FORMULÁRIO DE PROPOSTA DE PREÇOS

RAZÃO SOCIAL: **CLEAR TECNOLOGIA DA INFORMAÇÃO LTDA.**		
CNPJ: **30.088.923/001-08**	TELEFONES: **+55 92 3042-0123**	
EMAIL: **contato@clearit.com.br**		
ENDEREÇO: **RUA SALVADOR, 440 – SALA 1810 – EDIFICIO SOBERANE – TORRE CORPORATE ADRIANÓPOLIS – MANAUS – AMAZONAS – CEP: 69057-040**		
BANCO	AGÊNCIA:	CONTA CORRENTE:
BRADESCO	**0301**	**032817-0**

ITEM	DESCRIÇÃO	UND.	QTD.	VALOR UNITÁRIO	VALOR TOTAL
1	Licenciamento de software de proteção e resiliência de informações *Marca:* ***Veeam*** *Modelo:* ***Veeam Backup & Replication***	Un.	600	R$ 1.985,00	R$ 1.191.000,00
2.1	Unidade de armazenamento de informação – tipo 1 *Marca:* ***Exagrid*** *Modelo:* ***EX84***	Un.	2	R$ 1.328.400,00	R$ 2.656.800,00
2.2	Unidade de armazenamento de informação – tipo 2 *Marca:* ***Exagrid*** *Modelo:* ***EX36***	Un.	5	R$ 598.500,00	R$ 2.992.500,00
3.1	Fitoteca de armazenamento modular – base *Marca:* ***LENOVO/ IBM*** *Modelo:* ***SR630 / TS4300***	Un.	2	R$ 307.490,00	R$ 614.980,00
3.2	Fitoteca de armazenamento modular – expansão *Marca:* ***LENOVO/ IBM*** *Modelo:* ***TS4300 – Expansion Unit***	Un.	4	R$ 260.000,00	R$ 1.040.000,00
4	Serviço de instalação e configuração da solução *Marca:* ***ClearIT*** *Modelo:* --	Un.	100	R$ 350,00	R$ 35.000,00
5	Serviço de treinamento oficial *Marca:* ***Veeam*** *Modelo:* ***VMCE v11***	Un.	6	R$ 22.000,00	R$ 132.000,00
VALOR TOTAL:					R$ 8.662.280,00

Valor total por extenso da Proposta de Preços: R$ 8.662.280,00 (Oito milhões e seiscentos e sessenta e dois mil e duzentos e oitenta Reais)

Validade da proposta: 60 (sessenta) dias

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

Manaus, 18 de abril de 2022.

ROGERIO AUGUSTO FERREIRA:10134666836

Assinado de forma digital por ROGERIO AUGUSTO FERREIRA:10134666836
Dados: 2022.04.18 10:44:49 -03'00'

CLEAR TECNOLOGIA DA INFORMAÇÃO LTDA
30.088.923/0001-08
Rogério Augusto Ferreira
Diretor – Sócio Administrador
101.346.668-36

ANEXO – PROPOSTA TÉCNICA

DA GARANTIA E DO SUPORTE TÉCNICO

Conforme prescrito no Edital 33/2022, do Tribunal de Justiça do Estado do Amazonas – TJ/AM, nossa oferta contempla a garantia de 36 (trinta e seis) meses para os produtos apresentados nesta proposta comercial (conforme respectivas listas de Part-Numbers).

- E-mail para abertura de chamados (ClearIT): suporte@clearit.com.br
- Número para abertura de chamados (ClearIT): 0800 990 9000
- Website para abertura de chamados (ClearIT): https://clearit.odoo.com/helpdesk/

SERVIÇOS CORRELATOS

Declaramos que a oferta dos nossos serviços está em conformidade com os requisitos, premissas e níveis de qualidade descritos no Termo de Referência do Edital 33/2022, do Tribunal de Justiça do Estado do Amazonas – TJ/AM.

ACESSÓRIOS PARA INSTALAÇÃO

Declaramos que todos os acessórios e demais componentes necessários a integração, a adequada instalação física e a implantação lógica da nossa solução ofertada com o parque tecnológico do Tribunal de Justiça do Estado do Amazonas – TJ/AM, estão inclusos nessa proposta.

LOCAL E PRAZO DE ENTREGA OU APLICAÇÃO

Declaramos que estamos cientes e em conformidade com os prazos e locais de entrega preconizados no Termo de Referência do Edital 33/2022, do Tribunal de Justiça do Estado do Amazonas – TJ/AM.

VERSÕES DE FIRMWARE E SOFTWARE

Declaramos que todas as versões de licenças, de softwares, de firmwares etc., serão entregues na versão mais atual e estável disponibilizadas pelas respectivas fabricantes de nossa oferta.

Declaramos que na necessidade de conversão ou alteração de qualquer licenciamento, software e/ou firmware, correlatos ao objeto da contratação, para o perfeito funcionamento do parque computacional do Tribunal de Justiça do Estado do Amazonas com a nossa oferta destacada nesta proposta, será de nossa responsabilidade tal execução sem qualquer alteração ao escopo comercial e/ou técnico aqui descrito.

DESCRITIVO TÉCNICO

ITEM 01 – Licenciamento de software de proteção e resiliência de informações

Marca: Veeam

Modelo: Backup & Replication

Quantidade: 600

Lista de Part Numbers:

PART-NUMBER	DESCRIÇÃO	UNITÁRIO	TOTAL
P-VBRVUL-0I-SU3YP-00	Veeam Backup & Replication Universal Subscription License. Includes Enterprise Plus Edition features. 10 instance pack. 3 Years Subscription	1	60
228-04778	SQLSvrStd SNGL LicSAPk OLV NL 3Y AqY1 AP	1	1
359-01469	SQLCAL SNGL LicSAPk OLV NL 3Y AqY1 AP Usr CAL	2	2

ITEM 02.1 – Unidade de armazenamento de informação – tipo 1

Marca: Exagrid

Modelo: EX84

Quantidade: 02

Lista de Part Numbers:

PART-NUMBER	DESCRIÇÃO	UNITÁRIO	TOTAL
EX84	Disk Capacity: Raw: 192 TB, Useable: 168 TB. 84 TB Full Backup.	1	2
EX84-SW	168 TB Tiered Backup Storage Software - Software Only	1	2
EX-3YR-MS-P	Three year 7 x 24 Customer Support and product Maintenance	1	2
EX-40GIG-SFPF-A	40 Gigabit Ethernet Dual Port QSFP+ Optical Option for the EX52, EX52-SEC, EX84 and EX84-SEC only. Includes two qualified short-range QSFP+ Modules.	1	2
ACORD-DMM-8246	CORDAO | OM4 | DUPLEX | 15M | MM | LC/UPC - LC/UPC	2	4
ACE-CAT6-AZ	Patch Cord Cat6 Azul | 5M | RJ-45	1	2

ITEM 02.2 – Unidade de armazenamento de informação – tipo 2

Marca: Exagrid

Modelo: EX36

Quantidade: 05

Lista de Part Numbers:

PART-NUMBER	DESCRIÇÃO	UNITÁRIO	TOTAL
EX36	Disk Capacity: Raw: 96 TB, Useable: 72 TB. 36 TB Full Backup. Includes 1 10 Gigabit Add on Card	1	5
EX36-SW	72 TB Tiered Backup Storage Software -Software Only	1	5
EX-3YR-MS-P	Three year 7 x 24 Customer Support and product Maintenance	1	5
ACORD-DMM-8246	CORDAO | OM4 | DUPLEX | 15M | MM | LC/UPC - LC/UPC	2	10
ACE-CAT6-AZ	Patch Cord Cat6 Azul | 5M | RJ-45	1	5

ITEM 03.1 – Fitoteca de armazenamento modular - base

Marca: Lenovo / IBM

Modelo: SR630 / TS4300

Quantidade: 02

Lista de Part Numbers:

PART-NUMBER	DESCRIÇÃO	UNITÁRIO	TOTAL
Robô de Fitas			
6741L1U	Tape: IBM TS4300 3U Tape Library-Base Unit	1	2
B16H	Biblioteca de Fitas IBM TS4300 3U - Unidade de Base	1	2
B16X	Unidade LTO 7 HH SAS	1	2
AS10	Cabo HD-SAS para Mini-SAS	2	4

ATNZ	Cartuchos de Dados Ultrium 7 de 5 unidades	7	14
AS1E	Cartucho de Limpeza L1 UCC Ultrium	4	8
B16N	Kit de Montagem em Rack	1	2
B16K	Fonte de alimentação adicional	1	2
AS2U	Cabo de Alimentação de Rack 4,3 m, 10 A/100-250 V, C13 para IEC 320-C14	2	4
5WS7A03786	Advanced Service - 3Yr 24x7	1	2
ACE-CAT6-AZ	Patch Cord Cat6 Azul | 5M | RJ-45	1	2
Unidade Controladora / Servidor			
AUW0	Chassi ThinkSystem SR630 2,5 pol. com 8 compartimentos	1	2
B4HT	Processador Intel Xeon Silver 4208 8C 85 W 2,1 GHz	2	4
B4H2	RDIMM ThinkSystem de 16 GB TruDDR4 2.933 MHz (2R x 8 1,2 V)	2	4
AUWB	ThinkSystem SR530/SR630/SR570 2.5" SATA/SAS 8-Bay Backplane	1	2
AUNJ	Adaptador ThinkSystem RAID 930-8i de 12 Gb PCIe Flash 2 GB	1	2
AULV	HDD Hot Swap ThinkSystem 2,5 pol. 300 GB 15 K SAS 12 Gb 512n	4	8
AUMV	ThinkSystem M.2 com Kit de Ativação de Espelhamento	1	2
AUUV	SSD Não Hot Swap ThinkSystem M.2 128 GB SATA 6 Gbps	2	4
AUWC	Kit de Placa Riser 1 ThinkSystem SR530/SR570/SR630 x8/x16 PCIe LP+LP	1	2
AUWA	Kit de Placa Riser 2 ThinkSystem SR530/SR570/SR630 x16 PCIe LP	1	2
AUKJ	LOM SFP+ 2 portas 10 Gb ThinkSystem	1	2

AUNR	HBA ThinkSystem 430-8e SAS/SATA 12 Gb	1	2
AVW8	Fonte de alimentação ThinkSystem 550 W (230 V/115 V) Platinum hot-swap	2	4
6570	Cabo Jumper 2,0 m, 13 A/100-250 V, C13 a C14	2	4
AUPW	Atualização ThinkSystem XClarity Controller Padrão para Corporativo	1	2
AXCA	Trilho de Corrediça sem Ferramentas ThinkSystem	1	2
A1PK	Cabo IBM Passive DAC SFP+ de 5m	2	4
B0MJ	Recurso Habilitar TPM 1.2	1	2
B7Y0	Ativar IPMI-over-LAN	1	2
5WS7A68965	Advanced Service - 3Yr 24x7	1	2
ACE-CAT6-AZ	Patch Cord Cat6 Azul | 5M | RJ-45	1	2
7S05CTO3WW	Windows Server 2019	1	2
S09C	Windows Server 2019 Standard (16 core) - MultiLang (not preinstalled)	1	2
3523	Drop-in-the-Box	1	2
3444	Registration only	1	2
5372SWX	xSeries HIPO	1	2
S08H	Microsoft Windows Server 2019 Client Access License (1 User)	2	4
5372SWX	xSeries HIPO	1	2
S0AM	Windows Server 2019 Standard (16 core) - MultiLang (not preinstalled)	1	2
7S05CTO5WW	Windows Server 2019 CALs	1	2
S081	Microsoft Windows Server 2019 Client Access License (1 User)	2	4

ITEM 03.2 – Fitoteca de armazenamento modular - expansão

Marca: IBM / Lenovo

Modelo: TS4300 – Expansion Unit

Quantidade: 04

Lista de Part Numbers:

PART-NUMBER	DESCRIÇÃO	UNITÁRIO	TOTAL
6741L3U	Tape: IBM TS4300 3U Tape Library-Expansion Unit	1	4
B16J	Biblioteca de Fitas IBM TS4300 3U - Unidade de Expansão	1	4
B16X	Unidade LTO 7 HH SAS	1	4
AS10	Cabo HD-SAS para Mini-SAS	2	8
ATNZ	Cartuchos de Dados Ultrium 7 de 5 unidades	7	28
B16N	Kit de Montagem em Rack	1	4
B16L	Fonte de Alimentação Principal	1	4
B16K	Fonte de alimentação adicional	1	4
AS2U	Cabo de Alimentação de Rack 4,3 m, 10 A/100-250 V, C13 para IEC 320-C14	2	8
B5XM	Embalagem Externa 6741 CTO	1	4
5WS7A03947	Advanced Service - 3Yr 24x7	1	4

ITEM 04 – Serviço de instalação e configuração da solução

Marca: ClearIT (CLEAR TECNOLOGIA DA INFORMAÇÃO)

Modelo: N/A

Quantidade: 100

ITEM 05 – Serviço de treinamento oficial

Marca: Veeam

Modelo: VMCE v11

Quantidade: 06

Lista de Part Numbers:

PART-NUMBER	DESCRIÇÃO	UNITÁRIO	TOTAL
VMCEV11	Veeam Availability Suite v11: Configuration and Management	1	6

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

São Paulo
Av. Queiroz Filho, 1560 - sala 18 - Edifício Gaivota
Vila Leopoldina - CEP 05319-000 - SP
+55 11 4673-4999

COMPROVAÇÕES TÉCNICAS

ITEM	DOCUMENTAÇÃO	COMPROVAÇÃO	PÁGINA
5.3.1. Licenciamento de software de proteção e resiliência de informações	N/A	N/A	N/A
5.3.1.1. O licenciamento da solução de proteção e recuperação deverá ser baseado no modelo para o ambiente virtual, temporário ou por assinatura, baseando-se na quantidade de máquinas virtuais desde que não haja limitação de quantidades de uso de capacidade;	veeam_backup_11_0_user_guide_vsphere.pdf	Subscription license — license that expires at the end of the subscription term. The Subscription license term is normally 1-3 years from the date of license issue.	128
5.3.1.2. O licenciamento entregue deverá permitir a portabilidade das cargas de trabalho, garantindo proteção da informação independentemente de onde ela esteja, seja localmente, seja remotamente (em outro site, em nuvem pública, etc.);	https://www.veeam.com/licensing-policy.html	Veeam Backup & Replication Offered in packs of 10 Instances. Support and Maintenance is limited to Production level. Basic level is not offered. Applications included: Veeam Backup for AWS Veeam Backup for Microsoft Azure Veeam Backup for Nutanix AHV Veeam Agent for Windows Veeam Agent for Linux Veeam Agent for Oracle Solaris Veeam Agent for IBM AIX Veeam Agent for Mac Instance calculation for Veeam Universal License Veeam Universal License (available in both Subscription and Perpetual License Types) removes complexity from license consumption tracking, providing portability of the license across a breadth of product functionality. WORKLOAD TYPE VM or Cloud VM LICENSED OBJECT(S) 1 INSTANCE 1	N/A
5.3.1.3. Após o vencimento da assinatura deve existir um período mínimo de carência de 30 dias;	veeam_backup_11_0_user_guide_vsphere.pdf	Grace Period Duration 30 days	148

5.3.1.4. Para a medição do licenciamento do ambiente virtual, é a soma de todas as máquinas virtuais dos hypervisors (VMware vSphere ESXi e Acropolis Operating System), não importando a quantidade de núcleos deles, para esse licenciamento a solução não deverá ter limite de TB de entrada;	https://www.veeam.com/licensing-policy.html	Veeam Backup & Replication Offered in packs of 10 Instances. Support and Maintenance is limited to Production level. Basic level is not offered. Applications included: Veeam Backup for AWS Veeam Backup for Microsoft Azure Veeam Backup for Nutanix AHV Veeam Agent for Windows Veeam Agent for Linux Veeam Agent for Oracle Solaris Veeam Agent for IBM AIX Veeam Agent for Mac Instance calculation for Veeam Universal License Veeam Universal License (available in both Subscription and Perpetual License Types) removes complexity from license consumption tracking, providing portability of the license across a breadth of product functionality. WORKLOAD TYPE VM or Cloud VM LICENSED OBJECT(S) 1 INSTANCE 1	N/A
5.3.1.5. No licenciamento da solução, entendem-se que todas as funcionalidades descritas nesse termo estarão habilitadas e disponíveis para uso de forma total e irrestrita, na inteireza da capacidade licenciada, independentemente da quantidade ou tipo de agentes necessários, de acordo com a necessidade da CONTRATANTE, e, sem necessidade de aquisição de qualquer outro tipo de licença ou recurso adicional para execução de tais funcionalidades;	Conforme Lista de Part Numbers	N/A	N/A

5.3.1.6. A solução ofertada deve estar habilitada para permitir a instalação de quantos servidores de movimentação de dados e de gerência da solução, quanto forem necessários para configuração do ambiente a ser protegido, de acordo com as melhores práticas propostas pelo fabricante;	veeam_backup_11_0_user_guide_vsphere.pdf	Backup Proxy A backup proxy is an architecture component that sits between the backup server and other components of the backup infrastructure. While the backup server administers tasks, the proxy processes jobs and delivers backup traffic. Basic backup proxy tasks include the following: • Retrieving VM data from the production storage • Compressing • Deduplicating • Encrypting • Sending it to the backup repository (for example, if you run a backup job) or another backup proxy (for example, if you run a replication job) A backup proxy can operate as a data mover in NAS backup and transfer data from the source file share to the backup repository. For more information about using a backup proxy in NAS backup, see NAS Backup Support.	314
5.3.1.7. A versão ofertada deve ser a última versão disponível, não será aceita a utilização de versões anteriores para cobrir alguma especificação técnica;	Conforme Lista de Part Numbers	N/A	N/A
5.3.1.8. A solução ofertada deverá de maneira simples e objetiva mostrar a quantidade de licenças adquiridas e utilizadas;	veeam_backup_11_0_user_guide_vsphere.pdf	Viewing License Information • Instances — number of instances that you can use to protect workloads.	139
5.3.1.9. Caso a solução permita o consumo acima do que foi contratada, sem nenhuma trava, não será cobrado em hipótese nenhuma essa diferença, seja no licenciamento, seja em futuras renovações ou desistência da utilização do software;	veeam_backup_11_0_user_guide_vsphere.pdf	Exceeding License Limit	146
5.3.1.10. Ter suporte e garantia de 3 (três) anos, com atendimento 24 x 7 x 365 e até 2 (duas) horas para início do atendimento de chamados com prioridade mais alta, ou seja, quando o ambiente estiver em estado de indisponibilidade de uso;	https://www.veeam.com/support-policy.html	"4.4 Severity level descriptions	https://www.veeam.com/support-policy.html

5.3.1.11. Caso a solução ofertada necessite de algum banco de dados, o mesmo deverá ser fornecido devidamente licenciado sem nenhum custo extra para a CONTRATANTE.	https://bp.veeam.com/vbr/2_Design_Structures/D_Veeam_Components/D_VBR_DB/database.html	When protecting more than 500 VMs. The max database size allowed by Express Edition is usually sufficient, so do not consider this a constraint. Veeam Backup & Replication console and job processing may however slow down as a result of CPU and RAM constraints on the SQL Server Express instance.	N/A
5.3.2. Características da infraestrutura	N/A	N/A	N/A
5.3.2.1. Deve possuir arquitetura em múltiplas camadas ou arquitetura similar:	veeam_backup_11_0_user_guide_vsphere.pdf	Backup Infrastructure Veeam Backup & Replication is a modular solution that lets you build a scalable backup infrastructure for environments of different sizes and configuration.	244
5.3.2.1.1. Servidor de gerência de proteção;	veeam_backup_11_0_user_guide_vsphere.pdf	Backup Server The backup server is a Windows-based physical or virtual machine on which Veeam Backup & Replication is installed. It is the core component in the backup infrastructure that fills the role of the “configuration and control center”.	246
5.3.2.1.2. Servidores de movimentação de dados;	veeam_backup_11_0_user_guide_vsphere.pdf	Backup Proxy A backup proxy is an architecture component that sits between the backup server and other components of the backup infrastructure. While the backup server administers tasks, the proxy processes jobs and delivers backup traffic.	314
5.3.2.1.3. Clientes ou agentes de cópias.	veeam_backup_11_0_user_guide_vsphere.pdf	Veeam Agent Management To back up physical machines running Windows, Linux, Unix or macOS operating systems, Veeam Backup & Replication uses backup agents installed on each computer. Veeam Backup & Replication operates as a centralized control center for deploying and managing Veeam Agent for Microsoft Windows, Veeam Agent for Linux, Veeam Agent for IBM AIX, Veeam Agent for Oracle Solaris and Veeam Agent for Mac (Veeam Agents).	1947
5.3.3. O servidor de gerência de proteção deverá ter suporte para instalação no mínimo com os sistemas operacionais abaixo:	N/A	N/A	N/A
5.3.3.1.1. Microsoft Windows 2016;	veeam_backup_11_0_user_guide_vsphere.pdf	Microsoft Windows Server 2016	31

5.3.3.1.2. Microsoft Windows 2019.	veeam_backup_11_0_user_guide_vsphere.pdf	Microsoft Windows Server 2019	31
5.3.4. O servidor de movimentação de dados deverá ter suporte para instalação no mínimo com um dos sistemas abaixo:	N/A	N/A	N/A
5.3.4.1. Microsoft Windows 2019 ou superior;	veeam_backup_11_0_user_guide_vsphere.pdf	Microsoft Windows Server 2019	36
5.3.4.2. CentOS 8.x ou superior;	veeam_backup_11_0_user_guide_vsphere.pdf	CentOS 7 to 8.3, and 8.41, and 8.5	36
5.3.4.3. Ubuntu 18.04 ou superior;	veeam_backup_11_0_user_guide_vsphere.pdf	Ubuntu: 14.04 LTS, 16.04 LTS, 18.04 LTS, 19.10, 20.04 LTS, and 21.041, and 21.102	36
5.3.4.4. Debian 10.4 ou superior;	veeam_backup_11_0_user_guide_vsphere.pdf	Debian 9.0 to 10.8, and 11.0	36
5.3.4.5. Suse Linux Enterprise Server 15 SP2 ou superior;	veeam_backup_11_0_user_guide_vsphere.pdf	Debian 9.0 to 10.8, and 11.0	36
5.3.4.6. Red Hat Enterprise Linux 8.x ou superior.	veeam_backup_11_0_user_guide_vsphere.pdf	RHEL 6.0 to 8.3, and 8.41, and 8.52	36
5.3.5. Possuir um banco de dados ou catálogo interno, contendo informações sobre todos os arquivos e mídias onde os backups foram armazenados;	veeam_backup_11_0_enterprise_manager_user_guide.pdf	Veeam Backup Catalog Veeam Backup Catalog is a feature that stands for VM guest OS file indexing. Veeam Backup Catalog comprises Veeam Guest Catalog services that run on the following servers in the backup infrastructure: Veeam backup server and Veeam Backup Enterprise Manager server. • Veeam Guest Catalog service on the Veeam backup server works as a local catalog service. It collects index data for backup jobs on this specific Veeam backup server and stores this data locally in the Veeam Backup Catalog folder.	17
5.3.5.1. Não serão aceitas soluções que possuam catálogos distintos dentro da mesma arquitetura, ou seja, soluções onde caso uma tarefa seja executada em um certo módulo, exista a possibilidade dos dados e metadados não refletirem no mesmo catálogo único;	veeam_backup_11_0_enterprise_manager_user_guide.pdf	Veeam Backup Catalog Veeam Backup Catalog is a feature that stands for VM guest OS file indexing. Veeam Backup Catalog comprises Veeam Guest Catalog services that run on the following servers in the backup infrastructure: Veeam backup server and Veeam Backup Enterprise Manager server. • Veeam Guest Catalog service on the Veeam backup server works as a local catalog service. It collects index data for backup jobs on this specific Veeam backup server and stores this data locally in the Veeam Backup Catalog folder. • Veeam Guest Catalog service on Veeam Backup Enterprise Manager works as a federal	17

		catalog service. It communicates with Veeam Guest Catalog services on Veeam backup servers connected to Veeam Backup Enterprise Manager and performs the following tasks: o Replicates index data from Veeam backup servers to create a federal catalog o Maintains index data retention o Lets you search for machine guest OS files in backup files	
5.3.5.2. Para todas as funcionalidades aqui descritas, o catálogo deverá ser único, independentemente de aplicação e servidor protegido;	veeam_backup_11_0_enterprise_manager_user_guide.pdf	Veeam Backup Catalog Veeam Backup Catalog is a feature that stands for VM guest OS file indexing. Veeam Backup Catalog comprises Veeam Guest Catalog services that run on the following servers in the backup infrastructure: Veeam backup server and Veeam Backup Enterprise Manager server. • Veeam Guest Catalog service on the Veeam backup server works as a local catalog service. It collects index data for backup jobs on this specific Veeam backup server and stores this data locally in the Veeam Backup Catalog folder.	17
5.3.6. Caso a ferramenta faça uso de um software de banco de dados para armazenamento das informações, e este requeira uma licença para uso, essa licença deve ser fornecida em conjunto com a solução;	Ciente e de Acordo	N/A	N/A
5.3.7. Ser flexível e escalável, permitindo sua instalação, configuração e uso em sites remotos interligados ao site principal através de WAN. Além disso, a solução deve prover recursos de desduplicação na origem, desduplicação no destino, e compactação tanto no site principal como nos sites remotos na inteireza da capacidade previamente licenciada e sem necessidade de aquisição de qualquer outro tipo de licença ou recurso adicional para execução de tais operações;	veeam_backup_11_0_user_guide_vsphere.pdf	Distributed Deployment The distributed deployment scenario is recommended for large geographically dispersed virtual environments with multiple Veeam Backup & Replication servers installed across different sites. These backup servers are federated under Veeam Backup Enterprise Manager — an optional component that provides centralized management and reporting for these servers through a web interface. Veeam Backup Enterprise Manager collects data from backup servers and enables	609

		you to run backup and replication jobs across the entire backup infrastructure through a single web console, edit them and clone jobs using a single job as a template. It also provides reporting data for various areas (for example, all jobs performed within the last 24 hours or 7 days, all VMs engaged in these jobs and so on). Using indexing data consolidated on one server, Veeam Backup Enterprise Manager provides advanced capabilities to search for VM guest OS files in VM backups created on all backup servers (even if they are stored in backup repositories on different sites), and recover them in a single click. Search for VM guest OS files is enabled through Veeam Backup Enterprise Manager itself.	
5.3.7.1. As tecnologias de desduplicação e compressão de dados, devem ser independentes do tipo de repositório de armazenamento empregado na arquitetura, permitindo a ativação/desativação das tecnologias conforme o melhor cenário de implementação a ser utilizado;	veeam_backup_11_0_user_guide_vsphere.pdf	Data Compression and Deduplication Veeam Backup & Replication provides mechanisms of data compression and deduplication. Data compression and deduplication let you decrease traffic going over the network and disk space required for storing backup files and VM replica files.	757
	veeam_backup_11_0_user_guide_vsphere.pdf	Changing Data Compression Settings You can change data compression settings for existing jobs. After you change the settings, you will not need to create new full backups to use new settings. Veeam Backup & Replication will automatically apply the new compression level to the newly created backup files after you save the settings. Previously created backup files will not be affected.	757
	veeam_backup_11_0_user_guide_vsphere.pdf	You can change data deduplication settings for existing jobs. After you change the settings, you will not need to create new full backups to enable/disable the deduplication. Veeam Backup & Replication will automatically apply the change to the newly created backup	758

		files after you save the settings. Previously created backup files will not be affected.	
	veeam_backup_11_0_user_guide_vsphere.pdf	Local target (large blocks) Local target LAN target WAN target	759
5.3.8. Ter a funcionalidade para proteger localidades remotas, assegurando que a transmissão de dados através da WAN seja minimizada, provendo tanto desduplicação quanto replicação, enquanto possibilita recuperação granular de dados. A solução deve prover arquitetura flexível ao ponto de que a recuperação no escritório regional possa ser total (com todos os dados vindos do datacenter) ou parcial (com somente o envio dos dados que não estão em cache local);	veeam_backup_11_0_user_guide_vsphere.pdf	WAN Acceleration WAN acceleration is a Veeam technology that optimizes data transfer to remote locations. It is specific for off - site backup copy jobs and replication jobs. Off-site backup copy and replication always involve moving large volumes of data between remote sites. The most common problems that backup administrators encounter during off-site backup and replication are insufficient network bandwidth to support VM data traffic and transmission of redundant data. To solve these problems, Veeam Backup & Replication offers the WAN acceleration technology that combines: • Network traffic compression • Multistreaming upload • Global data deduplication • Variable block size deduplication. These technologies help optimize the data transfer and reduce the amount of data going over WAN	1749
5.3.8.1. Permite implementar um controle da quantidade de dados trafegados, seja limitando a quantidade de rede que a solução poderá utilizar, ou a vazão (throughput) máximo que a solução poderá utilizar para gravar no repositório.	veeam_backup_11_0_user_guide_vsphere.pdf	Enabling Traffic Throttling Traffic throttling setting in a network rule allows you to limit the impact of Veeam Backup & Replication tasks on network performance. Traffic throttling prevents jobs from utilizing the entire bandwidth available in your environment and makes sure that other network operations get enough traffic.	616
5.4. Funcionalidades de cópia e recuperação	N/A	N/A	N/A
5.4.1. Ser capaz de realizar cópia de arquivos abertos sem que a consistência deles seja comprometida;	veeam_backup_11_0_agent_management_guide.pdf	Application-Aware Processing If your computer runs VSS-aware applications, you can enable application-aware processing to create a transactionally consistent backup. The transactionally consistent backup guarantees	202

		proper recovery of applications without data loss.	
5.4.2. Possuir recursos avançados de agendamento de rotinas de proteção, para datas específicas, dias da semana recorrentes, dia do mês recorrente. Primeiro, segundo, terceiro e último dia do mês. Ser capaz de filtrar por mês e dia da semana;	veeam_backup_11_0_user_guide_vsphere.pdf	Define scheduling settings for the job: o To run the job at specific time daily, on defined week days or with specific periodicity, select Daily at this time. Use the fields on the right to configure the necessary schedule. o To run the job once a month on specific days, select Monthly at this time. Use the fields on the right to configure the necessary schedule. To run the job repeatedly throughout a day with a specific time interval, select Periodically every. In the field on the right, select the necessary time unit: Hours or Minutes. Click Schedule and use the time table to define the permitted time window for the job. In the Start time within an hour field, specify the exact time when the job must start.	862
5.4.3. Possuir a funcionalidade de paralelizar a gravação dos dados em dispositivos de armazenamento (funcionalidade conhecida como multiplexação);	veeam_backup_11_0_tapes_user_guide.pdf	Tape Parallel Processing You can use multiple drives simultaneously for writing data to tape. This option is useful if you have a lot of tape jobs running at the same time or you have a lot of data that must be written to tape in a limited backup window.	122
5.4.4. Ser capaz de enviar alertas através de e-mail com o objetivo de reportar eventos ocorridos na operação e configuração da solução;	veeam_backup_11_0_user_guide_vsphere.pdf	Specifying Email Notification Settings You can receive email notifications with results on jobs performed on the backup server.	652
5.4.5. A solução deverá permitir o transporte de dados de backup em infraestrutura de objetos, como S3;	veeam_backup_11_0_user_guide_vsphere.pdf	Capacity Tier Capacity tier is an additional tier of storage that can be attached to a scale-out backup repository. Data from the scale-out backup repository performance extents can be transported to the capacity tier for long-term storage. With capacity tier you can: • Move inactive backup chains to the capacity extent, as described in Moving Backups to Capacity Tier and Manually	525

		Moving Backups to Capacity Tier.	
5.4.5.1. A solução deverá estar licenciada para realizar o transporte dos dados para infraestruturas de objetos em nuvem pública e privada;	veeam_backup_11_0_user_guide_vsphere.pdf	Capacity Tier Capacity tier is an additional tier of storage that can be attached to a scale-out backup repository. Data from the scale-out backup repository performance extents can be transported to the capacity tier for long-term storage. With capacity tier you can: • Move inactive backup chains to the capacity extent, as described in Moving Backups to Capacity Tier and Manually Moving Backups to Capacity Tier. The capacity tier consists of a single capacity extent. The capacity extent can be either cloud -based object storage repository or on-premises object storage repository, such as: • S3-compatible object storage repository • Amazon S3 • Microsoft Azure Blob storage • Microsoft Azure Data Box • IBM Cloud Object Storage • Google Cloud Object Storage	525
5.4.5.1.1. Não se faz necessária a entrega dessa infraestrutura;	Ciente e de Acordo	N/A	N/A
5.4.5.2. Deverá ser compatível com, no mínimo, provedores de nuvem privada e pública, como:	N/A	N/A	N/A
5.4.5.2.1. Microsoft Azure;	veeam_backup_11_0_user_guide_vsphere.pdf	• S3-compatible object storage repository • Amazon S3 • Microsoft Azure Blob storage • Microsoft Azure Data Box • IBM Cloud Object Storage • Google Cloud Object Storage	525
5.4.5.2.2. AWS;	veeam_backup_11_0_user_guide_vsphere.pdf	• S3-compatible object storage repository • Amazon S3 • Microsoft Azure Blob storage • Microsoft Azure Data Box • IBM Cloud Object Storage • Google Cloud Object Storage	525
5.4.5.2.3. Nutanix Objects;	https://www.veeam.com/sys234	Alliance Product: Nutanix Objects	N/A

5.4.5.2.4. Dell EMC ECS;	https://www.veeam.com/sys219	Alliance Product: ECS	N/A
5.4.5.3. A solução deverá permitir a movimentação de dados para a nuvem (backup e restauração), de acordo com as políticas de backup implementadas. Não serão aceitas soluções que dependam de hardwares específicos para executar essa funcionalidade;	veeam_backup_11_0_user_guide_vsphere.pdf	Data Transfer Veeam Backup & Replication offers the following opportunities of data transfer to and from the capacity extent: • Copy policy: the backups are copied to the capacity extent automatically. • Move policy: the inactive backup chains can be transferred to the capacity extent. • Previously offloaded data can also be downloaded from the capacity tier back to the performance extents.	529
5.4.5.3.1. Todas as licenças necessárias à execução dessa funcionalidade deverão estar inclusas na solução;	Conforme Lista de Part Numbers	N/A	N/A
5.4.5.4. A solução deverá permitir a construção de um repositório de armazenamento de backup com escalabilidade horizontal, garantindo uma arquitetura híbrida entre nuvem privada e nuvem pública. Deverá permitir o uso simultâneo, com o propósito de criar uma entidade virtual de armazenamento, de storages, appliances de desduplicação e arquiteturas de nuvem;	veeam_backup_11_0_user_guide_vsphere.pdf	Scale-Out Backup Repository A scale-out backup repository is a repository system with horizontal scaling support for multi-tier storage of data. The scale-out backup repository consists of one or more backup repositories called performance tier, and can be expanded with object storage repositories for long-term and archive storage: capacity tier and archive tier. All the storage devices and systems inside the scale-out backup repository are joined into a system, with their capacities summarized. • It supports any backup target supported by Veeam: Windows or Linux servers with local or DAS storage, network shares, deduplicating storage appliances. All the features of any storage device or system are preserved. • It allows you to set up granular performance policy. For more information, see Backup File Placement. • It provides practically unlimited cloud-based storage capacity: you can instruct Veeam Backup & Replication to offload data from extents to the cloud for long-term storage. For details, see Capacity Tier.	516

5.4.5.4.1. A solução deverá permitir que o repositório de armazenamento escalável seja composto, concomitantemente, por armazenamento direto em Windows/Linux (SAN ou DAS), compartilhamentos de rede (NAS), equipamentos específicos para deduplicação (PBBA) e armazenamento de Objetos (S3 e HTTP) em nuvem pública e privada;	veeam_backup_11_0_user_guide_vsphere.pdf	Scale-Out Backup Repository It supports any backup target supported by Veeam: Windows or Linux servers with local or DAS storage, network shares, deduplicating storage appliances. All the features of any storage device or system are preserved. Capacity Tier Capacity tier is an additional tier of storage that can be attached to a scale-out backup repository. Data from the scale-out backup repository performance extents can be transported to the capacity tier for long-term storage. The capacity tier consists of a single capacity extent. The capacity extent can be either cloud -based object storage repository or on-premises object storage repository, such as: • S3-compatible object storage repository • Amazon S3 • Microsoft Azure Blob storage • Microsoft Azure Data Box • IBM Cloud Object Storage • Google Cloud Object Storage	516 525
5.4.5.4.2. A solução deverá permitir elencar, por características de desempenho dos repositórios, distintos níveis de armazenamento com o propósito de garantir estabilidade nos processos de backup e restauração de dados;	veeam_backup_11_0_user_guide_vsphere.pdf	A scale-out backup repository can comprise different tiers, or logical levels of storage. • Performance tier is the level used for the fast access to the data. It consists of one or more backup repositories called performance extents. • Capacity tier is an additional level for storing data that needs to be accessed less frequently. However, you still can restore your data directly from it. The capacity tier consists of one cloud -based or on-premises object storage repository called capacity extent. • Archive tier is an additional level for archive storage of infrequently accessed data. Applicable data from the capacity tier can be transported to the archive tier. For restore from the archive tier, data must undergo preparation process.	516 e 517

5.4.5.4.3. A solução deverá permitir a escolha do armazenamento contínuo das imagens de backup, de modo que um ciclo de backup possa estar presente em um único elemento da infraestrutura compartilhada (integralmente em um appliance de desduplicação), bem como em múltiplos elementos da infraestrutura compartilhada (backups completos em um appliance de deduplicação e backups incrementais em compartilhamentos NAS);	veeam_backup_11_0_user_guide_vsphere.pdf	Data Locality If you set the Data locality policy for a scale-out backup repository, all backup files that belong to the same backup chain are stored on the same extent of the scale-out backup repository. The Data locality policy does not put any limitations to backup chains. A new backup chain may be stored on the same extent or another extent. For example, if you create an active full backup, Veeam Backup & Replication may store the full backup file to another extent, and all dependent incremental backup files will be stored together with this full backup file. However, if you use a deduplicating storage appliance as an extent to the scale-out backup repository, Veeam Backup & Replication will attempt to place a new full backup (active or synthetic) to the extent where the full backup from the previous backup chain resides. Such behavior will help increase the data deduplication ratio.	521 e 522
	veeam_backup_11_0_user_guide_vsphere.pdf	Performance If you set the Performance policy for a scale-out backup repository, full backup files and incremental backup files that belong to the same backup chain are stored on different extents of the scale-out backup repository. If necessary, you can explicitly specify on which extents full backup files and incremental backup files must be stored.	522

