nutanix systems include an SNMP agent that provides interoperability with industry standard SNMP manager systems. Nutanix also provides a custom Management Information Base (MIB) for Nutanix-specific information.

Note: Nutanix recommends that you configure your SNMP manager to use the virtual IP address of the cluster to communicate with the Nutanix SNMP agent instead of the Controller VM IP address.

To configure SNMP, do the following:



Procedure

1. Click the gear icon in the main menu and then select **SNMP** in the **Settings** page. The SNMP Configuration dialog box appears.




Figure 226: SNMP Configuration Window

2. To enable SNMP for this cluster, check the **Enable SNMP** box. To disable SNMP, uncheck the box.

Note:

SNMP traps are sent by the Controller VM that functions as the Alert Manager leader. If you need to open your firewall to receive the traps, keep in mind that the Alert Manager leader can rotate during tasks like AOS or host upgrades. Therefore, it might be necessary to open all the Controller VM IP addresses to ensure that the traps are received.

3. To view the Nutanix MIB (**NUTANIX-MIB.txt**), click the **View MIB** link. To download **NUTANIX-MIB.txt**, right-click and select the appropriate download action for your browser and then copy **NUTANIX-MIB.txt** to your SNMP manager systems.
See your SNMP manager documentation for instructions on how to install the Nutanix MIB.
4. To add an SNMP transport, click the **Transports** tab and the **New Transport** button, and then do the following in the indicated fields. An SNMP transport is a combination of the transport protocol and port number on which you want the Nutanix SNMP agent to receive queries. SNMP transports enable you to combine transport protocols and port numbers

other than the default port number. The port numbers that are specified in SNMP transports are unblocked on the Controller VM, making them available to receive queries:



The image shows a web-based configuration dialog for SNMP transports. At the top, there are three tabs: 'Transports' (highlighted with a red box), 'Users', and 'Traps'. Below the tabs, there are two main input fields. The first is labeled 'PROTOCOL' and is a pull-down menu currently showing 'UDP'. The second is labeled 'PORT' and is an empty text input field. At the bottom of the dialog, there are three buttons: 'View MIB' on the left, 'Cancel' in the center, and 'Save' on the right (highlighted with a blue box).

Figure 227: SNMP Transport Fields

- a. **Protocol:** Select the protocol to use from the pull-down list.
The options are **TCP**, **TCP6**, **UDP**, and **UDP6**.
- b. **Port:** Enter the port number to use.
The standard SNMP port number is 161.
- c. When the fields are correct, click the **Save** button (lower right).
This saves the configuration and redisplay the dialog box with the new transport appearing in the list.

Note: To return to the SNMP Configuration window without saving, click the **Cancel** button.

5. To add an SNMP user entry, click the **Users** tab and the **New User** button and then do the following in the indicated fields:



The image shows a web-based configuration interface for adding an SNMP user. At the top, there are three tabs: 'Transports', 'Users', and 'Traps'. The 'Users' tab is highlighted with a red rectangular border. Below the tabs, the form contains the following fields:

- USERNAME**: A text input field.
- PRIV TYPE**: A pull-down menu with 'AES' selected.
- PRIV KEY**: A text input field.
- AUTH TYPE**: A pull-down menu with 'SHA' selected.
- AUTH KEY**: A text input field.

At the bottom of the form, there are three buttons: 'View MIB' (disabled), 'Cancel', and 'Save' (highlighted in blue).

Figure 228: SNMP User Fields

- Username:** Enter a user name.
- Priv Type:** Select the privacy encryption type from the pull-down list.
The only option is **AES** (Advanced Encryption Standard). In the nCLI, this setting is optional.
- Priv Key:** Enter a privacy key phrase (password) into this field.
The key phrase is AES encrypted when the user is created.
- Auth Type:** Select the authentication hash function type from the pull-down list.
The only option is **SHA** (Secure Hash Algorithm).
- Auth Key:** Enter an authentication key phrase (password) into this field.
The key phrase is SHA-1 encrypted when the user is created.
- When all the fields are correct, click the **Save** button (lower right).
This saves the configuration and redisplay the dialog box with the new user entry appearing in the list.

6. To add an SNMP trap receiver, click the **Traps** tab and the **New Trap Receiver** button, and then do the following in the indicated fields:



SNMP Configuration

Configure SNMP for the Nutanix software.

☒ Enable for Nutanix objects

Transports

Users

Traps

RECEIVER NAME

SNMP VERSION

☐ v2c

☒ v3

TRAP USERNAME

No Applicable Option

ADDRESS

PORT

162

ENGINE ID

INFORM

False

TRANSPORT PROTOCOL

UDP



View MIB

Prism | System Management | 604

Cancel

Save

- a. **Receiver Name:** Enter the receiver name.
- b. **SNMP Version:** Select (click the radio button for) the SNMP version, either v3 or v2c. (Version 1 is not supported.)
- c. **Trap Username:** Select a user from the pull-down list.
All users added previously (see step 5) appear in the pull-down list. You cannot add a trap receiver entry until at least one user has been added.
- 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 standard SNMP port number is 161.
- f. **Engine ID:** Optionally, enter an engine identifier value, which must be a hexadecimal string between 5 and 32 characters long.
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 pull-down list 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.

- h. **Transport Protocol:** Select the protocol to use from the pull-down list.
The options are **TCP**, **TCP6**, **UDP**, and **UDP6**.
- i. When all the fields are correct, click the **Save** button (lower right).
This saves the configuration and redisplay the dialog box with the new trap entry appearing in the list.
- j. To test all configured SNMP traps, click the **Traps** tab, and then click **Test All**.



SNMP Configuration

Configure SNMP for the Nutanix software.

☒ Enable for Nutanix objects

Transports · Users · Traps

+ New Trap Receiver

Test all


ADDRESS	PORT	ENGINE ID	USERNAME	VERSION	
192.168.1.100	162	0x8000A12F04a		v2c	 

View MIB

Figure 230: Test All Button

The Nutanix cluster sends test alerts to all the SNMP trap receivers configured on the cluster.

- 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 the **Save** button.

8. To delete an SNMP entry, click the appropriate tab (**Transports**, **Users**, or **Traps**) and then click the X icon for that entry in the list.

A window prompt appears to verify the delete action; click the **OK** button. The entry is removed from the list.

Nutanix MIB

Overview

The Simple Network Management Protocol (SNMP) enables administrators to monitor network-attached devices for conditions that warrant administrative attention. In the Nutanix SNMP implementation, information about entities in the cluster is collected and made available through the Nutanix MIB (**NUTANIX-MIB.txt**). The Nutanix enterprise tree is located at 1.3.6.1.4.1.41263.

The Nutanix MIB is divided into the following sections:

- Cluster information. Status information about the cluster as a whole.

- Software version information. Version information about the software packages that comprise the Controller VM.
- Service status information. Information about the status of essential services on each Controller VM.
- Hypervisor information. Information about each hypervisor instance.
- Virtual machine information. Information about hosted virtual machines.
- Disk information. Status information about the disks in the cluster.
- Controller VM resource information. Indicate how much CPU and memory capacity is available to a Controller VM.
- Storage pool information. Status information about the storage pools in the cluster.
- Storage Container information. Status information about the disks in the cluster.
- Alerts information. Information about generated alerts that can be captured through the SNMP trap (or inform) mechanism.

Table 918: Cluster Information Fields

Name	Description	Data Type
clusterName	Cluster name.	Display string
clusterVersion	Cluster version number. This is the Nutanix core package version expected on all the Controller VMs.	Display string
clusterStatus	Current status of the cluster. Possible values are "started" and "stopped".	Display string
clusterTotalStorageCapacity	Total storage capacity of the cluster in bytes.	Unsigned 64-bit integer
clusterUsedStorageCapacity	Storage used on the cluster, in bytes.	Unsigned 64-bit integer
clusterIops	Average I/O operations per second (IOPS) in the cluster.	Unsigned 64-bit integer
clusterLatency	Average I/O latency in the cluster, in milliseconds.	Unsigned 64-bit integer
clusterIOBandwidth	Cluster-wide I/O bandwidth in kilobytes per second (KBps).	Unsigned 64-bit integer

Table 919: Software Version Information Fields

Name	Description	Data Type
svtIndex	Unique index that is used to identify an entry in the software version information table.	Signed 32-bit integer
svtControllerVMId	Nutanix Controller VM identification number.	Display string



Name	Description	Data Type
svtNutanixBootstrap	Nutanix bootstrap software package version.	Display string
svtNutanixInfrastructure	Nutanix infrastructure software package version.	Display string
svtNutanixCore	Nutanix core software package version.	Display string
svtNutanixToolchain	Nutanix toolchain software package version.	Display string
svtNutanixServiceability	Nutanix serviceability software package version.	Display string
svtLinuxKernel	Linux kernel version currently installed.	Display string

Table 920: Service Status Information Fields

Name	Description	Data Type
cstIndex	Unique index that is used to identify an entry in the service status information table.	Signed 32-bit integer
cstControllerVMId	Nutanix Controller VM identification number.	Display string
cstControllerVMStatus	Status of the Nutanix node.	Display string
cstDataServiceStatus	Status of the core data services on the Controller VM.	Display string
cstMetadataServiceStatus	Status of the metadata services on the Controller VM.	Display string