5.4.5.4.4. A solução deverá validar diariamente, de modo automático, o estado dos distintos elementos de armazenamento que compõem o repositório compartilhado. A solução deverá validar o status de cada elemento, informando se eles estão online ou não, se os movimentadores de dados estão estáveis e qual o espaço de armazenamento remanescente no repositório compartilhado global;	veeam_backup_11_0_user_guide_vsphere.pdf	Receiving Scale-Out Backup Repository Reports Veeam Backup & Replication is capable of sending reports that contain information about processing results of your scale-out backup repositories data. • Reports are sent daily at time specified in global notification settings. • Reports are sent for all notification types selected in global notification settings, such as Success, Warning and Failure. • The title of a report is built up of "Scale-out Backup Repository" + a repository name. That said, if your scale-out backup repository name is Amazon, then the report title would be Scale-out Backup Repository Amazon. Description of Report Statuses The following table lists possible combinations of Warning and Error messages shown under the Status column of a report. Success Reports Warning Reports Error Reports	582 a 586
5.4.5.5. Ser capaz de enviar traps SNMP (Simple Network Management Protocol) com o objetivo de reportar eventos ocorridos na operação da solução;	veeam_backup_11_0_user_guide_vsphere.pdf	Specifying SNMP Settings You can receive SNMP traps with results on jobs performed on the backup server. You can use SNMP traps to feed data to other monitoring systems such as CA Unicenter, BMC Patrol, IBM Tivoli or HPE OpenView. SNMP traps can be sent to 5 different destinations.	656
5.4.5.6. A solução deverá permitir a restauração segura de imagens de backup, permitindo a criação de uma área específica, prévia à operação de recuperação, para a varredura de vírus ou malwares.	veeam_backup_11_0_user_guide_vsphere.pdf	How Secure Restore Works Veeam Backup & Replication uses the mount server as a staging server for scanning machine data with antivirus software. By default, the mount server role is assigned to one of the following machines: the Veeam backup server or a repository that stores machine backups. However, you can assign the mount server role to any 64-bit Microsoft Windows machine in your backup infrastructure.	1608

5.4.5.7. A solução deverá possuir um arquivo de configuração o qual deverá ser validado durante o processo de restauração para identificar qual software de varredura deverá ser ativado na análise de vírus ou malwares.	veeam_backup_11_0_user_guide_vsphere.pdf	Antivirus XML Configuration File The antivirus software that you plan to use for scanning backups is described in the AntivirusInfos.xml file. Veeam Backup & Replication creates this configuration file on every machine with the mount server role and stores the file in the %ProgramFiles%\Common Files\Veeam\Backup and Replication\MountService folder. During secure restore, Veeam Backup & Replication reads settings from the configuration file and triggers the antivirus to scan backup files. The settings in the file are already predefined for the following antivirus software: • Symantec Protection Engine • ESET • Windows Defender • Kaspersky Security 10	1609
5.4.5.8. Deverá ser compatível com fabricantes de varredura de vírus ou malwares como Symantec, ESET e Kaspersky.	veeam_backup_11_0_user_guide_vsphere.pdf	Antivirus XML Configuration File The antivirus software that you plan to use for scanning backups is described in the AntivirusInfos.xml file. Veeam Backup & Replication creates this configuration file on every machine with the mount server role and stores the file in the %ProgramFiles%\Common Files\Veeam\Backup and Replication\MountService folder. During secure restore, Veeam Backup & Replication reads settings from the configuration file and triggers the antivirus to scan backup files. The settings in the file are already predefined for the following antivirus software: • Symantec Protection Engine • ESET • Windows Defender • Kaspersky Security 10	1609

5.4.5.9. A console de gerenciamento da solução de backup deverá exibir os resultados da varredura efetuada pelo software terceiro de análise de vírus ou malwares.	veeam_backup_11_0_user_guide_vsphere.pdf	Viewing Malware Scan Results Results of the malware scan are available in restore session statistics. To view restore session statistics, do one of the following: • Open the Home view, in the inventory pane select Last 24 hours. In the working area, double-click the necessary restore session. Alternatively, you can select the session and click Statistics on the ribbon or right-click the session and select Statistics. • Open the History view, in the inventory pane select Restore. In the working area, double-click the necessary restore session. Alternatively, you can select the session and click Statistics on the ribbon or right-click the session and select Statistics.	1616
5.4.6. Possuir a funcionalidade de agendamento automático de tarefas de backup;	veeam_backup_11_0_user_guide_vsphere.pdf	Step 11. Define Job Schedule At the Schedule step of the wizard, select to run the backup job manually or schedule the job to run on a regular basis.	862
5.4.7. Para operações de dados gravadas em disco e fita, a solução de proteção deve possuir as seguintes funcionalidades:	N/A	N/A	N/A
5.4.7.1. Para um mesmo dado armazenado deve haver a possibilidade de configuração de diferentes períodos de retenção;	veeam_backup_11_0_user_guide_vsphere.pdf	Retention Policy for Backup Copy Jobs The retention policy of a backup copy job defines for how long Veeam Backup & Replication must retain copied restore points in the target backup repository. The retention policy of a backup copy job does not depend on retention policy of the source backup job. The backup copy job has its own retention policy settings. Short-Term Retention Policy The short-term retention policy allows to retain restore points created by backup copy jobs for a specified number of days or until the number of restore points reaches the specified number in the retention settings. Long-Term Retention Policy (GFS) The long-term or Grandfather-	1141 1141 1143

		Father-Son (GFS) retention policy allows you to store VM backups for long periods of time — for weeks, months and years. For this purpose, Veeam Backup & Replication creates synthetic or active full backup files and marks them with GFS flags.	
5.4.7.2. Para um dado armazenado deve haver a possibilidade de estender o período de retenção.	veeam_agent_windows_5_0_user_guide.pdf	Retention Policy for Outdated Backups In addition to specifying general retention policy settings, you can define retention policy for outdated backups. An outdated backup is a backup for which no new restore points were created within the last N days. Retention policy for outdated backups helps to avoid keeping redundant data in the target location: once the outdated backup retention period is over, an outdated backup is automatically removed from the target location.	51
5.4.7.3. Implementar a execução de cópias completas sintéticas ou similar, podendo implementar através de cópias do tipo eternamente incremental (Forever Incremental);	veeam_backup_11_0_user_guide_vsphere.pdf	Synthetic Full Backup In some situations, running active full backups periodically may not be an option. Active full backups are resource-intensive and consume considerable amount of network bandwidth. As an alternative, you can create synthetic full backups.	729
	veeam_backup_11_0_user_guide_vsphere.pdf	Forever Forward Incremental Backup Retention Policy If the number of days or restore points in forever forward incremental backup chains exceeds retention policy settings, Veeam Backup & Replication transforms the backup chain to make room for the most recent restore point.	734
5.4.7.3.1. Uma cópia completa sintética é gerada através de uma outra cópia completa tradicional (não sintetizada) anterior e de cópias incrementais ou diferenciais subsequentes ou de um	veeam_backup_11_0_user_guide_vsphere.pdf	How Synthetic Full Backup Works	731

backup incremental cumulativo. A cópia sintetizada deverá ser capaz de restaurar arquivos e diretórios da mesma maneira que um cliente faz a restauração de uma cópia tradicional;			
5.4.7.4. Permitir a gravação de cópias do tipo Disco-Para-Disco-Para-Unidade de Fita;	veeam_backup_11_0_tapes_user_guide.pdf	Linking Backup Jobs to Backup to Tape Jobs Veeam Backup & Replication provides two options for linking backup jobs to backup to tape jobs: • If you already have backup jobs configured, you can choose the necessary jobs in the Backup to Tape Job Wizard. For details, see Creating Backup to Tape Jobs. • Alternatively, you can point a backup job to an existing backup to tape job using the Backup Job wizard.	153
5.4.7.5. Ser compatível com bibliotecas auto-carregadoras de cartuchos de fitas magnéticas;	veeam_backup_11_0_tapes_user_guide.pdf	Supported Devices and Configuration Veeam Backup & Replication supports Linear Tape-Open tape libraries starting from generation 3 (LTO3) or later: • Physical libraries, standalone drives and virtual tape libraries. • Partitions of the physical or virtual tape libraries presented to the Veeam backup server. Within a library partition, only one generation of LTO drives are supported.	8
5.4.7.6. Possuir a funcionalidade de criar múltiplas cópias de backups armazenados, com a opção de recuperação dos dados de forma automática através da cópia secundária se a cópia primária não estiver mais disponível.	veeam_backup_11_0_user_guide_vsphere.pdf	Linking Backup Jobs to Backup Copy Jobs You can link backup jobs to backup copy jobs. This option lets you create a secondary target for the backup job and store backups created with the backup job in the secondary backup repository.	1203
	veeam_backup_11_0_user_guide_vsphere.pdf	About Backup Copy With backup copy, you can create several instances of the same backup file and copy them to secondary (target) backup repositories for long-term storage.	1127
5.5. Funcionalidades da console de gerenciamento, integração e alta-disponibilidade	N/A	N/A	N/A

5.5.1. Possuir interface que seja capaz de gerenciar e executar operações de proteção e recuperação dos sistemas operacionais Windows, Unix e Linux; ambientes de virtualização VMware e Acropolis Operating System; aplicações como Microsoft Active Directory e banco de dados Microsoft SQL Server, Oracle (Windows e Linux) e Oracle RAC (em Linux);	veeam_backup_11_0_enterprise_manager_user_guide.pdf	About Veeam Backup Enterprise Manager Veeam Backup Enterprise Manager (Enterprise Manager) is a management and reporting component that allows you to manage multiple Veeam Backup & Replication installations from a single web console.	10
	veeam_backup_11_0_enterprise_manager_user_guide.pdf	: • Manage jobs across multiple Veeam backup servers. • View on-going reporting data for all jobs running on these servers, set up email notifications to get information on the status of all jobs. • Search for machines, file shares, and guest files in backups and replicas. • Perform recovery operations for VMs and physical machines, including 1-Click restore, 1-click guest OS file restore and application items restore (for Microsoft Exchange mailboxes, Microsoft SQL Server databases and Oracle databases); perform 1-Click restore for file share backups. • Centrally manage and update licenses to ensure compliance. • Delegate permissions for restore operations to personnel in charge. • Manage VMware vCloud Director organizations and support their administrators with the Veeam SelfService Backup Portal. • Manage vSphere user accounts and support them with the vSphere Self-Service Backup Portal. • Install vSphere Web Client plug-in on vCenter servers. • Implement data encryption and decryption processes for the Veeam solutions. • Provide operation automation with Veeam Backup Enterprise Manager REST API	10

	veeam_backup_for_nutanix_ahv_3_0_user_guide.pdf	2. Nutanix AHV Plug-in: an installer component that deploys AHV Backup Proxy services on the Veeam Backup & Replication server. After you install Nutanix AHV Plug-in, you can add AHV clusters and AHV Backup Proxy instances to the Veeam Backup & Replication infrastructure. 3. AHV Backup Proxy: the main component that manages backup and restore processes. AHV Backup Proxy allows you to create VM backups, snapshots and restore VMs from backups to an AHV cluster. o AHV Backup Proxy web console: a web console that allows you to manage backup and restore operations, configure notifications and configure settings of AHV Backup Proxy.	11
5.5.2. O acesso administrativo ao console do servidor de gerenciamento da solução poderá ser feito através de ferramenta disponibilizada no próprio software (console gráfico) ou através de navegador Web;	veeam_backup_11_0_enterprise_manager_user_guide.pdf	Configuring Accounts and Roles Veeam Backup Enterprise Manager implements security based on user roles by limiting access to features and data. This empowers the administrator to delegate permissions in a granular way, on an as-needed basis. For example, the administrator can grant permissions to another user to recover files without being able to see the content of the files. Administrators grant users and groups access to Enterprise Manager by adding accounts. When adding an account, administrators assign a role to the account to provide it with permissions. Enterprise Manager offers the following roles: • Portal Administrator • Portal User • Restore Operator	100
5.5.3. Suportar cópia de segurança dos arquivos de catálogo e configuração, para promover recuperação dos serviços de gerenciamento no evento de falhas;	veeam_backup_for_nutanix_ahv_3_0_user_guide.pdf	Configuration Backup and Restore You can back up and restore the configuration database that Veeam Backup & Replication uses. If the backup server fails for some reason, you can re-install the backup server and quickly restore its configuration from the configuration backup.	676

5.5.4. Suportar unificação de autenticação (single sign on - SSO), permitindo a integração com o Microsoft Active Directory. A funcionalidade de integração com o Active Directory deverá permitir a definição granular das permissões administrativas aos recursos, objetos e servidores definidos naconfiguração do software;	veeam_backup_11_0_enterprise_manager_user_guide.pdf	Administrators can add accounts to Veeam Backup Enterprise Manager to grant users access to the website. Enterprise Manager offers the following account types: User, Group, External User and External Group. Type User Description Local or AD user How to Sign In By specifying a user name and password Name Format DOMAIN/Username Domain is optional	100
	veeam_backup_11_0_enterprise_manager_user_guide.pdf	SAML Authentication Support Veeam Backup Enterprise Manager supports single sign-on authentication based on the SAML 2.0 protocol. Enterprise organizations who use a single sign-on (SSO) service in their IT infrastructure can extend single signon capabilities to Veeam Backup Enterprise Manager. Once a user of the organization is logged in to the single sign-on service, the user can access Veeam Backup Enterprise Manager without the need to provide their credentials	29
5.5.5. A base de dados para armazenamento do catálogo deverá possuir mecanismo de proteção (backup) das informações armazenadas no catálogo e funcionalidades de recuperação rápida do catálogo em caso de desastre.	veeam_backup_for_nutanix_ahv_3_0_user_guide.pdf	Configuration Backup and Restore You can back up and restore the configuration database that Veeam Backup & Replication uses. If the backup server fails for some reason, you can re-install the backup server and quickly restore its configuration from the configuration backup.	676
5.6. Suporte à Criptografia:	veeam_backup_11_0_user_guide_vsphere.pdf	Data Encryption	1777
5.6.1. Implementar criptografia de dados na origem (cliente ou proxy de backup), de uma forma que seja garantido que o dado que trafegará na rede local ou na rede WAN seja criptografado;	veeam_backup_11_0_user_guide_vsphere.pdf	Veeam Backup & Replication uses the block cypher encryption algorithm. Encryption works at the source side. Veeam Backup & Replication reads VM or file data, encodes data blocks, transfers them to the target side in the encrypted format and stores the data to a file in the backup repository or archives the data to tape.	1777

5.6.2. Criptografia de dados no destino (servidor de backup);	veeam_backup_11_0_user_guide_vsphere.pdf	NOTE Veeam Backup & Replication will pass encryption keys to the target backup repository or cloud repository in the following cases: • If you run a backup copy job over WAN accelerators • If you perform health check for the encrypted backup files	1777
	veeam_backup_11_0_user_guide_vsphere.pdf	If you want to encrypt Veeam Agent backup files stored in the backup repository, select the Encrypt backups stored in this repository check box and choose the necessary password from the field below. If you have not specified a password beforehand, click Add on the right or the Manage passwords link to add a new password. Veeam Backup & Replication will encrypt files at the backup repository side using its built-in encryption mechanism.	475
5.6.3. Implementar no mínimo chaves de criptografia de 256 bits para cifrar os dados;	veeam_backup_11_0_user_guide_vsphere.pdf	Beside the job-level encryption, Veeam Backup & Replication allows you to encrypt network traffic going between the primary site and the disaster recovery site. Network traffic encryption is configured as part of global network traffic rules that are set for backup infrastructure components. For network traffic encryption, Veeam Backup & Replication uses the 256-bit Advanced Encryption Standard (AES).	1777
5.6.4. Implementar pares de chaves de criptografia de 4096 bits para recuperação de desastres;	veeam_backup_11_0_user_guide_vsphere.pdf	Enterprise Manager Keys • To generate Enterprise Manager keys required for data restore without a password, Veeam Backup & Replication uses the RSA algorithm with a 4096-bit key length. • To generate a request for data restore from a backup server, Veeam Backup & Replication uses the RSA algorithm with a 2048-bit key length.	1779
5.7. Suportar protocolos IPv4 e IPv6 para rotinas de backup;	veeam_backup_11_0_user_guide_vsphere.pdf	Network Consider the following requirements and limitations: • Domain names of all managed servers added to the Veeam backup infrastructure and machines you want to	28

		back up must be resolvable into IPv4 addresses.	
5.8. Integração com as seguintes aplicações para cópia e restauração	N/A	N/A	N/A
5.8.1. Realizar proteção e recuperação dos seguintes sistemas operacionais, aplicações, banco de dados e ambientes de virtualização:	N/A	N/A	N/A
5.8.1.1. Microsoft Windows 7 SP1, 8.1, 10, Server 2008 R2 SP1, 2012, 2012 R2, 2016 e 2019;	veeam_backup_11_0_agent_management_guide.pdf	• Microsoft Windows Server 2022 • Microsoft Windows Server 2019 • Microsoft Windows Server 2016 • Microsoft Windows Server Semi-Annual Channel (from version 1803 to version 20H2) • Microsoft Windows Server 2012 R2 • Microsoft Windows Server 2012 • Microsoft Windows Server 2008 R2 SP14 • Microsoft Windows 11 (version 21H2) • Microsoft Windows 10 Semi-Annual Channel (from version 1803 to version 21H2)5 • Microsoft Windows 10 Long-Term Servicing Channel (versions 1507, 1607 and 1809) • Microsoft Windows 8.1 • Microsoft Windows 7 SP1	61
5.8.1.2. Oracle Linux 6.x ou superiores;	veeam_backup_11_0_agent_management_guide.pdf	• Oracle Linux 6 – 8.5 (RHCK) • Oracle Linux 6 (starting from UEK R2) – Oracle Linux 8 (up to UEK R6 U2, kernel version 5.4.17-2102)	63
5.8.1.3. Red Hat Enterprise Linux 6.x ou superiores;	veeam_backup_11_0_agent_management_guide.pdf	• RHEL 6.0 – 8.5	63
5.8.1.4. Ubuntu 16.x ou superiores;	veeam_backup_11_0_agent_management_guide.pdf	• Ubuntu 14.04, 16.04, 18.04, 20.04, 20.10, 21.04, 21.10	63
5.8.1.5. Debian 8.x ou superiores;	veeam_backup_11_0_agent_management_guide.pdf	• Debian 9.0 – 11.1	63
5.8.1.6. Microsoft Active Directory 2012 ou superiores;	veeam_backup_11_0_agent_management_guide.pdf	Microsoft Active Directory Domain Controllers • Microsoft Windows Server 2019	74
5.8.1.7. Microsoft SQL Server 2012 ou superiores;	veeam_backup_11_0_agent_management_guide.pdf	Microsoft SQL Server • Microsoft SQL Server 2019	75
5.8.1.8. Oracle 11g R2 ou superiores (Linux ou Windows);	veeam_backup_11_0_agent_management_guide.pdf	Oracle Oracle Database versions 11g to 21c are supported for the	75

		following operating systems (32-bit and 64-bit architecture):	
5.8.1.9. MySQL 5.6.x ou superiores;	veeam_backup_11_0_agent_management_guide.pdf	MySQL • MySQL database system versions 5.6 to 8.0 are supported.	76
5.8.1.10. PostgreSQL 9.4 ou superiores;	veeam_backup_11_0_agent_management_guide.pdf	PostgreSQL • PostgreSQL database system versions 9.5 to 14.0 are supported.	76
5.8.1.11. VMware ESX/ESXi 6.0 ou superiores;	veeam_backup_11_0_user_guide_vsphere.pdf	• vSphere 7.0 (up to 7.0 U31)	23
5.8.1.12. Nutanix 5.10 ou superiores.	veeam_backup_for_nutanix_ahv_3_0_user_guide.pdf	Platform AHV Backup Proxy supports the following versions of the Nutanix AOS: • Nutanix AOS versions 5.5.X up to 6.0.X. Instant Recovery to Nutanix AHV is supported only for Nutanix AOS 6.0 and later.	13
5.9. Suporte ao Active Directory	veeam_backup_11_0_agent_management_guide.pdf	Specifying Active Directory Objects	114
5.9.1. Executar cópia em tempo de execução do Microsoft Active Directory;	veeam_backup_11_0_agent_management_guide.pdf	Specifying Active Directory Objects	
5.9.1.1. Possibilitar as seguintes opções de recuperação:	N/A	N/A	N/A
5.9.1.2. Recuperação de um objeto;	veeam_backup_11_0_explorers_user_guide.pdf	Veeam Explorer for Microsoft Active Directory allows you to restore and export Active Directory objects and containers from backups created in Veeam Backup & Replication.	12
5.9.1.3. Recuperação de um atributo;	veeam_backup_11_0_explorers_user_guide.pdf	Right-click an object and select View attributes to view object attributes. To copy necessary attributes to the clipboard, select an attribute and press CTRL+C on the keyboard. Multiple selection is also supported.	23
5.9.1.4. Recuperação de um atributo deletado de um objeto.	veeam_backup_11_0_explorers_user_guide.pdf	Right-click an object and select View attributes to view object attributes. To copy necessary attributes to the clipboard, select an attribute and press CTRL+C on the keyboard. Multiple selection is also supported.	23
5.10. Suporte a Oracle e Oracle RAC	N/A	N/A	N/A
5.10.1. Deverá executar proteção e recuperação de base da dados Oracle e Oracle RAC com as seguintes características nativas ou não:	N/A	N/A	N/A

5.10.1.1. Executar proteção e recuperação das bases de dados do Oracle/Oracle RAC via RMAN e sem parada do banco;	veeam_backup_11_0_plug-ins_user_guide.pdf	• Veeam Plug-in for Oracle RMAN: for transactionally-consistent RMAN-based backups of Oracle databases. Veeam Plug-in for Oracle RMAN Veeam Plug-in uses the backup and restore functionality of Oracle Recovery Manager (RMAN) and transfers backups to Veeam backup repositories. How Veeam Plug-in for Oracle RMAN Works Veeam Plug-in functions as an agent between Oracle RMAN and Veeam backup repository. By default, RMAN sends backups to a native RMAN location on disk (DEFAULT DEVICE TYPE TO DISK). When you configure the Veeam Plug-in, the default device type is changed to SBT_TAPE, which gives control over backup media management to Veeam Plug-in. Thus, after you deploy Veeam Plug-in on an Oracle server, you can perform all backup and restore operations in the Oracle RMAN console. Veeam Plug -in compresses, deduplicates database backups and transfers them to a backup repository connected to Veeam Backup & Replication.	98 99
5.10.1.2. Executar arquivamento do registro de eventos (log) possibilitando a criação de rotina de cópia para que ocorra com intervalos de 1 (uma) hora;	veeam_backup_11_0_user_guide_vsphere.pdf	8. To back up Oracle archived logs with Veeam Backup & Replication, select the Backup log every <N> minutes check box and specify the frequency for archived log backup. By default, archived logs are backed up every 15 minutes. The maximum log backup interval is 480 minutes.	857
5.10.1.3. Permitir a cópia do arquivamento de transações (archives logs) baseados na quantidade de arquivamento (archives);	veeam_backup_11_0_user_guide_vsphere.pdf	Retain Logs with Image-Level Backup By default, Veeam Backup & Replication retains archived log backups together with the corresponding imagelevel backup of the Oracle VM. Veeam Backup & Replication retains VM backup and log backups according to the short-term retention configured for VM backups. When Veeam Backup & Replication removes a restore point of the image-level backup from the backup chain,	801

		it also removes a chain of archived logs relating to this imagelevel backup.	
	veeam_backup_11_0_user_guide_vsphere.pdf	Log Shipping Servers For every Oracle VM whose archived logs you want to back up, Veeam Backup & Replication defines how to ship logs to the backup repository. Archived logs can be transported in the following ways: • Directly from the VM guest OS to the backup repository. This method is recommended — it does not involve additional resources and puts less load on the VM guest OS. • Via log shipping servers. If it is not possible to establish a direct connection between the VM guest OS and backup repository, you can configure Veeam Backup & Replication to use a log shipping server.	802
	veeam_backup_11_0_user_guide_vsphere.pdf	Log shipping servers are assigned per job session. When a new job session starts, Veeam Backup & Replication detects log shipping servers anew. Veeam Backup & Replication can also re-detect available servers during the job session. If a log shipping server becomes unavailable for some reason, Veeam Backup & Replication will fail over to another log shipping server.	802
	veeam_backup_11_0_user_guide_vsphere.pdf	Log Files At each start of the Oracle backup job ('parent'), a new .VLB is created to store log backups in the repository: • If the Use per-machine backup files option is selected for the repository, then Veeam will create a separate .VLB for each server processed by the job.	806
	veeam_backup_11_0_user_guide_vsphere.pdf	As described in the section above, during database log backup ('child') job session, log archiving is performed by native means of the Oracle server. Archived logs are stored to a temporary folder on the Oracle VM guest file system. Then Veeam copies	806

		archived log to the current .VLB in the repository. When the new 'parent' job session starts, another .VLB is created, and the archived log files that appear after that will be stored there during the 'child' job session.	
	veeam_backup_11_0_user_guide_vsphere.pdf	Backup Job Scheduling You can start backup jobs manually or schedule them to start automatically at specific time. Veeam Backup & Replication lets you configure the following settings for the job: • Scheduling settings • Job retry settings • Backup window settings	807
	veeam_backup_11_0_user_guide_vsphere.pdf	Short-Term Retention Policy Every successful backup job session creates a new restore point that lets you roll back VM data to an earlier point in time. To control the number of restore points in the backup chain, you must specify retention policy settings. The retention policy defines how many restore points you want to retain on disk and thus how 'far' you are able to roll back.	732
5.10.1.4. Permitir a configuração que após a cópia dos registros de transações (archives logs) os mesmos sejam mantidos ou deletados;	veeam_backup_11_0_user_guide_vsphere.pdf	Select Do not delete archived logs if you want Veeam Backup & Replication to preserve archived logs on the VM guest OS. Select Delete logs older than <N> hours or Delete logs over <N> GB if you want Veeam Backup & Replication to delete archived logs that are older than <N> hours or larger than <N> GB.	857
5.10.1.5. Além da proteção do Banco, a solução deverá proteger a área de catálogo, control file e sp file.	veeam_backup_11_0_plug-ins_user_guide.pdf	Backup of Control File and SPFILE If you set CONFIGURE CONTROLFILE AUTOBACKUP to ON, then RMAN automatically creates a control file and an SPFILE backup after you run the BACKUP command. For details, see the Oracle Documentation. Oracle Recovery Catalog You can back up the Oracle Recovery Catalog with Veeam Plug-in on the Recovery Catalog server according to the Oracle procedures. For details, see the Managing a Recovery	112 e 113

		Catalog section of the Database Backup and Recovery User's Guide.	
5.10.1.6. Possibilitar a recuperação com as seguintes características:	N/A	N/A	N/A
5.10.1.7. Recuperação completa da Base de dados no mesmo servidor	veeam_backup_11_0_plug-ins_user_guide.pdf	Database Recovery With the configured Veeam Plug-in for Oracle RMAN, you can perform all kinds of database restore operations available in Oracle RMAN. You can also restore from database backups in the Veeam Backup & Replication console. • Restore to Original Server • Restore to Another Server	154
5.10.1.8. Recuperação completa da Base de dados em outro servidor	veeam_backup_11_0_plug-ins_user_guide.pdf	Database Recovery With the configured Veeam Plug-in for Oracle RMAN, you can perform all kinds of database restore operations available in Oracle RMAN. You can also restore from database backups in the Veeam Backup & Replication console. • Restore to Original Server • Restore to Another Server	154
5.10.1.9. Recuperação de um datafile específico	veeam_backup_11_0_explorers_user_guide.pdf	When restoring with the Restore with different name and settings option, the location for the following files will be available for editing: o Control files o Data files o Log files o Temp files	262
5.10.1.10. Recuperação granular no nível de tabela	veeam_backup_11_0_explorers_user_guide.pdf	Data Publishing Publishing databases allows you to temporarily attach large Oracle databases to the target Oracle server without having to restore them. Publishing databases typically occurs faster than using standard restore features and might be convenient in certain cases, for example, when your time to perform disaster-recovery operations is limited. During publishing, Veeam Explorer for Oracle does the following: 1. Mounts VM disks from the backup file to the target Oracle	264

		server. Mounting is done to the C:\VeeamFLR directory for Windows-based servers and the /tmp directory for Linux-based servers. 2. Retrieves required database files. 3. Attaches associated databases directly to the target Oracle server so that you can perform required operations using Oracle tools.	
5.10.1.11. Recuperação em um momento do tempo específico;	veeam_backup_11_0_explorers_user_guide.pdf	Step 6. Specify Point in Time At this step of the wizard, do the following: • Select the Recover to the latest available option to restore data as of the latest available state. • Select the Recover to the following point in time option to select a state as of which you want to restore data: o Date. Select a date as of which you want to restore data. o SCN. Specify SCN (System Change Number). o Sequence. Specify the sequence.	261
5.11. Suporte a Microsoft SQL Server	veeam_backup_11_0_explorers_user_guide.pdf	Veeam Explorer for Microsoft SQL Server	60
5.11.1. Executar proteção e recuperação de base dos dados Microsoft SQL Server com as seguintes características nativas ou não:	veeam_backup_11_0_explorers_user_guide.pdf	Veeam Explorer for Microsoft SQL Server	60
5.11.1.1. Executar proteção e recuperação de bases de dados Microsoft SQL Server sem parada do banco;	veeam_backup_11_0_user_guide_vsphere.pdf	Application-Aware Processing In this article To create transactionally consistent backups or replicas of VMs that run Microsoft Active Directory, Microsoft SQL Server, Microsoft SharePoint, Microsoft Exchange or Oracle, you must enable application-aware processing in job settings. Application-aware processing is the Veeam technology based on Microsoft VSS. Microsoft VSS is responsible for quiescing applications on the VM and creating a consistent view of application data on the VM guest OS.	777

5.11.1.2. Executar cópia de registro de transações (transaction log) possibilitando a criação de rotina de cópia para que ocorra com intervalos de 1 (uma) hora;	veeam_backup_11_0_user_guide_vsphere.pdf	Transaction Log Truncation If you back up or replicate virtualized database systems that use transaction logs, for example, Microsoft Exchange or Microsoft SQL Server, you can instruct Veeam Backup & Replication to truncate transaction logs so that logs do not overflow the storage space on the VM. Veeam Backup & Replication provides the following options of transaction logs handling:	781
	veeam_backup_11_0_user_guide_vsphere.pdf	Sessions of Transaction Log Backup Jobs The transaction log backup job runs permanently in the background, shipping transaction logs to the backup repository at a specific time interval (by default, every 15 minutes). A sequence of time intervals between sessions of the parent backup job makes up a session of the transaction log backup job.	768
5.11.1.3. Permitir a configuração que após a cópia dos registros de transações (transaction log) os mesmos sejam mantidos ou deletados;	veeam_backup_11_0_user_guide_vsphere.pdf	Transaction Log Truncation If you back up or replicate virtualized database systems that use transaction logs, for example, Microsoft Exchange or Microsoft SQL Server, you can instruct Veeam Backup & Replication to truncate transaction logs so that logs do not overflow the storage space on the VM. Veeam Backup & Replication provides the following options of transaction logs handling: • Truncate logs • Do not truncate logs • Back up logs periodically	781
5.11.1.4. A solução deverá possibilitar a recuperação com as seguintes características:	N/A	N/A	N/A
5.11.1.5. Recuperação completa da base de dados no mesmo servidor	veeam_backup_11_0_explorers_user_guide.pdf	Browsing for Servers To browse for a server, do the following: • On the Local Servers tab, select a local SQL Server that is located on a machine with Veeam Explorer. • On the Network Server tab, select an SQL Server over the network.	92
5.11.1.6. Recuperação completa da base de dados em outro servidor	veeam_backup_11_0_explorers_user_guide.pdf	Browsing for Servers To browse for a server, do the following: • On the Local Servers tab, select a local SQL Server that is located on a machine with	92

		Veeam Explorer. • On the Network Server tab, select an SQL Server over the network.	
5.11.1.7. Recuperação de uma base específica	veeam_backup_11_0_explorers_user_guide.pdf	Restoring Single Database To restore a single Microsoft SQL Server database, do the following:	
5.11.1.8. Recuperação granular no nível de tabela	veeam_backup_11_0_explorers_user_guide.pdf	Data Publishing Publishing databases allows you to temporarily attach large SQL Server databases to the target Microsoft SQL Server without having to restore them. Publishing databases typically occurs faster than using standard restore features and might be convenient in certain cases, for example, when your time to perform disaster recovery operations is limited.	117
5.11.1.9. Recuperação em um momento do tempo específico;	veeam_backup_11_0_explorers_user_guide.pdf	Step 2. Specify Restore Point At this step of the wizard, select a state as of which you want to restore your database: • Select the Restore to the point in time of the selected image-level backup option to load database files as per date of the current restore point. • Select the Restore to a specific point in time option to load database files as per selected point in time. Use the slider control to choose a point you need. • Select the Perform restore to the specific transaction check box to load database files exactly as of the moment before undesired transactions.	89
5.12. Suporte a PostgreSQL	N/A	N/A	N/A
5.12.1. Executar proteção e recuperação de base da dados PostgreSQL Server com as seguintes características nativas ou não:	veeam_backup_11_0_agent_management_guide.pdf	PostgreSQL Processing Settings You can specify how Veeam Agent for Linux must process a PostgreSQL database: 1. At the Guest Processing step of the wizard, select the Enable application-aware processing	263
5.12.1.1. Cópia em tempo de execução do banco de dados;	veeam_backup_11_0_agent_management_guide.pdf	PostgreSQL Backup You can use Veeam Agent for Linux to create transactionally consistent backups of Veeam Agent machines that run the PostgreSQL database system.	42

5.12.1.2. Permitir a recuperação completa;	veeam_agent_linux_5_0_user_guide.pdf	Volume-Level Restore If data on a computer volume gets corrupted, you can restore this volume from the backup. For volume-level restore, you can use backups that were created at the volume level. File-level backups cannot be used for volume restore. When you perform volume-level restore, Veeam Agent for Linux restores the entire content of the volume. It retrieves from the backup data blocks pertaining to a specific volume and copies them to the necessary location. File-Level Restore If you have lost or modified files and directories on your computer by mistake, you can restore a copy of the necessary objects from the backup. For file-level restore, you can use a backup of any type: • Volume-level backup • File-level backup Veeam Agent for Linux does not simply extract files and folders from the backup file. During file-level restore, Veeam Agent for Linux performs the following operations: 1. Veeam Agent for Linux associates the backup file with a loop device, for example, /dev/loop0, to make the backup file accessible as a block device. 2. Veeam Agent for Linux mounts the loop device to the mount point directory in the computer's file system. o For file-level restore with the Veeam Agent for Linux control panel or Veeam Recovery Media, Veeam Agent for Linux mounts the backup content to the /mnt/backup directory. o For file-level restore with the command line interface, you can specify a directory in which Veeam Agent for Linux should mount the backup content. After the backup content is mounted, you can use Linux command line utilities or preferred file browser to work with restored files and directories. You can browse for	44 45

		files and directories in the mounted backup and copy them to their initial location or to a new location.	
	veeam_agent_linux_5_0_user_guide.pdf	Volume-Level Backup You can set up Veeam Agent for Linux to create volume-level backup. The volume-level backup captures the whole image of a data volume on your computer. You can use the volume-level backup to restore a computer volume, specific files and folders on the volume or perform bare metal recovery. File-Level Backup You can set up Veeam Agent for Linux to create file-level backup. The file-level backup captures only data of individual directories and files on the computer. You can use the file-level backup to restore files and directories that you have added to the backup scope.	15 17

5.12.1.3. Restaurar a base de dados ou seus arquivos no mesmo servidor em caminho diferente;	veeam_agent_linux_5_0_user_guide.pdf	Step 5. Specify Backup Location Settings Specify settings for the target storage that contains a backup file from which you plan to restore data: • Specify shared folder settings — if you have selected the Add shared folder option at the Select backup location step of the wizard. • Select local drive — if you have selected the Mount local disk option at the Select backup location step of the wizard. • Specify Veeam backup repository settings — if you have selected the Add VBR server option at the Select backup location step of the wizard. • Specify Veeam Cloud Connect repository settings — if you have selected the Add Cloud provider option at the Select backup location step of the wizard.	313
	veeam_agent_linux_5_0_user_guide.pdf	Step 5. Specify Backup Location Settings Specify settings for the target storage that contains a backup file from which you plan to restore data: • Specify shared folder settings — if you have selected the Add shared folder option at the Select backup location step of the wizard. • Select local drive — if you have selected the Mount local disk option at the Select backup location step of the wizard. • Specify Veeam backup repository settings — if you have selected the Add VBR server option at the Select backup location step of the wizard. • Specify Veeam Cloud Connect repository settings — if you have selected the Add Cloud provider option at the Select backup location step of the wizard.	273
5.12.1.4. Restaurar uma instância ou seus arquivos em um outro servidor.	veeam_agent_linux_5_0_user_guide.pdf	Step 5. Specify Backup Location Settings Specify settings for the target storage that contains a backup file from which you plan to restore data: • Specify shared folder settings	313

		— if you have selected the Add shared folder option at the Select backup location step of the wizard. • Select local drive — if you have selected the Mount local disk option at the Select backup location step of the wizard. • Specify Veeam backup repository settings — if you have selected the Add VBR server option at the Select backup location step of the wizard. • Specify Veeam Cloud Connect repository settings — if you have selected the Add Cloud provider option at the Select backup location step of the wizard.	
	veeam_agent_linux_5_0_user_guide.pdf	Step 5. Specify Backup Location Settings Specify settings for the target storage that contains a backup file from which you plan to restore data: • Specify shared folder settings — if you have selected the Add shared folder option at the Select backup location step of the wizard. • Select local drive — if you have selected the Mount local disk option at the Select backup location step of the wizard. • Specify Veeam backup repository settings — if you have selected the Add VBR server option at the Select backup location step of the wizard. • Specify Veeam Cloud Connect repository settings — if you have selected the Add Cloud provider option at the Select backup location step of the wizard.	273
5.13. Suporte a MySQL	N/A	N/A	N/A
5.13.1. Executar proteção e recuperação de base da dados MySQL Server com as seguintes características nativas ou não:	N/A	N/A	N/A
5.13.1.1. Cópia em tempo de execução do banco de dados;	veeam_backup_11_0_agent_management_guide.pdf	MySQL Backup You can use Veeam Agent for Linux to create transactionally consistent backups of Veeam Agent machines that run the MySQL database system.	40