Table 921: Hypervisor Information Fields

Name	Description	Data Type
hypervisorIndex	Number that is used to uniquely identify an entry in the hypervisor information table.	Signed 32-bit integer
hypervisorID	System-generated string that Nutanix uses to uniquely identify a hypervisor instance.	Display string
hypervisorName	Name of the hypervisor instance.	Display string
hypervisorVmCount	Number of VMs configured on the hypervisor instance.	Unsigned 32-bit integer
hypervisorCpuCount	Number of CPU cores available to the hypervisor instance.	Unsigned 32-bit integer
hypervisorCpuUsagePercent	Percentage of CPU resources in use by the hypervisor instance.	Unsigned 32-bit integer
hypervisorMemory	Total memory available to the hypervisor instance, in bytes.	Unsigned 64-bit integer



Name	Description	Data Type
hypervisorMemoryUsagePercent	Memory in use by the hypervisor instance, as a percentage of the total available memory.	Unsigned 64-bit integer
hypervisorReadIOPerSecond	Total number of read I/O operations per second (IOPS) being performed by the hypervisor.	Unsigned 32-bit integer
hypervisorWriteIOPerSecond	Total number of write I/O operations per second (IOPS) being performed by the hypervisor.	Unsigned 32-bit integer
hypervisorAverageLatency	Average I/O latency of the hypervisor in microseconds (μ s).	Unsigned 64-bit integer
hypervisorIOBandwidth	I/O bandwidth of the hypervisor in kilobytes per second (KBps).	Unsigned 64-bit integer
hypervisorRxBytes	Total number of bytes received by the hypervisor.	Unsigned 64-bit integer
hypervisorTxBytes	Total number of bytes transmitted by the hypervisor.	Unsigned 64-bit integer
hypervisorRxDropCount	Total number of packets dropped by the hypervisor when receiving data.	Unsigned 64-bit integer
hypervisorTxDropCount	Total number of packets dropped by the hypervisor when transmitting data.	Unsigned 64-bit integer

Table 922: Virtual Machine Information Fields

Name	Description	Data Type
vmIndex	Number that is used to uniquely identify an entry in the VM information table.	Signed 32-bit integer
vmId	System-generated string that Nutanix uses to uniquely identify a virtual machine.	Display string
vmName	Name of the VM.	Display string
vmHypervisorId	System-generated ID of the hypervisor on which the VM is provisioned.	Display string
vmPowerState	Power state of the VM.	Display string
vmCpuCount	Number of CPU cores available to the VM.	Unsigned 32-bit integer
vmCpuUsagePercent	Percentage of CPU resources in use by the VM.	Unsigned 32-bit integer
vmMemory	Total memory allocated to the VM, in bytes.	Unsigned 64-bit integer
vmMemoryUsagePercent	Memory in use by the VM, as a percentage of the total allocated memory.	Unsigned 64-bit integer



Name	Description	Data Type
vmReadIOPerSecond	Total number of read I/O operations per second (IOPS) being performed by the VM.	Unsigned 32-bit integer
vmWriteIOPerSecond	Total number of write I/O operations per second (IOPS) being performed by the VM.	Unsigned 32-bit integer
vmAverageLatency	Average I/O latency of the VM, in microseconds (μ s).	Unsigned 64-bit integer
vmIOBandwidth	I/O bandwidth of the VM in kilobytes per second (KBps).	Unsigned 64-bit integer
vmRxBytes	Total number of bytes received by the VM.	Unsigned 64-bit integer
vmTxBytes	Total number of bytes transmitted by the VM.	Unsigned 64-bit integer
vmRxDropCount	Total number of packets dropped by the VM when receiving data.	Unsigned 64-bit integer
vmTxDropCount	Total number of packets dropped by the VM when transmitting data.	Unsigned 64-bit integer

Table 923: Disk Information Fields

Name	Description	Data Type
dstIndex	Number that is used to uniquely identify an entry in the disk information table.	Signed 32-bit integer
dstDiskId	Disk identification number. The number is unique for each disk.	Display string
dstControllerVMId	Nutanix Controller VM identification number.	Display string
dstSerial	Disk serial number.	Display string
dstNumRawBytes	Physical storage capacity on the device, in terms of number of raw bytes.	Unsigned 64-bit integer
dstNumTotalBytes	Usable storage on the device through its file system, in terms of number of usable bytes.	Unsigned 64-bit integer
dstNumFreeBytes	Available storage on the device through its file system for non-root users, in terms of number of free bytes.	Unsigned 64-bit integer
dstNumTotalInodes	Total number of usable inodes on the device through its file system.	Unsigned 64-bit integer
dstNumFreeInodes	Total number of available (free) inodes on the device through its file system for non-root users.	Unsigned 64-bit integer
dstAverageLatency	Average I/O latency of the disk in microseconds (μ s).	Unsigned 64-bit integer



Name	Description	Data Type
dstIOBandwidth	I/O bandwidth of the disk in kilobytes per second (KBps).	Unsigned 64-bit integer
dstNumberIops	Current number of I/O operations per second (IOPS) for the disk.	Unsigned 64-bit integer
dstState	State of the disk.	Signed 32-bit integer

Table 924: Controller VM Resource Information Fields

Name	Description	Data Type
crtIndex	Number that is used to uniquely identify an entry in the Controller VM resource information table.	Signed 32-bit integer
crtControllerVMId	Nutanix Controller VM identification number.	Display string
crtMemory	Total memory available to the Controller VM in bytes.	Unsigned 64-bit integer
crtNumCpus	Total number of CPUs allocated to the Controller VM.	Signed 32-bit integer
crtName	Name of the Controller VM.	Display string

Table 925: Storage Pool Information Fields

Name	Description	Data Type
spitIndex	Number that is used to uniquely identify an entry in the storage pool information table.	Signed 32-bit integer
spitStoragePoolId	Storage pool identification number.	Display string
spitStoragePoolName	Storage pool name.	Display string
spitTotalCapacity	Total storage pool capacity in bytes.	Unsigned 64-bit integer
spitUsedCapacity	Used storage pool capacity in bytes.	Unsigned 64-bit integer
spitIOPerSecond	Current number of I/O operations per second (IOPS) for this storage pool.	Signed 32-bit integer
spitAvgLatencyUsecs	Average I/O latency for the storage pool in microseconds.	Unsigned 64-bit integer
spitIOBandwidth	I/O bandwidth of the storage pool in kilobytes per second (KBps).	Unsigned 64-bit integer

Table 926: Storage Container Information Fields

Name	Description	Data Type
citIndex	Number that is used to uniquely identify an entry in the storage container information table.	Signed 32-bit integer
citContainerId	Storage Container identification number.	Display string
citContainerName	Storage Container name.	Display string
citTotalCapacity	Total storage container capacity in bytes.	Unsigned 64-bit integer
citUsedCapacity	Used storage container storage in bytes.	Unsigned 64-bit integer
citIOPerSecond	Current number of I/O operations per second (IOPS) for this storage container.	Signed 32-bit integer
citAvgLatencyUsecs	Average I/O latency for the storage container in microseconds.	Unsigned 64-bit integer
citIOBandwidth	I/O bandwidth of the storage container in kilobytes per second (KBps).	Unsigned 64-bit integer

Trap Resolution

In addition to generating an SNMP trap when an alert condition is detected, Nutanix clusters generate a trap when an alert condition is resolved. The resolved trap, named `ntxTrapResolved`, is generated regardless of whether the alert condition is resolved manually or automatically. If a Prism Central alert is resolved in Prism Central, Prism Central sends a resolved trap. If a Nutanix cluster alert is resolved in either the Prism web console or Prism Central, both Prism Central and the cluster send a resolved trap.

In this section, the terms original alert and original trap are used to refer, respectively, to the alert and trap that are generated when the alert condition is detected, and the terms resolved alert and resolved trap are used to refer, respectively, to the alert and trap that are generated when the alert condition is resolved.

To enable you to associate a resolved trap with the original alert and original trap, the UUID of the original alert is included in the original trap, as part of the alert message (`ntxAlertDisplayMsg`), and in the resolved trap, as a MIB object (`ntxAlertUuid`).

A resolved trap has the following MIB objects:

- `ntxTrapName` (name of the original trap)
- `ntxAlertUuid` (UUID of the original alert)
- `ntxAlertResolvedTime`
- `ntxAlertDisplayMsg` (message of the original alert)
- `ntxAlertTitle` (title of the original alert)
- `ntxAlertSeverity` (severity of the original alert)

The fields are described in the rows that follow.



Table 927: Alerts Information Fields

Name	Description	Data Type
ntxAlertCreationTime	Time of alert creation. The value is the number of seconds since the UNIX epoch (01/01/1970).	Unsigned 64-bit integer
ntxAlertDisplayMsg	Alert message text.	Display string
ntxAlertTitle	Alert title.	Display string
ntxAlertSeverity	Alert severity, which has one of the following values: 1. Informational 2. Warning 3. Critical 4. Audit	Integer
ntxTrapName	Name of the trap that was generated when the alert condition was detected. This MIB object is included in resolved traps.	Display string
ntxAlertUuid	UUID of the alert that was generated when the alert condition was detected. This MIB object is included in resolved traps.	Display string
ntxAlertResolvedTime	Time at which the alert was resolved. The value is the number of seconds since the UNIX epoch (01/01/1970). This MIB object is included in resolved traps.	Unsigned 64-bit integer

Configuring a Banner Page

You have the option to create a welcome banner, which will be the first screen that appears when a user attempts to log into the web console. The content of the banner page is configurable, and it can include a custom message and graphics.

About this task

To configure a banner page, do the following:

Procedure

1. Click the gear icon in the main menu and then select **Welcome Banner** in the **Settings** page. The Edit Welcome Banner dialog box appears.
2. Enter (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.
3. Click the **Preview** button to display the banner in the pane on the right.
4. If the banner is not correct, update the HTML code as needed until the preview pane displays the desired message.