5.13.1.2. Permitir a recuperação completa;	veeam_agent_linux_5_0_user_guide.pdf	Volume-Level Restore If data on a computer volume gets corrupted, you can restore this volume from the backup. For volume-level restore, you can use backups that were created at the volume level. File-level backups cannot be used for volume restore. When you perform volume-level restore, Veeam Agent for Linux restores the entire content of the volume. It retrieves from the backup data blocks pertaining to a specific volume and copies them to the necessary location. File-Level Restore If you have lost or modified files and directories on your computer by mistake, you can restore a copy of the necessary objects from the backup. For file-level restore, you can use a backup of any type: • Volume-level backup • File-level backup Veeam Agent for Linux does not simply extract files and folders from the backup file. During file-level restore, Veeam Agent for Linux performs the following operations: 1. Veeam Agent for Linux associates the backup file with a loop device, for example, /dev/loop0, to make the backup file accessible as a block device. 2. Veeam Agent for Linux mounts the loop device to the mount point directory in the computer's file system. o For file-level restore with the Veeam Agent for Linux control panel or Veeam Recovery Media, Veeam Agent for Linux mounts the backup content to the /mnt/backup directory. o For file-level restore with the command line interface, you can specify a directory in which Veeam Agent for Linux should mount the backup content. After the backup content is mounted, you can use Linux command line utilities or preferred file browser to work with restored files and directories. You can browse for	44 45

		files and directories in the mounted backup and copy them to their initial location or to a new location.	
	veeam_agent_linux_5_0_user_guide.pdf	Volume-Level Backup You can set up Veeam Agent for Linux to create volume-level backup. The volume-level backup captures the whole image of a data volume on your computer. You can use the volume-level backup to restore a computer volume, specific files and folders on the volume or perform bare metal recovery. File-Level Backup You can set up Veeam Agent for Linux to create file-level backup. The file-level backup captures only data of individual directories and files on the computer. You can use the file-level backup to restore files and directories that you have added to the backup scope.	15 17

5.13.1.3. Restaurar a base de dados ou seus arquivos no mesmo servidor;	veeam_agent_linux_5_0_user_guide.pdf	Step 5. Specify Backup Location Settings Specify settings for the target storage that contains a backup file from which you plan to restore data: • Specify shared folder settings — if you have selected the Add shared folder option at the Select backup location step of the wizard. • Select local drive — if you have selected the Mount local disk option at the Select backup location step of the wizard. • Specify Veeam backup repository settings — if you have selected the Add VBR server option at the Select backup location step of the wizard. • Specify Veeam Cloud Connect repository settings — if you have selected the Add Cloud provider option at the Select backup location step of the wizard.	313
	veeam_agent_linux_5_0_user_guide.pdf	Step 5. Specify Backup Location Settings Specify settings for the target storage that contains a backup file from which you plan to restore data: • Specify shared folder settings — if you have selected the Add shared folder option at the Select backup location step of the wizard. • Select local drive — if you have selected the Mount local disk option at the Select backup location step of the wizard. • Specify Veeam backup repository settings — if you have selected the Add VBR server option at the Select backup location step of the wizard. • Specify Veeam Cloud Connect repository settings — if you have selected the Add Cloud provider option at the Select backup location step of the wizard.	273

5.13.1.4. Restaurar uma instância ou seus arquivos em um outro servidor.	veeam_agent_linux_5_0_user_guide.pdf	Step 5. Specify Backup Location Settings Specify settings for the target storage that contains a backup file from which you plan to restore data: • Specify shared folder settings — if you have selected the Add shared folder option at the Select backup location step of the wizard. • Select local drive — if you have selected the Mount local disk option at the Select backup location step of the wizard. • Specify Veeam backup repository settings — if you have selected the Add VBR server option at the Select backup location step of the wizard. • Specify Veeam Cloud Connect repository settings — if you have selected the Add Cloud provider option at the Select backup location step of the wizard.	313
	veeam_agent_linux_5_0_user_guide.pdf	Step 5. Specify Backup Location Settings Specify settings for the target storage that contains a backup file from which you plan to restore data: • Specify shared folder settings — if you have selected the Add shared folder option at the Select backup location step of the wizard. • Select local drive — if you have selected the Mount local disk option at the Select backup location step of the wizard. • Specify Veeam backup repository settings — if you have selected the Add VBR server option at the Select backup location step of the wizard. • Specify Veeam Cloud Connect repository settings — if you have selected the Add Cloud provider option at the Select backup location step of the wizard.	273
5.14. Suporte ao ambiente virtual (VMware e Acropolis Hypervisor)	N/A	N/A	N/A
5.14.1. Executar proteção e recuperação do Ambiente Virtual com as seguintes características:	N/A	N/A	N/A

5.14.1.1. Realizar recuperação da imagem completa da máquina virtual (ambientes VMware e Acropolis Hypervisor) e também de arquivos de maneira granular sem a necessidade de scripts, área temporária ou montagem dos arquivos vmdk, vhd;	veeam_backup_11_0_user_guide_vsphere.pdf	Data Recovery Veeam Backup & Replication offers the following types of recovery: • VM recovery — to restore entire VMs to different data protection environments: to VMware vSphere, Hyper-V, Amazon EC2 and so on. • Disk recovery — to recover and export disks. • Item recovery — to recover VM files, guest OS files and folders, and application items. • Secure restore — to scan data with antivirus software before restoring it to the production environment.	1367
	veeam_backup_for_nutanix_ahv_3_0_user_guide.pdf	Data Recovery Veeam Backup for Nutanix AHV offers the following recovery options for various disaster recovery scenarios: • Entire VM Restore: recover a VM from a backup file or from a snapshot to an AHV cluster. You can restore from backups of VMware vSphere, Hyper-V, AHV VMs or from backups created by Veeam Agent for Windows or Veeam Agent for Linux. • VM Disks Restore: recover a specific hard drive of a VM from backup or backup snapshot, and attach it to the original VM or to a new VM. • Instant VM Recovery: instantly start a VM directly from a backup file. • VM Disk Export: restore VM disks from backups created by AHV Backup Proxy and convert them to disks of the VMDK, VHD or VHDX format. • Guest OS File Recovery: recover individual guest OS files from Windows, Linux, Mac and other guest OS file systems. • Application Items Restore: restore application items (Microsoft Active Directory, Microsoft Exchange, Microsoft SharePoint, and Microsoft SQL Server). • Restore to Microsoft Azure: restore VMs from backups to Microsoft Azure. • Restore to Amazon EC2: restore VMs from a backup to Amazon EC2 as EC2 instances.	169

5.14.1.2. No caso da restauração granular, não há necessidade de se restaurar a Guest VM inteira;	veeam_backup_11_0_user_guide_vsphere.pdf	Guest OS File Recovery With guest OS file recovery, you can recover individual guest OS files and folders from backups, replicas, storage snapshots, Nutanix AHV snapshots and so on.	1565
	veeam_backup_for_nutanix_ahv_3_0_user_guide.pdf	VM Guest OS Files Restore You can use Veeam Backup & Replication to recover individual VM guest OS files and folders from VM backups and snapshots created by AHV Backup Proxy.	224
5.14.1.3. Permitir redirecionar a restauração de uma máquina virtual hospedada para uma pasta alternativa, outro volume de armazenamento;	veeam_backup_11_0_user_guide_vsphere.pdf	How Entire VM Restore Works A VM can be restored to its original location or to a new location.	1412
	veeam_backup_for_nutanix_ahv_3_0_user_guide.pdf	Step 4. Select Container The Storage Container step of the wizard is available if you have selected to change the location and settings of the restored VM. NOTE If you restore from a snapshot, you cannot select a different container. VM disks of the restored VM will be on the default container. To change the default target container, do the following: 1. In the Disk location field, expand the VM tree. 2. Click the Container button.	192
5.14.1.4. Incluir automaticamente máquinas virtuais novas criadas dentro de seleções de cópias anteriores;	veeam_backup_11_0_user_guide_vsphere.pdf	Instant Recovery to VMware vSphere With Instant Recovery to VMware vSphere, you can immediately recover different workloads (VMs, EC2 instances, physical servers and so on) as VMware vSphere VMs to your production environment by running them directly from compressed and deduplicated backup files. Instant Recovery helps improve recovery time objectives (RTO), minimize disruption and downtime of production workloads. The workloads are recovered in a matter of minutes.	1369
	veeam_backup_for_nutanix_ahv_3_0_user_guide.pdf	Adding PDs to Backup Job If you want to select a protection domain instead of individual VMs, do the following: 1. At the Virtual Machines step of the wizard, click Add on the toolbar. 2. In the Add Objects window,	117

		select the Dynamic mode check box. 3. Click + on the left of the cluster name to expand the list of protection domains.	
5.14.1.5. Permitir cópia completa (Full) e incremental para os servidores virtuais;	veeam_backup_11_0_user_guide_vsphere.pdf	The first backup job session always produces a full backup of the VM image. Subsequent backup job sessions are incremental — Veeam Backup & Replication copies only those data blocks that have changed since the last backup job session. To keep track of changed data blocks, Veeam Backup & Replication uses different approaches. For more information, see Changed Block Tracking.	715
	veeam_backup_for_nutanix_ahv_3_0_user_guide.pdf	Creating Active Full Backup When you create an active full backup, AHV Backup Proxy creates a new full backup file and adds it to the backup chain. The active full backup resets the backup chain. All subsequent incremental backups use the active full backup as a starting point.	127
5.14.1.6. Ser capaz de realizar cópias e restauração de servidores virtuais Linux e Windows, sejam elas estado de consistência ou aplicação;	veeam_backup_11_0_user_guide_vsphere.pdf	Guest Processing If you back up or replicate running VMs, you can enable guest processing options. Guest processing options are advanced tasks that require Veeam Backup & Replication to communicate with the VM guest OS. Veeam Backup & Replication offers the following guest processing options: • Application-aware processing. You can create transactionally consistent backups and replicas of VMs that run Microsoft Active Directory, Microsoft Exchange, Microsoft SharePoint, Microsoft SQL Server or Oracle Database. The transactionally consistent backup guarantees proper recovery of these applications without data loss. For information on system requirements for the applications, see System Requirements.	773

	veeam_backup_for_nutanix_ahv_3_0_user_guide.pdf	NOTE If you want to create application-consistent backups of AHV VMs, follow the guidelines of Nutanix. For more information, see the Application-Consistent Snapshot Guidelines section of the Prism Web Console Guide. Backup Job Functionality Limitations • AHV Backup Proxy creates application-consistent backups when the source VM meets Nutanix requirements described in the Application-Consistent Snapshot Guidelines section of the Prism Web Console Guide. If a VM does not meet these requirements, AHV Backup Proxy will create a crashconsistent backup.	108 110
5.14.1.7. Permitir que as tarefas de cópias e restauração sejam realizadas via interface gráfica;	veeam_backup_11_0_user_guide_vsphere.pdf	Backup & Replication Console The Veeam Backup & Replication console is a client-side component that provides access to the backup server. The console allows you to log in to Veeam Backup & Replication and perform all kinds of data protection and disaster recovery operations on the backup server.	279
	veeam_backup_for_nutanix_ahv_3_0_user_guide.pdf	Getting to Know UI of Web Console After you log in to the AHV Backup Proxy web console, the Dashboard tab opens. At the main panel of the web console, you can navigate through tabs to manage backup and restore operations and configure settings of AHV Backup Proxy.	82
5.14.1.8. O backup dos servidores virtuais deverá ser armazenado de maneira desduplicada;	veeam_backup_11_0_user_guide_vsphere.pdf	Data Compression and Deduplication Veeam Backup & Replication provides mechanisms of data compression and deduplication. Data compression and deduplication let you decrease traffic going over the network and disk space required for storing backup files and VM replica files.	757
	veeam_backup_for_nutanix_ahv_3_0_user_guide.pdf	Veeam Backup for Nutanix AHV uses agentless approach to back up Nutanix AHV VMs and volume groups. The solution creates image-based backups of VMs at a specific	9

		point in time. Veeam Backup for Nutanix AHV retrieves the VM data block by block, compresses, deduplicates it, and forwards it to a backup repository.	
5.14.1.9. Permitir orquestração de cópias de baixo nível da camada de armazenamento (Snapshot) de máquinas virtuais ou Domínios de Proteção no Nutanix AHV, com a retenção desses dados armazenados diretamente no cluster AHV.	veeam_backup_for_nutanix_ahv_3_0_user_guide.pdf	With the AHV Backup Proxy web console, you can: o Create backups of Nutanix AHV VMs and store them in Veeam backup repositories. o Create VM snapshots and protection domain snapshots. o Restore VMs from backups and snapshots to the original Nutanix AHV cluster. o Restore VM disks and attach them to VMs running on Nutanix AHV clusters.	9
5.14.1.10. Permitir a restauração granular de arquivos ou sistemas de arquivos a partir de cópias em disco. Em caso de backup armazenado em disco a recuperação granular poderá ser feito utilizando-se cópias que possam estar desduplicados;	veeam_backup_11_0_user_guide_vsphere.pdf	Guest OS File Recovery With guest OS file recovery, you can recover individual guest OS files and folders from backups, replicas, storage snapshots, Nutanix AHV snapshots and so on.	1565
	veeam_backup_for_nutanix_ahv_3_0_user_guide.pdf	VM Guest OS Files Restore You can use Veeam Backup & Replication to recover individual VM guest OS files and folders from VM backups and snapshots created by AHV Backup Proxy.	224
5.14.1.11. Possui capacidade de realizar a replicação de máquinas virtuais VMware localmente e remotamente em outro Cluster, realizando clones ou snapshots com proteção contínua dos dados por máquina virtual:	veeam_backup_11_0_user_guide_vsphere.pdf	Replication Replication is a technology that helps you protect mission-critical VMware virtual machines. When you replicate a VM, Veeam Backup & Replication creates an exact copy of the VM in the native VMware vSphere format on the target host. Veeam Backup & Replication maintains this copy in sync with the original VM. Replication provides minimum recovery time objective (RTO) in case a disaster strikes because VM replicas are in a ready -to-start state. We recommend you to replicate VMs for which recovery point objective (RPO) of hours is required. If you need RPO of seconds, consider continuous data protection (CDP).	910

5.14.1.11.1. Deverá suportar a replicação remota a fim de replicar os dados das máquinas virtuais entre soluções de armazenamento distintas, inclusive de diferentes fabricantes;	veeam_backup_11_0_user_guide_vsphere.pdf	Data Replication To replicate VMs, Veeam Backup & Replication leverages VMware vSphere snapshot capabilities. During replication, Veeam Backup & Replication requests VMware vSphere to create a VM snapshot. The VM snapshot can be thought of as a point-in-time copy of a VM that includes virtual disks, system state, configuration and so on. Veeam Backup & Replication uses the snapshot as a source of data for replication. Recovery If a disaster strikes and the production VM stops working properly, you can fail over to its replica. When you fail over to a replica, the replica takes over the role of the original VM. After your original VM is repaired, you can fail back to it and transfer all changes that occurred to replica to the original VM. If your original VM cannot be repaired, you can perform permanent failover, that is, permanently switch from the original VM to the VM replica and use this replica as the original VM. Source and Target Hosts The source host is the host where VMs that you plan to replicate are located. The target host is the host where VM replicas will be created and maintained in the ready-to-start state. The role of a target host can be assigned to a standalone ESXi host or ESXi cluster. If you assign a cluster or vCenter Server as a target, the replication process becomes more sustainable — the replication process will not fail if there is at least one available host in the cluster. To replicate data from and to hosts, they must be first added to the Veeam Backup & Replication infrastructure. For more information on how to add hosts, see the Adding VMware vSphere Servers section.	910 910 912

5.14.1.11.2. Suportar a orquestração de failover e failback das máquinas virtuais replicadas;	veeam_backup_11_0_user_guide_vsphere.pdf	Failover and Failback for Replication Failover and failback operations help you ensure that your business will function even if a disaster strikes your production site. Failover is a process of switching from the original VM on the source host to its VM replica on a host in the disaster recovery site. Failback is a process of returning from the VM replica to the original VM. Perform failover Perform failback	972 972 972
5.14.1.12. Permitir a execução de uma máquina virtual diretamente de uma imagem de backup desduplicada e comprimida;	veeam_backup_11_0_user_guide_vsphere.pdf	Instant Recovery to VMware vSphere With Instant Recovery to VMware vSphere, you can immediately recover different workloads (VMs, EC2 instances, physical servers and so on) as VMware vSphere VMs to your production environment by running them directly from compressed and deduplicated backup files. Instant Recovery helps improve recovery time objectives (RTO), minimize disruption and downtime of production workloads.	1369
	veeam_backup_for_nutanix_ahv_3_0_user_guide.pdf	Instant Recovery With the Instant Recovery feature of Veeam Backup & Replication, you can immediately restore different workloads as VMs to your production environment by running it directly from compressed and deduplicated backup files. Instant VM Recovery helps to improve recovery time objectives, minimize disruption and d owntime of production VMs.	207
5.14.1.12.1. Essa funcionalidade deverá permitir sua execução de modo agnóstico ao servidor e repositório de backup utilizado, seja para vSphere ou para Acropolis;	veeam_backup_11_0_user_guide_vsphere.pdf	Specifying Destination for One VM The Destination step of the wizard is available if you recover one workload and recover it to a new location or with different settings. At this step of the wizard, you configure destination settings such as the recovered VM name, target host, VM folder and so on. For workloads other than VMware vSphere VMs and vCloud Director VMs, you	1377

		must also configure a network mapping table.	
	veeam_backup_for_nutanix_ahv_3_0_user_guide.pdf	Step 4. Select Target Cluster At the Cluster step of the wizard, specify a Nutanix AHV cluster where you want to locate the recovered VMs: 1. In the VM location list, select one or more VMs and click Cluster. 2. In the Select Cluster window, select a standalone host or cluster where the selected VMs must be registered. To quickly find the necessary object, use the search field at the bottom of the Select Cluster window: a. Click the button on the left of the field to select the necessary type of object that should be searched for. b. Enter an object name or a part of it in the search field. c. Click the Start search button on the right or press [Enter]	214
5.14.1.12.1.1. Deverá permitir que a solução de virtualização empregada possa movimentar a máquina virtual para o ambiente de produção, posteriormente;	veeam_backup_11_0_user_guide_vsphere.pdf	Migration of Recovered VMs to Production Site To migrate the recovered VMs to the production storage, you can use one of the following relocation methods: • Use Storage vMotion to quickly migrate the recovered VM to the production storage without any downtime.	1370
5.14.1.12.1.2. A máquina virtual iniciada não deverá alterar os dados de backup existentes, ficando a encargo da solução tratar a área de armazenamento temporária da máquina;	veeam_backup_11_0_user_guide_vsphere.pdf	How Instant Recovery Works When Instant Recovery is performed, Veeam Backup & Replication uses the Veeam vPower technology to mount a workload image to an ESXi host directly from a compressed and deduplicated backup file. Since there is no need to extract the workload from the backup file and copy it to production storage, you can perform recovery from any restore point in a matter of minutes. The image of the workload remains in read-only state to avoid unexpected modifications. By default, all changes to virtual disks that take place while a recovered	1369 1370

		VM is running are logged to auxiliary redo log files residing on the NFS server (backup server or backup repository). These changes are discarded as soon as the recovered VM is removed, or merged if you migrate the VM to the production site.	
5.14.1.12.1.3. Permitir que uma máquina virtual Acropolis seja restaurada como máquina virtual VMware diretamente da imagem de backup;	veeam_backup_11_0_user_guide_vsphere.pdf	Instant Recovery to VMware vSphere With Instant Recovery to VMware vSphere, you can immediately recover different workloads (VMs, EC2 instances, physical servers and so on) as VMware vSphere VMs to your production environment by running them directly from compressed and deduplicated backup files. Supported Backup Types Backups of Nutanix AHV virtual machines created by Veeam Backup for Nutanix AHV	1369
	veeam_backup_for_nutanix_ahv_3_0_user_guide.pdf	Instant Recovery With the Instant Recovery feature of Veeam Backup & Replication, you can immediately restore different workloads as VMs to your production environment by running it directly from compressed and deduplicated backup files. Instant VM Recovery helps to improve recovery time objectives, minimize disruption and d owntime of production VMs.	207
	veeam_backup_for_nutanix_ahv_3_0_user_guide.pdf	Supported Workloads To recover machines to a Nutanix AHV cluster, you can use the following backups: • Backups of Nutanix AHV VMs created by Veeam Backup for Nutanix AHV • Backups of Microsoft Hyper-V and VMware vSphere VMs created by Veeam Backup & Replication	208

5.14.1.13. A solução deverá permitir a criação de uma área de testes isolada, compatível com VMware ou Acropolis, para depurar máquinas virtuais, testar upgrades de software e instalar novas aplicações nas máquinas virtuais;	veeam_backup_11_0_user_guide_vsphere.pdf	On-Demand Sandbox If you need to perform tests for production VMs, you can use an On-Demand SandboxTM. The On-Demand Sandbox is an isolated virtual environment where you can start one or more VMs from backups, VM replicas or VMs from storage snapshots. You can use the On-Demand Sandbox to perform the following tasks: • Troubleshoot problems with VMs • Test software patches and upgrades • Install new software and so on	1361
5.14.1.13.1. Quaisquer atualizações realizadas somente deverão ser aplicadas nas máquinas após a restauração completa dos dados no ambiente de produção;	veeam_backup_11_0_user_guide_vsphere.pdf	The On-Demand Sandbox uses a virtual lab — an isolated environment that is fully fenced off from the production environment. VMs started in the virtual lab remain in the read-only state. All changes made to VMs are written to redo logs (for VM backups and storage snapshots) or saved to delta files (for VM replicas). Redo logs and delta files are deleted after you finish working with the On-Demand Sandbox and power it off.	1361
5.14.1.13.2. Ações executadas no ambiente isolados deverão ocorrer em infraestrutura temporária, sendo descartadas ou rotacionadas caso a máquina não venha a ser restaurada;	veeam_backup_11_0_user_guide_vsphere.pdf	The On-Demand Sandbox uses a virtual lab — an isolated environment that is fully fenced off from the production environment. VMs started in the virtual lab remain in the read-only state. All changes made to VMs are written to redo logs (for VM backups and storage snapshots) or saved to delta files (for VM replicas). Redo logs and delta files are deleted after you finish working with the On-Demand Sandbox and power it off.	1361
5.15. Funcionalidade de desduplicação de cópia e arquivamento	N/A	N/A	N/A
5.15.1. Permitir uso da tecnologia de desduplicação de dados para toda a capacidade existente, não existindo limitações devido a licença empregada, eliminando blocos repetidos, para cópias e arquivamento em disco e movimentação de dados	veeam_backup_11_0_user_guide_vsphere.pdf	Deduplication Data deduplication decreases the size of files. With data deduplication enabled, Veeam Backup & Replication does not store to the resulting file identical data blocks and space that has been pre-allocated but not used.	758

desduplicados, independentemente de quantitativo de dispositivos de armazenamento que compõem a infraestrutura da CONTRATANTE.	veeam_backup_one_feature_comparison.pdf	Licensing details Veeam has previously offered socket-based licenses and have now pivoted to the feature-complete and portable VUL for full functionality across workloads and environments. This document will outline the features that are available by license type.	3
5.15.2. Implementar desduplicação a nível de blocos, não sendo aceita a técnica de Single-Instance Storage;	veeam_backup_11_0_user_guide_vsphere.pdf	Veeam Backup & Replication uses Veeam Data Movers to deduplicate VM data: • Veeam Data Mover in the source side deduplicates VM data at the level of VM disks. Before the source - side Veeam Data Mover starts processing a VM disk, it obtains digests for the previous restore point in the backup chain from Veeam Data Mover in the target side. The source-side Veeam Data Mover consolidates this information with CBT information from the hypervisor and filters VM disk data based on it. If some data block exists in the previous restore point for this VM, the source-side Veeam Data Mover does not transport this data block to the target. In addition to it, in case of thin disks the source-side Veeam Data Mover skips unallocated space. • Veeam Data Mover in the target side deduplicates VM data at the level of the backup file. It processes data for all VM disks of all VMs in the job. The target-side Veeam Data Mover uses digests to detect identical data blocks in transported data, and stores only unique data blocks to the resulting backup file.	758
5.15.3. Implementar desduplicação de blocos na origem (client-side deduplication), de forma que o cliente envie apenas novos blocos de dados criados e/ou modificados a partir da última cópia total completa;	veeam_backup_11_0_user_guide_vsphere.pdf	Veeam Backup & Replication uses Veeam Data Movers to deduplicate VM data: • Veeam Data Mover in the source side deduplicates VM data at the level of VM disks. Before the source - side Veeam Data Mover starts processing a VM disk, it obtains digests for the previous restore point in the backup chain from Veeam Data Mover in the target side. The source-side Veeam Data Mover consolidates this information with CBT	758

		information from the hypervisor and filters VM disk data based on it. If some data block exists in the previous restore point for this VM, the source-side Veeam Data Mover does not transport this data block to the target. In addition to it, in case of thin disks the source-side Veeam Data Mover skips unallocated space. • Veeam Data Mover in the target side deduplicates VM data at the level of the backup file. It processes data for all VM disks of all VMs in the job. The target-side Veeam Data Mover uses digests to detect identical data blocks in transported data, and stores only unique data blocks to the resulting backup file.	
5.15.4. Implementar desduplicação de blocos no destino (target-side deduplication), de forma que sejaresponsabilidade do servidor de transporte executar a tecnologia;	veeam_backup_11_0_user_guide_vsphere.pdf	Veeam Backup & Replication uses Veeam Data Movers to deduplicate VM data: • Veeam Data Mover in the source side deduplicates VM data at the level of VM disks. Before the source - side Veeam Data Mover starts processing a VM disk, it obtains digests for the previous restore point in the backup chain from Veeam Data Mover in the target side. The source-side Veeam Data Mover consolidates this information with CBT information from the hypervisor and filters VM disk data based on it. If some data block exists in the previous restore point for this VM, the source-side Veeam Data Mover does not transport this data block to the target. In addition to it, in case of thin disks the source-side Veeam Data Mover skips unallocated space. • Veeam Data Mover in the target side deduplicates VM data at the level of the backup file. It processes data for all VM disks of all VMs in the job. The target-side Veeam Data Mover uses digests to detect identical data blocks in transported data, and stores only unique data blocks to the resulting backup file.	758

5.15.5. Implementar desduplicação de dados em tarefas de cópia;	veeam_backup_11_0_user_guide_vsphere.pdf	Veeam Backup & Replication uses Veeam Data Movers to deduplicate VM data: • Veeam Data Mover in the source side deduplicates VM data at the level of VM disks. Before the source - side Veeam Data Mover starts processing a VM disk, it obtains digests for the previous restore point in the backup chain from Veeam Data Mover in the target side. The source-side Veeam Data Mover consolidates this information with CBT information from the hypervisor and filters VM disk data based on it. If some data block exists in the previous restore point for this VM, the source-side Veeam Data Mover does not transport this data block to the target. In addition to it, in case of thin disks the source-side Veeam Data Mover skips unallocated space. • Veeam Data Mover in the target side deduplicates VM data at the level of the backup file. It processes data for all VM disks of all VMs in the job. The target-side Veeam Data Mover uses digests to detect identical data blocks in transported data, and stores only unique data blocks to the resulting backup file.	758
5.15.6. Implementar desduplicação e compressão em uma mesma tarefa.	veeam_backup_11_0_user_guide_vsphere.pdf	Data Compression and Deduplication Veeam Backup & Replication provides mechanisms of data compression and deduplication. Data compression and deduplication let you decrease traffic going over the network and disk space required for storing backup files and VM replica files.	759
5.15.7. Implementar desduplicação em infraestrutura de armazenamento local (DAS) e vida rede (SAN e NAS);	veeam_backup_11_0_user_guide_vsphere.pdf	Local target 1024 KB Recommended for backup and replication to SAN, DAS or local storage. This option provides the fastest job performance but reduces the deduplication ratio, because with larger data blocks it is less likely to find identical blocks. LAN target 512 KB	759

		Recommended for backup and replication to NAS, and on-site backup and replication. This option provides a better deduplication ratio and reduces the size of a file because of reduced data block sizes.	
5.15.8. A tecnologia de desduplicação não deverá possuir limites quanto a quantidade de dados que serão lidos (front-end), nem limites para a quantidade de dados que serão escritos (back-end);	N/A	N/A	N/A
5.16. Reportes e alertas	N/A	N/A	N/A
5.16.1. Vir disponível com os seguintes relatórios e reportes:	N/A	N/A	N/A
5.16.1.1. Histórico de rotinas de proteção concluídos nas últimas 24 horas, nos últimos 30 dias e nos últimos 6 meses;	veeam_backup_11_0_user_guide_vsphere.pdf	Viewing Real-Time Statistics To view real-time statistics for a job, do one of the following: • Open the Home view, in the inventory pane select one of the following nodes: Jobs, Last 24 hours or Running. In the working area, right-click the job and select Statistics. • Open the Home view, in the inventory pane select one of the following nodes: Jobs, Last 24 hours or Running. In the working area, double-click the running job. The real-time statistics provides detailed data on job sessions: job progress, duration, processing rate, performance bottlenecks, amount of processed data, read and transferred data and details of the session performance, for example, warnings and errors that have occurred in the process of operation.	904
	veeam_backup_11_0_user_guide_vsphere.pdf	Viewing Job and Job Session Reports You can generate reports with details about all sessions of a job or a single session only.	908

	veeam_backup_11_0_enterprise_manager_user_guide.pdf	If you are using the Enterprise or Enterprise Plus edition, Veeam Backup Enterprise Manager will keep index files for backups that are currently stored on disk and for archived backups (for example, backups that were recorded to tape). Thus, you will be able to browse and search through backup contents even if the backup in repository is no longer available or it was removed by Remove from Backups or Remove from Disk command in Veeam Backup console. On the Session History tab, in the Guest file system catalog section, specify how long index files must be stored on the Veeam Backup Enterprise Manager server: a. Enter the desired number of months in the Retention period, months field. The default value is 3 months, the minimum allowed value is 1 month, and the maximum allowed value is 99 months.	96
5.16.1.2. Histórico de recuperações efetuadas nas últimas 24 horas, nos últimos 30 dias e nos últimos 6 meses;	veeam_backup_11_0_user_guide_vsphere.pdf	Viewing Real-Time Statistics To view real-time statistics for a job, do one of the following: • Open the Home view, in the inventory pane select one of the following nodes: Jobs, Last 24 hours or Running. In the working area, right-click the job and select Statistics. • Open the Home view, in the inventory pane select one of the following nodes: Jobs, Last 24 hours or Running. In the working area, double-click the running job. The real-time statistics provides detailed data on job sessions: job progress, duration, processing rate, performance bottlenecks, amount of processed data, read and transferred data and details of the session performance, for example, warnings and errors that have occurred in the process of operation.	904
	veeam_backup_11_0_user_guide_vsphere.pdf	Viewing Job and Job Session Reports You can generate reports with details about all sessions of a job or a single session only.	908

	veeam_backup_11_0_enterprise_manager_user_guide.pdf	If you are using the Enterprise or Enterprise Plus edition, Veeam Backup Enterprise Manager will keep index files for backups that are currently stored on disk and for archived backups (for example, backups that were recorded to tape). Thus, you will be able to browse and search through backup contents even if the backup in repository is no longer available or it was removed by Remove from Backups or Remove from Disk command in Veeam Backup console. On the Session History tab, in the Guest file system catalog section, specify how long index files must be stored on the Veeam Backup Enterprise Manager server: a. Enter the desired number of months in the Retention period, months field. The default value is 3 months, the minimum allowed value is 1 month, and the maximum allowed value is 99 months.	96
5.16.1.3. Reportes de rotinas de backup concluídos com sucesso, com erro ou não concluídos;	veeam_backup_11_0_user_guide_vsphere.pdf	Viewing Job and Job Session Reports You can generate reports with details about all sessions of a job or a single session only. Veeam Backup & Replication displays the following counters in reports: • The Success, Warning and Error counters show how many workloads were processed with the Success, Warning	908
5.16.1.4. Taxa de desduplicação por rotina de backup;	veeam_backup_11_0_user_guide_vsphere.pdf	• The Backup size counter shows the resulting backup file size. • The Dedupe counter shows the deduplication level. • The Compression counter shows the compression level.	909
5.16.1.5. Possuir relatórios com as seguintes características:	N/A	N/A	N/A
5.16.1.5.1. Horário de início e término de uma rotina de backup;	veeam_backup_11_0_user_guide_vsphere.pdf	• The Duration counter shows time from the job start till the current moment or job end.	909
5.16.1.5.2. Tempo de duração de uma rotina de backup;	veeam_backup_11_0_user_guide_vsphere.pdf	• The Duration counter shows time from the job start till the current moment or job end.	909
5.16.1.5.3. Status do backup (situação):	veeam_backup_11_0_user_guide_vsphere.pdf	• The Success, Warning and Error counters show how many	908

		workloads were processed with the Success, Warning and Error statuses.	
5.16.1.5.3.1. Relação dos objetos incluídos na rotina de backup;	veeam_backup_11_0_user_guide_vsphere.pdf	• The Status box shows information about the job results. This box informs how many tasks have completed with the Success, Warning and Error statuses (1 task per 1 VM). • The pane at the lower left corner shows a list of objects included in the job. • The pane at the lower right corner shows a list of operations performed during the job. To see a list of operations for a specific object included in the job, click the object in the pane on the left. To see a list of operations for the whole job, click anywhere on the blank area in the left pane.	905
5.16.1.5.3.2. Horário de início e término do backup de cada objeto;	veeam_backup_11_0_user_guide_vsphere.pdf	• The Status box shows information about the job results. This box informs how many tasks have completed with the Success, Warning and Error statuses (1 task per 1 VM). • The pane at the lower left corner shows a list of objects included in the job. • The pane at the lower right corner shows a list of operations performed during the job. To see a list of operations for a specific object included in the job, click the object in the pane on the left. To see a list of operations for the whole job, click anywhere on the blank area in the left pane.	905
5.16.1.5.3.3. Tempo de duração do backup de cada objeto;	veeam_backup_11_0_user_guide_vsphere.pdf	• The Status box shows information about the job results. This box informs how many tasks have completed with the Success, Warning and Error statuses (1 task per 1 VM). • The pane at the lower left corner shows a list of objects included in the job. • The pane at the lower right corner shows a list of operations performed during the job. To see a list of operations for a specific object included in the job, click the object in the pane on the left. To see a list of operations for	905

		the whole job, click anywhere on the blank area in the left pane.	
5.16.1.5.3.4. Volume de dados na origem durante a rotina de backup;	veeam_backup_11_0_user_guide_vsphere.pdf	The Data read counter shows the amount of data read from the datastore before compression and deduplication.	909
5.16.1.5.3.5. Volume de dados trafegados durante a rotina de backup;	veeam_backup_11_0_user_guide_vsphere.pdf	The Transferred counter shows the amount of data transferred from the source side to the target side after applying compression and source-side deduplication.	909
5.16.1.5.3.6. Volume de dados com compressão e desduplicação;	veeam_backup_11_0_user_guide_vsphere.pdf	The Backup size counter shows the resulting backup file size.	909
5.16.1.5.3.7. Taxa de desduplicação de dados;	veeam_backup_11_0_user_guide_vsphere.pdf	• The Backup size counter shows the resulting backup file size. • The Dedupe counter shows the deduplication level. • The Compression counter shows the compression level.	909
5.16.1.5.3.8. Taxa de compressão de dados;	veeam_backup_11_0_user_guide_vsphere.pdf	• The Backup size counter shows the resulting backup file size. • The Dedupe counter shows the deduplication level. • The Compression counter shows the compression level.	909
5.16.1.6. A solução ofertada deverá enviar os seguintes alertas via e-mail:	veeam_backup_11_0_user_guide_vsphere.pdf	Configuring Job Notification Settings	654
5.16.1.7. Rotina de backup finalizada com sucesso;	veeam_backup_11_0_user_guide_vsphere.pdf	Notify on success	655
5.16.1.8. Rotina de backup finalizada com erro;	veeam_backup_11_0_user_guide_vsphere.pdf	Notify on warning	655
5.16.1.9. Rotina de backup com problema;	veeam_backup_11_0_user_guide_vsphere.pdf	Notify on error	655
5.16.1.10. Alerta para utilização de licenciamento.	veeam_backup_11_0_user_guide_vsphere.pdf	Support Contract Expiration Notification By default, Veeam Backup & Replication informs email recipients specified in global notification settings about the support expiration date in every email notification. Veeam Backup & Replication starts sending such notifications 14 days before the expiration date. Expiration information is also shown on the splash screen and on the License Information window (to display the License Information window, select	661

		Help > License from the main menu).	
	veeam_backup_11_0_user_guide_vsphere.pdf	Exceeding License Limit For Veeam Universal Licenses, Veeam Backup & Replication allows you to protect more workloads than covered by the number of instances specified in the license. An increase in the number of protected workloads is allowed throughout the duration of the contract (license key).	146

ITEM	DOCUMENTAÇÃO	COMPROVAÇÃO	PÁGINA
5.17. Unidade de armazenamento de informação	N/A	N/A	N/A
5.17.1. Características Gerais	N/A	N/A	N/A
5.17.1.1. Ser homologada pelo software de proteção ofertada;	https://www.exagrid.com/exagrid-products/supported-data-backup-applications/veeam-backup/	Veeam is an ExaGrid Technology Partner. ExaGrid's Tiered Backup Storage is changing the economics of Veeam backup. ExaGrid's cost-effective scale-out growth model has a lower cost up front and a lower cost over time compared to standard disk solutions and traditional deduplication storage solutions.	N/A
	https://helpcenter.veeam.com/docs/backup/vsphere/deduplicating_appliance_exgrid.html?ver=110	ExaGrid You can use ExaGrid appliances as backup repositories.	N/A
5.17.1.2. Prover infraestrutura de armazenamento, voltados para a proteção de dados do ambiente de hiperconvergência ou nuvem privada;	ExaGrid-Detailed-Product-Description_DS-1-1.pdf	ExaGrid Basic Concept ExaGrid appliances work seamlessly with over 25 industry-leading backup applications and database utilities by presenting themselves as standard NAS shares (CIFS or NFS), a Veeam Data Mover target, or Veritas NetBackup OST target ExaGrid appliances are easily integrated into existing backup environments, as illustrated below	2

	ExaGrid-Detailed-Product-Description_DS-1-1.pdf	ExaGrid Appliance Models The product line's multiple appliance models can be combined into a single scale-out system with a confguration of up to 32 appliances in a single system for a total of 6 14PB raw capacity, 5 37PB of usable capacity, and supports a full backup of up to 2 69PB, total backup logical data capacity of 43PB	4
5.17.1.3. Corresponder a um módulo de armazenamento de backup em disco, com o propósito específico de ingestão dos dados de backup com compactação, desduplicação e replicação (local e remota/nuvem) dos dados desduplicados;	ExaGrid-Detailed-Product-Description_DS-1-1.pdf	Data Deduplication ExaGrid's innovative approach to data deduplication minimizes the amount of data to be stored by using zone-level data deduplication across all received backups	7
	ExaGrid-Detailed-Product-Description_DS-1-1.pdf	The ExaGrid system performs deduplication in parallel with backups by using idle system cycles Termed "adaptive deduplication," this approach keeps deduplication out of the backup path, resulting in faster backups and a shorter backup window Adaptive Deduplication gives full system resources to backups for the shortest backup window and uses idle cycles to perform deduplication and offsite replication, allowing data to be replicated offsite for a strong offsite recovery point (RPO) for update to data for disaster recoveries	7
	ExaGrid-Detailed-Product-Description_DS-1-1.pdf	SPLITTER COMPUTE HASH HASH COMPARE BUCKETING COMPRESSION	7
5.17.1.4. Ser novo, de primeiro uso e estar em linha de fabricação na data da abertura da licitação. Não serão aceitos equipamentos usados, remanufaturados, de demonstração ou gateways;	Conforme Lista de Part Numbers	N/A	N/A
5.17.1.5. Constar no site do fabricante (documento oficial e público) como um sistema de armazenamento de backup em disco, em linha de produção;	ExaGrid-Detailed-Product-Description_DS-1-1.pdf	ExaGrid Appliance Models EX84 EX36	4