- When the preview is correct, check the **Enable Banner** box (lower left) and then the **Save** button.

You can disable (or enable) the banner at any time by unchecking (checking) the **Enable Banner** box.

Example

The following figure is a sample banner page configuration.

Note: A live banner page includes an "Accept terms and conditions" bar at the bottom. Clicking on this bar sends the user to the login page.




Figure 231: Welcome Banner Window

VM Management through Prism Element (ESXi)

From AOS 5.0 or later releases, you can perform your core VM management operations directly from Prism without using any hypervisor management interface (for example, vCenter Server). The VM Management through Prism for ESXi feature provides a unified management interface for all of the ESXi hypervisors. For this functionality to work, you need to register vCenter Server with the Prism Element or multiple vCenter Servers with the Prism Central. For more information about registering vCenter Server to your cluster, see [Registering a vCenter Server](#) on page 617.

Note:

By using this feature you can perform following operations directly through Prism.

- Create, clone, update, and delete VMs.
- Create and delete NICs.
- Attach and delete disks.
- Power operations: Power on or off, reset, suspend, resume, guest shutdown, and guest restart.
- Open and launch VM console.
- Manage VM guest tools (mounting VMware guest tools, mounting NGT).

Note:

- You can perform the power operations and launching of VM console even when 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 is managed by a single vCenter Server.
- Ensure that DRS is enabled on the vCenter Server.
- Ensure that you are running ESXi and vCenter Server 5.5 or later releases.
- Ensure that you have homogeneous network configuration. For example, network should have either 1G or 10G NICs.
- Ensure that you unregister the vCenter Server from the cluster before changing the IP address of the vCenter Server. After you change the IP address of the vCenter Server, you must register the vCenter Server again with the new IP address.
- The vCenter Server 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 vCenter Server. In this case, there will be new vCenter entry with host connection status as Connected and you need to register to this vCenter Server. For more information about registering vCenter Server again, see [Managing vCenter Server Registration Changes](#) on page 619.

CAUTION: If multiple vCenter Servers are managing the hosts, you will not be able to perform the VM management operations. Move all the hosts into one vCenter Server.

Requirements and Limitations

- SCSI, IDE, and SATA disks are supported. PCI disks are not supported.
- The E1000, E1000e, PCnet32, VMXNET, VMXNET 2, VMXNET 3 network adapter types (NICs) are 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 gets deleted.
- Network configuration (creation of port groups or VLANs) is not supported.

Registering a vCenter Server

To perform core VM management operations directly from Prism without switching to vCenter Server, you need to register your cluster with the vCenter Server.

Before you begin

Ensure that you have vCenter Server Extension privileges as these privileges provide permissions to perform vCenter registration for the Nutanix cluster.

About this task

Following are some of the important points about registering vCenter Server.



- Nutanix does not store vCenter Server credentials.
- Whenever a new node is added to cluster, vCenter Sever registration for the new node is automatically performed.

Procedure

1. Log into the Prism web console.
2. Click the gear icon in the main menu and then select **vCenter Registration** in the **Settings** page.
The vCenter Server that is managing the hosts in the cluster is auto-discovered and displayed.
3. Click the **Register** link.
The IP address is auto-populated in the **Address** field. The port number field is also auto-populated with 443. Do not change the port number.
4. Type the administrator user name and password of the vCenter Server in the **Admin Username** and **Admin Password** fields.


A screenshot of the 'vCenter Registration' dialog box. The dialog has a black header bar with the title 'vCenter Registration' on the left, and a question mark icon and a close (X) icon on the right. Below the header, there are four input fields: 'IP ADDRESS' with the value '10.4.234.1', 'PORT' with the value '443', 'ADMIN USERNAME' (empty), and 'ADMIN PASSWORD' (empty). At the bottom right of the dialog, there are two buttons: 'Cancel' and 'Register'.

Figure 232: vCenter Registration Figure 1

5. Click **Register**.
During the registration process a certificate is generated to communicate with the vCenter Server. If the registration is successful, relevant message is displayed in the **Tasks** dashboard.

The Host Connection field displays as Connected, which implies that all the hosts are being managed by the vCenter Server that is registered.




Figure 233: vCenter Registration Figure 2

Unregistering a Cluster from the vCenter Server

To unregister the vCenter Server from your cluster, perform the following procedure.

About this task

- Ensure that you unregister the vCenter Server from the cluster before changing the IP address of the vCenter Server. After you change the IP address of the vCenter Server, you should register the vCenter Server again with the new IP address with the cluster.
- The vCenter Server 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 vCenter Server. In this case, there will be new vCenter entry with host connection status as Connected and you need to register to this vCenter Server.

Procedure

1. Log into the Prism web console.
2. Click the gear icon in the main menu and then select **vCenter Registration** in the **Settings** page.
A message that cluster is already registered to the vCenter Server is displayed.
3. Type the administrator user name and password of the vCenter Server in the **Admin Username** and **Admin Password** fields.
4. Click **Unregister**.

If the credentials are correct, the vCenter Server is unregistered from the cluster and a relevant message is displayed in the **Tasks** dashboard.

Managing vCenter Server Registration Changes

After you change the IP address of the vCenter Server or if current vCenter Server is not managing the hosts then the Host Connection field changes to **Not Connected** (may occur if current vCenter Server is not managing the hosts), you must register your cluster with the new vCenter Server. Perform the following procedure to re-register the vCenter Server.

Procedure

1. Log into the Prism web console.
2. Click the gear icon in the main menu and then select **vCenter Registration** in the **Settings** page.

The vCenter Server that is managing the hosts in the cluster is displayed along with the vCenter Server that you must unregister.




Figure 234: vCenter Registration Management

3. Unregister the existing vCenter Server instance by clicking **Unregister**.
For more information about unregistering vCenter Server, click [Unregistering a Cluster from the vCenter Server](#) on page 619.
4. Register the new vCenter Server instance by clicking **Register**.
For more information about registering vCenter Server, click [Registering a vCenter Server](#) on page 617.

In-Place Hypervisor Conversion

The In-place hypervisor conversion feature provides you with an option to convert your existing ESXi cluster to an AHV cluster. All the VMs that are running in the ESXi cluster are converted so that they can run on the AHV cluster. The VM conversion is independent of the host that the VM is running on. The conversion process occurs in such a way that the VM down time is minimised. You can also convert the cluster back to ESXi. This conversion from AHV to ESXi is a file system restore of the host. Hence the cluster will be similar to what it was before the conversion.

Note:

- This feature converts your existing ESXi cluster to an AHV cluster. You cannot start the conversion process on the AHV cluster.
- Do not remove the hosts from the vCenter Server if you want to perform the reverse conversion process (AHV to ESXi).

- This feature is supported on the clusters with multiple hypervisors (combination of ESXi nodes with one AHV node or multiple AHV nodes).
- When you are converting the cluster from ESXi to AHV, remove the unused adapters displayed under **Teaming and Failover** in vCenter Server and keep only those adapters that are active and of the same speeds as the standby and active adapters.

Following are enhancements to In-Place Hypervisor Conversion feature.

- Decreased VM downtime: With the implementation of the new workflow of converting the nodes in the cluster in a rolling manner, the VM downtime is reduced to only shut down time and conversion time required for that particular VM. Approximately, the VM downtime has been reduced from 3 to 4 hours to less than 5 minutes.
- Prism state: The Prism console is responsive during the conversion process. However, the Prism goes into Read-Only state.
- State of the VM: The current state of the VM is preserved and the VM is brought back into the same state post conversion. For example, the VM is automatically powered on post conversion if that VM was powered on before you started the conversion.
- Preservation of the MAC addresses of the VM NICs after conversion: After the cluster is converted to AHV or ESXi, MAC addresses of the VM NICs are preserved. As a result of the MAC address preservation, both the static and dynamically configured IP addresses are preserved after conversion.

Note: IP addresses are not preserved for Windows VMs.

- VMs that were created after conversion to AHV are retained post conversion back to ESXi.

Requirements and Limitations for In-Place Hypervisor Conversion

Following are the general prerequisites along with requirements and limitations for the in-place hypervisor conversion feature.

Prerequisites

- HA and DRS must be enabled.
- All the hosts should be managed by the same vCenter Server.
- The vCenter Server VM that is managing the cluster must not be on the same cluster which is being converted.
- Install NGT on all the VMs. See [Enabling NGT and Mounting the NGT Installer in a VM](#) on page 299 for more information.
- All the platforms are supported.
- Ensure that you have saved all the ESXi ISOs (if you have more than one version of ESXi running in your cluster) at the **foundation/isos/hypervisor/esx/** location. This ISO image or images are used to bring the cluster back to its pre-conversion state if you abort the cluster conversion operation from ESXi to AHV.

General Limitations

- This feature is not supported if you have metro availability protection domains in your environment.



- Conversion back from ESXi to AHV fails if you have converted the cluster using free ESXi license.
- Cluster conversion does not proceed if you have Nutanix Files deployed in your cluster.
- Affinity policies (if any configured) are not honored post conversion to AHV.
- This feature is not supported for single-node clusters and two-node clusters.

ESXi Supported Versions

Table 928: Supported ESXi Versions

Version	Support
5.5	All the versions that are included in the ISO Whitelist (see <code>~/foundation/config/iso_whitelist.json</code> file) that is bundled with AOS are supported. If you have upgraded foundation by using standalone foundation, the ISO Whitelist is invalidated. In this case, you need to see the original Whitelist that was present before you performed the upgrade. To verify the json file that was bundled with AOS, see the <code>nutanix-packages.json</code> file.
6.0	
6.5	
6.7	

ESXi Network Requirements and Limitations

Table 929: ESXi Network Requirements and Limitations

Component	Description
NIC	Minimum of two or more homogenous NICs of 1 GB, 10 GB, or 40 GB are supported.
vSwitch	• Each host should have only one external vSwitch. If you have more than one external vSwitch conversion validation fails. • Ensure you have one active and multiple passive failover configuration of NICs on vSwitch. • Active/active load balancing policies are not supported and is converted to active/passive on AHV. • For a standard switch configuration it is recommended that all the port groups are present on all the hosts because there might be a possibility that the conversion process might move the VMs to a different host.
Distributed vSwitch	Supported
Internal vSwitch	Internal vSwitch apart from Nutanix vSwitch is not supported.
LACP	Not supported.



- The vSwitch is converted to an open vSwitch on the AHV side. If you perform any configuration changes on the AHV, these changes may not be maintained after converting the cluster back to ESXi.
- Networks with no VMs prior to the cluster conversion may not be converted.
- All the VMs may end up on a single host during reverse conversion. The network configuration that you might have defined for the VMs may not be maintained.
- MAC address of the VMs are not preserved because the MAC address is provided by the hypervisor (ESXi and AHV) and both the hypervisors follow different conventions.
- After migration to AHV, NIC adapter may not have same type of virtualized hardware device (for example VMXNET3 may be E1000 after conversion).
- Serial ports and virtual graphics processing unit (vGPU) configurations may be lost after conversion.

Virtual Machine Requirements and Limitations

- Only VMs with flat disks are supported. The delta disks are not supported.
- IDE/SCSI disks are only supported. SATA and PCI disks are not supported.
- Set the SAN policy to OnlineAll for all the Windows VMs for all the non-boot SCSI disks so that they can be automatically brought online. For more information about setting SAN policy, see “” () .
- Folders with NFS mounts and VMware tags in vSphere for categorisation may not be supported.
- Ensure that you do not have VMs with UEFI boot firmware enabled.
- Virtual machines with attached volume groups or shared virtual disks are not supported.
- After conversion of the cluster back to ESXi from AHV, the VMs are converted to the maximum hardware version that is supported by that specific ESXi version.
- Guest OS network interface on Linux VMs may change to generic type during ESXi to AHV conversion (for example, in RHEL 7 network interface enoXXXX changes to eth0). You have to reconfigure the network settings according to changes post conversion.
- Guest OS type for the Linux VMs may change to a more generic type (for example RHEL 7 may change to Other Linux 64-bit) during the conversion back from AHV to ESXi.

Supported Operating Systems

Table 930: Supported Operating Systems

Operating System	Versions	Limitations
Windows	• Windows 2008 R2 or later versions • Windows 7 or later versions	• Only 64-bit operating system is supported.



Operating System	Versions	Limitations
Linux	- CentOS 6.5 and 7.0 - Red Hat Enterprise Linux (RHEL) 6.5 or later and RHEL 7.0 or later. - Oracle Linux 6.5 and 7.0 - Ubuntu 14.04	SUSE Linux Enterprise Server (SLES) operating system is not supported.

In-Place Hypervisor Conversion Process

ESXi to AHV

After you start the conversion process, all the nodes in the cluster are converted in a rolling manner to AHV. The conversion happens one node at a time instead of happening on all the nodes at the same time. Once you start the conversion process, the first node is placed in the maintenance mode and all the VMs that are running on the node is migrated to other ESXi nodes in the cluster by using the HA and DRS feature. After the migration process is completed, the node is converted to AHV. After node is successfully converted, all VMs that were migrated to ESXi are converted and placed one at a time on AHV. Once the first node is converted to AHV, similar steps are performed for the rest of the nodes in the ESXi cluster.

Note: The VMs that are running on the last ESXi node in the cluster are converted and migrated to the available AHV hosts and then the ESXi host goes into conversion.

If any error occurs in the VM conversion, appropriate alerts or error messages are displayed. The source VM disk of the VM is not modified. Any errors during imaging, storing, or restoring of configuration on the host are fatal errors and conversion will be stopped and you will be prompted to abort the conversion.

Note: After conversion, do not remove the host from the vCenter Server because it may impact the reverse conversion process.

Port Groups and VLAN Transformation

All the ESXi port groups, any VLAN IDs, and virtual machines that belong to a particular group are captured, a corresponding Acropolis virtual network for every unique ESXi port group on the cluster is created, a corresponding VLAN ID is assigned, and then the VMs are transferred into the right virtual network. If the ESXi host has a VLAN set on the management port group, after conversion the Acropolis management interface and the Controller VM public interface are placed in that same VLAN.

AHV to ESXi

During the reverse conversion (AHV to ESXi), the process of conversion is similar. Additionally, if the cluster does not have the ESXi ISO stored on the cluster, you need to provide the ESXi ISO image during the conversion process.

Note:

- The image that you provide should be of the same major ESXi version that you have used during ESXi to AHV conversion.

- If new nodes were added to the cluster after conversion to AHV, then these nodes need to be removed first before starting the reverse conversion process.

After conversion to ESXi, all the hosts are automatically registered to the vCenter Server.

Converting Cluster (ESXi to AHV)

Perform the following procedure to convert the cluster from ESXi to AHV.


Before you begin

Verify that you have met all the networking and virtual machine requirements as described in [Requirements and Limitations for In-Place Hypervisor Conversion](#) on page 621.

About this task

Procedure

1. Log into the web console.
2. Click the gear icon in the main menu and then select **Convert Cluster** in the **Settings** page.
3. From the **Available Hypervisor** drop-down menu, select **AHV**.



Convert Cluster ? X

To provide a smooth hypervisor conversion experience you need to fulfill the following [requirements](#).

AVAILABLE HYPERVISORS

AHV Validate

VM BOOT OPTIONS

Preserve power state of user VMs

Cancel Convert Cluster

Figure 235: Convert Cluster

4. Select the state of the VMs that you want post conversion from the **VM Boot Options** drop-down menu.
 - » **Preserve power state of the user VMs:** Select this option if you want to keep the original power state of the VMs. For example, if you want the VMs to be in a running state post conversion automatically, select this option.
 - » **Power Off User VMs:** Select this option if you want power off all the VMs running on ESXi cluster before you start the conversion process. After conversion, these VMs will be in the powered off state.
5. Click **Validate** to enter vCenter Server credentials and to verify whether you have met all the requirements.
6. Enter the IP address of the vCenter Server in the **vCenter Sever IP Address** along with the administrator user name and password of the vCenter Server in the **Username** and **Password** fields.
7. Click **Yes**.

A validation that you have met all the requirements is performed. Once the validation is successful the conversion process proceeds. If validation fails (for any reason), a relevant message to take appropriate action is displayed.
8. Click **Convert Cluster**.

A confirmation message is displayed.

Note: The cluster changes to Read Only mode. A message **Oops - server Error** may be displayed when the Prism node is undergoing conversion. Wait for the conversion process to complete. You can access other Controller VM for Read-Only Prism operations.

The entire conversion process may take 3 to 4 hours depending on the nodes that are present in your cluster. However, the VM downtime will be less than 5 minutes because all the nodes in the cluster are converted in a rolling manner. You can also track the progress of the conversion by logging again into the web console.

Converting Cluster (AHV to ESXi)

Perform the following procedure to convert the cluster back to ESXi from AHV.

About this task

Perform this procedure only if you have converted your cluster from ESXi to AHV. You cannot start the conversion process on the AHV cluster.

Before you begin

Verify that you have met all the networking and virtual machine requirements as described in [Requirements and Limitations for In-Place Hypervisor Conversion](#) on page 621 topic.

Procedure

1. Log into the web console.
2. Click the gear icon in the main menu and then select **Convert Cluster** in the **Settings** page.
3. From the **Available Hypervisor** drop-down menu, select **ESXi**.



4. Select the state of the VMs that you want post conversion from the **VM Boot Options** drop-down menu.
 - » **Preserve power state of the user VMs:** Select this option if you want to keep the original power state of the VMs. For example, if you want the VMs to be in a running state post conversion automatically, select this option.
 - » **Power Off User VMs:** Select this option if you want power off all the VMs running on ESXi cluster before you start the conversion process. After conversion, these VMs will be in the powered off state.
5. Click **Validate** to enter vCenter Server credentials and to verify whether you have met all the requirements.
6. Enter the IP address of the vCenter Server in the **vCenter Sever IP Address** along with the administrator user name and password of the vCenter Server in the **Username** and **Password** fields.
7. Click **Yes**.

A validation that you have met all the requirements is performed. Once the validation is successful the conversion process proceeds. If validation fails (for any reason), a relevant message to take appropriate action is displayed.
8. Click **Convert Cluster** button to start the conversion process.
9. (Optional) If you have not saved the ESXi ISOs at the **foundation/isos/hypervisor/esx/** location, click **Choose File** and select the ESXi ISO.




Figure 236: Software Upload

Note: If you have different versions of ESXi running in your cluster, you have to perform this step for every version of ESXi ISO.

10. Click **Convert**.

All cluster operations stops and you will not be able to perform any operations on the cluster while the conversion is in progress. A message to this effect is displayed.

11. Click **Yes** to proceed with the conversion.

The conversion begins. The time taken to finish the conversion is dependent on your environment. For example, it might take between 2 to 3 hours for the conversion to finish for a 3 or 4 nodes cluster. However, the VM downtime will be less than 5 minutes.

Aborting Cluster Conversion

If any issues occur with the cluster conversion process, a message to abort the cluster conversion or to continue with the conversion process is displayed. Depending on the type of message, you can appropriate action.