5.17.1.6. Não serão aceitas soluções definidas por Software (Virtual Appliance);	ExaGrid-Detailed-Product-Description_DS-1-1.pdf	ExaGrid appliances are comprised of processors, memory, RAID6, and a hot spare drive per appliance, using enterprise-class SAS drives and ExaGrid software See the ExaGrid Technical Specifcations data sheet for more details	3
5.17.1.7. O hardware do módulo de armazenamento de cópias em disco não poderá ser compartilhado com nenhum outro software para operar;	ExaGrid-Detailed-Product-Description_DS-1-1.pdf	Each appliance plugs in and is virtualized into a shared system with a single user interface, global deduplication, and automatic load balancing The media server is connected to the same network and sees the appliances as one or more NAS shares, Veeam Data Mover targets, or Veritas NetBackup OST targets Since each appliance includes the appropriate amount of processor, memory, Landing Zone disk, deduplicated repository disk, and bandwidth for the rated data size, performance increases as more appliances are added to the system	3
5.17.1.8. Ser do tipo agnóstico, ou seja, possuir compatibilidade com diversas soluções de software de proteção. Não serão aceitas soluções proprietárias (“lock in”) ou seja, aqueles que só funcionam com um software de backup específico;	ExaGrid-Detailed-Product-Description_DS-1-1.pdf	Support for Heterogeneous Backup Applications and Utilities Customer environments are made up of many backup approaches, backup applications, and utilities and different disk-based backup approaches support these in different ways Customers may have any number of backups occurring in their environment, including traditional backup applications such as Veeam, Veritas NetBackup, Commvault, IBM Spectrum Protect (TSM), Dell EMC NetWorker, HYCU, Zerto, and many others as well as direct-to-disk SQL dumps, Oracle RMAN direct backups, and specifc UNIX utilities such as UNIX TAR ExaGrid supports over 25 backup applications and utilities	9

5.17.1.8.1. Deverá possuir compatibilidade com softwares e aplicações de backup comuns de mercado, como, no mínimo, Arcserve, Backup Exec, Commvault, DellEMC Networker, Hycu, IBM TSM, Oracle RMAN, Microsoft SQL, NetBackup e Veeam;	ExaGrid-Detailed-Product-Description_DS-1-1.pdf	Support for Heterogeneous Backup Applications and Utilities Customer environments are made up of many backup approaches, backup applications, and utilities and different disk-based backup approaches support these in different ways Customers may have any number of backups occurring in their environment, including traditional backup applications such as Veeam, Veritas NetBackup, Commvault, IBM Spectrum Protect (TSM), Dell EMC NetWorker, HYCU, Zerto, and many others as well as direct-to-disk SQL dumps, Oracle RMAN direct backups, and specifc UNIX utilities such as UNIX TAR ExaGrid supports over 25 backup applications and utilities	9
5.17.1.9. Estar licenciada para toda sua capacidade e funcionalidade, incluindo replicação;	Conforme Lista de Part Numbers	N/A	N/A
5.17.1.10. Permitir o particionamento lógico da área de armazenamento, sem prejuízo às características de desduplicação solicitadas neste certame;	ExaGridAdminGuide.pdf	Backup Application Specific Shares Application specific shares are tailored to specific backup applications. ExaGrid provides detailed step-by-step instructions on how to use your backup application with these shares: Universal Backup Shares ExaGrid Universal Backup Shares are designed to accommodate those backup applications that do not have an ExaGrid Share designed specifically for their use. Universal Backup Shares are deduplicated and can be replicated and migrated like application specific ExaGrid Shares.	39 40
5.17.1.11. Todos os valores de capacidade de armazenamento devem ser calculados considerando o sistema de cálculo BASE 2, ou seja, 1 Terabyte (TB) é igual a 1024 Gigabytes (GB);	Conforme Lista de Part Numbers	N/A	N/A

5.17.1.12. Possuir recursos de tolerância a falhas de, pelo menos, discos, fontes de alimentação e ventiladores. Os discos rígidos deverão ser hot-pluggable e hot-swappable permitindo substituição sem necessidade interrupção do funcionamento da solução;	ExaGrid-Detailed-Product-Description_DS-1-1.pdf	RAID6 protection with a hot swappable spare Redundant hot swappable power supplies	15
	ExaGrid-Technical_Specifications_DS-3.pdf	Redundant fans	1
5.17.1.13. Possuir mecanismos que protejam contra a inconsistência dos dados mesmo em casos de interrupção abrupta ou desligamento acidental;	ExaGrid-Technical_Specifications_DS-3.pdf	Flash-Backed RAID Cache The industry-leading PCI RAID controller has onboard writeback cache backed up by a super-cap powered flash memory Unexpected loss of appliance power does not result in backup data loss because any in-process writes to any disk are preserved until power is restored	19
5.17.1.14. Ser entregue com arranjos de discos rígidos do tipo RAID-6 configurado de tal modo a tolerar a falha de até 2 (dois) discos rígidos, para os discos destinados ao armazenamento de dados de backup;	ExaGrid-Technical_Specifications_DS-3.pdf	RAID6 Internal Storage with Consistency Checking All ExaGrid internal storage is accessed using an industry-leading PCI RAID controller at the RAID6 level of disk protection with a global “hot spare” disk	19
5.17.1.14.1. Deverá possuir, no mínimo, 1 (um) disco configurado como hot-spare;	ExaGrid-Technical_Specifications_DS-3.pdf	RAID6 Internal Storage with Consistency Checking All ExaGrid internal storage is accessed using an industry-leading PCI RAID controller at the RAID6 level of disk protection with a global “hot spare” disk	19
5.17.1.15. Possuir funcionalidade de desduplicação dos dados em nível de bytes ou blocos, com capacidade de eliminação de dados redundantes para racionalizar a utilização do espaço de armazenamento;	https://www.exagrid.com/exagrid-products/data-deduplication/	ExaGrid has taken a more innovative path. ExaGrid uses zone-level deduplication, which breaks data into larger “zones” and then compares at the byte level. This approach allows for the best of all worlds. First, the tracking table is 1,000th the size of the block-level approach and allows for full appliances in a scale-out solution.	N/A
5.17.1.16. Implementar desduplicação global para o módulo de armazenamento de backup em disco, considerando todos os dados retidos, sendo capaz de identificar dados duplicados de backups de diferentes origens dentro de um mesmo conjunto de armazenamento de modo a maximizar a taxa de desduplicação e garantindo que os dados sejam gravados uma única vez;	ExaGrid-Detailed-Product-Description_DS-1-1.pdf	Deduplication efciency across the entire system with global deduplication	2

5.17.1.17. Suportar simultaneamente acessos de leitura e gravação pelos protocolos CIFS e NFS;	ExaGrid-Detailed-Product-Description_DS-1-1.pdf	ExaGrid Basic Concept ExaGrid appliances work seamlessly with over 25 industry-leading backup applications and database utilities by presenting themselves as standard NAS shares (CIFS or NFS), a Veeam Data Mover target, or Veritas NetBackup OST target	2
5.17.1.18. Suportar a implementação do OpenStorage Technology;	ExaGrid-Detailed-Product-Description_DS-1-1.pdf	ExaGrid Basic Concept ExaGrid appliances work seamlessly with over 25 industry-leading backup applications and database utilities by presenting themselves as standard NAS shares (CIFS or NFS), a Veeam Data Mover target, or Veritas NetBackup OST target	2
5.17.1.19. Permitir a criação de backups sintéticos completos, do software de proteção ofertado, internamente no appliance;	VeeamLatestGuide.pdf	In addition, the entire Veeam synthetic full operation occurs on the ExaGrid server eliminating the need to move data between the Veeam backup server and backup storage. This greatly reduces the time to complete a synthetic full.	6

5.17.1.20. Permitir a execução de processos de backup e restore em paralelo;	ExaGrid-Detailed-Product-Description_DS-1-1.pdf	Data Backup Performance The ExaGrid system performs deduplication in parallel with backups by using idle system cycles Termed “adaptive deduplication,” this approach keeps deduplication out of the backup path, resulting in faster backups and a shorter backup window Data Restore Performance Because ExaGrid’s appliance allows each full backup to frst land on the disk-cache Landing Zone before deduplication, the system maintains the most recent backup in its full, undeduplicated form for fast restores, Instant VM recoveries in seconds to minutes, and fast offsite tape copies Since over 90% of restores and 100% of instant VM recoveries and tape copies are done from the most recent backup, this approach avoids the overhead incurred from “rehydrating” data during critical restores As a result, restore, recovery, and copy times from an ExaGrid system are an order of magnitude faster than solutions that only store deduplicated data	7 e 8
5.17.1.21. Deverá implementar tecnologia que detém dos serviços de movimentação de dados compatível com o Software de Proteção ofertado, removendo a necessidade de utilização de servidores gateways, servidores intermediários, servidores auxiliares ou similares para o emprego de tais serviços;	VeeamLatestGuide.pdf	The ExaGrid-Veeam Accelerated Agent is integrated with ExaGrid appliances and allows backups, restores and recoveries to complete faster. The Veeam backup server more efficiently inter-operates with its own Veeam data mover using optimized Veeam communications versus generic CIFS	6
	https://helpcenter.veeam.com/docs/backup/vsphere/deduplicating_appliance_exgrid.html?ver=110#exagrid-deployment	To communicate with ExaGrid, Veeam Backup & Replication uses two Veeam Data Movers that are responsible for data processing and transfer: Veeam Data Mover on the backup proxy Veeam Data Mover on the ExaGrid appliance	N/A

5.17.1.21.1. Caso o appliance não implemente internamente os serviços de movimentação de dados do software de proteção ofertado, será aceita a entrega de servidor físico adicional com recurso suficiente para comportar o tráfego de, ao menos, 170 TB (cento e setenta terabytes líquidos) de dados protegidos durante uma janela de 10 (dez) horas;	N/A	N/A	N/A
5.17.1.22. Integração entre o software de proteção ofertado e o módulo de armazenamento:	N/A	N/A	N/A
5.17.1.22.1. O módulo de armazenamento deverá permitir a inicialização de máquinas virtuais, através do software ofertado, diretamente da sua área de armazenamento, desde que o virtualizador suporte a funcionalidade;	VeeamLatestGuide.pdf	Instant VM Recovery from an ExaGrid System In the event of a primary storage outage or to prove your backups are recoverable, Veeam can instantly recover a virtual machine and run it directly from an ExaGrid Server's landing space. You then have the option to migrate the running VM to your production storage.	4
5.17.1.22.2. Deverá permitir a restauração de máquinas virtuais, arquivos únicos e objetos/arquivos específicos de aplicações diretamente do repositório de armazenamento;	VeeamLatestGuide.pdf	Instant File-Level and Application Item Recovery Quickly recover an entire VM, an individual file from an image-level backup, or an Application Item Recovery. Veeam instant file-level recovery enables you to quickly restore individual files from your backups and replicas, taking advantage of ExaGrid's landing zone architecture to achieve the fastest restores and recoveries. Veeam Explorers can be used to restore application specific items such as Exchange mailboxes, folders, Active Directory Users and Groups, databases and database objects.	4

5.17.1.23. Possuir funcionalidade para replicação de cópias em equipamento similar e do mesmo fabricante de forma assíncrona, utilizando recursos de desduplicação e reduzindo consumo do link de comunicação, através de rede IP (WAN/LAN);	ExaGridAdminGuide.pdf	Share Replication Once an ExaGrid Share has been created (See “Creating Shares” on page 47), if it resides in a multi-site ExaGrid System, you can replicate that share to one or two ExaGrid Sites. A copy of a share that has been replicated is known as a replica. Replicating copies to other ExaGrid Sites provides an additional layer of protection for your data. ExaGrid Spoke Sites can replicate shares to their ExaGrid Hub Site. If there is another Spoke Site, the shares replicated to the Hub Site can then be replicated to the other Spoke Site. Spoke Sites cannot directly replicate shares between each other. ExaGrid Hub Sites can replicate shares directly to one or two Spoke Sites:	51
	ExaGridAdminGuide.pdf	All interfaces configured for Any Site traffic must use an address that is reachable from all such addresses on all ExaGrid servers in the ExaGrid System - i.e. all hub Sties and all Spoke sites. An error will be returned if an interface set to All Sites cannot logically reach an interface in another site that is set to All Sites. This is a logical evaluation only based on IP networking rules - there is no testing done of actual physical connectivity.	9
5.17.1.24. Possuir arquitetura baseada em camadas que permita a proteção contra “ransomware attack”, independente do software de backup.	ExaGrid-Detailed-Product-Description_DS-1-1.pdf	ExaGrid Retention Time-Lock for Ransomware Recovery ExaGrid’s unique approach ransomware recovery is called Retention Time-Lock It prevents hackers from deleting the backups and allows for retention points to be purged The result is a strong data protection and recovery solution at a very low cost of storage	16

5.17.1.24.1. Entende-se por equipamento multi-camadas àqueles onde as camadas de armazenamento são nativas (não podem ser criadas ou removidas) e onde pelo menos uma das camadas não pode ser acessada diretamente pelo software de backup para escrita. Além disso, deve possuir pelo menos uma camada isolada do acesso externo com funcionalidade de atraso de deleção, onde os dados retidos ao longo do tempo devem ser armazenados no formato imutável e não podem ser imediatamente deletados por comando do software de backup.	ExaGrid-Detailed-Product-Description_DS-1-1.pdf	ExaGrid provides Tiered Backup Storage with a front-end disk-cache Landing Zone and separate Retention Tier containing all retention data Data is written directly to the "network facing" ExaGrid disk-cache Landing Zone Then it is tiered into a "non-network facing" long-term retention repository where it is stored as deduplicated data objects to reduce the storage cost of long-term retention data	1
5.17.1.24.1.1. Tal atraso de deleção deve ser configurável em dias, proporcionado ao menos o atraso por 15 dias.	ExaGridAdminGuide.pdf	Retention Time-Lock Overview ExaGrid's Tiered Backup Storage architecture has: ■ A high-speed, disk cache that makes up ExaGrid's front-end Landing Zone Tier where the latest backup data resides. ■ A separate Retention Tier that contains all deduplicated and retained data.	20
	ExaGridAdminGuide.pdf	Because the Retention Tier is not network facing and not accessible to ransomware and other malicious activity, a time-lock can be set on the ExaGrid Retention Tier that blocks deletes for up to 30 days.	20
5.17.1.24.1.2. Para desativação ou modificação desse recurso deve ser possível requerer escalação e duplo fator de autenticação.	ExaGridAdminGuide.pdf	Required security officer approval for changes to time-lock settings	20
5.17.1.24.2. Caso a solução ofertada não possua arquitetura multi-camadas e não possua todas as características solicitadas no item anterior, deve ser entregue com um segundo equipamento (para cada unidade contratada no certame), possuindo as mesmas características do equipamento primário, juntamente com um mecanismo que realize o filtro da replicação de dados entre eles para isolar os dados replicados do repositório primário e impedir a	N/A	N/A	N/A

propagação do ataque de ransomware no momento da sincronização entre os sistemas (Air-Gap, Atraso de Sincronismo). Todos os componentes para o funcionamento dessa proteção devem ser fornecidos com a solução.			
5.17.1.25. O appliance deverá possuir baterias, supercapacitores ou tecnologia similar, para proteger a cache de escrita, evitando a perda de dados em eventos de falha elétrica;	ExaGrid-Detailed-Product-Description_DS-1-1.pdf	Flash-Backed RAID Cache The industry-leading PCI RAID controller has onboard writeback cache backed up by a super-cap powered flash memory Unexpected loss of appliance power does not result in backup data loss because any in-process writes to any disk are preserved until power is restored	19
5.17.1.26. O appliance deverá implementar mecanismos de validação da consistência dos dados desduplicados armazenados, garantindo que eles estejam íntegros durante backups, restaurações e replicações. A tecnologia deverá reparar, automaticamente, dados que não estejam consistentes com as rotinas executadas;	ExaGrid-Detailed-Product-Description_DS-1-1.pdf	Deduplicated Metadata Transactional Consistency Metadata that tracks all of the deduplicated data is kept in a database and on internal storage Software techniques are used to ensure transactional integrity of all metadata changes, including flushing flesystem pages into the flash-backed RAID onboard cache The data flow of deduplicated backup data is protected end-to-end by the combination of checksums (above) and metadata transactional consistency	19
5.17.1.27. Deverá possuir integração com Microsoft Active Directory para autenticação de usuários quanto ao acesso a interface de gerência da solução;	ExaGridAdminGuide.pdf	ExaGrid Servers can be members of an Active Directory Domain. ExaGrid roles can be mapped to users or groups in their server's Active Directory Domain. For details see " Managing Active Directory" on page 29.	19

5.17.1.28. Deverá permitir a implementação de topologias de replicação, como 1 para 1, 1 para N e o cascateamento de equipamentos. A solução deverá permitir a criação de topologias de nuvem privada e híbrida;	ExaGridAdminGuide.pdf	Share Replication Once an ExaGrid Share has been created (See “Creating Shares” on page 47), if it resides in a multi-site ExaGrid System, you can replicate that share to one or two ExaGrid Sites. A copy of a share that has been replicated is known as a replica. Replicating copies to other ExaGrid Sites provides an additional layer of protection for your data. ExaGrid Spoke Sites can replicate shares to their ExaGrid Hub Site. If there is another Spoke Site, the shares replicated to the Hub Site can then be replicated to the other Spoke Site. Spoke Sites cannot directly replicate shares between each other. ExaGrid Hub Sites can replicate shares directly to one or two Spoke Sites:	51
	ExaGridAdminGuide.pdf	Note – Replication to a Cloud ExaGrid Site is always encrypted.	58
5.17.1.29. Possuir recursos para monitoramento remoto pelo fabricante, tal como notificação do tipo Call-Home ou Email-Home, para verificação proativa de componentes de hardware em situação de falha ou pré-falha.	ExaGridAdminGuide.pdf	As detailed in “ExaGrid Reports” on page 91 the ExaGrid system allows you to manually log in and view the status of your ExaGrid System. In addition, the ExaGrid System also provides several automated tools that monitor and report status.	97
5.17.1.30. Ser montado em rack padrão 19” e deve ser entregue com todos os trilhos, cabos, conectores, manuais de operação e quaisquer outros componentes que sejam necessários à instalação, customização e plena operação;	Conforme Proposta Comercial	N/A	N/A
5.17.1.31. Ter suporte e garantia de 3 (três) anos, com atendimento 24 x 7 x 365 e até 2 (duas) para início do atendimento de chamados com prioridade mais alta, ou seja, quando o equipamento estiver em estado de indisponibilidade de uso. É troca de peça no próximo dia útil;	ExaGrid-Support-and-Maintenance-Americas-Data-Sheet.pdf	Requirement: Current annual customer maintenance and support renewal Support Hours: 8:00 a.m. to 5:00 p.m. (Eastern standard time), Monday – Friday (optional 7x24 support available for an additional fee) Methods: Phone or email support Response Time: 80% of phone calls and emails will be responded to in less than an hour	1

5.17.1.31.1. Permitir abertura de chamados diretamente com a equipe do fabricante de engenharia nível 2 (técnicos especializados para atuar na investigação do problema, sugestão de ajustes e correção, coleta e avaliação de logs). Além disso, esse engenheiro deve estar disponível para implementar atualizações e correções, revisar as configurações do ambiente e sugerir ajustes de acordo com melhores práticas, mesmo sem a ocorrência de problemas ou indisponibilidade na solução.	Conforme Lista de Part Numbers	N/A	N/A
5.18. Unidade de armazenamento de informação - tipo 1	N/A	N/A	N/A
5.18.1. Possuir 149 TB (cento e quarenta e nove terabytes) de área útil;	ExaGrid-Detailed-Product-Description_DS-1-1.pdf	ExaGrid Appliance Models EX84 Usable Capacity (TB) 168 TB = 152,795 TiB	4
5.18.2. Possuir alguma das seguintes arquiteturas para o módulo de armazenamento:	N/A	N/A	N/A
5.18.2.1. Scale-up: Soluções com arquitetura tradicional (crescimento veritical) baseada em uma ou duas controladoras interconectadas a um ou mais gabinetes de discos, onde a ampliação do armazenamento é realizada com a adição de gavetas de disco e está limitada à capacidade das controladoras e a desduplicação é do tipo em linha (in-line) e global para o volume de armazenamento gerenciado por essas controladoras. Nesse caso a solução:	N/A	N/A	N/A
5.18.2.1.1. Entregar duas controladoras, no mínimo no modelo ativo-passivo, com discos sólidos (SSD) para aceleração de desduplicação e reconstrução de dados (rehidratação);	N/A	N/A	N/A
5.18.2.1.2. Deverá possuir interfaces de rede redundantes e dedicadas a interconexão de alta disponibilidade da solução, empregando interfaces 10G Ethernet SFP+, incluindo transceivers Short-Range e fibras OM4, multi-modo, de 1.0m;	N/A	N/A	N/A

5.18.2.1.3. Permitir desempenho de, no mínimo, 25 TB/h (vinte e cinco terabytes por hora) para tarefas de backup. O desempenho deve ser possível sem considerar desduplicação na origem, compressão ou componentes de software e hardware externos;	N/A	N/A	N/A
5.18.2.1.4. Deve permitir desduplicação global quando associado a, pelo menos, dois outros equipamentos similares do mesmo fabricante, com objetivo de constituir uma única área de armazenamento lógica. Se não possuir essa capacidade, deve ser fornecido com área de armazenamento 37,5% maior, considerando um ganho futuro de 1,6:1;	N/A	N/A	N/A
5.18.2.1.5. Deverá ser entregue habilitado para desduplicar sua máxima capacidade suportada.	N/A	N/A	N/A
5.18.2.2. Scale-out: Soluções com arquitetura hiperconvergente (crescimento horizontal), que possuem em seu módulo computacional processador, memória, interfaces de rede e discos associados e permita a agregação de vários módulos em um mesmo cluster onde a ampliação do armazenamento é realizada com asimples adição de módulos e a desduplicação é global entre eles. Nesse caso a solução:	ExaGrid-Detailed-Product-Description_DS-1-1.pdf	Scalable Architecture – Scale-Out	5
5.18.2.3. Permitir desempenho de, no mínimo, 13 TB/h (treze terabytes por hora) para tarefas de backup e restore, sem contabilizar o uso externo de softwares e hardwares;	ExaGrid-Detailed-Product-Description_DS-1-1.pdf	Max Backup Thruput (TB/hr) 15,25 15,25 TB = 13,86 TiB	4

5.18.2.4. Permitir desduplicação global quando associado a, pelo menos, dois outros módulos similares do mesmo fabricante, com objetivo de constituir uma única área de armazenamento lógica;	ExaGridAdminGuide.pdf	Adding Additional ExaGrid Servers When you add an ExaGrid Server, you not only increase the amount of storage space in your ExaGrid System, you also increase the system's ability to take in and process backup data. More ExaGrid Servers mean more memory, CPUs and ports over which backup data can flow. When a new ExaGrid Server is added, the ExaGrid System will automatically redistribute retained backup versions to balance the storage load across the ExaGrid Servers. ExaGrid 5.0 and later has global deduplication. Migrated moved jobs/clients will be deduplicated against their previous share/location.	110 e 111
5.18.2.5. Deverá ser entregue habilitado para desduplicar sua máxima capacidade suportada.	Conforme Lista de Part Numbers	N/A	N/A
5.18.3. Ser fornecido com portas Ethernet de 10Gbps do tipo SFP+, com suas respectivas GBICS, na quantidade suficiente para que o desempenho especificado seja alcançado;	Conforme Lista de Part Numbers	N/A	N/A
5.18.4. Deverá possuir, no mínimo, 16 (dezesseis) núcleos de processamento, com o dobro de threads, por controladora entregue;	ExaGrid-Technical_Specifications_DS-3.pdf	1 CPU 16 hyper-threaded cores	1
5.18.5. Deverá ser entregue com o máximo de memória suportado pelo equipamento, conforme divulgado em documentação oficial da fabricante;	ExaGrid-Detailed-Product-Description_DS-1-1.pdf	In contrast, ExaGrid's scalable approach with global deduplication adds full servers—including memory, processor, and bandwidth as well as disk The fgure below shows the differences between how the two different architectures cope with data growth over time	5
	ExaGrid-Technical_Specifications_DS-3.pdf	Memory 128 GiB	1
5.18.6. A solução deverá ser escalável a, no mínimo, 1000 TB (mil terabytes líquidos), seja através da adição de gavetas de discos ou de equipamentos similares em uma arquitetura scale-out;	ExaGrid-Detailed-Product-Description_DS-1-1.pdf	Scales to 32 appliances in a single scale-out system (2.69PB full backup @ 488TB/hr. ingest rate)	4

5.18.7. A solução deverá estar licenciada para receber imagens de backup desduplicadas na origem;	VeeamLatestGuide.pdf	ExaGrid Systems recommends configuring Veeam to deduplicate data within a job. Typically Veeam will achieve a 2:1 deduplication ratio. Once the deduplicated Veeam backup data has landed, the ExaGrid System will further deduplicate the Veeam data anywhere from 3:1 to 7:1.	71
5.18.8. Ter pelo menos 1 (um) Porta IPMI, 2 (Duas) Portas 1GB Ethernet e 2 (Duais) portas 10GB SFP+.	ExaGrid-Technical_Specifications_DS-3.pdf	Two 1GbE; two optional 10GbE, 25GbE or 40GbE	1
	ExaGridAdminGuide.pdf	The IPMI port is located in the back of your ExaGrid Server.	67

ITEM	DOCUMENTAÇÃO	COMPROVAÇÃO	PÁGINA
5.19. Unidade de armazenamento de informação - tipo 2	N/A	N/A	N/A
5.19.1. Possuir 62 TB (sessenta e dois terabytes) de área útil;	ExaGrid-Detailed-Product-Description_DS-1-1.pdf	EX36 Usable Capacity (TB) 72 TB = 65,4836 TiB	4
5.19.2. Possuir alguma das seguintes arquiteturas para o módulo de backup:	N/A	N/A	N/A
5.19.2.1. Scale-up: Soluções com arquitetura tradicional (crescimento veritical) baseada em uma ou duas controladoras interconectadas a um ou mais gabinetes de discos, onde a ampliação do armazenamento é realizada com a adição de gavetas de disco e está limitada à capacidade das controladoras e a desduplicação é do tipo em linha (in-line) e global para o volume de armazenamento gerenciado por essas controladoras. Nesse caso a solução:	N/A	N/A	N/A
5.19.2.1.1. Entregar duas controladoras, no mínimo no modelo ativo-passivo, com discos sólidos (SSD) para aceleração de desduplicação e reconstrução de dados (rehidratação);	N/A	N/A	N/A
5.19.2.1.2. Deverá possuir interfaces de rede redundantes e dedicadas a interconexão de alta disponibilidade da solução, empregando interfaces 10G Ethernet SFP+, incluindo transceivers Short-Range e fibras OM4, multi-modo, de 1.0m;	N/A	N/A	N/A

5.19.2.1.3. Permitir desempenho de, no mínimo, 10 TB/h (dez terabytes por hora) para tarefas de backup. O desempenho deve ser possível sem considerar desduplicação na origem, compressão ou componentes de software e hardware externos;	N/A	N/A	N/A
5.19.2.1.4. Permitir desduplicação global quando associado a, pelo menos, dois outros equipamentos similares do mesmo fabricante, com objetivo de constituir uma única área de armazenamento lógica. Se não possuir essa capacidade, deve ser fornecido com área de armazenamento 37,5% maior, considerando um ganho futuro de 1,6:1;	N/A	N/A	N/A
5.19.2.1.5. Deverá ser entregue habilitado para desduplicar sua máxima capacidade suportada.	N/A	N/A	N/A
5.19.2.2. Scale-out: Soluções com arquitetura hiperconvergente (crescimento horizontal), que possuem em seu módulo computacional processador, memória, interfaces de rede e discos associados e permita a agregação de vários módulos em um mesmo cluster onde a ampliação do armazenamento é realizada com a simples adição de módulos e a desduplicação é global entre eles. Nesse caso a solução:	ExaGrid-Detailed-Product-Description_DS-1-1.pdf	Scalable Architecture – Scale-Out	5
5.19.2.2.1. Permitir desempenho de, no mínimo, 5,5 TB/h (cinco e meio terabytes por hora) para tarefas de backup e restore, sem contabilizar o uso externo de softwares e hardwares;	ExaGrid-Detailed-Product-Description_DS-1-1.pdf	Max Backup Thruput (TB/hr) 7,78 tb = 7,075869 TiB	4

5.19.2.2.2. Permitir desduplicação global quando associado a, pelo menos, dois outros módulos similares do mesmo fabricante, com objetivo de constituir uma única área de armazenamento lógica.	ExaGridAdminGuide.pdf	Adding Additional ExaGrid Servers When you add an ExaGrid Server, you not only increase the amount of storage space in your ExaGrid System, you also increase the system's ability to take in and process backup data. More ExaGrid Servers mean more memory, CPUs and ports over which backup data can flow. When a new ExaGrid Server is added, the ExaGrid System will automatically redistribute retained backup versions to balance the storage load across the ExaGrid Servers. ExaGrid 5.0 and later has global deduplication. Migrated moved jobs/clients will be deduplicated against their previous share/location.	110 e 111
5.19.2.2.3. Deverá ser entregue habilitado para desduplicar sua máxima capacidade suportada.	Conforme Lista de Part-Numbers	N/A	N/A
5.19.2.3. Ser fornecido com portas Ethernet de 10Gbps do tipo SFP+, com suas respectivas GBICS, na quantidade suficiente para que o desempenho especificado seja alcançado;	Conforme Lista de Part-Numbers	N/A	N/A
5.19.2.4. Deverá possuir, no mínimo, 6 (seis) núcleos de processamento, com o dobro de threads, por controladora entregue;	ExaGrid-Technical_Specifications_DS-3.pdf	1 CPU 6 hyper-threaded cores	1
5.19.2.5. Deverá ser entregue com o máximo de memória suportado pelo equipamento, conforme divulgado em documentação oficial da fabricante;	ExaGrid-Technical_Specifications_DS-3.pdf	Memory 64 GB	1
5.19.2.6. A solução deverá ser escalável a, no mínimo, 500 TB (quinhentos terabytes líquidos), seja através da adição de gavetas de discos ou de equipamentos similares em uma arquitetura scale-out;	ExaGrid-Detailed-Product-Description_DS-1-1.pdf	Scales to 32 appliances in a single scale-out system (2.69PB full backup @ 488TB/hr. ingest rate)	4
5.19.2.7. A solução deverá estar licenciada para receber imagens de backup desduplicadas na origem;	VeeamLatestGuide.pdf	ExaGrid Systems recommends configuring Veeam to deduplicate data within a job. Typically Veeam will achieve a 2:1 deduplication ratio. Once the deduplicated Veeam backup data has landed, the ExaGrid System will further deduplicate the Veeam data anywhere from 3:1 to 7:1.	71

5.19.2.8. Deve ter pelo menos 1 (um) Porta IPMI, 2 (Duas) Portas 1GB Ethernet e 2 (Duais) portas 10GB SFP+.	ExaGrid-Technical_Specifications_DS-3.pdf	Four 1GbE; two optional 10GbE or 25GbE	1

ITEM	DOCUMENTAÇÃO	COMPROVAÇÃO	PÁGINA
5.20. Fitoteca de armazenamento modular	N/A	N/A	N/A
5.20.1. Características Gerais	N/A	N/A	N/A
5.20.1.1. Ser composto por todos os equipamentos e acessórios necessários para plena instalação e funcionamento;	Conforme Lista de Part Numbers	N/A	N/A
5.20.1.2. Ser do mesmo fabricante dos equipamentos ofertados para solução de software de proteção, ou estar homologado por ele, estando presente na lista de compatibilidade de hardware desses equipamentos;	lp0773.pdf	Veeam Backup and Retention	17
	https://helpcenter.veeam.com/docs/backup/vsphere/tape_supported_devices.html?ver=110	Veeam Backup & Replication supports Linear Tape-Open tape libraries starting from generation 3 (LTO3) or later:	N/A
5.20.1.3. Gabinete para rack, com tamanho máximo de 3U, acompanhado de:	lp0773.pdf	Base module: 3U rack-mount or stand-alone	6
5.20.1.3.1. Cabo de alimentação compatível com as PDU's do Rack;	Conforme Lista de Part Numbers	N/A	N/A
5.20.1.3.2. Trilhos e demais elementos de fixação necessários para a instalação em rack de 19 polegadas, do próprio fabricante dos equipamentos;	Conforme Lista de Part Numbers	N/A	N/A
5.20.1.3.3. Possuir fonte de alimentação redundante 110/220Vac;	Conforme Lista de Part Numbers	N/A	N/A
5.20.1.3.4. Ser modular permitindo expandir a capacidade de drives e slots através de módulos de expansão para no mínimo 21 unidades de operação e 272 slots;	lp0773.pdf	Per 7-module library: From up to 7x full-high and 7x half-high tape drives to up to 21x half-high tape drives	6
	lp0773.pdf	Per 7-module library: 280	6
5.20.1.3.5. Ter suporte e garantia de 3 (três) anos, com atendimento 24 x 7 x 365 e até 2 (duas) horas para início do atendimento de chamados com prioridade mais alta, ou seja, quando o equipamento estiver em estado de indisponibilidade de uso. É	lp0773.pdf	Advanced Service: 24x7 service coverage with 2-hour onsite response or 6-hour committed repair	14
	lp1049.pdf	Advanced Service: 24x7 service coverage with 2-hour onsite response or 6-hour committed repair	79

troca de peça no próximo dia útil.			
5.20.1.3.6. Ter acesso direto ao engenheiro de nível 2 do fabricante nomeado para a CONTRATANTE, que possibilite a execução das seguintes atividades durante toda a vigência da garantia, sem limite de quantidade:	Conforme Lista de Part Numbers	N/A	N/A
5.20.1.3.6.1. Atualização de software e/ou aplicação de correções;	Conforme Lista de Part Numbers	N/A	N/A
5.20.1.3.6.2. Revisão do ambiente para validação das configurações e, se necessário, realizar os ajustes com as melhores práticas indicadas;	Conforme Lista de Part Numbers	N/A	N/A
5.20.1.3.6.3. Abertura de chamados de suporte direto com engenheiro de nível 2, sem necessidade de triagem de nível 1;	Conforme Lista de Part Numbers	N/A	N/A
5.20.1.3.6.4. Reinstalações ou reconfigurações que se fizerem necessárias, mesmo que não sejam decorrentes de problemas de suporte.	Ciente e de Acordo	N/A	N/A
5.21. Fitoteca de Armazenamento – base	N/A	N/A	N/A
5.21.1. Suportar múltiplos caminhos e particionamento lógico;	lp0773.pdf	Multi-path feature to allow sharing of the library robotics across heterogeneous applications independently of each other. The library can be partitioned into up to 21 logical libraries (each partition must have at least one drive). It also can provide each logical library its own separate and distinct drives, storage slots, and control paths.	2

5.21.2. Suportar funcionalidade de alta disponibilidade de caminhos que garanta o uso de um caminho de comunicação redundante quando o caminho principal falha;	lp0773.pdf	Multi-path feature to allow sharing of the library robotics across heterogeneous applications independently of each other. The library can be partitioned into up to 21 logical libraries (each partition must have at least one drive). It also can provide each logical library its own separate and distinct drives, storage slots, and control paths.	2
5.21.3. Suportar até 3 unidades LTO 6, 7 ou 8;	lp0773.pdf	Per module: Up to 1x full-high and 1x half-high tape drives, or up to 3x half-high tape drives	6
5.21.4. Ser entregue com 1 unidade LTO-7 (Linear Tape-Open geração 7), com capacidade de gravação mínima de 6TB em cada cartucho, sem o uso de compressão;	Conforme Lista de Part Numbers	N/A	N/A
5.21.5. Ser acompanhada por no mínimo 2 (duas) portas SAS de 6 Gb;	lp0773.pdf	Half High 6 Gb SAS LTO 7 HH SAS Drive	9
5.21.6. Ter compatibilidade de leitura e escrita com o padrão LTO-6, e de leitura com o padrão LTO-5;	lp0773.pdf	LTO 7 tape drives can read and write LTO 7 and LTO 6 data cartridges, and can read LTO 5 data cartridges.	11
5.21.7. Cada unidade de leitura e gravação deverá possuir taxa de transferência de no mínimo 300 Mbps, sem o uso de compressão;	lp0773.pdf	LTO 8 HH, LTO7: Up to 300 MB/s	7
5.21.8. Cada unidade deverá ser acompanhada por no mínimo 2 (dois) cabos Mini-SAS para SAS de 1,5 m;	Conforme Lista de Part Numbers	N/A	N/A
5.21.9. O equipamento deverá ser do tipo “library”, com capacidade de armazenamento mínima para 40 (quarenta) cartuchos LTO-7;	lp0773.pdf	Number of cartridge slots Per module: 40 Per 7-module library: 280	6
5.21.10. Contar com interface Ethernet dedicada para gerenciamento, através de redes TCP/IP, compatível com os protocolos HTTP e SNMP;	lp0773.pdf	2x 10/100/1000 Mb Ethernet ports (UTP, RJ-45) for remote management: Web-based GUI; SNMP and email notifications	7