Before you begin

Ensure that you have saved the ESXi ISO at the **foundation/isos/hypervisor/esx/** location. This ISO image is used to bring the cluster back to its pre-conversion state if you abort the cluster conversion operation from ESXi to AHV.

About this task

Issues in the cluster conversion process might occur because of the following reasons.

- Imaging Issue: If this issue occurs, you can only abort the cluster conversion process and a relevant message is displayed. Aborting the conversion reverts the cluster back to the its original state.
- VM Conversion Issue: If this issue occurs, you can either abort the conversion or can continue with the process. If you decide to continue with the process, cluster conversion is completed keeping the current state of the VM. After conversion is completed, you must perform appropriate actions on the VMs to bring back the VMs.

Procedure

Click **Abort** when a message to abort the cluster conversion is displayed.




Figure 237: Abort Operation

The cluster goes back to its original state. You may need to manually power on the VMs depending on its state after you perform this operation.

Internationalization (i18n)

The following table lists all the supported and unsupported entities in UTF-8 encoding.

Table 931: Internationalization Support

Supported Entities	Unsupported Entities
Cluster name	Acropolis file server
Storage Container name	Share path
Storage pool	Internationalized domain names
VM name	E-mail IDs
Snapshot name	Hostnames
Volume group name	Integers
Protection domain name	Password fields
Remote site name	Any Hardware related names (for example, vSwitch, iSCSI initiator, vLAN name)
User management	
Chart name	

CAUTION: The creation of none of the above entities are supported on Hyper-V because of the DR limitations.

Entities Support (ASCII or non-ASCII) for the Active Directory Server

- In the New Directory Configuration, **Name** field is supported in non-ASCII.
- In the New Directory Configuration, **Domain** field is not supported in non-ASCII.
- In Role mapping, **Values** field is supported in non-ASCII.
- User names and group names are supported in non-ASCII.

Localization (L10n)

Nutanix localizes the user interface in Simplified Chinese and Japanese language. All the static screens are translated to the selected locale language. All the dashboards (including tool tips) and menus of the Prism Element are localized.

You have an option to change the language settings of the cluster from English (default) to Simplified Chinese or Japanese. For more information, see [Changing the Language Settings](#) on page 630.

If the Prism Central instance is launched from the Prism Element, language settings of the Prism Central takes precedence over Prism Element.

You can also create new users with the specified language setting. For more information, see [Creating New Users](#).



Guidelines and Limitations

- Logical entities that do not have a contextual translation available in the localized language are not localized.
- 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.

Changing the Language Settings

Perform the following procedure to set the Prism language settings to a chosen language (English, simplified Chinese, or Japanese). You also have an option to set the calendar, date, and time format to a selected region.

Procedure

1. Click the gear icon in the main menu and then select **Language Settings** in the **Settings** page.
The Language Settings dialog box appears.




Figure 238: Language Settings Window

2. To change the language, select the desired language from the **Languages** field pull-down menu.

The **English** language is selected by default, but you can change that to either **Simplified Chinese** or **Japanese**.

3. To change the locale settings (date, time, calendar), select the appropriate region from the **Region** field drop-down menu.

A default locale is set based on the language setting. However, you can change the region to display the date, time, and calendar in some other format. This format for date, time, and calendar is applied to the entire cluster.

4. When the settings are correct, click the **Save** button.

The language and locale settings are changed according to the selection. For example, in the following screenshot, the language setting for the cluster is changed to Chinese and the locale setting is changed to Russian.

Language Settings ?

Language

简体中文 (Simplified Chinese)

Region

Russia

Long Time	22 декабря 2015 г., 16:00:00
Short Time	22.12.15, 16:00:00
Number	1 234 567,89

Discard Changes Save

Figure 239: Locale Settings

For more information on the entities that are supported, see [Internationalization \(i18n\)](#) on page 629. For more information about localization, see [Localization \(L10n\)](#) on page 629.

Hyper-V Setup

Joining the Cluster and Hosts to a Domain

After completing foundation of the cluster, you need to join the cluster and its constituent hosts to the Active Directory domain. The joining of cluster and hosts to the domain facilitates centralized administration and security through the use of other Microsoft services such as Group Policy and enables administrators to manage the distribution of updates and hotfixes.

Before you begin

- If you have a VLAN segmented network, verify that you have assigned the VLAN tags to the Hyper-V hosts and Controller VMs. To configure VLANs for the Controller VM, see the *Advanced Setup Guide*.
- Ensure that you have valid credentials of the domain account that has the privileges to create a new or modify an existing computer accounts in the Active Directory domain. Active Directory domain created by using non-ASCII text may not be supported. For more

information about usage of ASCII or non-ASCII text in Active Directory configuration, see *Internationalization (i18n)* section.

Procedure

1. Log into the Web console by using one of the Controller VM IP address or by using cluster virtual IP address.
2. Click the gear icon in the main menu and select **Join Cluster and Hosts to the Domain** from the **Settings** page.

Join Cluster and Hosts to the Domain ? X

Here you can join the cluster and its constituent hosts to an Active Directory domain. This step requires a reboot.

FULL DOMAIN NAME

GREAT.COMPANY.COM

NAME SERVER IP ADDRESS

CLUSTER NAME

allgood-smb

CLUSTER VIRTUAL IP ADDRESS

10.66.22.33

Host details
All hosts can be renamed with the prefix. Hosts can also be individually renamed in the Advanced View.

PREFIX

Advanced View

Credentials ⓘ
Credentials will not be stored.

LOGON NAME (DOMAIN\USERNAME)

PASSWORD

Cancel Join

Figure 240: Join Cluster and Hosts to the Domain

3. Type the fully qualified name of the domain that you want to join the cluster and its constituent hosts to in the **Full Domain Name** text box.
4. Type the IP address of the name server in the **Name Server IP Address** text box that can resolve the domain name that you have entered in the **Full Domain Name** text box.

5. Type a name for the cluster in the **Cluster Name** text box.
The maximum length of the cluster name should not be more than 15 characters and it should be a valid NetBIOS name.
6. Type the virtual IP address of the cluster in the **Cluster Virtual IP Address** text box.
If you have not already configured the virtual IP address of the cluster, you can configure it by using this field.
7. Type the prefix that should be used to name the hosts (according to your convention) in the **Prefix** text box.
 - The prefix name should not end with a period.
 - The maximum length of the prefix name should not be more than 11 characters.
 - Should be a valid NetBIOS name.

For example, if you enter prefix name as Tulip, the hosts are named as Tulip-1, Tulip-2, and so on, in the increasing order of the external IP address of the hosts.

If you do not provide any prefix, the default name of NTN**X-block-number** is used. Click the **Advanced** link to see the expanded view of all the hosts in all the blocks of the cluster and to rename them individually.
8. In the **Credentials** field, type the logon name and password of the domain account that has the privileges to create a new or modify an existing computer accounts in the Active Directory domain.
This logon name must be in the format **DOMAIN\USERNAME**. The credentials are required to join the cluster and its constituent hosts to the domain. Nutanix does not store the credentials.
9. When all the information is correct, click **Join**.
The cluster is joined to the domain. All the hosts are renamed and joined to the domain and are restarted. Wait for about 5 to 10 minutes before proceeding to allow the hosts and Controller VMs time to start. After the cluster is ready, the logon page is displayed.

What to do next

Create a Microsoft failover cluster. For more information, see [Creating a Failover Cluster for Hyper-V](#) on page 633.

Creating a Failover Cluster for Hyper-V

Before you begin

Perform the following tasks before you create a failover cluster:

- Join the hosts to the domain as described in "Joining the Cluster and Hosts to a Domain" in the *Hyper-V Administration for Acropolis Guide* or the *Prism Web Console Guide*.
- Ensure that the Windows failover cluster IP address and hypervisor or host IP address are in the same subnet.

Perform the following procedure to create a failover cluster that includes all the hosts in the cluster.

Procedure

1. Log on to the Prism Element web console by using one of the Controller VM IP addresses or by using the cluster virtual IP address.



- Click the gear icon in the main menu and select **Configure Failover Cluster** from the **Settings** page.

Configure Failover Cluster

This step allows you to setup a Hyper-V failover cluster.

FAILOVER CLUSTER NAME
allgood.great.company.com

FAILOVER CLUSTER IP ADDRESS
10.66.22.34

CREDENTIALS (i)
Credentials will not be stored.

LOGON NAME (DOMAIN\USERNAME)

PASSWORD

Cancel Create Cluster

Figure 241: Configure Failover Cluster

- Type the failover cluster name in the **Failover Cluster Name** text box.
The maximum length of the failover cluster name must not be more than 15 characters and must be a valid NetBIOS name.
- Type an IP address for the Hyper-V failover cluster in the **Failover Cluster IP Address** text box.
This address is for the cluster of Hyper-V hosts that are currently being configured. It must be unique, different from the cluster virtual IP address and from all other IP addresses assigned to the hosts and Controller VMs. It must be in the same network range as the Hyper-V hosts.
- In the **Credentials** field, type the logon name and password of the domain account that has the privileges to create a new account or modify existing accounts in the Active Directory domain.
The logon name must be in the format **DOMAIN\USERNAME**. The credentials are required to create a failover cluster. Nutanix does not store the credentials.
- Click **Create Cluster**.
A failover cluster is created by the name that has been provided and it includes all the hosts in the cluster.

Enabling Kerberos for Hyper-V

If you are running Windows Server 2012 R2, perform the following procedure to configure Kerberos to secure the storage. You do not have to perform this procedure for Windows Server 2016 because Kerberos is enabled automatically during failover cluster creation.

Before you begin

- Join the hosts to the domain as described in [Joining the Cluster and Hosts to a Domain](#) on page 631.
- Verify that you have configured a service account for delegation. For more information on enabling delegation, see the *Microsoft documentation*.

Procedure

1. Log on to the web console by using one of the Controller VM IP addresses or by using the cluster virtual IP address.
2. Click the gear icon in the main menu and select **Kerberos Management** from the **Settings** page.




Figure 242: Configure Failover Cluster

3. Set the **Kerberos Required** option to enabled.
4. In the **Credentials** field, type the logon name and password of the domain account that has the privileges to create and modify the virtual computer object representing the cluster in Active Directory. The credentials are required for enabling Kerberos.
The logon name must be in the format **DOMAIN\USERNAME**. Nutanix does not store the credentials.
5. Click **Save**.

Configuring the Hyper-V Computer Object by Using Kerberos

About this task

Perform the following procedure to complete the configuration of the Hyper-V Computer Object by using Kerberos and SMB signing (for enhanced security).

Note: Nutanix recommends you to configure Kerberos during a maintenance window to ensure cluster stability and prevent loss of storage access for user VMs.

Procedure

1. Log on to Domain Controller and perform the following for each Hyper-V host computer object.
 - a. Right-click the host object, and go to **Properties**. In the **Delegation** tab, select the **Trust this computer for delegation to specified services only** option, and select **Use any authentication protocol**.
 - b. Click **Add** to add the **cifs** of the Nutanix storage cluster object.





Figure 243: Adding the cifs of the Nutanix storage cluster object

2. Check the Service Principal Name (SPN) of the Nutanix storage cluster object.

```
> Setspn -l name_of_cluster_object
```

Replace *name_of_cluster_object* with the name of the Nutanix storage cluster object.

Output similar to the following is displayed.



```
Administrator: Windows PowerShell
PS C:\Users\Administrator>
PS C:\Users\Administrator>
PS C:\Users\Administrator>
PS C:\Users\Administrator>
PS C:\Users\Administrator>
PS C:\Users\Administrator>
PS C:\Users\Administrator>
PS C:\Users\Administrator> setspn -l virat
Registered ServicePrincipalNames for CN=VIRAT,CN=Computers,DC=hyperv16,DC=qa,DC=com:
cifs/VIRAT.HYPERV16.QA.COM
cifs/VIRAT
PS C:\Users\Administrator>
```

Figure 244: SPN Registration

If the SPN is not registered for the Nutanix storage cluster object, create the SPN by running the following commands.

```
> Setspn -S cifs/name_of_cluster_object name_of_cluster_object
> Setspn -S cifs/FQDN_of_the_cluster_object name_of_cluster_object
```

Replace *name_of_cluster_object* with the name of the Nutanix storage cluster object and *FQDN_of_the_cluster_object* with the domain name of the Nutanix storage cluster object.

Example

```
> Setspn -S cifs/virat virat
> Setspn -S cifs/virat.sre.local virat
```

3. [Optional] To enable SMB signing feature, log on to each Hyper-V host by using RDP and run the following PowerShell command to change the **Require Security Signature** setting to **True**.

```
> Set-SMBClientConfiguration -RequireSecuritySignature $True -Force
```

CAUTION: The SMB server will only communicate with an SMB client that can perform SMB packet signing, therefore if you decide to enable the SMB signing feature, it must be enabled for all the Hyper-V hosts in the cluster.

SECURITY AND USER MANAGEMENT

Prism provides several mechanisms to maintain security in a cluster and control user access.

Note: The following links are to topics in the "Security Management using Prism Element (PE)" chapter in the *Nutanix Security Guide*.

- Set the user authentication method (see [""](#) ())
- Add, edit, or delete local user accounts (see [""](#) ())
- Install or replace an SSL certificate (see [""](#) ())
- Control SSH access to the cluster (see [""](#) ())
- Enable data-at-rest encryption (see [""](#) ())
- Enable network segmentation (see [""](#) ())
- Review authentication best practices (see [""](#) ()) and firewall requirements (see [""](#) ())

SUPPORT SERVICES

Nutanix provides support services in several ways.

- Nutanix technical support can monitor your clusters and provide assistance when problems occur (see [Controlling Remote Connections](#) on page 645, [Configuring HTTP Proxy](#) on page 646, and [Configuring Pulse](#) on page 642).
- Nutanix technical support maintains a portal that you can access to request assistance, download AOS updates, or view documentation (see [Accessing the Nutanix Support Portal](#) on page 648).
- Nutanix supports a REST API that allows you to request information or run administration scripts for a Nutanix cluster (see [Accessing the REST API Explorer](#) on page 652).

Pulse Health Monitoring

The feature known as Pulse provides diagnostic system data to Nutanix Support to deliver proactive, context-aware support for Nutanix solutions. Pulse unobtrusively collects this information with no effect on system performance. Pulse shares only basic system-level information necessary for monitoring the health and status of a Nutanix cluster. Information includes:

- system alerts
- current Nutanix software version
- Nutanix processes and Controller VM information
- hypervisor details such as type and version

When Pulse is enabled, it sends a summary email of the cluster configuration to a Nutanix Support server daily by default. Pulse also collects the most important data like system-level statistics and configuration information more frequently to automatically detect issues and help make troubleshooting easier. With this information, Nutanix Support can apply advanced analytics to optimize your implementation and to address potential problems.

Note: Pulse sends messages through ports 80/8443/443 or through your mail server (see [Pulse Access Requirements](#) on page 644).

Note: When logging into Prism the first time after installation or an upgrade, the system checks whether Pulse is enabled. If it is not, a screen appears to enable Pulse (see [Logging Into the Web Console](#) on page 38).

Pulse Transport Methods

Configure one of the following Pulse transport methods (in order of preference):



- [first choice] Enable Pulse and use Prism Central as a proxy for the Pulse data transmitted by each node (for clusters registered with Prism Central). Advantages: The configuration is automatic (as described in "Prism Central Proxy for Pulse Data" in the *Prism Central Guide*), and no new firewall configurations are required when you add a node to the cluster or remove a node from the cluster.
- [second choice] Enable Pulse and configure an HTTP proxy server (see [Configuring HTTP Proxy](#) on page 646). Advantage: No new firewall configurations are required when you add a node to the cluster or remove a node from the cluster.
- [third choice] Enable Pulse and configure your firewall. Enable Pulse by using each Controller VM IP address in each managed cluster (see [Pulse Access Requirements](#) on page 644). 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 is a service that enables Nutanix Support to request granular diagnostic information from Pulse-enabled clusters. Pulse streams configuration data, metrics, alerts, events, and select logs back to Nutanix Support, but this information is aggregate in nature. The Pulse information is designed to provide a high-level state representation of the cluster. When the aggregate data is not detailed enough to diagnose a specific issue, Nutanix Support often needs to collect more diagnostic data from the cluster. Remote Diagnostics allows Nutanix Support to remotely collect the following data (and only the following data):

- Nutanix services logs
- Custom gflags being set for any Nutanix service
- Activity traces for Nutanix services
- Hypervisor logs
- Hypervisor config
- Cluster configuration
- System statistics like memory usage
- Nutanix NCC health check reports

Each time Remote Diagnostics triggers a collection, an entry is added to the audit trail for the cluster. There are always two entries, the start (initiation) and finish (termination) of the diagnostics collection.

Remote Diagnostics is enabled by default for every cluster where Pulse is enabled. If your security policy (or other consideration) does not allow Nutanix Support cluster access for remote diagnostics collection, you can disable Remote Diagnostics without turning off Pulse. Nutanix Support will still provide seamless and proactive support based on the Pulse data.

- To check the Remote Diagnostics status, SSH into a Controller VM and enter the following command:

```
nutanix@cvm$ zkcat /appliance/logical/nusights/collectors/kCommand/override_config
```

- To disable Remote Diagnostics, SSH into a Controller VM and enter the following command:

```
nutanix@cvm$ /home/nutanix/ncc/bin/nusights/set_remote_diagnostics_status --enable=false --reason="text"
```

The `--reason` argument is optional. Use it to enter a text string describing the reason for disabling Pulse.



- To enable Remote Diagnostics, SSH into a Controller VM and enter the following command:

```
nutanix@cvm$ /home/nutanix/ncc/bin/nusights/set_remote_diagnostics_status --enable=true --reason="text"
```

The `--reason` argument is optional. Use it to enter a text string describing the reason for enabling Pulse.

Configuring Pulse

About this task

- See [Pulse Health Monitoring](#) on page 640 for Pulse configuration recommendations.
- Pulse sends messages through ports 80/8443/443, or if this is not allowed, through your mail server (see *Configuring an SMTP Server in the Prism Web Console Guide*).
- When logging in to Prism Central the first time after installation or an upgrade, the system checks whether Pulse is enabled. If it is not, a message appears recommending that you enable Pulse. To enable Pulse, click the **Continue** button in the message and follow the prompts; to continue without enabling Pulse, check the **Disable Pulse (not recommended)** box and then click the **Continue** button.

You can enable (or disable) Pulse at any time. To configure Pulse, do the following:

Procedure

- Click the gear icon in the main menu and then select **Pulse** in the **Settings** page.




Figure 245: Pulse Window

- To enable (disable) this feature, check (uncheck) the **Enable** box. Nutanix recommends enabling Pulse to allow Nutanix Support to receive system data and deliver proactive and context-aware support.

Note: Nutanix does not collect any personally-identifiable information (PII) through Pulse.

3. Click the **Additional Info** button.
Additional connection and recipient information appears.

Pulse?

support in the event of an issue. Nutanix does not collect personally-identifiable information.