5.21.11. Possuir dispositivo que permita a identificação dos cartuchos por código de barras;	lp0773.pdf	Robotics: The base library will include scaling robotics with a barcode reader for identifying, locating, and moving LTO cartridges between storage locations and drives.	3
5.21.12. Estar acompanhado de 35 (trinta e cinco) cartuchos normais de fita, no padrão LTO-7, e quatro cartuchos de limpeza;	Conforme Lista de Part Numbers	N/A	N/A
5.21.13. Os cartuchos já devem ser acompanhados das respectivas etiquetas de código de barras;	Conforme Lista de Part Numbers	N/A	N/A
5.21.14. Ser compatível, e estar homologado, com os sistemas operacionais:	N/A	N/A	N/A
5.21.14.1. Microsoft Windows Server 2016 ou superior;	https://www.ibm.com/docs/en/ts4300-tape-library?topic=instructions-windows-server-2008-2012-2016-2019	This section describes how to install, remove, and uninstall the Windows tape and medium changer device drivers on Windows Server 2008, 2012, 2016, and 2019.	N/A
5.21.14.2. Red Hat Enterprise Linux 7.6 ou superior;	https://delivery04.dhe.ibm.com/sar/CMA/STA/08y3t/1/lin_tape.fixlist	- Support for TS4300 - Initial support for RHEL 7.6 (min. kernel version 3.10.0-957.1.3 only)	N/A
5.21.14.3. SUSE Linux Enterprise Server (SLES) 15 ou superior.	https://delivery04.dhe.ibm.com/sar/CMA/STA/08y3t/1/lin_tape.fixlist	- Support for TS4300 - Initial support for SLES 15 SP1 (min. kernel version 4.12.14-197.1)	N/A
5.21.15. Vir acompanhado também da unidade de controle, que deve possuir no mínimo:	N/A	N/A	N/A
5.21.15.1. Sistema operacional base licenciado, Microsoft Windows Server 2016 ou superior. Não serão aceitos softwares que não possuam suporte da fabricante;	Conforme Lista de Part Numbers	N/A	N/A
5.21.15.1.1. O sistema operacional deverá ser instalado em SSD redundante, configurado em RAID-1 e com, ao menos, 80 GiB de área de armazenamento líquida total;	Conforme Lista de Part Numbers	N/A	N/A
5.21.15.2. Entregue com redundância de CPU, com no mínimo 8 cores e hyperthread. Deverão ser da última geração de	Conforme Lista de Part Numbers	N/A	N/A

processadores ofertados pela fabricante do mesmo;			
5.21.15.3. Suporte a CPUs com memória base de 2.400, 2.666 e 2.933 MHz;	lp1049.pdf	4208 | 2400 MHz 5215 | 2666 MHz 5222 | 2933 MHz	23 e 24
5.21.15.4. Entregue com 32 GB de memória RAM;	Conforme Lista de Part Numbers	N/A	N/A
5.21.15.4.1. O servidor ofertado deverá suportar uma quantidade idêntica de DIMMs de memória por processador instalado, não sendo aceitas ofertas onde os processadores podem ser configurados com quantidades distintas de DIMMs por soquete;	Ciente e de Acordo	N/A	N/A
5.21.15.5. Entregue com 2 (duas) interfaces de rede de 10Gbps e cabos passivos de conexão direta e 5,0m;	Conforme Lista de Part Numbers	N/A	N/A
5.21.15.6. Respeitando as seguintes características de armazenamento:	N/A	N/A	N/A
5.21.15.6.1. Placa de hardware RAID com 2GB de cache;	lp0652.pdf	RAID 930-8i 2GB (Standard)	4
5.21.15.6.2. Suporte aos níveis RAID 0, 1, 10, 5, 50, 6, 60;	lp0652.pdf	0, 1, 10, 5, 50, 6, 60	4
5.21.15.6.3. Suporte a discos HDD, SSD e SED;	lp0652.pdf	HDD, SED, SSD	4
5.21.15.6.4. Com 2 (duas) interfaces Mini-SAS de 12Gb/s;	lp0652.pdf	2x Mini-SAS HD x4 (SFF-8643)	4
5.21.15.6.4.1. Deverá possuir, no mínimo, um slot extra para expansão de HBAs;	lp1049.pdf	Up to four slots. Slot 4 is the fixed slot on the system planar, and the remaining slots depend on the riser cards installed. The slots are as follows:	8
5.21.15.7. Respeitando as seguintes características de gerenciamento:	N/A	N/A	N/A
5.21.15.7.1. Controlar o consumo energético do módulo;	lp1049.pdf	Capping power usage	67
5.21.15.7.2. Permitir o gerenciamento remoto da solução;	lp1049.pdf	Remotely controlling server power (Power on, Power off, Restart)	66
5.21.15.7.3. Permitir o gerenciamento IPMI-over-LAN;	lp1049.pdf	Redirecting serial console via IPMI	66
5.21.15.7.4. Permitir o mapeamento de imagens através de compartilhamentos HTTPS, SFTP, CIFS e NFS;	lp1049.pdf	Mounting the remote ISO and image files via HTTPS, SFTP, CIFS, and NFS	67

5.21.15.7.5. Permitir o uso concomitante da interface de gerência por, no mínimo, 6 (seis) usuários;	lp1049.pdf	Collaborating across up to six users of the virtual console	67
5.21.15.7.6. Permitir o controle de consumo de banda de rede;	lp1049.pdf	Controlling quality and bandwidth usage	67
5.21.15.8. Possuir fontes e ventiladores hot-swap e redundantes, com tensão bivolt;	lp1049.pdf	Up to two redundant hot-swap 550 W, 750 W, or 1100 W (100 - 240 V) High Efficiency Platinum, or 750 W (200 - 240 V) High Efficiency Titanium AC power supplies. HVDC support (PRC only). Five (one processor) or seven (two processors) hot-swap dual-rotor system fans with N+1 redundancy	8

ITEM	DOCUMENTAÇÃO	COMPROVAÇÃO	PÁGINA
5.22. Fitoteca de Armazenamento – Expansão	N/A	N/A	N/A
5.22.1. Suportar até 3 unidades LTO 6, 7 ou 8;	lp0773.pdf	Per module: Up to 1x full-high and 1x half-high tape drives, or up to 3x half-high tape drives	6
5.22.2. Possuir 1 unidade LTO-7 (Linear Tape-Open geração 7), com capacidade de gravação mínima de 6TB em cada cartucho, sem o uso de compressão;	lp0773.pdf	Native cartridge capacity of up to 18 TB (Ultrium 9), 12 TB (Ultrium 8), 9 TB (Ultrium 7 initialized LTO 8 Type M), 6 TB (Ultrium 7).	2
5.22.3. Cada unidade LTO-7 deverá possuir conectividade SAS, de no mínimo, 6Gb com 2 portas por unidade LTO;	Conforme Lista de Part-Numbers	N/A	N/A
5.22.4. Cada unidade deverá ser acompanhada por no mínimo 2 (dois) cabos Mini-SAS para SAS de 1,5m;	Conforme Lista de Part-Numbers	N/A	N/A
5.22.5. Ter compatibilidade de leitura e escrita com o padrão LTO-6, e de leitura com o padrão LTO-5;	lp0773.pdf	LTO 7 tape drives can read and write LTO 7 and LTO 6 data cartridges, and can read LTO 5 data cartridges.	11
5.22.6. Cada unidade de leitura e gravação deverá possuir taxa de transferência de no mínimo 300 Mbps, sem o uso de compressão;	lp0773.pdf	Native data transfer rate of up to 400 MB/s (Ultrium 9 FH), 300 MBp/s (Ultrium 9 HH), 360 MB/s (Ultrium 8 FH), 300 MB/s (Ultrium 8 HH or Ultrium 7), or up to 160 MB/s (Ultrium 6) per tape drive.	2
5.22.7. O equipamento deverá ser do tipo “Expansion module”, com capacidade de armazenamento mínima para 40 cartuchos LTO-7;	lp0773.pdf	The following figure shows the front view of the TS4300 Tape Library expansion module	5

5.22.8. Estar acompanhado de 35 (trinta e cinco) cartuchos normais de fita, no padrão LTO-7;	Conforme Lista de Part-Numbers	N/A	N/A
5.22.9. Ter suporte e garantia de 3 (três) anos, com atendimento 24 x 7 x 365 e até 2 (duas) para início do atendimento de chamados com prioridade mais alta, ou seja, quando o equipamento estiver em estado de indisponibilidade de uso. É troca de peça no próximo dia útil;	Conforme Lista de Part-Numbers	N/A	N/A
5.22.90. Os cartuchos já devem ser acompanhados das respectivas etiquetas de código de barras.	Conforme Lista de Part-Numbers	N/A	N/A

Lenovo ThinkSystem RAID 930 Series Internal RAID Adapters

Product Guide

The ThinkSystem RAID 930 family of internal 12 Gbps SAS RAID controllers are high-performance RAID-on-chip (ROC) adapters. These adapters support RAID levels 0/1/10/5/50/6/60 as well as JBOD, and include an extensive list of RAS and management features.

The family is comprised of four adapters:

- The ThinkSystem RAID 930-4i supports up to four internal SAS and SATA drives
- The ThinkSystem RAID 930-8i supports up to eight internal SAS and SATA drives
- The ThinkSystem RAID 930-16i supports up to 16 internal SAS and SATA drives
- The ThinkSystem RAID 930-24i supports up to 24 internal SAS and SATA drives

The ThinkSystem RAID 930-8i is shown in the following figure.

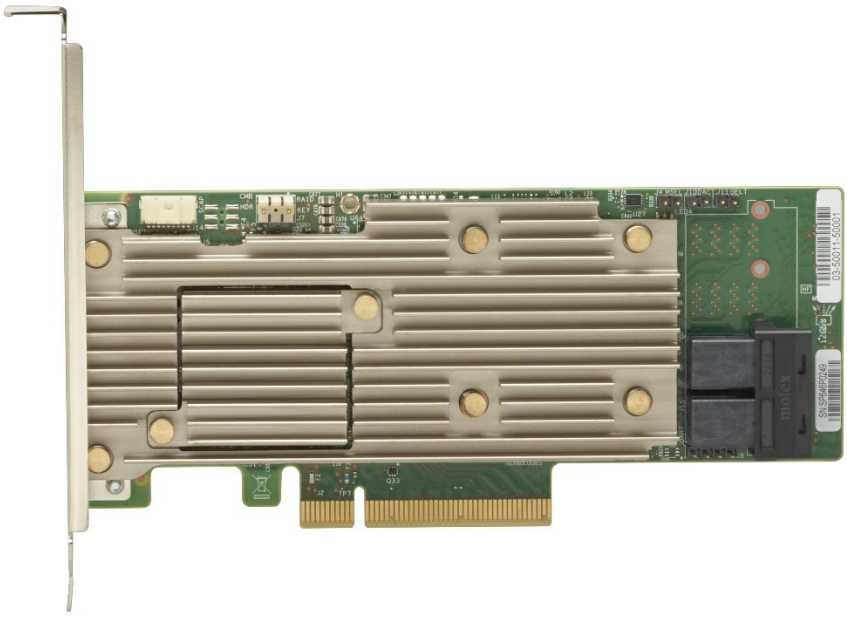


Figure 1. ThinkSystem RAID 930-8i 2GB Flash PCIe 12Gb Adapter

Did you know?

RAID on Chip-based controllers such as the RAID 930 adapters have a dedicated processor that offloads all RAID functions from the server's CPU. With hardware acceleration for RAID 5 and 6 operations plus dedicated memory for caching, the 930-8i, 16i and 24i offer the ultimate performance for ThinkSystem servers.

Click here to check for updates

Part number information

The following table provides the ordering part numbers for the adapters.

Table 1. Part numbers and feature codes

Part number	Feature code	Description
7Y37A01086	AUV1	ThinkSystem RAID 930-24i 4GB Flash PCIe 12Gb Adapter
4Y37A09721	B31E	ThinkSystem RAID 930-16i 8GB Flash PCIe 12Gb Adapter
7Y37A01085	AUNK	ThinkSystem RAID 930-16i 4GB Flash PCIe 12Gb Adapter
7Y37A01084	AUNJ	ThinkSystem RAID 930-8i 2GB Flash PCIe 12Gb Adapter
4Y37A16227	B5DP	ThinkSystem SR670 RAID 930-8i 2GB Flash PCIe 12Gb Adapter
7M27A03917	AUYS	ThinkSystem RAID 930-4i-2GB 2 Drive Adapter Kit for SN550
7M17A03933	AVED	ThinkSystem RAID 930-4i-2GB 4 Drive Adapter Kit for SN850

The following figure shows the ThinkSystem RAID 930-16i.

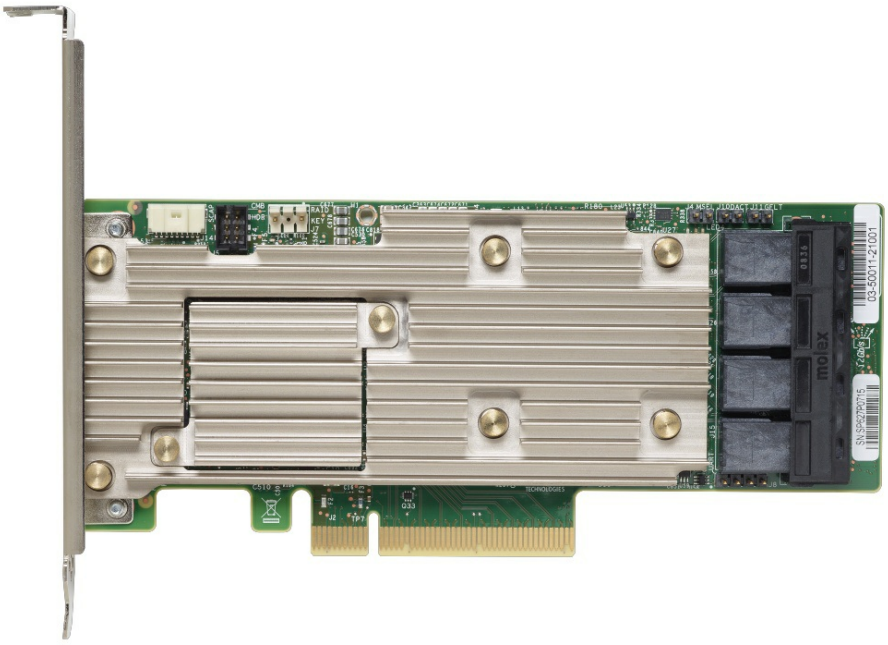


Figure 2. ThinkSystem RAID 930-16i PCIe 12Gb Adapter

Technical specifications

The ThinkSystem RAID 930 Series internal RAID adapters have the following specifications:

- PCIe 3.0 x8 host interface
- 12 Gbps SAS/SATA RAID controllers, based on the Broadcom MegaRAID 9460 adapter family
- 2GB, 4GB or 8GB of integrated flash-backed cache
- Connectivity for up to 2, 8, 16 or 24 internal SAS or SATA drives
- Support for intermixing SAS and SATA HDDs and SSDs. Mixing SAS and SATA drives in the same array is not supported. Mixing of HDDs and SSDs in the same array is not supported.
- Support for intermixing of 12 Gbps and 6 Gbps drives.

- Support for RAID 0, 1, 10, 5, 50, 6 and 60 standard (The 930-4i supports only RAID 0, 1 in the SN550 and RAID 0, 1, 5, 6 in the SN850 because those blade servers only support two and four drives respectively)
- Support for JBOD (non-RAID) drive state
- Support for up to 64 virtual disks, up to 128 arrays, up to 16 virtual disks per array
- Support for logical drive sizes greater than 2 TB.
- Configurable stripe size from 64 KB up to 1 MB
- Supports 512e, 512n and 4K sector formatted drives
- Supports UNMAP command for SAS SSDs (JBOD and RAID volumes)
- Supports TRIM command for SATA SSDs (JBOD)
- Compliant with Disk Data Format (DDF) configuration on disk (CoD).
- S.M.A.R.T. support.
- Configuration through
    - XClarity Provisioning Manager UEFI interface
    - XClarity Controller web interface
    - XClarity Administrator Configuration Patterns
    - StorCLI command-line interface
    - LSI Storage Authority (LSA) GUI interface
    - UEFI Human Interface Infrastructure (HII)

Note: CacheCade is not supported by these adapters

The following table compares the specifications of the RAID 930 adapters.

Table 2. Specifications

Feature	RAID 930-4i	RAID 930-8i	RAID 930-16i	RAID 930-24i
Part number	7M27A03917 7M17A03933	7Y37A01084 4Y37A16227	7Y37A01085 4Y37A09721	7Y37A01086
Form factor	Custom	Low profile	Low profile	Full height half length
Controller chip	LSI SAS3504	LSI SAS3508	LSI SAS3516	LSI SAS3508 + LSI SAS35x36R
Host interface	PCIe 3.0 x8	PCIe 3.0 x8	PCIe 3.0 x8	PCIe 3.0 x8
Port interface	12 Gb SAS	12 Gb SAS	12 Gb SAS	12 Gb SAS
Number of ports	2 or 4*	8	16	24
Port connectors	1x Mini-SAS HD x4 (SFF-8643)	2x Mini-SAS HD x4 (SFF-8643)	4x Mini-SAS HD x4 (SFF-8643)	6x Mini-SAS HD x4 (SFF-8643)
Drive interface	SAS, SATA	SAS, SATA	SAS, SATA	SAS, SATA
Drive type	HDD, SED, SSD	HDD, SED, SSD	HDD, SED, SSD	HDD, SED, SSD
Hot-swap drives	Yes	Yes	Yes	Yes
Max devices	2	8	16	24
RAID levels	0, 1, 5, 6*	0, 1, 10, 5, 50, 6, 60	0, 1, 10, 5, 50, 6, 60	0, 1, 10, 5, 50, 6, 60
JBOD drive state	Yes	Yes	Yes	Yes
Cache	2GB (Standard)	2GB (Standard)	4GB (Standard) 8GB (Standard)	4GB (Standard)
CacheVault cache protection with flash	Yes	Yes	Yes	Yes
Performance Accelerator (FastPath)	Yes	Yes	Yes	Yes
SED support (SafeStore)	Yes	Yes	Yes	Yes

* The 930-4i for use in the SN550 blade server only supports RAID 0 and RAID 1 because the server has only two drives. The 930-4i for use in the SN850 blade server only supports RAID 0, 1, 5 and 6 because the server has only four drives.

Features

The ThinkSystem RAID 930 adapters have the following standard features:

- MegaRAID FastPath SSD performance acceleration

  MegaRAID FastPath software provides high-performance I/O acceleration for SSD-based virtual drives by using a low latency I/O path to increase the maximum I/O per second (IOPS) capability of the controller. This feature boosts the performance of applications with a highly random data storage access pattern, such as transactional databases.

- MegaRAID flash cache protection

  MegaRAID flash cache protection uses NAND flash memory, which is powered by a CacheVault Power Module supercapacitor, to protect data that is stored in the controller cache. This module eliminates the need for a lithium-ion battery, which is commonly used to protect DRAM cache memory on PCI RAID controllers. To avoid the possibility of data loss or corruption during a power or server failure, flash cache protection technology transfers the contents of the DRAM cache to NAND flash using power from the offload power module. After the power is restored to the RAID controller, the content of the NAND flash is transferred back to the DRAM, which is flushed to disk.

- Auto-resume on array rebuild or array reconstruction after the loss of system power

  Auto-resume uses non-volatile RAM (NVRAM) to save the rebuild progress during a host reboot or power failure to automatically resume from the last checkpoint. Auto-resume ensures that data integrity is maintained throughout the process. The card supports a number of features that can be implemented without rebooting the server. Applications, such as email and web server, benefit from avoiding downtime during the transition.

- Online Capacity Expansion

  Online Capacity Expansion (OCE) allows the capacity of a virtual disk to be expanded by adding new physical disks or making use of unused space on existing disks, without requiring a reboot.

- Online RAID Level Migration

  Online RAID Level Migration (RLM), which is also known as logical drive migration, can migrate a virtual disk from any RAID level to any other RAID level without requiring a reboot. System availability and application functionality remain unaffected.

- Fast initialization for quick array setup

  Fast initialization quickly writes zeros to the first and last sectors of the virtual drive. This feature allows you to immediately start writing data to the virtual drive while the initialization is running in the background.

- Consistency check for background data integrity

  Consistency check verifies that all stripes in a virtual disk with a redundant RAID level are consistent. The consistency check mirrors data when an inconsistent stripe is detected for RAID 1 and re-creates the parity from the peer disks for RAID 5 or RAID 6. Consistency checks can be scheduled to take place periodically.

- Extensive online configuration options and advanced monitoring and event notification

  Management tools provide convenience for the configuration of logical volumes and alerting when errors have occurred or are about to occur.

- Patrol read for media scanning and repairing

  Patrol read is a background sentry service that pro-actively discovers and corrects media defects (bad sectors) that arise normally as a disk drive ages. The service issues a series of verify commands, and if a bad block is discovered, the card's firmware uses RAID algorithms to re-create the missing data and remap the sector to a good sector. The task is interruptible based on controller activity and host operations. The firmware also provides an interface where the patrol read task can be initiated, set up for continuous operation, and terminated from a management application. Patrol read can be activated by a manual command or automatically.

- Global and dedicated hot spare with revertible hot spare support

  A hot spare rebuilds data from all virtual disks within the disk group in which it is configured. You can define a physical disk as a hot spare to replace a failed drive. Hot spares can be configured as either global or dedicated. A global hot spare allows any physical drive to be designated as a hot spare. A dedicated hot spare allows the user to assign a hot spare drive to a particular array of the same drive type.

- Drive roaming

  Drive roaming occurs when the physical disks are changed to different ports on the same controller. When the drives are placed on different channels, the controller detects the RAID configuration from the configuration data on the drives.

- MegaRAID SafeStore support for self-encrypting drive (SED) services

  MegaRAID SafeStore encryption services offer instant secure erase and local key management for self-encrypting drives. This technology represents a step forward in securing data on a disk drive from any unauthorized access or modification resulting from theft, loss, or repurposing of drives. Instant secure erase permanently removes data when repurposing or decommissioning SEDs. SafeStore local key management provides the necessary management and protection of SEDs by using a simple pass phrase, security key identifier, and security key file that can be set and applied to all SEDs that are assigned to a RAID adapter. This feature removes the complexity of managing each SED's unique encryption key, and it essentially relieves the administrator of most of the daily tasks of securing data.

- XClarity Provisioning Manager for pre-boot array configuration and management

  Provisioning Manager is the ThinkSystem UEFI-based application that includes a RAID setup wizard to help you configure drive groups and virtual disks before installing or booting the operating system.

- XClarity Controller web interface for remote storage management

  XClarity Controller (XCC) is the systems management processor integrated in all ThinkSystem servers. The XCC web interface allows you to perform storage inventory, create and edit virtual disks, view events, import a new controller configuration, and perform firmware updates on the adapter.

Note: MegaRAID CacheCade and MegaRAID Storage Manager are not supported.

The ThinkSystem RAID 930-24i is shown in the following figure.

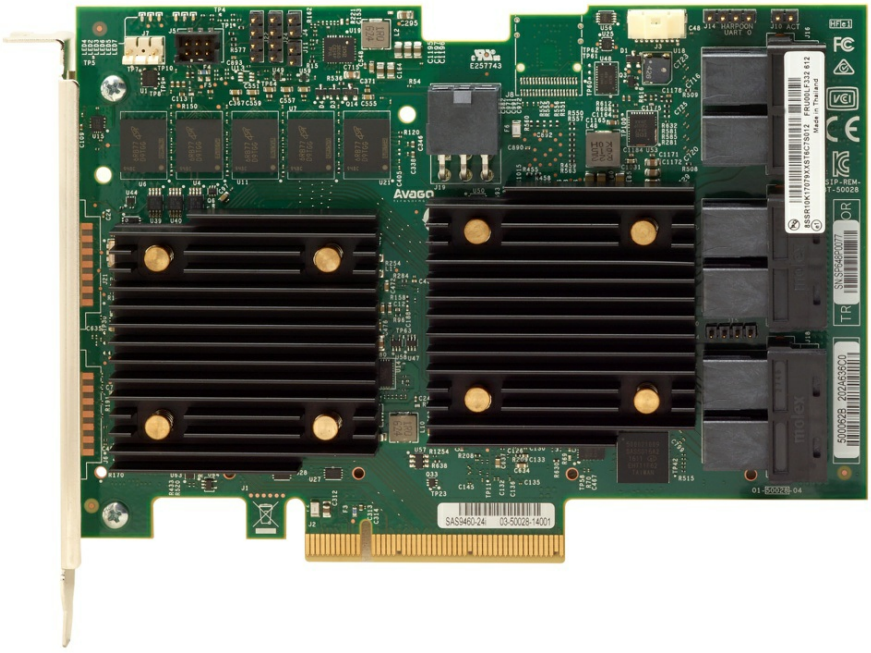


Figure 3. ThinkSystem RAID 930-24i 4GB Flash PCIe 12Gb Adapter

Server support

The following tables list the ThinkSystem servers that are compatible.

Table 3. Server support (Part 1 of 2)

Part Number	Description	Edge		1S Intel V2			2S Intel V2				AMD				Dense V2				4S V2		8S
		SE350 (7Z46 / 7D1X)	SE450 (7D8T)	ST50 V2 (7D8K / 7D8J)	ST250 V2 (7D8G / 7D8F)	SR250 V2 (7D7R / 7D7Q)	ST650 V2 (7Z75 / 7Z74)	SR630 V2 (7Z70 / 7Z71)	SR650 V2 (7Z72 / 7Z73)	SR670 V2 (7Z22 / 7Z23)	SR635 (7Y98 / 7Y99)	SR655 (7Y00 / 7Z01)	SR645 (7D2Y / 7D2X)	SR665 (7D2W / 7D2V)	SD630 V2 (7D1K)	SD650 V2 (7D1M)	SD650-N V2 (7D1N)	SN550 V2 (7Z69)	SR850 V2 (7D31 / 7D32)	SR860 V2 (7Z59 / 7Z60)	SR950 (7X11 / 7X12)
7Y37A01086	ThinkSystem RAID 930-24i 4GB Flash PCIe 12Gb Adapter	N	N	N	N	N	N	N	N	N	N	Y	N	N	N	N	N	N	N	N	N
4Y37A09721	ThinkSystem RAID 930-16i 8GB Flash PCIe 12Gb Adapter	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	Y
7Y37A01085	ThinkSystem RAID 930-16i 4GB Flash PCIe 12Gb Adapter	N	N	N	N	N	Y	Y	Y	N	Y	Y	Y	Y	N	N	N	N	Y	Y	Y
7Y37A01084	ThinkSystem RAID 930-8i 2GB Flash PCIe 12Gb Adapter	N	N	N	N	N	Y	Y	Y	Y	Y	Y	Y	Y	N	N	N	N	Y	Y	Y
4Y37A16227	ThinkSystem SR670 RAID 930-8i 2GB Flash	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N
7M27A03917	ThinkSystem RAID 930-4i-2GB 2 Drive Adapter Kit for SN550	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N
7M17A03933	ThinkSystem RAID 930-4i-2GB 4 Drive Adapter Kit for SN850	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N

Table 4. Server support (Part 2 of 2)

Part Number	Description	1S Intel V1				2S Intel V1								Dense V1				4S V1		
		ST50 (7Y48 / 7Y50)	ST250 (7Y45 / 7Y46)	SR150 (7Y54)	SR250 (7Y52 / 7Y51)	ST550 (7X09 / 7X10)	SR530 (7X07 / 7X08)	SR550 (7X03 / 7X04)	SR570 (7Y02 / 7Y03)	SR590 (7X98 / 7X99)	SR630 (7X01 / 7X02)	SR650 (7X05 / 7X06)	SR670 (7Y36 / 7Y37)	SD530 (7X21)	SD650 (7X58)	SN550 (7X16)	SN850 (7X15)	SR850 (7X18 / 7X19)	SR850P (7D2F / 2D2G)	SR860 (7X69 / 7X70)
7Y37A01086	ThinkSystem RAID 930-24i 4GB Flash PCIe 12Gb Adapter	N	N	N	N	Y	N	N	N	N	N	Y	N	N	N	N	N	N	N	N
4Y37A09721	ThinkSystem RAID 930-16i 8GB Flash PCIe 12Gb Adapter	N	N	N	N	Y	Y	Y	Y	Y	Y	Y	N	N	N	N	N	Y	Y	Y
7Y37A01085	ThinkSystem RAID 930-16i 4GB Flash PCIe 12Gb Adapter	N	Y	N	Y	Y	Y	Y	Y	Y	Y	Y	N	N	N	N	N	Y	Y	Y
7Y37A01084	ThinkSystem RAID 930-8i 2GB Flash PCIe 12Gb Adapter	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	N	N	N	N	N	Y	Y	Y
4Y37A16227	ThinkSystem SR670 RAID 930-8i 2GB Flash	N	N	N	N	N	N	N	N	N	N	N	Y	N	N	N	N	N	N	N
7M27A03917	ThinkSystem RAID 930-4i-2GB 2 Drive Adapter Kit for SN550	N	N	N	N	N	N	N	N	N	N	N	N	N	N	Y	N	N	N	N
7M17A03933	ThinkSystem RAID 930-4i-2GB 4 Drive Adapter Kit for SN850	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N	Y	N	N	N

The ThinkSystem RAID 930-4i is shown in the following figure.

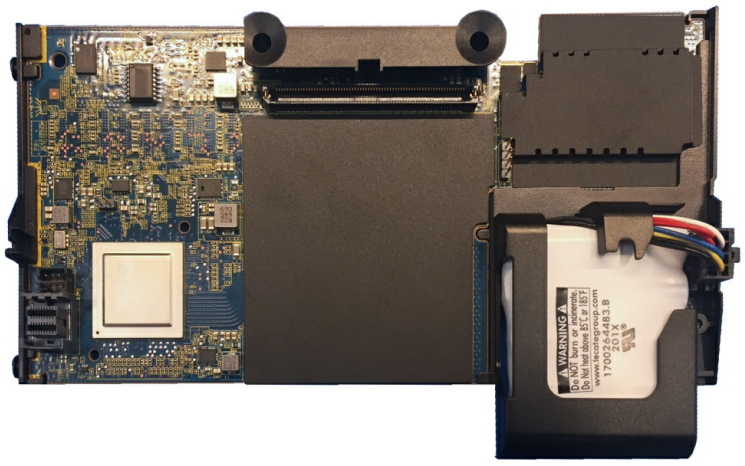


Figure 4. ThinkSystem RAID 930-4i-2GB 2 Drive Adapter Kit for SN550 (included heatsink not shown)

Operating system support

The following tables list the supported operating systems for the adapters:

- ThinkSystem RAID 930-8i 2GB Flash PCIe 12Gb Adapter, 7Y37A01084
- ThinkSystem RAID 930-16i 4GB Flash PCIe 12Gb Adapter, 7Y37A01085
- ThinkSystem RAID 930-16i 8GB Flash PCIe 12Gb Adapter, 4Y37A09721
- ThinkSystem RAID 930-24i 4GB Flash PCIe 12Gb Adapter, 7Y37A01086

Table 5. Operating system support for ThinkSystem RAID 930-8i 2GB Flash PCIe 12Gb Adapter, 7Y37A01084 (Part 1 of 3)

Operating systems	SR150	SR250	ST250	ST50	SR630 V2	SR650 V2	SR670 V2	SR850 V2	SR860 V2	ST650 V2	SR635	SR645	SR655	SR665
Microsoft Windows Server 2012 R2	N	N	N	N	N	N	N	N	N	N	N	N	N	N
Microsoft Windows Server 2016	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Microsoft Windows Server 2019	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Microsoft Windows Server 2022	N	N	N	N	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Microsoft Windows Server version 1709	N	N	N	N	N	N	N	N	N	N	N	N	N	N
Microsoft Windows Server version 1803	N	N	N	N	N	N	N	N	N	N	N	N	N	N
Red Hat Enterprise Linux 6 Server x64 Edition	N	N	N	N	N	N	N	N	N	N	N	N	N	N
Red Hat Enterprise Linux 6.10	N	N	N	N	N	N	N	N	N	N	N	N	N	N
Red Hat Enterprise Linux 6.9	N	N	N	N	N	N	N	N	N	N	N	N	N	N
Red Hat Enterprise Linux 7.3	N	N	N	N	N	N	N	N	N	N	N	N	N	N
Red Hat Enterprise Linux 7.4	N	N	N	N	N	N	N	N	N	N	N	N	N	N
Red Hat Enterprise Linux 7.5	Y	Y	Y	Y	N	N	N	N	N	N	N	N	N	N
Red Hat Enterprise Linux 7.6	N	N	N	N	N	N	N	N	N	N	Y	Y	Y	Y
Red Hat Enterprise Linux 7.7	Y	Y	Y	Y	N	N	N	N	N	N	Y	Y	Y	Y
Red Hat Enterprise Linux 7.8	Y	Y	Y	Y	N	N	N	N	N	N	Y	Y	Y	Y
Red Hat Enterprise Linux 7.9	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.0	Y	Y	Y	Y	N	N	N	N	N	N	Y	N	Y	N
Red Hat Enterprise Linux 8.1	Y	Y	Y	Y	N	N	N	N	N	N	Y	Y	Y	Y
Red Hat Enterprise Linux 8.2	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.3	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.4	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.5	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 11 SP4	N	N	N	N	N	N	N	N	N	N	N	N	N	N
SUSE Linux Enterprise Server 11 SP4 with Xen	N	N	N	N	N	N	N	N	N	N	N	N	N	N
SUSE Linux Enterprise Server 12 SP2	N	N	N	N	N	N	N	N	N	N	N	N	N	N
SUSE Linux Enterprise Server 12 SP2 with Xen	N	N	N	N	N	N	N	N	N	N	N	N	N	N
SUSE Linux Enterprise Server 12 SP3	Y	Y	Y	N	N	N	N	N	N	N	N	N	N	N
SUSE Linux Enterprise Server 12 SP3 with Xen	Y	Y	Y	N	N	N	N	N	N	N	N	N	N	N
SUSE Linux Enterprise Server 12 SP4	N	N	N	N	N	N	N	N	N	N	Y	N	Y	N
SUSE Linux Enterprise Server 12 SP4 with Xen	N	N	N	N	N	N	N	N	N	N	Y	N	Y	N

Operating systems	SR150	SR250	ST250	ST50	SR630 V2	SR650 V2	SR670 V2	SR850 V2	SR860 V2	ST650 V2	SR635	SR645	SR655	SR665
SUSE Linux Enterprise Server 12 SP5	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 12 SP5 with Xen	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15	Y	Y	Y	N	N	N	N	N	N	N	N	N	N	N
SUSE Linux Enterprise Server 15 SP1	Y	Y	Y	Y	N	N	N	N	N	N	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP1 with Xen	Y	Y	Y	Y	N	N	N	N	N	N	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP2	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP2 with Xen	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP3	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP3 with Xen	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 with Xen	Y	Y	Y	N	N	N	N	N	N	N	N	N	N	N
Ubuntu 18.04 LTS	N	N	N	N	Y	Y	Y	N	N	Y	N	N	N	N
Ubuntu 20.04 LTS	N	N	N	N	Y	Y	N	N	N	N	N	N	N	N
VMware vSphere Hypervisor (ESXi) 6.0 U3	N	N	N	N	N	N	N	N	N	N	N	N	N	N
VMware vSphere Hypervisor (ESXi) 6.5	N	N	N	N	N	N	N	N	N	N	N	N	N	N
VMware vSphere Hypervisor (ESXi) 6.5 U1	N	N	N	N	N	N	N	N	N	N	N	N	N	N
VMware vSphere Hypervisor (ESXi) 6.5 U2	Y	Y	Y	Y	N	N	N	N	N	N	N	N	N	N
VMware vSphere Hypervisor (ESXi) 6.5 U3	Y	Y	Y	Y	N	N	N	N	N	N	Y	N	Y	N
VMware vSphere Hypervisor (ESXi) 6.7	Y	Y	Y	N	N	N	N	N	N	N	N	N	N	N
VMware vSphere Hypervisor (ESXi) 6.7 U1	N	N	N	N	N	N	N	N	N	N	N	N	N	N
VMware vSphere Hypervisor (ESXi) 6.7 U2	Y	Y	Y	Y	N	N	N	N	N	N	N	N	N	N
VMware vSphere Hypervisor (ESXi) 6.7 U3	Y	Y	Y	Y	Y	Y	Y	N	N	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 7.0	Y	Y	Y	Y	N	N	N	N	N	N	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 7.0 U1	Y	Y	Y	Y	N	N	N	Y	Y	N	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 7.0 U2	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 7.0 U3	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y

Table 6. Operating system support for ThinkSystem RAID 930-8i 2GB Flash PCIe 12Gb Adapter, 7Y37A01084 (Part 2 of 3)

Operating systems	SR530 (Gen 2)	SR550 (Gen 2)	SR570 (Gen 2)	SR590 (Gen 2)	SR630 (Gen 2)	SR650 (Gen 2)	SR670 (Gen 2)	SR850 (Gen 2)	SR850P	SR860 (Gen 2)	SR950 (Gen 2)	ST550 (Gen 2)
Microsoft Windows Server 2012 R2	N	N	N	N	N	N	N	N	N	N	N	N
Microsoft Windows Server 2016	Y	Y	Y	Y	Y	Y	N	Y	Y	Y	Y	Y
Microsoft Windows Server 2019	Y	Y	Y	Y	Y	Y	N	Y	Y	Y	Y	Y
Microsoft Windows Server 2022	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Microsoft Windows Server version 1709	N	N	N	N	N	N	N	N	N	N	N	N
Microsoft Windows Server version 1803	N	N	N	N	N	N	N	N	N	N	N	N

Operating systems	SR530 (Gen 2)	SR550 (Gen 2)	SR570 (Gen 2)	SR590 (Gen 2)	SR630 (Gen 2)	SR650 (Gen 2)	SR670 (Gen 2)	SR850 (Gen 2)	SR850P	SR860 (Gen 2)	SR950 (Gen 2)	ST550 (Gen 2)
Red Hat Enterprise Linux 6 Server x64 Edition	N	N	N	N	N	N	N	N	N	N	N	N
Red Hat Enterprise Linux 6.10	N	N	N	N	N	N	N	N	N	N	N	N
Red Hat Enterprise Linux 6.9	N	N	N	N	N	N	N	N	N	N	N	N
Red Hat Enterprise Linux 7.3	N	N	N	N	N	N	N	N	N	N	N	N
Red Hat Enterprise Linux 7.4	N	N	N	N	N	N	N	N	N	N	N	N
Red Hat Enterprise Linux 7.5	N	N	N	N	N	N	Y	N	N	N	N	N
Red Hat Enterprise Linux 7.6	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 7.7	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 7.8	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 7.9	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.0	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.1	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.2	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.3	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.4	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.5	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 11 SP4	N	N	N	N	N	N	N	N	N	N	N	N
SUSE Linux Enterprise Server 11 SP4 with Xen	N	N	N	N	N	N	N	N	N	N	N	N
SUSE Linux Enterprise Server 12 SP2	N	N	N	N	N	N	N	N	N	N	N	N
SUSE Linux Enterprise Server 12 SP2 with Xen	N	N	N	N	N	N	N	N	N	N	N	N
SUSE Linux Enterprise Server 12 SP3	N	N	N	N	N	N	N	N	Y	N	N	N
SUSE Linux Enterprise Server 12 SP3 with Xen	N	N	N	N	N	N	N	N	Y	N	N	N
SUSE Linux Enterprise Server 12 SP4	Y	Y	Y	Y	Y	Y	N	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 12 SP4 with Xen	Y	Y	Y	Y	Y	Y	N	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 12 SP5	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 12 SP5 with Xen	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15	Y	Y	Y	Y	Y	Y	N	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP1	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP1 with Xen	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP2	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP2 with Xen	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP3	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP3 with Xen	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 with Xen	Y	Y	Y	Y	Y	Y	N	Y	Y	Y	Y	Y
Ubuntu 18.04 LTS	N	N	N	N	N	N	N	N	N	N	N	N
Ubuntu 20.04 LTS	N	N	N	N	N	N	N	N	N	N	N	N
VMware vSphere Hypervisor (ESXi) 6.0 U3	N	N	N	N	N	N	N	N	N	N	N	N