PULSE CONNECTION STATUS	STATE
Status	SUCCESS
Last Checked Time	09/25/18, 9:38:08 PM

EMAIL CONNECTION STATUS	STATE
Mode	Default Nutanix Tunnel
Status	● DISABLED

Email Recipients

☐ Nutanix Support (nos-asups@nutanix.com)

Comma Separated Email List

Additional email recipients

Additional Info

Save

Figure 246: Pulse Window: Additional Info

4. Check the **Email Connection Status** field for connection information.

This field appears only when the feature is enabled, and it provides the following information:

- **Mode:** Displays the mode of e-mail transport, which is either Default Nutanix Tunnel or SMTP Server.
- **Status:** Displays the transport mechanism status, which is **Success** when the default tunnel or an SMTP server is active or **Failure** when neither is available (or **Unknown** when the status is not known).
- **Last Attempted:** Time of the last summary e-mail attempt.

5. Check the **Pulse Connection Status** field for connection information.

This field provides the following information:

- **Status:** Displays the transport mechanism status, which is **Success** when the HTTP client can access the Pulse server successfully or **Failure** 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.

6. To add (remove) Nutanix customer support as a recipient of the cluster information, check (uncheck) the box next to **Nutanix Support (nos-asups@nutanix.com)** in the **Email Recipients** field.

Cluster information is e-mailed automatically to Nutanix customer support at **nos-asups@nutanix.com** when this feature is enabled. If you do not want Nutanix customer support to receive the cluster information, you can uncheck this box at any time (or restart the e-mails by checking it again).

7. To send the cluster information to additional recipients, enter one or more e-mail addresses in the **Comma Separated Email List** field.

In addition to (or instead of) sending the cluster information to Nutanix customer support, you can e-mail the information to yourself or others by entering recipient e-mail addresses in this field. Use a comma to separate multiple addresses.

8. Click the **Save** button to save the new settings.

Pulse Access Requirements

In order to successfully send Pulse and alert messages from a cluster to the Nutanix support servers, Pulse requires the following access:

- Messages are sent from the Zeus (cluster configuration manager) leader, so the firewall must allow the Zeus leader IP address. Because the Zeus leader can change, it is recommended that the IP addresses for all Controller VMs in the cluster be open in the firewall.
- Required firewall rules configured to allow traffic from the cluster to Nutanix support servers. See Firewall Requirements in the *Security Guide*.

Pulse data is collected through the following methods:



- By default, emails are sent daily through port 8443 via nsc01.nutanix.net or nsc01.nutanix.net to nos-asups@nutanix.com.
- If a SMTP server is configured on a cluster, emails will be sent via the SMTP server to nos-asups@nutanix.com.
- Throughout the day, each Controller VM in the cluster sends Pulse data to insights.nutanix.com over port 443 using the HTTPS REST endpoint.

Note:

Firewall port requirements are automatically met if you have a Prism Central deployment, see "Prism Central Proxy for Pulse Data" in the *Prism Central Guide*.

Controlling Remote Connections

About this task

Nutanix technical support can remotely assist with problems by logging into the cluster through an SSH connection. This remote connection is disabled by default, but it may be enabled during the site installation procedure or at a later date through the web console.

You must configure rules in your external firewall to allow traffic from the cluster to Nutanix remote support servers. See Firewall Requirements in the *Security Guide*.

To enable (or disable) Nutanix technical support remote access to your cluster through this connection, do the following:

Procedure

1. Click the gear icon in the main menu and then select **Remote Support** in the **Settings** page. The Remote Support dialog box appears.
2. Select (click the radio button for) the desired access state.
 - To allow remote access (temporarily) for Nutanix support, select **Enable** and enter the number and duration (hours or minutes) in the two boxes. When remote support is enabled temporarily, an icon appears in the screen header that includes a dynamic countdown until remote support is disabled.
3. Click the **Save** button to save the new setting and close the window.




Figure 247: Remote Support Window

Configuring HTTP Proxy

About this task

If the customer site cannot send traffic to a Nutanix service center directly, an HTTP proxy is required. To configure an HTTP proxy, do the following:

Procedure

1. Click the gear icon in the main menu and then select **HTTP Proxy** in the **Settings** page. The HTTP Proxy dialog box appears.





Figure 248: HTTP Proxy Window

2. To add an HTTP proxy, click the **New Proxy** button and do the following in the displayed fields:

Note: Only one HTTP proxy can be configured at a time. If one exists currently, you must first delete it 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:** Check the box(es) for the protocol to proxy (HTTP, HTTPS, or both).
- g. When all the fields are correct, click the **Save** button (lower right).
This saves the configuration and redisplay the dialog box with the new HTTP proxy entry appearing in the list.

Note: To return to the HTTP Proxy window without saving, click the **Cancel** button.



The image shows a 'Create HTTP Proxy' dialog box. It has a title bar with the text 'Create HTTP Proxy' and a question mark icon. The dialog contains the following fields and controls:

- Name:** A text input field.
- Address:** A text input field.
- Port:** A text input field.
- Username:** A text input field.
- Password:** A text input field.
- Protocols:** A section with two checkboxes:
 - ☒ HTTP
 - ☐ HTTPS
- Buttons:** 'Cancel' and 'Save' buttons at the bottom right.

Figure 249: Create HTTP Proxy Window

3. To edit an HTTP proxy entry, click the pencil icon on the line for that entry, update one or more of displayed field entries as desired, and then click the **Save** button.

The Update HTTP Proxy dialog box appears with the same fields as the Create HTTP Proxy dialog box plus the option (below the protocol check boxes) to add whitelist entries. To configure HTTP proxy whitelist entries, do the following:

- To add a whitelist target, click the **+ Create** link. This opens a line to enter a target address. Enter the target IP address and then click the **Save** link in that field.
- To edit a whitelist target, click the pencil icon for that target and update as needed.
- To delete a whitelist target, click the X icon for that target.

PROTOCOLS

☒ HTTP

☒ HTTPS

Whitelist [+ create](#)

TARGET	ACTIONS
10.10.10.10	
Target	Save Cancel

[Cancel](#) [Save](#)

Figure 250: Whitelist Targets

4. To delete an HTTP proxy entry, click the X icon for that entry. A window prompt appears to verify the action; click the **OK** button. The entry is removed from the HTTP proxy list.


Accessing the Nutanix Support Portal

About this task

Nutanix provides a variety of support services and materials through its support portal.

Procedure

1. To access the Nutanix support portal, select **Support Portal** from the question mark icon



pull-down list.
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 pull-down lists 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 932: Main Menu Options

Category	Option	Description
Documentation	Product Documentation	Displays a page from which you can view the Nutanix product 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.



Category	Option	Description
	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.
	Acropolis Upgrade Paths	Displays a table of 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 & Forums	Open Case	Displays a form to create a support case.
	View Cases	Displays a page from which you can view your current support cases.
	.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.
Downloads	AOS (NOS)	Displays a page from which you can download AOS releases.
	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.
	Tools & Firmware	Displays a table of tools that can be downloaded, including the Nutanix Cluster Check (NCC) and Prism Central VM.
	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.



Category	Option	Description
	Licenses	Displays a table of your product licenses along with buttons to add or upgrade licenses for your clusters.




Figure 251: Nutanix Support Portal

Nutanix REST API

The Nutanix REST APIs allow you to create scripts that run system administration commands against the Nutanix cluster. The API enables the use of HTTP requests to get information about the cluster as well as make changes to the configuration. Output from the commands are returned in JSON format.

There are two versions of the Nutanix REST API.

- v1: The original Nutanix REST API.
- v2: An update of the v1 API. Users of the v1 API are encouraged to migrate to v2.

A complete list of REST API functions and parameters is available in the REST API Explorer section of the Nutanix web console. In addition, the complete reference for the v2 Nutanix API, including code samples in multiple languages, and tutorials are available at <http://developer.nutanix.com/>.

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.

The v1 and v2 APIs can both be viewed in the REST API Explorer.



Procedure

1. Open the explorer for the desired version of the API.
 - » v1: Connect to the Nutanix web console, click the user icon in the upper-right corner of the web console, and select **REST API Explorer**. In the explorer, select **Version 1** from the menu.




Figure 252: User menu

- » v2: Connect to the Nutanix web console, click the user icon in the upper-right corner of the web console, and select **REST API Explorer**. In the explorer, select **Version 2** from the menu.




Figure 253: API version selection

The REST API Explorer displays a list of the cluster objects that can be managed by the API. Each line has four options:

- **Show/Hide:** Expand or reduce the detail shown for the object
- **List Operations:** Show all operations that can be run on this object
- **Expand Operations:** Show the detailed view of the operations that can be run on this object




Figure 254: REST API Explorer

Tip: The objects are listed by a relative path that is appended to the base URL [https://management_ip_addr:9440/PrismGateway/services/rest/v\[1,2,3\]/api](https://management_ip_addr:9440/PrismGateway/services/rest/v[1,2,3]/api), where *management_ip_addr* is the IP address of any Nutanix Controller VM in the cluster.

- Find the line for the object you want to explore and click **Expand Operations**. For this example, you will operate on a storage pool.




Figure 255: Cluster Operations

3. Click **GET** on the first line to show the details for this API call.




Figure 256: Get Cluster Details

The explorer displays the parameters that can be passed when this action is used.

4. Click **Try it out!** to test the API call when used with your cluster.




Figure 257: Get Cluster Details Response

The test displays the request URL required for scripting, as well as sample responses expected from the API call.