Operating systems	SR530 (Gen 2)	SR550 (Gen 2)	SR570 (Gen 2)	SR590 (Gen 2)	SR630 (Gen 2)	SR650 (Gen 2)	SR670 (Gen 2)	SR850 (Gen 2)	SR850P	SR860 (Gen 2)	SR950 (Gen 2)	ST550 (Gen 2)
VMware vSphere Hypervisor (ESXi) 6.5	N	N	N	N	N	N	N	N	N	N	N	N
VMware vSphere Hypervisor (ESXi) 6.5 U1	N	N	N	N	N	N	N	N	N	N	N	N
VMware vSphere Hypervisor (ESXi) 6.5 U2	Y	Y	Y	Y	Y	Y	N	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 6.5 U3	Y	Y	Y	Y	Y	Y	N	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 6.7	N	N	N	N	N	N	N	N	N	N	N	N
VMware vSphere Hypervisor (ESXi) 6.7 U1	Y	Y	Y	Y	Y	Y	N	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 6.7 U2	Y	Y	Y	Y	Y	Y	N	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 6.7 U3	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 7.0	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 7.0 U1	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 7.0 U2	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 7.0 U3	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y

Table 7. Operating system support for ThinkSystem RAID 930-8i 2GB Flash PCIe 12Gb Adapter, 7Y37A01084 (Part 3 of 3)

Operating systems	SR530 (Gen 1)	SR550 (Gen 1)	SR570 (Gen 1)	SR590 (Gen 1)	SR630 (Gen 1)	SR650 (Gen 1)	SR670 (Gen 1)	SR850 (Gen 1)	SR860 (Gen 1)	SR950 (Gen 1)	ST550 (Gen 1)
Microsoft Windows Server 2012 R2	Y	Y	Y	Y	Y	Y	N	Y	Y	Y	Y
Microsoft Windows Server 2016	Y	Y	Y	Y	Y	Y	N	Y	Y	Y	Y
Microsoft Windows Server 2019	Y	Y	Y	Y	Y	Y	N	Y	Y	Y	Y
Microsoft Windows Server 2022	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Microsoft Windows Server version 1709	Y	Y	Y	Y	Y	Y	N	Y	Y	Y	Y
Microsoft Windows Server version 1803	N	N	N	N	Y	Y	N	Y	Y	Y	N
Red Hat Enterprise Linux 6 Server x64 Edition	N	Y	N	N	N	N	N	N	N	N	N
Red Hat Enterprise Linux 6.10	Y	Y	Y	Y	Y	Y	N	Y	Y	Y	Y
Red Hat Enterprise Linux 6.9	Y	Y	Y	Y	Y	Y	N	Y	Y	Y	Y
Red Hat Enterprise Linux 7.3	Y	Y	N	N	Y	Y	N	Y	N	Y	Y
Red Hat Enterprise Linux 7.4	Y	Y	Y	Y	Y	Y	N	Y	Y	Y	Y
Red Hat Enterprise Linux 7.5	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 7.6	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 7.7	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 7.8	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 7.9	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.0	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.1	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y

Operating systems	SR530 (Gen 1)	SR550 (Gen 1)	SR570 (Gen 1)	SR590 (Gen 1)	SR630 (Gen 1)	SR650 (Gen 1)	SR670 (Gen 1)	SR850 (Gen 1)	SR860 (Gen 1)	SR950 (Gen 1)	ST550 (Gen 1)
Red Hat Enterprise Linux 8.2	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.3	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.4	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.5	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 11 SP4	Y	Y	Y	Y	Y	Y	N	Y	Y	Y	Y
SUSE Linux Enterprise Server 11 SP4 with Xen	Y	Y	Y	Y	Y	Y	N	Y	Y	Y	Y
SUSE Linux Enterprise Server 12 SP2	Y	Y	N	N	Y	Y	N	Y	N	Y	Y
SUSE Linux Enterprise Server 12 SP2 with Xen	Y	Y	N	N	Y	Y	N	Y	N	Y	Y
SUSE Linux Enterprise Server 12 SP3	Y	Y	Y	Y	Y	Y	N	Y	Y	Y	Y
SUSE Linux Enterprise Server 12 SP3 with Xen	Y	Y	Y	Y	Y	Y	N	Y	Y	Y	Y
SUSE Linux Enterprise Server 12 SP4	Y	Y	Y	Y	Y	Y	N	Y	Y	Y	Y
SUSE Linux Enterprise Server 12 SP4 with Xen	Y	Y	Y	Y	Y	Y	N	Y	Y	Y	Y
SUSE Linux Enterprise Server 12 SP5	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 12 SP5 with Xen	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15	Y	Y	Y	Y	Y	Y	N	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP1	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP1 with Xen	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP2	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP2 with Xen	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP3	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP3 with Xen	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 with Xen	Y	Y	Y	Y	Y	Y	N	Y	Y	Y	Y
Ubuntu 18.04 LTS	N	N	N	N	N	N	N	N	N	N	N
Ubuntu 20.04 LTS	N	N	N	N	N	N	N	N	N	N	N
VMware vSphere Hypervisor (ESXi) 6.0 U3	Y	Y	Y	Y	Y	Y	N	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 6.5	Y	Y	N	N	Y	Y	N	Y	N	Y	Y
VMware vSphere Hypervisor (ESXi) 6.5 U1	Y	Y	Y	Y	Y	Y	N	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 6.5 U2	Y	Y	Y	Y	Y	Y	N	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 6.5 U3	Y	Y	Y	Y	Y	Y	N	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 6.7	Y	Y	Y	Y	Y	Y	N	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 6.7 U1	Y	Y	Y	Y	Y	Y	N	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 6.7 U2	Y	Y	Y	Y	Y	Y	N	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 6.7 U3	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 7.0	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 7.0 U1	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 7.0 U2	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 7.0 U3	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y

Table 8. Operating system support for ThinkSystem RAID 930-16i 4GB Flash PCIe 12Gb Adapter, 7Y37A01085 (Part 1 of 3)

Operating systems	SR250	ST250	SR630 V2	SR650 V2	SR850 V2	SR860 V2	ST650 V2	SR635	SR645	SR655	SR665
Microsoft Windows Server 2012 R2	N	N	N	N	N	N	N	N	N	N	N
Microsoft Windows Server 2016	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Microsoft Windows Server 2019	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Microsoft Windows Server 2022	N	N	Y	Y	Y	Y	Y	Y	Y	Y	Y
Microsoft Windows Server version 1709	N	N	N	N	N	N	N	N	N	N	N
Microsoft Windows Server version 1803	N	N	N	N	N	N	N	N	N	N	N
Red Hat Enterprise Linux 6 Server x64 Edition	N	N	N	N	N	N	N	N	N	N	N
Red Hat Enterprise Linux 6.10	N	N	N	N	N	N	N	N	N	N	N
Red Hat Enterprise Linux 6.9	N	N	N	N	N	N	N	N	N	N	N
Red Hat Enterprise Linux 7.3	N	N	N	N	N	N	N	N	N	N	N
Red Hat Enterprise Linux 7.4	N	N	N	N	N	N	N	N	N	N	N
Red Hat Enterprise Linux 7.5	Y	Y	N	N	N	N	N	N	N	N	N
Red Hat Enterprise Linux 7.6	N	N	N	N	N	N	N	Y	Y	Y	Y
Red Hat Enterprise Linux 7.7	Y	Y	N	N	N	N	N	Y	Y	Y	Y
Red Hat Enterprise Linux 7.8	Y	Y	N	N	N	N	N	Y	Y	Y	Y
Red Hat Enterprise Linux 7.9	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.0	Y	Y	N	N	N	N	N	Y	N	Y	N
Red Hat Enterprise Linux 8.1	Y	Y	N	N	N	N	N	Y	Y	Y	Y
Red Hat Enterprise Linux 8.2	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.3	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.4	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.5	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 11 SP4	N	N	N	N	N	N	N	N	N	N	N
SUSE Linux Enterprise Server 11 SP4 with Xen	N	N	N	N	N	N	N	N	N	N	N
SUSE Linux Enterprise Server 12 SP2	N	N	N	N	N	N	N	N	N	N	N
SUSE Linux Enterprise Server 12 SP2 with Xen	N	N	N	N	N	N	N	N	N	N	N
SUSE Linux Enterprise Server 12 SP3	Y	Y	N	N	N	N	N	N	N	N	N
SUSE Linux Enterprise Server 12 SP3 with Xen	Y	Y	N	N	N	N	N	N	N	N	N
SUSE Linux Enterprise Server 12 SP4	N	N	N	N	N	N	N	Y	N	Y	N
SUSE Linux Enterprise Server 12 SP4 with Xen	N	N	N	N	N	N	N	Y	N	Y	N
SUSE Linux Enterprise Server 12 SP5	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 12 SP5 with Xen	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15	Y	Y	N	N	N	N	N	N	N	N	N
SUSE Linux Enterprise Server 15 SP1	Y	Y	N	N	N	N	N	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP1 with Xen	Y	Y	N	N	N	N	N	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP2	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y

Operating systems	SR250	ST250	SR630 V2	SR650 V2	SR850 V2	SR860 V2	ST650 V2	SR635	SR645	SR655	SR665
SUSE Linux Enterprise Server 15 SP2 with Xen	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP3	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP3 with Xen	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 with Xen	Y	Y	N	N	N	N	N	N	N	N	N
Ubuntu 18.04 LTS	N	N	Y	Y	N	N	Y	N	N	N	N
Ubuntu 20.04 LTS	N	N	Y	Y	N	N	N	N	N	N	N
VMware vSphere Hypervisor (ESXi) 6.0 U3	N	N	N	N	N	N	N	N	N	N	N
VMware vSphere Hypervisor (ESXi) 6.5	N	N	N	N	N	N	N	N	N	N	N
VMware vSphere Hypervisor (ESXi) 6.5 U1	N	N	N	N	N	N	N	N	N	N	N
VMware vSphere Hypervisor (ESXi) 6.5 U2	Y	Y	N	N	N	N	N	N	N	N	N
VMware vSphere Hypervisor (ESXi) 6.5 U3	Y	Y	N	N	N	N	N	Y	N	Y	N
VMware vSphere Hypervisor (ESXi) 6.7	Y	Y	N	N	N	N	N	N	N	N	N
VMware vSphere Hypervisor (ESXi) 6.7 U1	N	N	N	N	N	N	N	N	N	N	N
VMware vSphere Hypervisor (ESXi) 6.7 U2	Y	Y	N	N	N	N	N	N	N	N	N
VMware vSphere Hypervisor (ESXi) 6.7 U3	Y	Y	Y	Y	N	N	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 7.0	Y	Y	N	N	N	N	N	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 7.0 U1	Y	Y	N	N	Y	Y	N	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 7.0 U2	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 7.0 U3	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y

Table 9. Operating system support for ThinkSystem RAID 930-16i 4GB Flash PCIe 12Gb Adapter, 7Y37A01085 (Part 2 of 3)

Operating systems	SR530 (Gen 2)	SR550 (Gen 2)	SR570 (Gen 2)	SR590 (Gen 2)	SR630 (Gen 2)	SR650 (Gen 2)	SR850 (Gen 2)	SR850P	SR860 (Gen 2)	SR950 (Gen 2)	ST550 (Gen 2)
Microsoft Windows Server 2012 R2	N	N	N	N	N	N	N	N	N	N	N
Microsoft Windows Server 2016	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Microsoft Windows Server 2019	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Microsoft Windows Server 2022	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Microsoft Windows Server version 1709	N	N	N	N	N	N	N	N	N	N	N
Microsoft Windows Server version 1803	N	N	N	N	N	N	N	N	N	N	N
Red Hat Enterprise Linux 6 Server x64 Edition	N	N	N	N	N	N	N	N	N	N	N
Red Hat Enterprise Linux 6.10	N	N	N	N	N	N	N	N	N	N	N
Red Hat Enterprise Linux 6.9	N	N	N	N	N	N	N	N	N	N	N
Red Hat Enterprise Linux 7.3	N	N	N	N	N	N	N	N	N	N	N
Red Hat Enterprise Linux 7.4	N	N	N	N	N	N	N	N	N	N	N
Red Hat Enterprise Linux 7.5	N	N	N	N	N	N	N	N	N	N	N

Operating systems	SR530 (Gen 2)	SR550 (Gen 2)	SR570 (Gen 2)	SR590 (Gen 2)	SR630 (Gen 2)	SR650 (Gen 2)	SR850 (Gen 2)	SR850P	SR860 (Gen 2)	SR950 (Gen 2)	ST550 (Gen 2)
Red Hat Enterprise Linux 7.6	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 7.7	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 7.8	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 7.9	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.0	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.1	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.2	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.3	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.4	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.5	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 11 SP4	N	N	N	N	N	N	N	N	N	N	N
SUSE Linux Enterprise Server 11 SP4 with Xen	N	N	N	N	N	N	N	N	N	N	N
SUSE Linux Enterprise Server 12 SP2	N	N	N	N	N	N	N	N	N	N	N
SUSE Linux Enterprise Server 12 SP2 with Xen	N	N	N	N	N	N	N	N	N	N	N
SUSE Linux Enterprise Server 12 SP3	N	N	N	N	N	N	N	Y	N	N	N
SUSE Linux Enterprise Server 12 SP3 with Xen	N	N	N	N	N	N	N	Y	N	N	N
SUSE Linux Enterprise Server 12 SP4	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 12 SP4 with Xen	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 12 SP5	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 12 SP5 with Xen	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP1	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP1 with Xen	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP2	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP2 with Xen	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP3	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP3 with Xen	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 with Xen	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Ubuntu 18.04 LTS	N	N	N	N	N	N	N	N	N	N	N
Ubuntu 20.04 LTS	N	N	N	N	N	N	N	N	N	N	N
VMware vSphere Hypervisor (ESXi) 6.0 U3	N	N	N	N	N	N	N	N	N	N	N
VMware vSphere Hypervisor (ESXi) 6.5	N	N	N	N	N	N	N	N	N	N	N
VMware vSphere Hypervisor (ESXi) 6.5 U1	N	N	N	N	N	N	N	N	N	N	N
VMware vSphere Hypervisor (ESXi) 6.5 U2	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 6.5 U3	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 6.7	N	N	N	N	N	N	N	N	N	N	N
VMware vSphere Hypervisor (ESXi) 6.7 U1	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y

Operating systems	SR530 (Gen 2)	SR550 (Gen 2)	SR570 (Gen 2)	SR590 (Gen 2)	SR630 (Gen 2)	SR650 (Gen 2)	SR850 (Gen 2)	SR850P	SR860 (Gen 2)	SR950 (Gen 2)	ST550 (Gen 2)
VMware vSphere Hypervisor (ESXi) 6.7 U2	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 6.7 U3	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 7.0	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 7.0 U1	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 7.0 U2	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 7.0 U3	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y

Table 10. Operating system support for ThinkSystem RAID 930-16i 4GB Flash PCIe 12Gb Adapter, 7Y37A01085 (Part 3 of 3)

Operating systems	SR530 (Gen 1)	SR550 (Gen 1)	SR570 (Gen 1)	SR590 (Gen 1)	SR630 (Gen 1)	SR650 (Gen 1)	SR850 (Gen 1)	SR860 (Gen 1)	SR950 (Gen 1)	ST550 (Gen 1)
Microsoft Windows Server 2012 R2	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Microsoft Windows Server 2016	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Microsoft Windows Server 2019	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Microsoft Windows Server 2022	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Microsoft Windows Server version 1709	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Microsoft Windows Server version 1803	N	N	N	N	Y	Y	Y	Y	Y	N
Red Hat Enterprise Linux 6 Server x64 Edition	N	Y	N	N	N	N	N	N	N	N
Red Hat Enterprise Linux 6.10	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 6.9	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 7.3	Y	Y	N	N	Y	Y	Y	N	Y	Y
Red Hat Enterprise Linux 7.4	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 7.5	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 7.6	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 7.7	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 7.8	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 7.9	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.0	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.1	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.2	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.3	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.4	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.5	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 11 SP4	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 11 SP4 with Xen	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y

Operating systems	SR530 (Gen 1)	SR550 (Gen 1)	SR570 (Gen 1)	SR590 (Gen 1)	SR630 (Gen 1)	SR650 (Gen 1)	SR850 (Gen 1)	SR860 (Gen 1)	SR950 (Gen 1)	ST550 (Gen 1)
SUSE Linux Enterprise Server 12 SP2	Y	Y	N	N	Y	Y	Y	N	Y	Y
SUSE Linux Enterprise Server 12 SP2 with Xen	Y	Y	N	N	Y	Y	Y	N	Y	Y
SUSE Linux Enterprise Server 12 SP3	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 12 SP3 with Xen	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 12 SP4	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 12 SP4 with Xen	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 12 SP5	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 12 SP5 with Xen	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP1	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP1 with Xen	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP2	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP2 with Xen	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP3	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP3 with Xen	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 with Xen	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Ubuntu 18.04 LTS	N	N	N	N	N	N	N	N	N	N
Ubuntu 20.04 LTS	N	N	N	N	N	N	N	N	N	N
VMware vSphere Hypervisor (ESXi) 6.0 U3	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 6.5	Y	Y	N	N	Y	Y	Y	N	Y	Y
VMware vSphere Hypervisor (ESXi) 6.5 U1	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 6.5 U2	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 6.5 U3	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 6.7	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 6.7 U1	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 6.7 U2	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 6.7 U3	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 7.0	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 7.0 U1	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 7.0 U2	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 7.0 U3	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y

Table 11. Operating system support for ThinkSystem RAID 930-16i 8GB Flash PCIe 12Gb Adapter, 4Y37A09721 (Part 1 of 2)

Operating systems	SR530 (Gen 2)	SR550 (Gen 2)	SR570 (Gen 2)	SR590 (Gen 2)	SR630 (Gen 2)	SR650 (Gen 2)	SR850 (Gen 2)	SR850P	SR860 (Gen 2)	SR950 (Gen 2)	ST550 (Gen 2)
Microsoft Windows Server 2012 R2	N	N	N	N	N	N	N	N	N	N	N
Microsoft Windows Server 2016	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Microsoft Windows Server 2019	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Microsoft Windows Server 2022	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Microsoft Windows Server version 1803	N	N	N	N	N	N	N	N	N	N	N
Red Hat Enterprise Linux 6.10	N	N	N	N	N	N	N	N	N	N	N
Red Hat Enterprise Linux 6.9	N	N	N	N	N	N	N	N	N	N	N
Red Hat Enterprise Linux 7.3	N	N	N	N	N	N	N	N	N	N	N
Red Hat Enterprise Linux 7.4	N	N	N	N	N	N	N	N	N	N	N
Red Hat Enterprise Linux 7.5	N	N	N	N	N	N	N	N	N	N	N
Red Hat Enterprise Linux 7.6	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 7.7	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 7.8	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 7.9	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.0	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.1	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.2	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.3	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.4	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.5	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 11 SP4	N	N	N	N	N	N	N	N	N	N	N
SUSE Linux Enterprise Server 12 SP2	N	N	N	N	N	N	N	N	N	N	N
SUSE Linux Enterprise Server 12 SP3	N	N	N	N	N	N	N	Y	N	N	N
SUSE Linux Enterprise Server 12 SP4	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 12 SP4 with Xen	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 12 SP5	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 12 SP5 with Xen	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP1	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP1 with Xen	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP2	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP2 with Xen	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP3	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP3 with Xen	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 with Xen	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y

Operating systems	SR530 (Gen 2)	SR550 (Gen 2)	SR570 (Gen 2)	SR590 (Gen 2)	SR630 (Gen 2)	SR650 (Gen 2)	SR850 (Gen 2)	SR850P	SR860 (Gen 2)	SR950 (Gen 2)	ST550 (Gen 2)
VMware vSphere Hypervisor (ESXi) 6.0 U3	N	N	N	N	N	N	N	N	N	N	N
VMware vSphere Hypervisor (ESXi) 6.5 U1	N	N	N	N	N	N	N	N	N	N	N
VMware vSphere Hypervisor (ESXi) 6.5 U2	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 6.5 U3	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 6.7	N	N	N	N	N	N	N	N	N	N	N
VMware vSphere Hypervisor (ESXi) 6.7 U1	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 6.7 U2	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 6.7 U3	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 7.0	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 7.0 U1	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 7.0 U2	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 7.0 U3	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y

Table 12. Operating system support for ThinkSystem RAID 930-16i 8GB Flash PCIe 12Gb Adapter, 4Y37A09721 (Part 2 of 2)

Operating systems	SR530 (Gen 1)	SR550 (Gen 1)	SR570 (Gen 1)	SR590 (Gen 1)	SR630 (Gen 1)	SR650 (Gen 1)	SR850 (Gen 1)	SR860 (Gen 1)	SR950 (Gen 1)	ST550 (Gen 1)
Microsoft Windows Server 2012 R2	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Microsoft Windows Server 2016	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Microsoft Windows Server 2019	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Microsoft Windows Server 2022	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Microsoft Windows Server version 1803	N	N	N	N	Y	Y	Y	Y	Y	N
Red Hat Enterprise Linux 6.10	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 6.9	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 7.3	Y	Y	N	N	Y	Y	Y	N	Y	Y
Red Hat Enterprise Linux 7.4	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 7.5	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 7.6	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 7.7	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 7.8	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 7.9	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.0	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.1	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.2	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.3	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y

Operating systems	SR530 (Gen 1)	SR550 (Gen 1)	SR570 (Gen 1)	SR590 (Gen 1)	SR630 (Gen 1)	SR650 (Gen 1)	SR850 (Gen 1)	SR860 (Gen 1)	SR950 (Gen 1)	ST550 (Gen 1)
Red Hat Enterprise Linux 8.4	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.5	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 11 SP4	Y	N	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 12 SP2	Y	N	N	N	Y	Y	Y	N	Y	Y
SUSE Linux Enterprise Server 12 SP3	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 12 SP4	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 12 SP4 with Xen	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 12 SP5	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 12 SP5 with Xen	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP1	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP1 with Xen	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP2	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP2 with Xen	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP3	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP3 with Xen	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 with Xen	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 6.0 U3	Y	N	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 6.5 U1	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 6.5 U2	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 6.5 U3	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 6.7	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 6.7 U1	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 6.7 U2	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 6.7 U3	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 7.0	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 7.0 U1	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 7.0 U2	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 7.0 U3	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y

Table 13. Operating system support for ThinkSystem RAID 930-24i 4GB Flash PCIe 12Gb Adapter, 7Y37A01086

Operating systems	SR655	SR650 (Gen 2)	ST550 (Gen 2)	SR650 (Gen 1)	ST550 (Gen 1)
Microsoft Windows Server 2012 R2	N	N	N	Y	Y

Operating systems	SR655	SR650 (Gen 2)	ST550 (Gen 2)	SR650 (Gen 1)	ST550 (Gen 1)
Microsoft Windows Server 2016	Y	Y	Y	Y	Y
Microsoft Windows Server 2019	Y	Y	Y	Y	Y
Microsoft Windows Server 2022	Y	Y	Y	Y	Y
Microsoft Windows Server version 1709	N	N	N	Y	Y
Microsoft Windows Server version 1803	N	N	N	Y	N
Red Hat Enterprise Linux 6.10	N	N	N	Y	Y
Red Hat Enterprise Linux 6.9	N	N	N	Y	Y
Red Hat Enterprise Linux 7.3	N	N	N	Y	Y
Red Hat Enterprise Linux 7.4	N	N	N	Y	Y
Red Hat Enterprise Linux 7.5	N	N	N	Y	Y
Red Hat Enterprise Linux 7.6	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 7.7	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 7.8	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 7.9	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.0	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.1	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.2	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.3	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.4	Y	Y	Y	Y	Y
Red Hat Enterprise Linux 8.5	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 11 SP4	N	N	N	Y	Y
SUSE Linux Enterprise Server 11 SP4 with Xen	N	N	N	Y	Y
SUSE Linux Enterprise Server 12 SP2	N	N	N	Y	Y
SUSE Linux Enterprise Server 12 SP2 with Xen	N	N	N	Y	Y
SUSE Linux Enterprise Server 12 SP3	N	N	N	Y	Y
SUSE Linux Enterprise Server 12 SP3 with Xen	N	N	N	Y	Y
SUSE Linux Enterprise Server 12 SP4	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 12 SP4 with Xen	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 12 SP5	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 12 SP5 with Xen	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15	N	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP1	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP1 with Xen	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP2	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP2 with Xen	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP3	Y	Y	Y	Y	Y
SUSE Linux Enterprise Server 15 SP3 with Xen	Y	Y	Y	Y	Y

Operating systems	SR655	SR650 (Gen 2)	ST550 (Gen 2)	SR650 (Gen 1)	ST550 (Gen 1)
SUSE Linux Enterprise Server 15 with Xen	N	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 6.0 U3	N	N	N	Y	Y
VMware vSphere Hypervisor (ESXi) 6.5	N	N	N	Y	Y
VMware vSphere Hypervisor (ESXi) 6.5 U1	N	N	N	Y	Y
VMware vSphere Hypervisor (ESXi) 6.5 U2	N	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 6.5 U3	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 6.7	N	N	N	Y	Y
VMware vSphere Hypervisor (ESXi) 6.7 U1	N	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 6.7 U2	N	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 6.7 U3	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 7.0	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 7.0 U1	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 7.0 U2	Y	Y	Y	Y	Y
VMware vSphere Hypervisor (ESXi) 7.0 U3	Y	Y	Y	Y	Y

Warranty

The adapters carry a 1-year limited warranty. When installed in a supported ThinkSystem server, the adapter assumes the server's base warranty and any warranty upgrades.

Operating environment

The ThinkSystem RAID 930 Series adapters are supported in the following environment:

- Operating:
  - Temperature: 10°C to 55°C (50°F to 131°F)
  - Relative humidity: 20% to 90% (non-condensing)
- Storage
  - Temperature with package: -40°C to 70°C (-40°F to 158°F)
  - Relative humidity: 5% to 95% (non-condensing)

Agency approvals

The ThinkSystem RAID 930 adapters have the following agency approvals:

- FCC Part 15 Class A
- Australia/New Zealand (AS/NZS CISPR 22)
- Canada (ICES-003 Class B)
- Europe (EN55022/EN55024)
- Japan (VCCI V-3)
- Korea (RRA no 2013-24 & 25)
- RoHS compliant
- EN/IEC/UL 60950
- Taiwan (CNS 13438)
- USA (FCC 47 CFR part 15 Subpart B class B)
- WEEE

Related publications and links

For more information, see the following documents:

- Lenovo ThinkSystem product publications:
  http://thinksystem.lenovofiles.com/help/index.jsp
- ServerProven hardware compatibility:
  http://www.lenovo.com/us/en/serverproven
- Lenovo RAID Management Tools and Resources:
  https://lenovopress.com/lp0579-lenovo-raid-management-tools-and-resources
- Lenovo RAID Introduction
  https://lenovopress.com/lp0578-lenovo-raid-introduction

Related product families

Product families related to this document are the following:

- RAID Adapters

Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area. Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service. Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

Lenovo (United States), Inc.
8001 Development Drive
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION ”AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary. Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk. Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

© Copyright Lenovo 2022. All rights reserved.

This document, LP0652, was created or updated on February 28, 2021.

Send us your comments in one of the following ways:

- Use the online Contact us review form found at:
  https://lenovopress.com/LP0652
- Send your comments in an e-mail to:
  comments@lenovopress.com

This document is available online at https://lenovopress.com/LP0652.

Trademarks

Lenovo and the Lenovo logo are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. A current list of Lenovo trademarks is available on the Web at https://www.lenovo.com/us/en/legal/copytrade/.

The following terms are trademarks of Lenovo in the United States, other countries, or both:
Lenovo®
ServerProven®
ThinkSystem
XClarity®

The following terms are trademarks of other companies:

Intel® is a trademark of Intel Corporation or its subsidiaries.

Linux® is the trademark of Linus Torvalds in the U.S. and other countries.

Microsoft®, Windows Server®, and Windows® are trademarks of Microsoft Corporation in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

IBM TS4300 Tape Library
Product Guide

IBM TS4300 Tape Library (Machine Type 6741) combines IBM enterprise tape automation and reliability with open system affordability. It is a high-density, highly scalable, easy-to-manage solution designed to keep data securely stored long-term, while helping reduce the costs associated with data center space and utilities.

The TS4300 utilizes modular design to meet growing long-term storage requirements. Starting with a single 3U base module, customers can scale the tape library by adding up to six 3U expansion modules for the overall capacity of up to 280 LTO cartridges. Each 3U base or expansion module can house up to 40 cartridges in two removable magazines.

Each 3U module can accommodate up to three half-height tape drives or a combination of one half-height tape drive and one full-height tape drive. The choice of the tape drives includes IBM LTO Ultrium 9, 8, 7, or 6 Tape Drives with SAS or Fibre Channel host connectivity interfaces.



Figure 1. IBM TS4300 Tape Library: Base module (bottom) with one expansion module (top)

Did you know?

Featuring the latest LTO Ultrium 9 technology, the TS4300 Tape Library provides up to 5.04 PB of native backup storage.

Lenovo can be your trusted partner that offers "one stop shop" and single point of contact for delivery of leading edge technologies and innovations from Lenovo and other leading IT vendors. These offerings can satisfy the wide range of your end-to-end IT infrastructure needs, including end-user devices, servers, storage, networking, services, management software, and financing.

Click here to check for updates

Key features

The TS4300 tape library offers the following features:

- Scalable tape library that is designed for high capacity, high performance, and high reliability.
- Adherence to LTO specifications.
- Support for half-high (HH) or full-high (FH) LTO Ultrium 9, 8, 7, or 6 tape drives
- Support for up to three half-high tape drives or a combination of one full-high tape drive and one half-high tape drive in each 3U module. The maximum number of drives in the tape library ranges from up to 7 full-high and up to 7 half-high to up to 21 half-high tape drives.
- Native data transfer rate of up to 400 MB/s (Ultrium 9 FH), 300 MBp/s (Ultrium 9 HH), 360 MB/s (Ultrium 8 FH), 300 MB/s (Ultrium 8 HH or Ultrium 7), or up to 160 MB/s (Ultrium 6) per tape drive.
- Native cartridge capacity of up to 18 TB (Ultrium 9), 12 TB (Ultrium 8), 9 TB (Ultrium 7 initialized LTO 8 Type M), 6 TB (Ultrium 7).
- 12Gb SAS, 6 Gb SAS or 8 Gb Fibre Channel attachments for connection to Lenovo servers that run Microsoft Windows or Linux server operating systems.
- Support for LTO 9 media specification tape cartridge compressed capacity of up to 45 TB, 30 TB compressed capacity for LTO8, and 15 TB compressed capacity for LTO7 with 2.5 to 1 compression per cartridge
- Data security and regulatory compliance through support for LME and WORM media on LTO 9, 8, and 7 tape drives
- Two removable magazines in each 3U module that can house up to 20 data cartridges each for a total of up to 280 cartridges in the tape library to enable quick population of the tape library and ease the storage needs for media.
- Flexibility in configuring a 5-slot I/O Station for each module, or one I/O Station can be accessed by several modules to help support continuous library operations, while importing and exporting media with little or no disruption to library operation.
- A standard high-speed bar code reader with which you can operate the TS4300 in random access mode. Sequential mode is not currently supported.
- Mixed media support, including IBM LTO Ultrium 9, 8, 7 .
- Remote management capabilities to allow for remote administration of the TS4300 through a web interface.
- Multi-path feature to allow sharing of the library robotics across heterogeneous applications independently of each other. The library can be partitioned into up to 21 logical libraries (each partition must have at least one drive). It also can provide each logical library its own separate and distinct drives, storage slots, and control paths.
- Optional Path Failover that enables the host device driver to transparently switch to an alternative path if there is a communication path failure to the application.
- Extra power supplies for redundant operations.
- Support from an extensive ecosystem of backup software vendors.

The features in the IBM LTO 9 tape drive:

- Servo and track layout technology: There are 6,656 data tracks in LTO 9 versus 3,584 data tracks in LTO 8. The high-bandwidth servo system features a low-mass servo to help more effectively track servo bands and improve data throughput with damaged media in less-than-optimal shock and vibration environments.
- Surface Control Guiding Mechanism: Designed to guide the tape along the tape path in the LTO 9, 8, 7, and 6 tape drive. This method is designed to use the surface of the tape, rather than the edges, to control tape motion to reduce tape damage (especially to the edges of the tape) and tape debris, which comes from the damaged edges and accumulation in the head area.

- Robust drive components optimized for automation environments: With the most robust components available, steel ball bearings in loader, robust leader block design, and single circuit card, the drives offer enhanced reliability and prolonged life.
- Separate writing of multiple filemarks: This is designed to cause any write command of two or more filemarks to cause a separate data set to be written containing all filemarks after the first.
- LTO Data Compression (LTO-DC): The LTO 9 uses LTO-DC, which is an implementation of a Lempel-Ziv class 1 (LZ-1) data compression algorithm. LTO-DC is an extension of Adaptive Lossless Data Compression (ALDC) and an improvement over previous IBM lossless compression algorithms. Patented "Scheme-Swapping" compression from IBM is designed to look ahead at incoming data and determine the most efficient storage method (either ALDC or pass-through mode) to help optimize data compression and increase data throughput. The compressed ratio is up to 2.5 to 1 for LTO 9.
- LTO Cartridge Memory (LTO-CM): Contained within the LTO data cartridge is the LTO-CM, which is a passive, contactless silicon storage device that is physically a part of the cartridge. The LTO-CM is designed to hold information about that specific cartridge, the media in the cartridge, and the data on the media. The storage capacity of the Generation 9 LTO-CM is 16,320 bytes. Communication between the drive and the LTO-CM is through a low-level RF field transmitted by the drive to the cartridge.
- Statistical Analysis and Reporting System (SARS): The LTO 9 tape drive uses SARS to help isolate failures between media and hardware. SARS uses the cartridge performance history saved in the Cartridge Memory module and the drive performance history kept in the drive flash to help determine the likely cause of failure. SARS is designed to cause the drive to request a cleaner tape, to mark the media as degraded, and to indicate that the hardware has degraded.
- Highly integrated electronics using IBM-engineered copper technology: This technology is designed to reduce the total number of components in the drive, help lower chip temperatures, and reduce power requirements, resulting in a more reliable drive. The ninth-generation drive electronics are designed to provide error correction of soft errors in the memory arrays in data and control paths.

Base Module and Expansion

TS4300 base model has the following characteristics:

- Base module can be installed and operate on its own, providing capacity for 40 LTO cartridges.
- Cartridge capacity: configured to hold two removable magazines, one on the left side and one on the right side, holding 20 slots each for a total of 40 data cartridges. This enables quick population of the tape library, as well as ease of storage for media
- I/O station: The full group of five slots will be configurable as either import/export elements or as storage elements. This enables continuous library operation during import or export of data cartridges.
- Media options: Uses LTO Ultrium® 9 Data Cartridge (5-pack)
- Robotics: The base library will include scaling robotics with a barcode reader for identifying, locating, and moving LTO cartridges between storage locations and drives.
- Power supplies: Power supplies are rated 80 PLUS Silver. Each module can run with a single power supply for nonredundant operation. Power supplies will be hot swappable when two are installed in a single module.

TS4300 expansion module has the following characteristics:

The capacity of a base library can be increased by adding one or more expansion modules. Each expansion will add 40 slots of LTO cartridge capacity and can include an additional three half-height LTO drives or an additional one full-height drive plus one half-height LTO drive.

- Cartridge capacity: expansion modules are configured to hold two removable magazines, one on the left side and one on the right side, holding 20 slots each for a total of 40 data cartridges. You may order additional magazines as optional features to ease import/export tape cartridges on a regular

basis.

- I/O station: The full group of five slots will be configurable as either import/export elements or as storage elements. This enables continuous library operation during import or export of data cartridges.
- Media options: Uses LTO Ultrium® 9 Data Cartridge (5-pack)
- Power supplies: Each module can run with a single power supply for nonredundant operation. An optional power supply will provide redundant operation. Power supplies will be hot swappable when two are installed in a single module.

Components and connectors

The following figure shows the front view of the TS4300 Tape Library base module

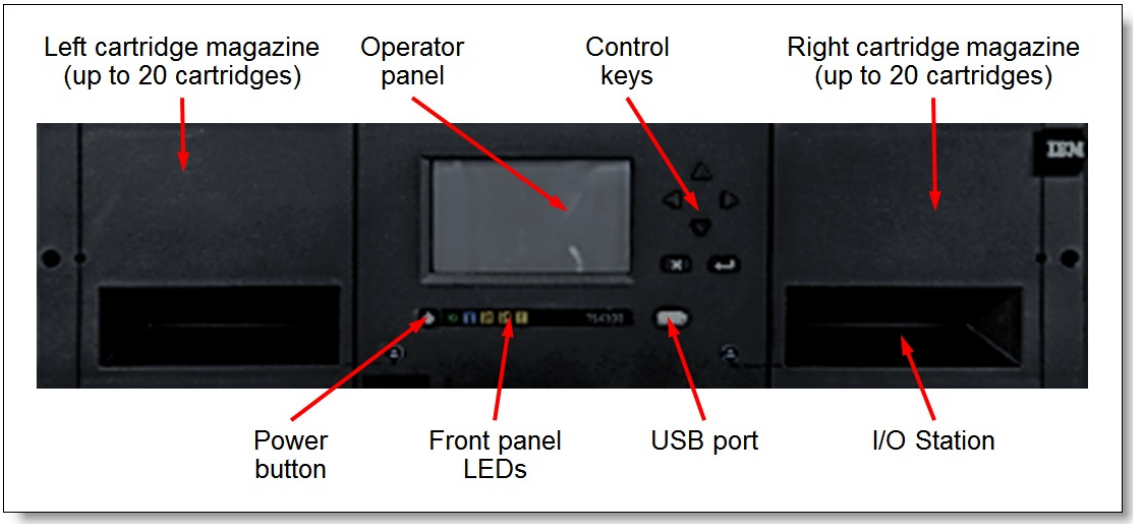


Figure 2. Front view of the TS4300 Tape Library base module

The following figure shows the front view of the TS4300 Tape Library expansion module



Figure 3. Front view of the TS4300 Tape Library expansion module

The following figure shows the rear view of the TS4300 Tape Library base module

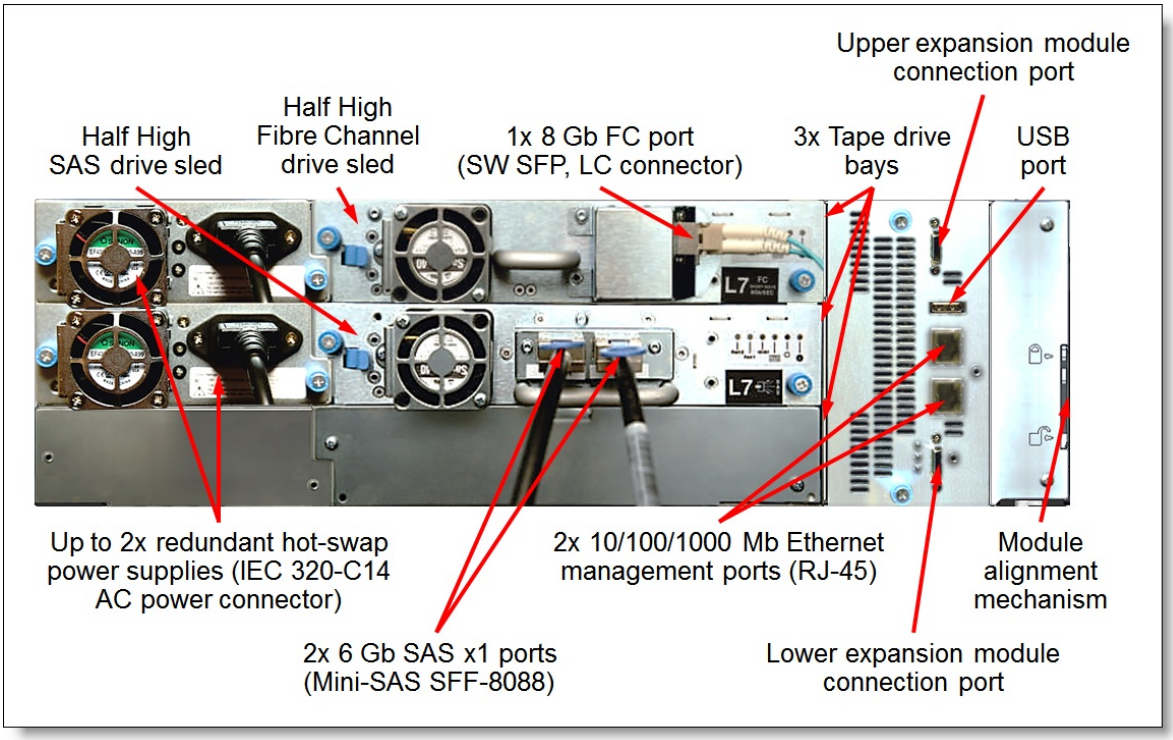


Figure 4. Rear view of the TS4300 Tape Library base module

The following figure shows the rear view of the TS4300 Tape Library expansion module.

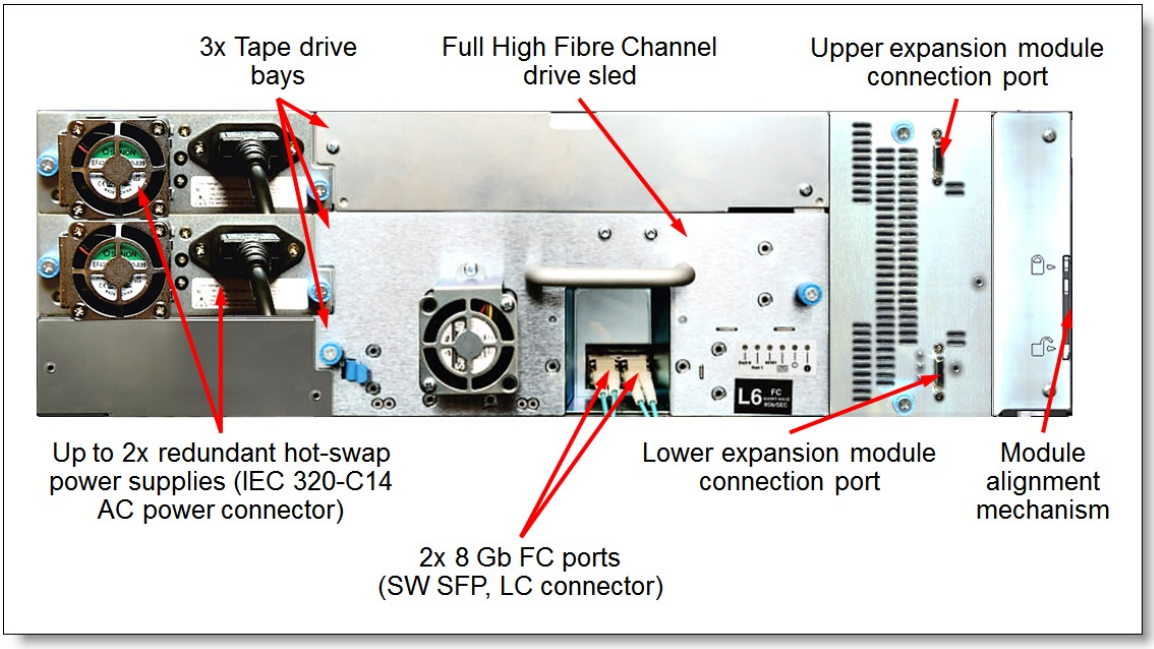


Figure 5. Rear view of the TS4300 Tape Library expansion module

System specifications

The following table lists the TS4300 Tape Library system specifications.

Table 1. System specifications

Attribute	Specification
Form factor	• Base module: 3U rack-mount or stand-alone • Expansion module: 3U rack-mount; up to 6 expansion modules • 7-module library: 21U rack-mount
Drive technology	LTO Ultrium 9, 8, 7, and 6 Full-High (FH) and Half-High (HH) tape drives
Drive interface	Ultrium LTO9 12 Gbps SAS Half and Full High drives Ultrium LTO8 6 Gbps SAS Half-High drive Ultrium LTO7 6 Gbps SAS Half-High drive Ultrium LTO9 8 Gbps Fibre Channel Half and Full High drives Ultrium LTO8 8 Gbps Fibre Channel Half and Full High drives Ultrium LTO7 8 Gbps Fibre Channel Half and Full High drives
Maximum number of drives	• Per module: Up to 1x full-high and 1x half-high tape drives, or up to 3x half-high tape drives • Per 7-module library: From up to 7x full-high and 7x half-high tape drives to up to 21x half-high tape drives
Number of cartridge slots	• Per module: 40 • Per 7-module library: 280
Configurable I/O station slots	• Per module: 5 • Per 7-module library: 35
Native cartridge capacity	• LTO 9 (L9): 18 TB • LTO 8 (L8): 12 TB • LTO 7 initialized LTO 8 Type M (M8): 9 TB • LTO 7 (L7): 6 TB

Attribute	Specification
Native backup storage capacity	Per module: • LTO 9 (L9): Up to 680 TB • LTO 8 (L8): Up to 480 TB • LTO 7 (M8): Up to 360 TB • LTO 7 (L7): Up to 240 TB Per 7-module library: • LTO 9 (L9) Up to 5.00 PB • LTO 8 (L8): Up to 3.36 PB. • LTO 7 (M8): Up to 2.52 PB. • LTO 7 (L7): Up to 1.68 PB.
Native data transfer rate	Per drive: • LTO 9 FH: Up to 400 MB/s • LTO 8 FH: Up to 360 MB/s • LTO 8 HH, LTO7: Up to 300 MB/s
Software features	Library partitioning, Path Failover (optional), Library Managed Encryption (optional)
Security features	Secure Socket Layer (SSL), user level security, LDAP, LTO encryption (optional)
Management interfaces (Base module)	• Operator panel • 2x 10/100/1000 Mb Ethernet ports (UTP, RJ-45) for remote management: Web-based GUI; SNMP and email notifications • 2x USB ports (service technician use only)
Cooling	Fixed fans on the tape drives and power supplies
Power supplies	Up to two redundant hot-swap 230 W AC 80 PLUS Silver power supplies (IEC 320-C14 power connector): • Base module: The first power supply is included; the second power supply is optional. • Expansion module: The first power supply is optional; the second power supply is optional. The first power supply is required if the expansion module contains a tape drive.
Hot-swap parts	SAS and Fibre Channel drive sleds, power supplies.
Dimensions	Base or expansion module: • Width: 446 mm (17.6 in.) • Depth: 873 mm (34.4 in.) • Height: 133 mm (5.2 in.)
Weight	• Base module: 21 kg (46.3 lb) • Expansion module: 13 kg (28.7 lb)
Warranty	3 year customer-replaceable unit (CRU) limited warranty with 9x5 next business day parts delivered
Service and support	Optional warranty service upgrades are available through Lenovo Services: onsite coverage, 24x7 coverage, 2-hour or 4-hour response time, 6-hour or 24-hour committed service repair (select countries), 1-year or 2-year warranty extensions
Host operating systems*	Microsoft Windows Server, Red Hat Enterprise Linux (RHEL), SUSE Linux Enterprise Server (SLES)

Models

The following table lists the ordering information for the TS4300 Tape Library models.

Table 2. TS4300 models

Description	Part number	Machine Type-Model	Feature code
IBM TS4300 3U Tape Library-Base Unit	6741A1F	6741L1U	B16H
IBM TS4300 3U Tape Library-Expansion Unit	6741A3F	6741L3U	B16J

The TS4300 Tape Library models include the following items:

- One TS4300 module with the following components:
  - Power supply (base module only)
  - Bar code reader (base module only)
  - Expansion cable (expansion module only)
- Two removable 20-cartridge magazines
- Documentation package

Notes:

- Tape drives are not included; at least one tape drive must be ordered with the TS4300 unit (See Tape drives for details).
- A cleaning cartridge is not included. See Tape media for ordering information.
- A power cable is not included. A rack power cable or a line cord must be ordered with the TS4300 unit (See Power supplies and cables for details).
- Rack mount kit is not included; it should be ordered with the TS4300 unit, if required (see the following table for details). Rack mount kit is required for each module if the configuration contains more than one module.

The following table lists the rack mount kit for the TS4300 units.

Table 3. Rack mount kit

Description	Part number	Feature code	Maximum quantity per module
Rack Mount Kit	01KP930	B16N	1

The following figure shows the connectivity topology.

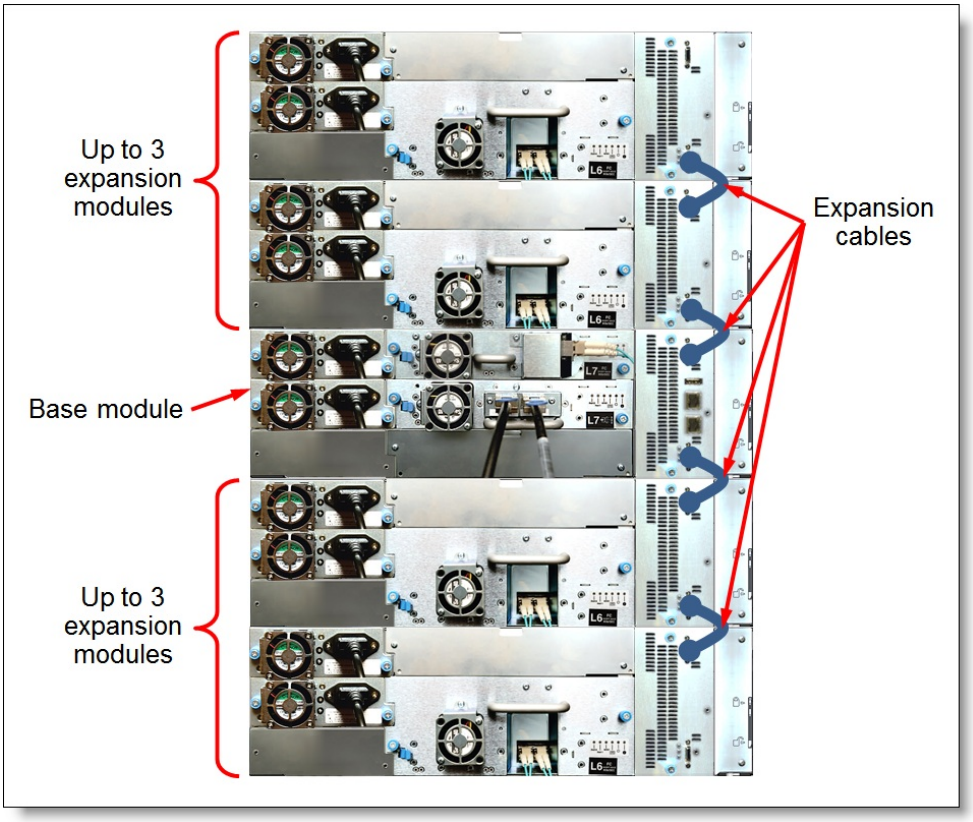


Figure 6. TS4300 tape library connectivity topology

Tape drives

Each TS4300 module (base or expansion) can accommodate up to three LTO Ultrium half-height tape drives or a combination of one LTO Ultrium half-height tape drive and one LTO Ultrium full-height tape drive.

The following table lists the supported tape drives for the TS4300.

Table 4. Tape drive options

Description	Part number	Feature code	Maximum quantity per module
Full Height 8 Gb Fibre Channel			
LTO 9 FH Fibre Channel Drive	02JH837	BLW0	1
LTO 8 FH Fibre Channel Drive	01KP954	B1ET	1
LTO 7 FH Fibre Channel Drive	01KP938	B16Y	1
Half Height 8 Gb Fibre Channel			
LTO 9 HH Fibre Channel Drive	02JH835	BLVZ	3
LTO 8 HH Fibre Channel Drive	01KP952	B1ER	3
LTO 7 HH Fibre Channel Drive	01KP936	B16W	3
Full Height 12 Gb SAS			
LTO 9 FH SAS Drive	02JH838	BLW6	3
Half Height 12 Gb SAS			
LTO 9 HH SAS Drive	02JH836	BLW7	3

Description	Part number	Feature code	Maximum quantity per module
Half High 6 Gb SAS			
LTO 8 HH SAS Drive	01KP953	B1ES	3
LTO 7 HH SAS Drive	01KP937	B16X	3
Tape drive support (Wrap tool)			
Fibre Tape Drive Support (Wrap Tool)	01KP939	B16R	7
SAS Tape Drive Support (Wrap Tool)	01KP940	B16S	7

Configuration note: Tape drive support options provide wrap tools (loopback connectors) to isolate drive sled issues. The Fibre Tape Drive Support is required when at least one Fibre Channel tape drive is installed in the tape library, and the SAS Tape Drive Support is required when at least one SAS tape drive is installed in the tape library.

A SAS cable or Fibre Channel cable is required to attach a tape drive in the TS4300 Tape Library to a host.

The following table lists supported cable options for LTO9 connectivity.

Table 5. SAS cables for LTO9

Description	Part number	Feature code	Maximum quantity
SAS connectivity cables: Mini-SAS HD (HBA SFF-8644 host) to Mini-SAS HD (HBA SFF-8644 tape drive)			
3 m Mini-SAS HD/Mini-SAS HD 1x Cable (from HBA with SFF-8644 to tape drive (LTO-9) with SFF-8644)	02JH830	BLW3	1
1.5 m Mini-SAS HD/Mini-SAS HD 1x Cable (from HBA with SFF-8644 to tape drive (LTO-9) with SFF-8644)	02JH831	BLW4	1
3 m Mini-SAS HD/Mini-SAS HD (Y) 2X Cable	02JH832	BLW5	1

The following table lists supported cable options for LTO8 and LTO7 connectivity.

Table 6. Cable options for LTO 8 & LTO 7

Description	Part number	Feature code	Maximum quantity per drive
SAS connectivity cables - Mini-SAS (SFF-8088 host) to Mini-SAS (SFF-8088 tape)			
Mini-SAS/Mini-SAS 4x Interposer (from HBA with SFF-8088 to max 4 cables with SFF-8088)	00NA077	AS0Y	1*
2 m Mini-SAS/Mini-SAS 1x Cable (host SFF-8088 to target SFF-8088)	00NA009	AS0Z	2
SAS connectivity cables - Mini-SAS HD (SFF-8644 host) to Mini-SAS (SFF-8088 tape)			
HD-SAS Cable to Mini-SAS (4 m)	00NV419	AS10	2
Y SAS HD to Mini-SAS Cable (3 m)	00NV420	AS12	2**
Fibre Channel connectivity cable8s			
10 m LC-LC Fibre Cable	01KP929	B16M	2

* The interposer is used on the host side to connect up to four SAS tape drives to a single Mini-SAS x4 port on a SAS HBA with up to four 2 m Mini-SAS to Mini-SAS cables (part number 00NA009).
** The Y SAS HD cable is used to connect up to two SAS tape drives to a single Mini-SAS HD x4 port on a SAS HBA.

Tape media

The TS4300 3U module supports up to 40 data cartridges and comes standard with two removable tape magazines with 20 cartridge slots each. More removable tape magazines can be ordered for the TS4300, if needed.

Note: The first release of the TS4300 tape library supported up to 32 cartridge slots, including up to four I/O station slots, in the single-module configurations and in the lowest module in the multi-module configurations. The second release of the TS4300 tape library supports up to 40 cartridge slots, including up to five I/O station slots, in the single-module configurations and in the lowest module in the multi-module configurations. Please contact a Lenovo service representative in your area for additional details.

The following table lists tape media options and extra magazines that can be ordered for the TS4300.

Table 7. Tape media and extra magazines

Description	Part number	Feature code	Supported tape drive types
Cleaning cartridge			
Ultrium Cleaning Cartridge (universal)	00NA017	AS1E	All drives
Data cartridges			
Ultrium 9 Data Cartridges 5-Pack	02JH833	BLW1	LTO 9
Ultrium 9 Data Cartridge (single)	02JH834	BLW2	LTO 9
Ultrium 8 Data Cartridges 5-Pack	01KP955	B1EW	LTO 9, LTO 8
LTO M8 Uninitialized Data Cartridges 5-Pack	02JH005	B5XL	LTO 8
Ultrium 7 Data Cartridges 5-Pack	00WF771	ATNZ	LTO 8, LTO 7
Optional magazines			
Right Side Magazine	01KP931	B16P	All drives
Left Side Magazine	01KP932	B16Q	All drives

LTO Ultrium tape drives support the following backward compatibility for previous generations of tape media:

- LTO 9 Tape Drives can read and write LTO 9 and LTO 8 data cartridges.
- LTO 8 tape drives can read and write LTO 8, LTO 7, and LTO 7 initialized LTO 8 Type M data cartridges.
- LTO 7 tape drives can read and write LTO 7 and LTO 6 data cartridges, and can read LTO 5 data cartridges.

Configuration notes:

- A 5-Pack of the Ultrium 6, 7, 8 or 9 Data Cartridges includes five unattached barcode labels.
- An LTO 7 initialized LTO 8 Type M data cartridge is a new (unused) LTO 7 data cartridge labelled and initialized to support 9 TB native storage capacity in the LTO 8 tape drive. Once initialized, an LTO 8 Type M cartridge cannot be used in the LTO 7 tape drives, and it cannot be changed back to a 6 TB LTO 7 cartridge.
- LTO M8 Uninitialized Data Cartridges 5-Pack (02JH005) includes five uninitialized LTO 7 data cartridges and five M8 barcode labels. These cartridges will be automatically initialized by the tape library firmware while under the control of ISV applications that recognize the M8 barcode label.

Software

The functionality of the TS4300 can be expanded with the following optional software features:

- Path Failover
- Library Managed Encryption

The TS4300 offers an optional path failover feature that enables the host device driver to resend the command to an alternative control path for the same logical library to ensure smooth library operations if there is a communication error. With control path failover, the alternative control path can include another HBA, SAN, or library control path drive. The device driver starts error recovery and continues the operation on the alternative control path without interrupting the application.

Ultrium 9, 8, and 7 tape drives are encryption-capable, and they support the following modes of encryption management:

- Library Managed: Available for Linux and Windows; requires the IBM Security Key Lifecycle Manager (SKLM).
- Application Managed: Available with third-party software such as IBM Spectrum Protect.

Library Managed Encryption requires an optional encryption license. Application Managed Encryption does not require a license.

The following table lists ordering information for optional software features for the TS4300.

Table 8. Optional software features for TS4300

Description	Part number	Feature code	Maximum quantity per library
6741 Path Failover	01KP984	B21P	1
6741 LTO Library Managed Encryption	01KP985	B21Q	1

Power supplies and cables

The TS4300 modules support up to two redundant hot-swap 230 W AC power supplies. The base module comes with one power supply, and the second power supply can be ordered for redundancy. The expansion module comes without power supplies, and the first power supply is required when at least one tape drive is installed in the expansion module.

The following table lists the power supply options for the TS4300.

Table 9. Power supply options for TS4300

Description	Part number	Feature code	Maximum quantity per module
First Power Supply	01KP928	B16L	1
Additional Power Supply	01KP927	B16K	1

The TS4300 modules and the power supply options ship without power cables. The power cables for the TS4300 are listed in the following table.

Table 10. Power cables

Description	Part number	Feature code	Maximum quantity per module
Line cords			
2.8m, 10A/120V, C13 to NEMA 5-15P (US) Line Cord	00NA027	AS29	2
2.8m, 10A/230V, C13 to CEE7-VII (Europe) Line Cord	00NA029	AS2A	2
2.8m, 10A/230V, C13 to DK2-5a (Denmark) Line Cord	00NA031	AS2B	2
2.8m, 10A/230V, C13 to BS 1363/A (UK) Line Cord	00NA033	AS2C	2
2.8m, 10A/230V, C13 to SI 32 (Israel) Line Cord	00NA035	AS2D	2
2.8m, 10A/230V, C13 to SEV 1011-S24507 (Sws) Line Cord	00NA037	AS2E	2
2.8m, 10A/230V, C13 to SABS 164 (South Africa) Line Cord	00NA039	AS2F	2
2.8m, 10A/230V, C13 to CEI 23-16 (Italy) Line Cord	00NA041	AS2G	2
2.8m, 10A/230V, C13 to AS/NZS 3112 (Aus/NZ) Line Cord	00NA043	AS2H	2
2.8m, 10A/208V, C13 to NEMA 6-15P (US) Line Cord	00NA045	AS2J	2
2.8m, 10A/220V, C13 to IRAM 2073 (Argentina) Line Cord	00NA047	AS2K	2
2.8m, 10A/220V, C13 to CNS 10917-3 (Taiwan) Line Cord	00NA049	AS2L	2
2.8m, 10A/220V, C13 to GB 2099.1 (China) Line Cord	00NA051	AS2M	2
2.8m, 10A/110V, C13 to CNS 10917-3 (Taiwan) Line Cord	00NA053	AS2N	2
2.8m, 12A/100V, C13 to JIS C-8303 (Japan) Line Cord	00NA055	AS2P	2
2.8m, 12A/220V, C13 to KSC 8305 (S. Korea) Line Cord	00NA057	AS2Q	2
2.8m, 10A/240V, C13 to IS 6538 (India) Line Cord	00NA059	AS2R	2
2.8m, 10A/220V, C13 to NBR 6147 (Brazil) Line Cord	00NA061	AS2T	2
Rack power cord			
4.3m, 10A/100-250V, C13 to IEC 320-C14 Rack Power Cable	00NA063	AS2U	2

Physical specifications

The 3U module of the TS4300 Tape Library has the following dimensions and weight (approximate):

- Width: 446 mm (17.6 in.)
- Depth: 873 mm (34.4 in.)
- Height: 133 mm (5.2 in.)
- Weight:
    - Base module: 21 kg (46.3 lb)
    - Expansion module: 13 kg (28.7 lb)

Operating environment

The TS4300 Tape Library is supported in the following environment:

- Air temperature:
  - Operating: 15 °C - 32 °C (59 °F - 90 °F)
  - Non-operating: 5 °C - 45 °C (41 °F - 113 °F)
  - Maximum altitude: 3,050 m (10,000 ft)
- Humidity:
  - Operating: 20% - 80%
  - Non-operating: 8% - 80%
- Electrical power:
  - 100 V AC - 240 V AC (5.0 A to 3.5 A)
  - Frequency: 50 Hz / 60 Hz
  - Power consumption: 350 watts

Warranty services and upgrades

The TS4300 Tape Library has a three-year Customer Replaceable Unit (CRU) warranty.

The standard warranty terms are customer-replaceable unit (CRU) and onsite (for field-replaceable units FRUs only) with standard call center support during normal business hours and 9x5 Next Business Day Parts Delivered.

Lenovo's additional support services provide a sophisticated, unified support structure for your data center, with an experience consistently ranked number one in customer satisfaction worldwide. Available offerings include:

- **Premier Support**

  Premier Support provides a Lenovo-owned customer experience and delivers direct access to technicians skilled in hardware, software, and advanced troubleshooting, in addition to the following:

  - Direct technician-to-technician access through a dedicated phone line
  - 24x7x365 remote support
  - Single point of contact service
  - End to end case management
  - Third-party collaborative software support
  - Online case tools and live chat support
  - On-demand remote system analysis
- **Warranty Upgrade (Preconfigured Support)**

  Services are available to meet the on-site response time targets that match the criticality of your systems.

  - 3, 4, or 5 years of service coverage
  - 1-year or 2-year post-warranty extensions
  - **Foundation Service**: 9x5 service coverage with next business day onsite response
  - **Essential Service:** 24x7 service coverage with 4-hour onsite response or 24-hour committed repair (available only in select markets)
  - **Advanced Service:** 24x7 service coverage with 2-hour onsite response or 6-hour committed repair (available only in select markets)

- **Managed Services**

  Lenovo Managed Services provides continuous 24x7 remote monitoring (plus 24x7 call center availability) and proactive management of your data center using state-of-the-art tools, systems, and practices by a team of highly skilled and experienced Lenovo services professionals.

  Quarterly reviews check error logs, verify firmware & OS device driver levels, and software as needed. We'll also maintain records of latest patches, critical updates, and firmware levels, to ensure you systems are providing business value through optimized performance.

- **Technical Account Management (TAM)**

  A Lenovo Technical Account Manager helps you optimize the operation of your data center based on a deep understanding of your business. You gain direct access to your Lenovo TAM, who serves as your single point of contact to expedite service requests, provide status updates, and furnish reports to track incidents over time. In addition, your TAM will help proactively make service recommendations and manage your service relationship with Lenovo to make certain your needs are met.

- **Enterprise Server Software Support**

  Enterprise Software Support is an additional support service providing customers with software support on Microsoft, Red Hat, SUSE, and VMware applications and systems. Around the clock availability for critical problems plus unlimited calls and incidents helps customers address challenges fast, without incremental costs. Support staff can answer troubleshooting and diagnostic questions, address product comparability and interoperability issues, isolate causes of problems, report defects to software vendors, and more.

- **Health Check**

  Having a trusted partner who can perform regular and detailed health checks is central to maintaining efficiency and ensuring that your systems and business are always running at their best. Health Check supports Lenovo-branded server, storage, and networking devices, as well as select Lenovo-supported products from other vendors that are sold by Lenovo or a Lenovo-Authorized Reseller.

Examples of region-specific warranty terms are second or longer business day parts delivery or parts-only base warranty.

If warranty terms and conditions include onsite labor for repair or replacement of parts, Lenovo will dispatch a service technician to the customer site to perform the replacement. Onsite labor under base warranty is limited to labor for replacement of parts that have been determined to be field-replaceable units (FRUs). Parts that are determined to be customer-replaceable units (CRUs) do not include onsite labor under base warranty.

If warranty terms include parts-only base warranty, Lenovo is responsible for delivering only replacement parts that are under base warranty (including FRUs) that will be sent to a requested location for self-service. Parts-only service does not include a service technician being dispatched onsite. Parts must be changed at customer's own cost and labor and defective parts must be returned following the instructions supplied with the spare parts.

Lenovo Service offerings are region-specific. Not all preconfigured support and upgrade options are available in every region. For information about Lenovo service upgrade offerings that are available in your region, refer to the following resources:

- Service part numbers in Lenovo Data Center Solution Configurator (DCSC):
  http://dcsc.lenovo.com/#/services
- Lenovo Services Availability Locator
  http://lenovolocator.com/

For service definitions, region-specific details, and service limitations, please refer to the following documents:

- Lenovo Statement of Limited Warranty for Infrastructure Solutions Group (ISG) Servers and System Storage
  http://pcsupport.lenovo.com/us/en/solutions/ht503310
- Lenovo Data Center Services Agreement
  http://support.lenovo.com/us/en/solutions/ht116628

Regulatory compliance

The TS4300 Tape Library conforms to the following regulations:

- FCC: Verified to comply with Part 15 of the FCC Rules, Class A
- Canada ICES-003/NMB-03, Class A
- CISPR 32, Class A
- EN 55032, Class A
- Japan VCCI Class A statement
- People's Republic of China Class A Electronic Emission statement
- Taiwan Class A compliance statement
- Korea KCC Class A statement
- Russia Electromagnetic Interference (EMI) Class A Statement

Interoperability

Lenovo provides end-to-end storage compatibility testing to deliver interoperability throughout the network. The TS4300 Tape Library supports attachment to Lenovo servers by using SAS or Fibre Channel connectivity.

Note: End-to-end storage configuration support (including adapters, network switches, and operating systems) *must* be verified through the System Storage Interoperation Center (SSIC) for the TS4300

Fibre Channel SAN switches

Lenovo offers the ThinkSystem DB Series of Fibre Channel SAN switches for high-performance storage expansion. See the DB Series product guides for models and configuration options:

- ThinkSystem DB Series SAN Switches:
  https://lenovopress.com/storage/switches/rack#rt=product-guide

Operating systems

For the list of supported operating systems, refer to the IBM System Storage Interoperation Center (SSIC): https://www-03.ibm.com/systems/support/storage/ssic/interoperability.wss

Note: The LTO tape drives are supported when attached and run from the host operating system, however, in virtualization environments (for example, Hyper-V, KVM, Xen), the tape drives are not supported in the guest operating systems.

ISV backup software compatibility

The Independent Software Vendors (ISVs) indicate support for the LTO Ultrium Tape Drives for the following ISV backup software:

- Acronis Backup Advanced
- Arcserve Backup
- ASG Time Navigator
- CommVault Simpana
- Dell/EMC NetWorker
- Echoleaf Systems Virtual Drive
- Gemalto SafeNet
- HPE/Micro Focus Data Protector
- IBM High Performance Storage System
- IBM RMS
- IBM Security Key Lifecycle Manager
- IBM Spectrum Archive Drive Edition
- IBM Spectrum Archive Library Edition
- IBM Spectrum Archive Enterprise Edition
- IBM Spectrum Protect
- Microsoft System Center Data Protection Manager
- NetApp Data Storage
- Oracle DIVA
- Oracle OSB
- Oracle SAM-QFS-HSM
- Point Archival Gateway
- Quest NetVault
- Qstar
- Veritas Backup Exec
- Veritas NetBackup
- Veeam Backup and Retention
- Versity Storage Manager

For more information, see the ISV matrix for IBM LTO that is available at this website: https://www.ibm.com/downloads/cas/ORLZ7EMG

Lenovo Financial Services

Lenovo Financial Services reinforces Lenovo's commitment to deliver pioneering products and services that are recognized for their quality, excellence, and trustworthiness. Lenovo Financial Services offers financing solutions and services that complement your technology solution anywhere in the world.

We are dedicated to delivering a positive finance experience for customers like you who want to maximize your purchase power by obtaining the technology you need today, protect against technology obsolescence, and preserve your capital for other uses.

We work with businesses, non-profit organizations, governments and educational institutions to finance their entire technology solution. We focus on making it easy to do business with us. Our highly experienced team of finance professionals operates in a work culture that emphasizes the importance of providing outstanding customer service. Our systems, processes and flexible policies support our goal of providing customers with a positive experience.

We finance your entire solution. Unlike others, we allow you to bundle everything you need from hardware and software to service contracts, installation costs, training fees, and sales tax. If you decide weeks or months later to add to your solution, we can consolidate everything into a single invoice.

Our Premier Client services provide large accounts with special handling services to ensure these complex transactions are serviced properly. As a premier client, you have a dedicated finance specialist who manages your account through its life, from first invoice through asset return or purchase. This specialist develops an in-depth understanding of your invoice and payment requirements. For you, this dedication provides a high-quality, easy, and positive financing experience.

For your region-specific offers, please ask your Lenovo sales representative or your technology provider about the use of Lenovo Financial Services. For more information, see the following Lenovo website:

https://www.lenovo.com/us/en/landingpage/lenovo-financial-services/

Related publications and links

For more information, see the following resources:

- Storage Options for ThinkSystem Servers
  https://lenovopress.com/lp0761-storage-options-for-thinksystem-servers
- Lenovo Data Center Solution Configurator (DCSC)
  http://dcsc.lenovo.com/
- IBM System Storage Interoperation Center (SSIC)
  http://ibm.com/systems/support/storage/ssic
- IBM Redbooks publication, IBM Tape Library Guide for Open Systems
  http://www.redbooks.ibm.com/abstracts/sg245946.html?Open

Related product families

Product families related to this document are the following:

- Tape Autoloaders and Libraries
- Backup Units

Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area. Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service. Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

Lenovo (United States), Inc.
8001 Development Drive
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION ”AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary. Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk. Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

© Copyright Lenovo 2022. All rights reserved.

This document, LP0773, was created or updated on March 18, 2022.

Send us your comments in one of the following ways:

- Use the online Contact us review form found at:
  https://lenovopress.com/LP0773
- Send your comments in an e-mail to:
  comments@lenovopress.com

This document is available online at https://lenovopress.com/LP0773.

Trademarks

Lenovo and the Lenovo logo are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. A current list of Lenovo trademarks is available on the Web at https://www.lenovo.com/us/en/legal/copytrade/.

The following terms are trademarks of Lenovo in the United States, other countries, or both:
Lenovo®
Lenovo Services
ThinkSystem

The following terms are trademarks of other companies:

Linux® is the trademark of Linus Torvalds in the U.S. and other countries.

Hyper-V®, Microsoft®, Windows Server®, and Windows® are trademarks of Microsoft Corporation in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Lenovo ThinkSystem SR630 Server (Xeon SP Gen 1 / Gen 2)

Product Guide

Lenovo ThinkSystem SR630 is an ideal 2-socket 1U rack server for small businesses up to large enterprises that need industry-leading reliability, management, and security, as well as maximizing performance and flexibility for future growth. The SR630 server is designed to handle a wide range of workloads, such as databases, virtualization and cloud computing, virtual desktop infrastructure (VDI), infrastructure security, systems management, enterprise applications, collaboration/email, streaming media, web, and HPC.

Featuring the second generation of the Intel Xeon Processor Scalable Family (Xeon SP Gen 2), the SR630 server offers scalable performance and storage capacity. The SR630 server supports up to two processors, up to 3 TB of memory capacity with TruDDR4 DIMMs or up to 7.5 TB of memory capacity with a combination of TruDDR4 DIMMs and Intel DC persistent memory modules (DCPMMs), up to 12x 2.5-inch or 4x 3.5-inch drive bays with an extensive choice of NVMe PCIe SSDs, SAS/SATA SSDs, and SAS/SATA HDDs, and flexible I/O expansion options with a LOM slot, a dedicated storage controller slot, and up to 3x PCIe slots.

The following figure shows the Lenovo ThinkSystem SR630 with 2.5-inch hot-swap drives.



Figure 1 Lenovo ThinkSystem SR630 with 2.5-inch drive bays

Did you know?

The SR630 server features a unique AnyBay design that allows a choice of drive interface types in the same drive bay: SAS drives, SATA drives, or U.2 NVMe PCIe drives.

The SR630 server offers onboard NVMe PCIe ports that allow direct connections to the U.2 NVMe PCIe SSDs, which frees up I/O slots and helps lower NVMe solution acquisition costs.

The SR630 server delivers outstanding memory performance with Performance+ 2933 MHz DIMMs, which is achieved by supporting two-DIMMs-per-channel configurations at speeds up to 10% faster than the Intel specification defines, while still maintaining world-class reliability.

Click here to check for updates

Key features

Combining performance and flexibility, the SR630 server is a great choice for small and medium businesses up to the large enterprise. It can provide outstanding uptime to keep business-critical applications and cloud deployments running safely. Ease of use and comprehensive systems management tools help make deployment easier. Outstanding reliability, availability, and serviceability (RAS) and high-efficiency design improve your business environment and help save operational costs.

Scalability and performance

The SR630 server offers numerous features to boost performance, improve scalability, and reduce costs:

- Improves productivity by offering superior system performance with the second generation of the Intel Xeon Processor Scalable Family with up to 28-core processors, up to 38.5 MB of last level cache (LLC), up to 2933 MHz memory speeds, and up to 10.4 GT/s Ultra Path Interconnect (UPI) links.
    - Support for up to two processors, 56 cores, and 112 threads allows to maximize the concurrent execution of multithreaded applications.
    - Intelligent and adaptive system performance with energy efficient Intel Turbo Boost 2.0 Technology allows CPU cores to run at maximum speeds during peak workloads by temporarily going beyond processor thermal design power (TDP).
    - Intel Hyper-Threading Technology boosts performance for multithreaded applications by enabling simultaneous multithreading within each processor core, up to two threads per core.
    - Intel Virtualization Technology integrates hardware-level virtualization hooks that allow operating system vendors to better utilize the hardware for virtualization workloads.
    - Intel Speed Select Technology provides improvements in server utilization and guaranteed per-core performance service levels with more granular control over processor performance.
    - Intel Deep Learning Boost (Vector Neural Network Instruction set [VNNI]) is designed to deliver significant, more efficient Deep Learning (Inference) acceleration for high-performance Artificial Intelligence (AI) workloads.
    - Intel Advanced Vector Extensions 512 (AVX-512) enable acceleration of enterprise-class and high performance computing (HPC) workloads.
- Helps maximize system performance for data intensive applications with up to 2933 MHz memory speeds and up to 3 TB of memory capacity with 3DS RDIMMs.
- Boosts the performance of data-intensive applications and delivers consistent service levels at scale for virtualized and cloud environments by using the innovative persistent memory technology that provides a unique combination of affordable large memory capacity and non-volatility for up to 7.5 TB of total server memory capacity, including 3DS RDIMMs and DCPMMs (DC persistent memory modules).
- Offers flexible and scalable internal storage in a 1U rack form factor with up to 12x 2.5-inch drives for performance-optimized configurations or up to 4x 3.5-inch drives for capacity-optimized configurations, providing a wide selection of SAS/SATA HDD/SSD and PCIe NVMe SSD types and capacities.
- Provides flexibility to use SAS, SATA, or NVMe PCIe drives in the same drive bays with a unique AnyBay design.
- Provides I/O scalability with the LOM slot, PCIe 3.0 slot for an internal storage controller, and up to three PCI Express (PCIe) 3.0 I/O expansion slots in a 1U rack form factor.
- Reduces I/O latency and increases overall system performance with Intel Integrated I/O Technology that embeds the PCI Express 3.0 controller into the Intel Xeon Processor Scalable Family.

Availability and serviceability

The SR630 server provides many features to simplify serviceability and increase system uptime:

- Offers protection in the event of a non-correctable memory failure with Single Device Data Correction (SDDC, also known as Chipkill, requires x4-based DIMMs), Adaptive Double Device Data Correction (ADDDC, also known as Redundant Bit Steering [RBS], requires x4-based DIMMs and Intel Xeon Gold or Platinum processors), memory mirroring, and memory rank sparing.
- Provides easy access to upgrades and serviceable parts (such as processors, memory DIMMs, and adapter cards) with tool-less cover removal.
- Offers data protection and greater system uptime with hot-swap drives supporting basic or advanced RAID redundancy.
- Provides availability for business-critical applications with redundant hot-swap power supplies and redundant hot-swap fans.
- Simplifies servicing, speeds up problem resolution, and helps improve system availability with light path diagnostics.
- Allows preventive actions in advance of possible failure, thereby increasing server uptime and application availability with Proactive Platform Alerts (including PFA and SMART alerts) for processors, voltage regulators, memory, internal storage (SAS/SATA HDDs and SSDs, NVMe SSDs, M.2 storage, flash storage adapters), fans, power supplies, RAID controllers, and server ambient and sub-component temperatures.
- Continuously monitors system parameters, triggers alerts, and performs recovery actions in case of failure to minimize downtime with Built-in XClarity Controller (XCC).
- Provides quick access to system status, firmware, network, health, and alerts information via Virtual Operator Panel from the XClarity Mobile App running on the Android or iOS mobile device that is connected to the front USB port with XClarity Controller access.
- Speeds up troubleshooting tasks to reduce service time with diagnostics built into the XClarity Provisioning Manager.