HELP RESOURCES

There are several information sources that you can access at any time when you need help:

- Context-sensitive help documentation (see [Accessing Online Help](#) on page 656).
- Health dashboard tutorial (see Menu Options in Health Dashboard in the *Prism Web Console Guide*).
- Customer support portal (see [Accessing the Nutanix Support Portal](#) on page 648).
- Nutanix community forum (see [Accessing the Nutanix Next Community](#) on page 660).
- REST API explorer (see [Accessing the REST API Explorer](#) on page 652).
- Glossary of terms (see [Glossary](#) on page 661).

Accessing Online Help

About this task

The web console 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 question mark icon pull-down list of the main menu (see Main Menu Options in Cluster Management in the *Prism Web Console Guide*):
 - » Select **Help with this page** to display help documentation that describes the current screen.

Note: In a task window click the question mark icon in the upper right to display the help documentation for that window.

- » Select **General Help** 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 258: Help Organization Page

2. To select a topic from the table of contents, click the collapse menu icon (also known as a "hamburger" button) in the upper left.

A table of contents pane appears on the left. Click a topic in the table of contents to display that topic.





Figure 259: Table of Contents Pane

3. To display all the help contents as a single document (*Web Console Guide*), click the **epub** or **pdf** button in the upper right.

You can view the *Web Console Guide* in either ePUB or PDF format by selecting the appropriate button. If your browser does not support the selected format, you can download the PDF or ePUB file.

4. To search for a topic, click the



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 260: 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, select **Nutanix Next Community** from the question mark icon  pull-down list of the main menu (see Main Menu Options in Cluster Management in the Prism Web Console Guide).

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 261: Next Community Screen

Glossary

aCLI

Acropolis command-line interface.

Acropolis

The Nutanix converged software fabric for virtualization and storage management. It consists of the Acropolis base software, Acropolis Distributed Storage Fabric, AHV, App Mobility Fabric, Prism, and Acropolis APIs.

Acropolis App Mobility Fabric

Provides virtualization management, volume management, and other distributed management functions for Acropolis.

AHV

Nutanix-specific hypervisor host providing core server virtualization and optimized for Nutanix cluster and guest VM use.

block

A set of Nutanix nodes contained in a single enclosure.



block fault tolerance

When certain conditions are met, redundant copies of any data required to serve I/O are placed on nodes that are not in the same block. In the case where multiple nodes in a block fail, the cluster can continue to run because cluster configuration data exists on other blocks. Guest VMs can continue to run because redundant copies of guest VM data and metadata exist on other blocks.

clone

A writeable copy of a vDisk.

Cloud Connect

A feature that enables you to back up and restore copies of virtual machines and files to and from an on-premise cluster and a Nutanix Controller VM located on the Amazon Web Service (AWS) or Microsoft Azure cloud. The cloud-based cluster is managed as a remote site from the Prism Web Console or nCLI.

cluster

A group of nodes contained in one or more Nutanix blocks.

cold data

Data that did not have many rewrites or overwrites for a long time. For example, snapshots, file servers, archives, or backups.

compression

An optional storage container setting that enables one of two types of compression.

compression, inline

Data is compressed as it is written. This type of compression is recommended for workloads that perform batch processing.

compression, post-process

Data is compressed after it is written. The delay time between write and compression is configurable. Because every workload has a different I/O profile, Nutanix has no recommended delay value. This type of compression is recommended for most workloads.

consistency group

A subset of VMs in a protection domain. All VMs within a consistency group for that protection domain will be snapshotted in a crash-consistent manner. For all VMs in a consistency group, a snapshot creates one snapshot for all VMs in the group.

Controller VM

A Nutanix VM that manages storage and other cluster functions on a node.

data resiliency

A measure of the types of failures a cluster can withstand; affected by block awareness and redundancy factor.

datastore

A logical storage container for files necessary for VM operations.



deduplication

The sharing of identical guest VM data on premium tiers (RAM and Flash) for improved performance or on capacity tiers (HDD) for storage space savings. Enabled by properties of a storage container or vDisk.

Distributed Storage Fabric

All storage functionality, including snapshots and clones, data protection, disaster recovery, data path redundancy, replication factors, deduplication, compression, erasure coding, and so on.

Enterprise Cloud Platform

The Nutanix solution that natively converges compute, virtualization and storage into a resilient, software-defined solution with rich machine intelligence

erasure coding

Optional algorithm included in the Acropolis base software to help reduce the storage used for fault tolerance. It helps to increase the effective or usable capacity on a cluster, depending on cluster size and data coldness.

failback

Planned failover initiated from recovery site.

failover

Moving VMs from a primary site to a recovery site.

failover, disaster

Failover when the primary site is down.

failover, planned

Failover when both sites are up.

guest VM

A VM running on a Nutanix cluster that executes a workload, such as VDI or Exchange, as opposed to a VM that is involved in cluster operations, such as a Controller VM.

host

An instance of the hypervisor that runs on a Nutanix node.

image service

A workflow in the Prism web console that enables a user to upload ISO or disk images (in ESXi or Hyper-V format) to a Nutanix AHV cluster by specifying a remote repository URL or by uploading a file from a local machine.

local replication

Multiple copies of data within a storage container. These copies exist for fault tolerance: if a physical disk fails, the cluster can recover data from another copy. The cluster manages the replicated data, and the copies are not visible to the user.

local snapshot

Snapshots stored on the same cluster where they originated.



nCLI

Nutanix command-line interface.

node

A physical server contained in a Nutanix block; runs a hypervisor host.

oplog

A write cache on a faster, more expensive storage tier.

Prism

Web-based management interface for managing Nutanix clusters.

Prism Central

Centralized management tool that runs as a separate VM configured as a single-node cluster to monitor and manage multiple clusters through a single web console.

Prism Element

A single cluster being managed by and available through the Prism Central web console.

protection domain

A group of VMs to be backed up locally on a cluster or replicated on the same schedule to one or more clusters. Protection domains may be associated with remote sites. Protection domain names must be unique across sites. A VM can be in at most one protection domain.

protection domain, active

A protection domain that manages live VMs and makes, replicates, and expires snapshots.

protection domain, inactive

A protection domain that receives snapshots from a remote cluster.

redundancy factor

The number of nodes plus 1 that the cluster can tolerate being down at one time. By default, Nutanix clusters have a redundancy factor of 2, which means that they can tolerate 1 node being down. They are configurable to redundancy factor 3 to enable tolerating 2 nodes being down.

remote replication, one-to-one

Replicates a production cluster with one idle cluster as a DR target.

remote replication, reciprocal

Cross replication within running (non-idle) production clusters.

remote site

A pair of clusters that can replicate data to each other. A configured remote site can also be located in the cloud (based on Amazon AWS or Microsoft Azure, for example).

remote snapshot

A snapshot copied asynchronously from one cluster to another.



replication factor

The number of copies of data kept by a storage container. The default is 2. Storage Containers on clusters with redundancy factor of 3 can have replication factor of 3.

reserved capacity

A property of a storage container or vDisk that guarantees that a certain amount of storage space is available.

schedule

A property of a protection domain that specifies the intervals to take snapshots and how long the snapshots should be retained. A schedule optionally specifies which remote site to replicate to.

Self-Service Restore

Allows a user to restore a file within a virtual machine from the Nutanix protected snapshot with minimal Nutanix administrator intervention.

shadow clone

A cache of a vDisk on all the nodes in the cluster. When a vDisk is read by multiple VMs (such as the base image for a VDI clone pool), the cluster creates shadow clones of the vDisk.

snapshot

A read-only copy of the state and data of a VM at a point in time. Snapshots for a VM are crash consistent, which means that the VMDK on-disk images are consistent with a single point in time. That is, the snapshot represents the on-disk data as if the VM crashed. The snapshots are not, however, application consistent, meaning that application data is not quiesced at the time of snapshot, unless the protection domain is configured to use application-consistent snapshots.

storage container

A subset of available storage within a storage pool.

storage pool

A group of physical disks from one or more tiers.

storage replication adapter

A Nutanix-provided module that allows VMware Site Replication Manager (SRM) to use native remote replication.

tier

A type of physical storage in a Nutanix node.

vDisk

Data associated with a VM represented as a set of files on a datastore.

VM high availability

In virtualization management VM high availability, when a node becomes unavailable, VMs that are running on that node are restarted on another node in the same cluster.



VM mobility

The ability to export your existing VMs from one non-AHV cluster to an AHV cluster. This option requires that you install the Nutanix VM Mobility installer on all the VMs. The Nutanix VM Mobility installer deploys the drivers that are required at the destination AHV cluster. After you prepare the source VMs, they can be exported to the AHV cluster.

vStore

A separate mount point within a storage container which has its own NFS namespace. This namespace maps to a protection domain. Each vStore is exported as a device through the Nutanix SRA.



COPYRIGHT

Copyright 2020 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.

License

The provision of this software to you does not grant any licenses or other rights under any Microsoft patents with respect to anything other than the file server implementation portion of the binaries for this software, including no licenses or any other rights in any hardware or any devices or software that are used to communicate with or in connection with this software.

Conventions

Convention	Description
<code>variable_value</code>	The action depends on a value that is unique to your environment.
<code>ncli> command</code>	The commands are executed in the Nutanix nCLI.
<code>user@host\$ command</code>	The commands are executed as a non-privileged user (such as nutanix) in the system shell.
<code>root@host# command</code>	The commands are executed as the root user in the vSphere or Acropolis host shell.
<code>> command</code>	The commands are executed in the Hyper-V host shell.
<code>output</code>	The information is displayed as output from a command or in a log file.

Version

Last modified: May 15, 2020 (2020-05-15T10:56:44+05:30)