Manageability and security

Powerful systems management features simplify local and remote management of the SR630 server and deliver enterprise-class data protection:

- Provides advanced service processor control, monitoring, and alerting functions with XClarity Controller, a next generation service processor.
- Improves Unified Extensible Firmware Interface (UEFI) system setup, configuration, updates, simplified error handling, and operating system deployment with the embedded XClarity Provisioning Manager.
- Offers XClarity Essentials software tools that can help you set up, use, and maintain the server.
- Increases uptime, reduces costs, and improves productivity through advanced server management capabilities with Lenovo XClarity Administrator that provides comprehensive hardware management.
- Provides on-the-go monitoring and management of devices in XClarity Administrator from anywhere with the Lenovo XClarity mobile app, which can help improve efficiency and reduce downtime risks.
- Centralizes infrastructure resource management with Lenovo XClarity Integrators for VMware vCenter and Microsoft System Center, extending XClarity Administrator features to virtualization management software tools and enabling users to deploy and manage infrastructure end-to-end.
- Offers advanced cryptographic functionality (such as digital signatures and remote attestation) with an integrated Trusted Platform Module (TPM) or optional Trusted Cryptographic Module (TCM) or Nationz TPM (available only in PRC).
- Keeps user data safe with Lenovo Business Vantage, a security software tool suite designed to work with the Trusted Cryptographic Module (available only in PRC).
- Offers enterprise-class data protection with advanced RAID and optional self-encrypting drives.
- Provides faster, stronger encryption with industry-standard AES NI support.
- Helps prevent certain classes of malicious buffer overflow attacks with Intel Execute Disable Bit

functionality, when combined with a supporting operating system.

- Enhances security through hardware-based resistance to malicious software attacks with Intel Trusted Execution Technology, allowing an application to run in its own isolated space, protected from all other software running on a system.

Energy efficiency

The SR630 server offers the following energy-efficiency features to save energy, reduce operational costs, increase energy availability, and contribute to the green environment:

- Delivers impressive compute power per watt, featuring 80 PLUS Titanium and Platinum redundant power supplies.
- Enables customers to lower energy costs with design to meet ASHRAE A4 standards in select configurations.
- Reduces power drawn with Intel Intelligent Power Capability that powers individual processor elements on and off as needed.
- Helps reduce power consumption with variable speed fans.
- Helps achieve lower heat output and reduced cooling needs with Lenovo XClarity Energy Manager that provides advanced data center power notification, analysis, and policy-based management.

Components and connectors

The following figure shows the front of the SR630 server with four 3.5-inch drive bays.

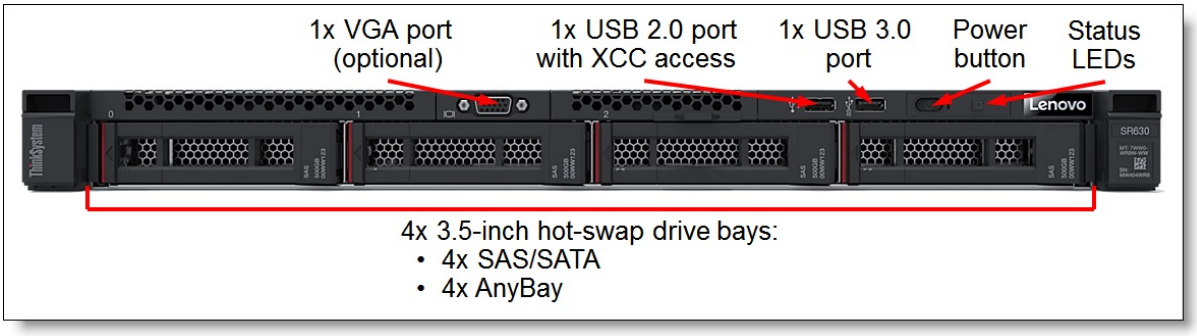


Figure 2. Front view of the SR630: 4x 3.5-inch drive bays

The following figure shows the front of the SR630 server with eight 2.5-inch drive bays.

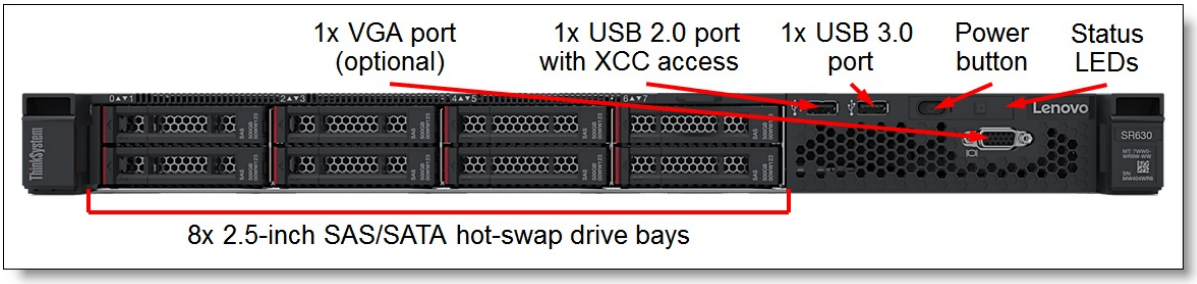


Figure 3. Front view of the SR630: 8x 2.5-inch drive bays

The following figure shows the front of the SR630 server with ten 2.5-inch drive bays.

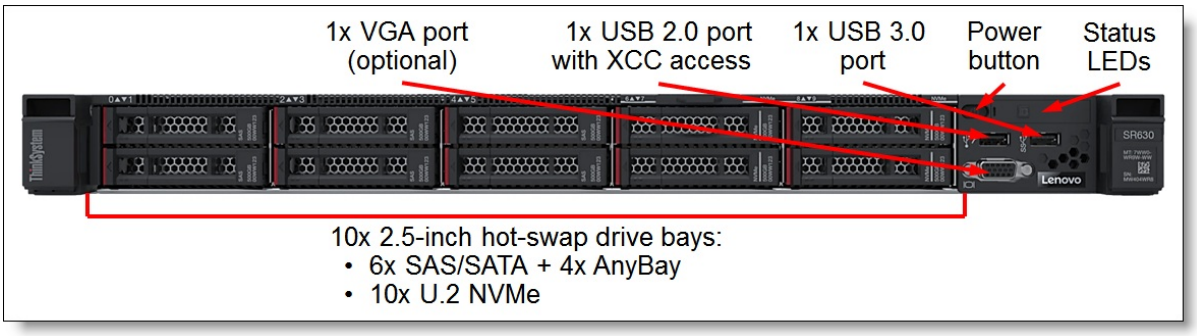


Figure 4. Front view of the SR630: 10x 2.5-inch drive bays

The front of the SR630 server includes the following components:

- 4x 3.5-inch, or 8x 2.5-inch, or 10x 2.5-inch hot-swap drive bays.
- One VGA port (optional).
- One USB 3.0 port.
- One USB 2.0 port with XClarity Controller access.
- Power button.
- Status LEDs.

The following figure shows the rear of the SR630 server with three PCIe low profile slots.

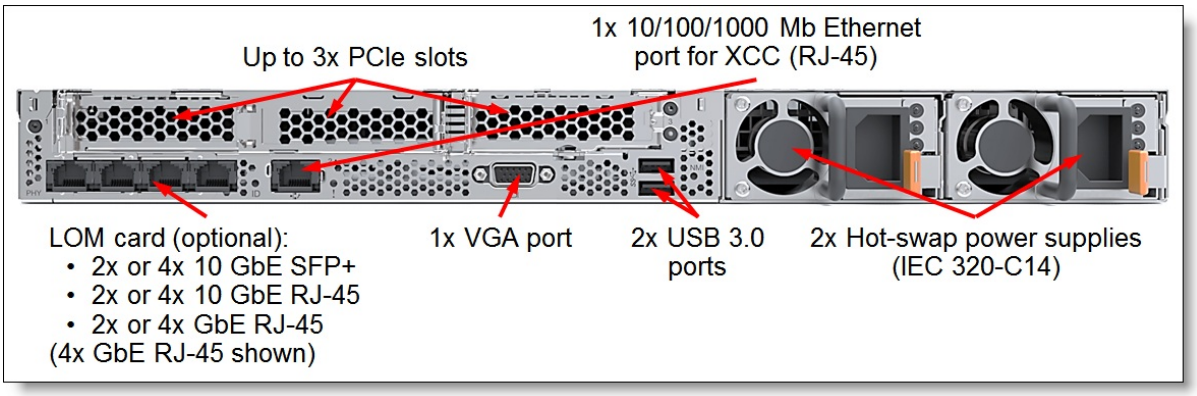


Figure 5. Rear view of the SR630

The rear of the SR630 server includes the following components:

- Up to three PCIe expansion slots (depending on the riser cards selected).
- One LOM card slot.
- One 1 GbE port for XClarity Controller.
- One VGA port.
- Two USB 3.0 ports.
- Up to two hot-swap power supplies.

The following figure shows the locations of key components inside the SR630 server.

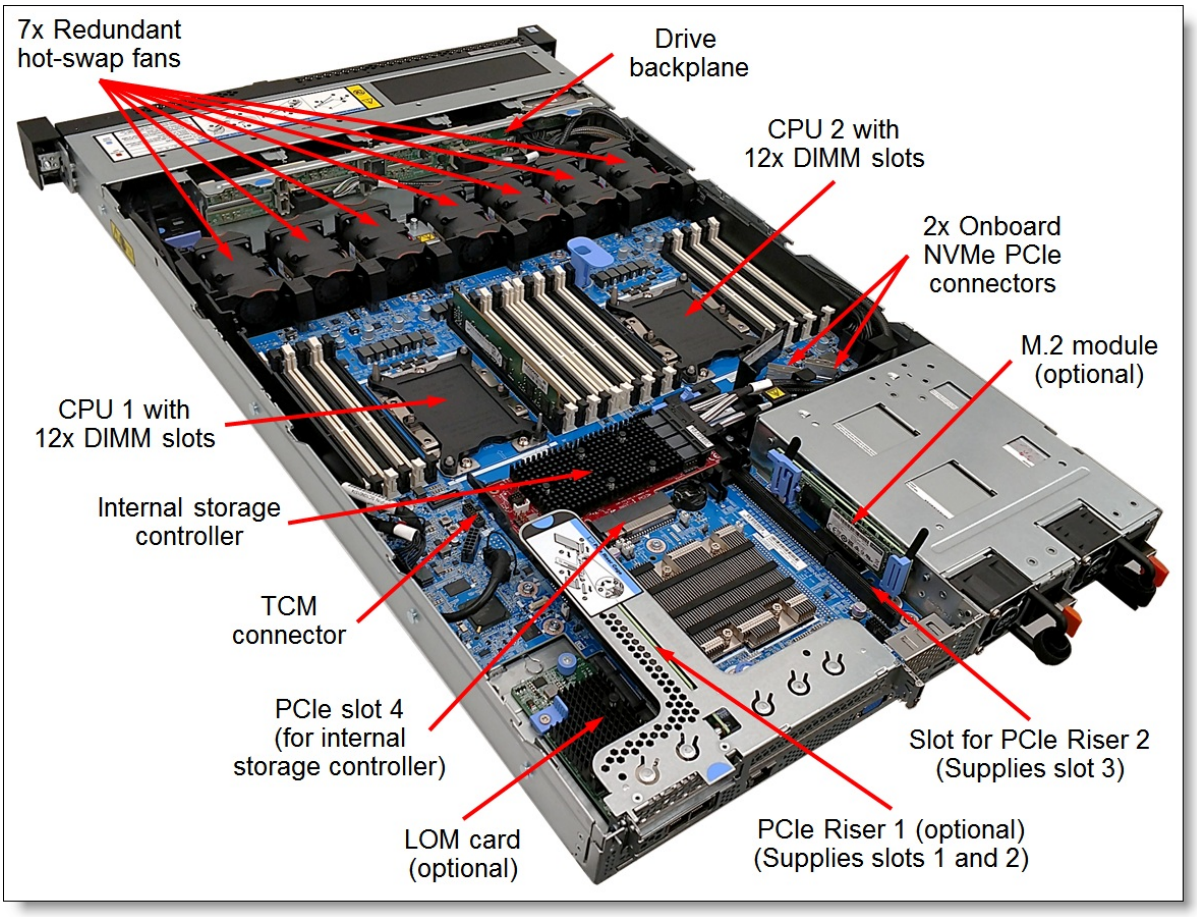


Figure 6. Internal view of the SR630

The following key components are located inside the SR630 server:

- Up to two processors.
- 24 DIMM slots (12 DIMM slots per processor).
- Drive backplanes.
- Two onboard NVMe PCIe connectors.
- One M.2 module connector.
- One LOM card connector.
- One onboard PCIe slot 4.
- Two slots for PCIe riser cards.
- One TCM connector.
- Five (one processor) or seven (two processors) hot-swap system fans.

System specifications

The following table lists the system specifications for the SR630 server.

Table 1. SR630 system specifications

Attribute	Specification
Machine types	7X01 - 1 year warranty 7X02 - 3 year warranty
Form factor	1U rack-mount

Attribute	Specification
Processor	Up to two 2nd Gen Intel Xeon Bronze, Silver, Gold, or Platinum processors: • Up to 28 cores (2.7 GHz core speeds). • Up to 3.8 GHz core speeds (4 cores). • UPI links up to 10.4 GT/s (2 UPI links used). • Up to 38.5 MB cache. • Up to 2933 MHz memory speed. 1st Gen Intel Xeon processors are also supported.
Chipset	Intel C624
Memory	Up to 24 DIMM sockets (12 DIMMs per processor; six memory channels per processor with two DIMMs per channel). Support for RDIMMs, LRDIMMs (1st Gen processors only), or 3DS RDIMMs. Memory speed up to 2933 MHz depending on the processor selected. Memory types cannot be intermixed.
Persistent memory	Up to 12x TruDDR4 2666 MHz DCPMMs in the DIMM slots. Not supported with 1st Gen Intel Xeon SP processors.

Memory capacity	• Memory DIMMs only: Up to 3 TB with up to 24x 128 GB 3DS RDIMMs (Up to 1.5 TB per processor).
	• Memory DIMMs and persistent memory modules:
	  ◦ App Direct Mode: Up to 7.5 TB with up to 12x 128 GB 3DS RDIMMs and up to 12x 512 GB DCPMMs (Up to 3.75 TB per processor).
	  ◦ Memory Mode: Up to 6 TB with up to 12x 512 GB DCPMMs (Up to 3 TB per processor).
	Note: Server configurations with more than 1 TB of memory capacity per socket (including DCPMMs and RDIMMs or 3DS RDIMMs) require processors that support up to 4.5 TB (L-suffix) per socket.

Memory protection	• Processor's integrated memory controllers: Error correction code (ECC), SDDC (for x4-based memory DIMMs), ADDDC (for x4-based memory DIMMs, requires Intel Xeon Gold or Platinum processors), memory mirroring, memory rank sparing, patrol scrubbing, and demand scrubbing. • DCPMM's onboard memory controllers: ECC, SDDC, DDDC, patrol scrubbing, and demand scrubbing. **Note:** In the configurations with DCPMMs, memory mirroring is supported only in the App Direct mode (other DCPMM modes do not support memory mirroring) and applies only to the RDIMMs or 3DS RDIMMs (DCPMMs are not mirrored). Memory sparing is not supported in the configurations with DCPMMs.
Drive bays	• Up to 6 SAS/SATA hot-swap drive bays: 4x 3.5" (front) + 2x 2.5" (rear) • 4 LFF AnyBay hot-swap drive bays: 4x 3.5" (front) • Up to 10 SFF SAS/SATA hot-swap drive bays: 8x 2.5" (front) + 2x 2.5" (rear) • Up to 12 SFF hot-swap drive bays: 6x 2.5" SAS/SATA & 4x 2.5" AnyBay (front) + 2x 2.5" SAS/SATA (rear) • 10 SFF U.2 NVMe PCIe SSD hot-swap drive bays

Attribute	Specification

Internal storage capacity	• 2.5-inch drives:
	  ◦ 184.32TB using 12x 15.36TB 2.5-inch SAS/SATA SSDs
	  ◦ 153.6TB using 10x 15.36TB 2.5-inch NVMe SSDs
	  ◦ 28.8TB using 12x 2.4TB 2.5-inch HDDs
	• 3.5-inch drives:
	  ◦ 72TB using 4x 18TB 3.5-inch HDDs
	  ◦ 30.72TB using 4x 7.68TB 3.5-inch SAS/SATA SSDs
	  ◦ 30.72TB using 4x 7.68TB 3.5-inch NVMe SSDs

Storage controller	• 12 Gb SAS/SATA RAID adapters with up to 8GB flash-backed cache • 12 Gb SAS/SATA HBA (non-RAID) • Onboard PCIe NVMe (with Intel VROC NVMe RAID support for Intel SSDs and optionally non-Intel SSDs) • NVMe Switch Adapters (with Intel VROC NVMe RAID support for Intel SSDs and optionally non-Intel SSDs)
Optical drive bays	None. Support for an external USB DVD RW Optical Disk Drive (SeeOptical drives).

Network interfaces	• Onboard LOM slot for up to 4x 1/10 Gb Ethernet ports:
	  ◦ 2x or 4x 1 GbE RJ-45 ports (no 10/100 Mb support)
	  ◦ 2x or 4x 10 GbE RJ-45 ports (no 10/100 Mb support)
	  ◦ 2x or 4x 10 GbE SFP+ ports (no 10/100 Mb support)
	• Optional Mezzanine LOM (ML2) slot for dual-port 10 GbE cards with SFP+ or RJ-45 connectors or single- or dual-port 25 GbE cards with SFP28 connectors.
	• 1x RJ-45 10/100/1000 Mb Ethernet systems management port.

I/O expansion slots	Up to four slots. Slot 4 is the fixed slot on the system planar, and the remaining slots depend on the riser cards installed. The slots are as follows: • Slot 1: PCIe 3.0 x8, ML2 x8, or ML2 x16; low profile • Slot 2: PCIe 3.0 x16 or x8; low profile or full-height, half-length • Slot 3: PCIe 3.0 x16; low profile • Slot 4: PCIe 3.0 x8 (dedicated to an internal storage controller) Slot 3 requires the second processor to be installed.
Ports	• Front: 1x USB 2.0 port with XClarity Controller access and 1x USB 3.0 port; optional 1x VGA port. • Rear: 2x USB 3.0 ports and 1x VGA port. Optional 1x DB-9 serial port.
Cooling	Five (one processor) or seven (two processors) hot-swap dual-rotor system fans with N+1 redundancy.
Power supply	Up to two redundant hot-swap 550 W, 750 W, or 1100 W (100 - 240 V) High Efficiency Platinum, or 750 W (200 - 240 V) High Efficiency Titanium AC power supplies. HVDC support (PRC only).
Video	Matrox G200e with 16 MB memory integrated into the XClarity Controller. Maximum resolution is 1920x1200 at 60 Hz with 32 bits per pixel.
Hot-swap parts	Drives, power supplies, and fans.
Systems management	XClarity Controller (XCC) Standard, Advanced, or Enterprise (Pilot 4 chip), proactive platform alerts, light path diagnostics, XClarity Provisioning Manager, XClarity Essentials, XClarity Administrator, XClarity Integrators for VMware vCenter and Microsoft System Center, XClarity Energy Manager, Capacity Planner.
Security features	Power-on password, administrator's password, secure firmware updates, Trusted Platform Module (TPM) 1.2 or 2.0 (configurable UEFI setting). Optional lockable front bezel. Optional Trusted Cryptographic Module (TCM) or Nationz TPM (available only in PRC). Optional Lenovo Business Vantage security software (available only in PRC).

Attribute	Specification
Operating systems	Microsoft Windows Server, Red Hat Enterprise Linux, SUSE Linux Enterprise Server, VMware ESXi. See the Operating systems section for specifics.
Warranty	One-year (Machine Type 7X01) or three-year (Machine Type 7X02) customer-replaceable unit (CRU) and onsite limited warranty with 9x5 Next Business Day Parts Delivered.
Service and support	Optional service upgrades are available through Lenovo Services: 2-hour or 4-hour response time, 6-hour or 24-hour committed service repair (select areas), warranty extension up to 5 years, 1-year or 2-year post-warranty extensions, YourDrive Your Data, Enterprise Software Support, and Basic Hardware Installation Services.
Dimensions	Width: 435 mm (17.1 in.), height: 43 mm (1.7 in.), depth: 750 mm (29.5 in.). See Physical specifications for details.
Weight	Minimum configuration: 11.9 kg (26.2 lb), maximum: 18.8 kg (41.4 lb)

Models

ThinkSystem SR630 models can be configured by using the Lenovo Data Center Solution Configurator (DCSC).

Configure-to-order (CTO) models are used to create models with factory-integrated server customizations. For CTO models, two base CTO models are available for the SR630 as listed in the following table, CTO1WW and CTOLWW:

- The CTO1WW base CTO model is for general business and is selectable by choosing **General Purpose** mode in DCSC.
- The CTOLWW base model is intended for High Performance Computing (HPC) and Artificial Intelligence (AI) configurations and solutions, including configurations for Lenovo Scalable Infrastructure (LeSI), and is enabled using either the **HPC & AI LeSI Solutions** mode or **HPC & AI ThinkSystem Hardware** mode in DCSC. CTOLWW configurations can also be built using System x and Cluster Solutions Configurator (x-config).

Preconfigured server models may also be available for the SR630, however these are region-specific; that is, each region may define their own server models, and not all server models are available in every region.

The following table lists the base CTO models of the ThinkSystem SR630 server.

Table 2. Base CTO models

Description	Machine Type/Model General purpose	Machine Type/Model for HPC and AI
ThinkSystem SR630 - 3 year Warranty	7X02CTO1WW	7X02CTOLWW
ThinkSystem SR630 - 1 year Warranty	7X01CTO1WW	7X01CTOLWW

For CTO orders, the following table lists the base chassis feature codes for the server.

There are currently two base feature codes for each of the 2.5-inch and 3.5-inch chassis. The "v2" bases include the new SR570 Air Duct Kit v2 which is required if a RAID 9350 adapter is to be configured. See the RAID flash power module (supercap) support section for details. The non-v2 bases can be selected if any other RAID adapter or HBA is selected.

Table 3. Base chassis for CTO models

Feature code	Description
Base feature codes - suitable for all configurations except those with a RAID 9350 adapter	
AUW0	ThinkSystem SR630 2.5" Chassis with 8 Bays
AUW1	ThinkSystem SR630 2.5" Chassis with 10 Bays
AUW2	ThinkSystem SR630 3.5" Chassis with 4 Bays
Base feature codes - suitable for all configurations including ones with a RAID 9350 adapter (includes the SR630 Air Duct Kit v2)	
BNPP	ThinkSystem SR630 2.5" Chassis with 8 Bays v2
BNPN	ThinkSystem SR630 2.5" Chassis with 10 Bays v2
BNPQ	ThinkSystem SR630 3.5" Chassis with 4 Bays v2

The following tables list the available models, grouped by region.

- Models for Australia and New Zealand
- Models for South East Asian countries (ASEAN)
- Models for Brazil
- Models for EMEA region
- Models for Hong Kong, Taiwan, Korea (HTK)
- Models for India
- Models for Japan
- Models for Latin American countries (except Brazil)
- Models for USA and Canada

Refer to the Specifications section for information about standard features of the server.

Common to all models:

- All models indicated as having the 750W power supply are using the Platinum power supply

Models for Australia and New Zealand

Table 4. Models for Australia and New Zealand

Model	Intel Xeon Scalable processor†	Memory	RAID	Front drive bays and drives	LOM	Rear Slots	Power supply	Front VGA	XCC	Rail kit
TopSeller models with a 3-year warranty (machine type 7X02)										
7X02A0BSAU	1x Bronze 3204 6C 85W 1.9G	1x 16GB 2Rx8 2933	530-8i	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide CMA
7X02A0BTAU	1x Silver 4208 8C 85W 2.1G	1x 16GB 2Rx8 2933	530-8i	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide CMA
7X02A0BXAU	1x Silver 4208 8C 85W 2.1G	1x 16GB 2Rx8 2933	930-8i	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide CMA
7X02A0HHAU	1x Silver 4208 8C 85W 2.1G	1x 16GB 2Rx8 2933	Option	Option 2.5"/8, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide CMA
7X02A0BPAU	1x Silver 4210 10C 85W 2.2G	1x 16GB 2Rx8 2933	930-8i	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide CMA
7X02A0BQAU	1x Silver 4210 10C 85W 2.2G	1x 32GB 2933	530-8i	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide CMA
7X02A0BRAU	1x Silver 4210 10C 85W 2.2G	1x 32GB 2933	930-8i	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide CMA
7X02A0BYAU	1x Silver 4210 10C 85W 2.2G	1x 16GB 2Rx8 2933	530-8i	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide CMA
7X02A0HJAU	1x Silver 4210 10C 85W 2.2G	1x 32GB 2933	Option	Option 2.5"/8, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide CMA
7X02A0F2AU	1x Silver 4214 12C 85W 2.2G	1x 16GB 2Rx8 2666	730-8i 2GB	8x 2.5" SAS, Open bay	Open	Open	1x 750W	Opt	Ent	Slide CMA
7X02A0F3AU	1x Silver 4214 12C 85W 2.2G	1x 16GB 2Rx8 2666	930-8i	8x 2.5" SAS, Open bay	Open	Open	1x 750W	Opt	Ent	Slide CMA
7X02A0BNAU	1x Silver 4215 8C 85W 2.5G	1x 16GB 2Rx8 2933	930-8i	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide CMA
7X02A0C0AU	1x Silver 4215 8C 85W 2.5G	1x 32GB 2933	930-8i	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide CMA
7X02A0FYAU	1x Silver 4215R 8C 130W 3.2G	1x 32GB 2933	930-8i	8x 2.5" SAS, Open bay	Open	Open	1x 750W	Yes	Ent	Slide
7X02A0BVAU	1x Silver 4216 16C 100W 2.1G	1x 32GB 2933	930-8i	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide CMA
7X02A0BZAU	1x Silver 4216 16C 100W 2.1G	1x 16GB 2Rx8 2933	930-8i	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide CMA
7X02A0FWAU	1x Silver 4216 16C 100W 2.1G	1x 32GB 2933	930-8i	8x 2.5" SAS, Open bay	Open	Open	1x 750W	Yes	Ent	Slide
7X02A0BWAU	1x Gold 5217 8C 115W 3.0G	1x 32GB 2933	930-8i	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide CMA
7X02A0BUAU	1x Gold 5218 16C 125W 2.3G	1x 32GB 2933	930-8i	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide CMA
7X02A0FVAU	1x Gold 6226R 16C 150W 2.9G	1x 64GB 2933	Option	Option 2.5"/8, Open bay	Open	Open	1x 750W	Yes	Ent	Slide
7X02A0FXAU	1x Gold 6226R 16C 150W 2.9G	1x 32GB 2933	930-8i	8x 2.5" SAS, Open bay	Open	Open	1x 750W	Yes	Ent	Slide
7X02A0FZAU	1x Gold 6226R 16C 150W 2.9G	1x 32GB 2933	Option	Option 2.5"/8, Open bay	Open	Open	1x 750W	Yes	Ent	Slide
7X02A0HKAU	1x Gold 6226R 16C 150W 2.9G	1x 32GB 2933	Option	Option 2.5"/8, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide CMA

† Processor description: Processor model, number of cores, thermal design power (TDP), core frequency

Models for South East Asian countries (ASEAN)

Table 5. Models for South East Asian countries (ASEAN)

Model	Intel Xeon Scalable processor†	Memory	RAID	Front drive bays and drives	LOM	Rear Slots	Power supply	Front VGA	XCC	Rail kit
TopSeller models with a 3-year warranty (machine type 7X02)										
7X02A09DSG	1x Bronze 3204 6C 85W 1.9G	1x 16GB 2Rx8 2933	930-16i 4GB	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A0ACSG	1x Bronze 3204 6C 85W 1.9G	1x 16GB 2Rx8 2933	930-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A08CSG	1x Silver 4208 8C 85W 2.1G	1x 16GB 2Rx8 2933	930-16i 4GB	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A08KSG	1x Silver 4208 8C 85W 2.1G	1x 16GB 2Rx8 2933	930-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A087SG	1x Silver 4210 10C 85W 2.2G	1x 16GB 2Rx8 2933	930-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A095SG	1x Silver 4210 10C 85W 2.2G	1x 16GB 2Rx8 2933	930-16i 4GB	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A08PSG	1x Silver 4214 12C 85W 2.2G	1x 16GB 2Rx8 2933	930-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A09WSG	1x Silver 4214 12C 85W 2.2G	1x 16GB 2Rx8 2933	930-16i 4GB	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A08ESG	1x Silver 4215 8C 85W 2.5G	1x 16GB 2Rx8 2933	930-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A0B2SG	1x Silver 4215 8C 85W 2.5G	1x 16GB 2Rx8 2933	930-16i 4GB	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A0A5SG	1x Silver 4216 16C 100W 2.1G	1x 16GB 2Rx8 2933	930-16i 4GB	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A0AZSG	1x Silver 4216 16C 100W 2.1G	1x 16GB 2Rx8 2933	930-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A08USG	1x Gold 5215 10C 85W 2.5G	1x 16GB 2Rx8 2933	930-16i 4GB	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A09VSG	1x Gold 5215 10C 85W 2.5G	1x 16GB 2Rx8 2933	930-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A07ZSG	1x Gold 5217 8C 115W 3.0G	1x 16GB 2Rx8 2933	930-16i 4GB	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A08TSG	1x Gold 5217 8C 115W 3.0G	1x 16GB 2Rx8 2933	930-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A099SG	1x Gold 5218 16C 125W 2.3G	1x 16GB 2Rx8 2933	930-16i 4GB	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A09XSG	1x Gold 5218 16C 125W 2.3G	1x 16GB 2Rx8 2933	930-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A08QSG	1x Gold 5220 18C 125W 2.2G	1x 16GB 2Rx8 2933	930-16i 4GB	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A0AUSG	1x Gold 5220 18C 125W 2.2G	1x 16GB 2Rx8 2933	930-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A08FSG	1x Gold 6230 20C 125W 2.1G	1x 16GB 2Rx8 2933	930-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A08HSG	1x Gold 6230 20C 125W 2.1G	1x 16GB 2Rx8 2933	930-16i 4GB	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A0AYSG	1x Gold 6240 18C 150W 2.6G	1x 16GB 2Rx8 2933	930-16i 4GB	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt

Model	Intel Xeon Scalable processor†	Memory	RAID	Front drive bays and drives	LOM	Rear Slots	Power supply	Front VGA	XCC	Rail kit
7X02A0BLSG	1x Gold 6240 18C 150W 2.6G	1x 16GB 2Rx8 2933	930-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A0ATSG	1x Gold 6242 16C 150W 2.8G	1x 16GB 2Rx8 2933	930-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A0BJSG	1x Gold 6242 16C 150W 2.8G	1x 16GB 2Rx8 2933	930-16i 4GB	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A091SG	1x Gold 6248 20C 150W 2.5G	1x 16GB 2Rx8 2933	930-16i 4GB	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A0BMSG	1x Gold 6248 20C 150W 2.5G	1x 16GB 2Rx8 2933	930-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A082SG	1x Gold 6252 24C 150W 2.1G	1x 16GB 2Rx8 2933	930-16i 4GB	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A09TSG	1x Gold 6252 24C 150W 2.1G	1x 16GB 2Rx8 2933	930-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt

† Processor description: Processor model, number of cores, thermal design power (TDP), core frequency

Models for Brazil

Table 6. Models for Brazil

Model	Intel Xeon Scalable processor†	Memory	RAID	Front drive bays and drives	LOM	Rear Slots	Power supply	Front VGA	XCC	Rail kit
Standard models with a 3-year warranty (machine type 7X02)										
7X02A0GHBR	1x Gold 5218R 20C 125W 2.1G	1x 32GB 2933	930-16i 4GB	8x 2.5" SAS, Open bay	2x10Gb RJ45	x8 LP, x16 FH	2x 750W	Yes	Ent	Slide
7X02A0GKBR	1x Gold 5218R 20C 125W 2.1G	1x 32GB 2933	930-16i 4GB	8x 2.5" SAS, Open bay	2x10Gb RJ45	x8 LP, x16 FH	2x 750W	Yes	Ent	Slide
7X02A0GGBR	1x Gold 6240R 24C 165W 2.4G	2x 32GB 2933	930-16i 4GB	8x 2.5" SAS, Open bay	4x10Gb RJ45	x8 LP, x16 FH	2x 750W	Yes	Ent	Slide
7X02A0GJBR	1x Gold 6240R 24C 165W 2.4G	2x 32GB 2933	930-16i 4GB	8x 2.5" SAS, Open bay	4x10Gb RJ45	x8 LP, x16 FH	2x 750W	Yes	Ent	Slide
TopSeller models with a 3-year warranty (machine type 7X02)										
7X02A0GSBR	1x Bronze 3204 6C 85W 1.9G	1x 32GB 2933	730-8i 2GB	4x 3.5" SAS, 2x 4TB SATA HDD	4x1Gb	x8 LP, x16 LP	2x 750W	Yes	Std	Slide
7X02A0F5BR	1x Silver 4208 8C 85W 2.1G	1x 32GB 2933	730-8i 2GB	8x 2.5" SAS, Open bay	4x1Gb	x8 LP, x16 FH	2x 750W	Opt	Ent	Slide
7X02A0F9BR	1x Silver 4208 8C 85W 2.1G	1x 32GB 2933	730-8i 2GB	8x 2.5" SAS, Open bay	4x1Gb	x8 LP, x16 FH	2x 750W	Opt	Std	Slide
7X02A0G3BR	1x Silver 4208 8C 85W 2.1G	1x 32GB 2933	730-8i 2GB	4x 3.5" SAS, Open bay	4x1Gb	x8 LP, x16 LP	2x 750W	Yes	Std	Slide
7X02A0G7BR	1x Silver 4208 8C 85W 2.1G	1x 32GB 2933	730-8i 2GB	4x 3.5" SAS, 1x 2TB SATA HDD	4x1Gb	x8 LP, x16 LP	2x 750W	Yes	Std	Slide
7X02A0GTBR	1x Silver 4208 8C 85W 2.1G	1x 32GB 2933	730-8i 2GB	4x 3.5" SAS, 2x 4TB SATA HDD	4x1Gb	x8 LP, x16 LP	2x 750W	Yes	Std	Slide
7X02A0F8BR	1x Silver 4210 10C 85W 2.2G	1x 32GB 2933	730-8i 2GB	8x 2.5" SAS, Open bay	4x1Gb	x8 LP, x16 FH	2x 750W	Opt	Std	Slide
7X02A0FSBR	1x Silver 4210R 10C 100W 2.4G	1x 32GB 2933	730-8i 2GB	8x 2.5" SAS, Open bay	4x1Gb	x8 LP, x16 FH	2x 750W	Opt	Std	Slide
7X02A0GUBR	1x Silver 4210R 10C 100W 2.4G	1x 32GB 2933	730-8i 2GB	4x 3.5" SAS, 2x 4TB SATA HDD	4x1Gb	x8 LP, x16 LP	2x 750W	Yes	Std	Slide

† Processor description: Processor model, number of cores, thermal design power (TDP), core frequency

Models for EMEA region

Table 7. Models for EMEA region

Model	Intel Xeon Scalable processor†	Memory	RAID	Front drive bays and drives	LOM	Rear Slots	Power supply	Front VGA	XCC	Rail kit
SAP HANA Business One models with a 3-year warranty (machine type 7X02)										
7X02A0HLEA	2x Silver 4210R 10C 100W 2.4G	12x 16GB 1Rx4 2933	940-8i 4GB	8x 2.5" SAS, 3x 960GB S4620	4x1Gb	x8 LP, x16 LP, x16 LP	2x 750W	Yes	Ent	Slide CMA
Standard models with a 3-year warranty (machine type 7X02)										
7X02A0EPEA	1x Bronze 3206R 8C 85W 1.9G	1x 32GB 2933	530-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide
7X021019EA	1x Silver 4208 8C 85W 2.1G	1x 32GB 2933	930-8i	8x 2.5" SAS, 2x 480GB MV SSD	Open	Open	2x 750W	Opt	Ent	Slide
7X02101AEA	1x Silver 4208 8C 85W 2.1G	2x 32GB 2933	930-8i	8x 2.5" SAS, 2x 480GB MV SSD	4x1Gb	x8 LP, x16 LP	2x 750W	Opt	Ent	Slide
7X02A0A9EA	1x Silver 4208 8C 85W 2.1G	1x 16GB 2Rx8 2666	930-8i	8x 2.5" SAS, Open bay	Open	Open	1x 750W	Opt	Ent	Slide
7X02A0F1EA	1x Silver 4208 8C 85W 2.1G	1x 32GB 2933	930-8i	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide
7X02A0FTEA	1x Silver 4208 8C 85W 2.1G	1x 32GB 2933	930-8i	8x 2.5" SAS, 2x 480GB MV SSD	Open	x8 LP, x16 LP	2x 750W	Opt	Ent	Slide
7X02A0G8EA	1x Silver 4208 8C 85W 2.1G	1x 32GB 2933	940-8i 4GB	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide
7X02A0GPEA	1x Silver 4208 8C 85W 2.1G	1x 32GB 2933	930-8i	8x 2.5" SAS, 2x 240GB S4510	Open	Open	2x 750W	Opt	Ent	Slide
7X02A0GXEA	1x Silver 4208 8C 85W 2.1G	1x 32GB 2933	940-8i 4GB	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide
7X02A0HFEA	1x Silver 4208 8C 85W 2.1G	1x 32GB 2933	9350-8i	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide
7X02A0HQEA	1x Silver 4208 8C 85W 2.1G	1x 32GB 2933	9350-8i	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide
7X02A088EA	1x Silver 4210 10C 85W 2.2G	1x 32GB 2666	930-8i	8x 2.5" SAS, Open bay	Open	Open	1x 750W	Opt	Ent	Slide
7X02A0AGEA	1x Silver 4210 10C 85W 2.2G	1x 16GB 2Rx8 2666	930-8i	8x 2.5" SAS, Open bay	Open	Open	1x 750W	Opt	Ent	Slide
7X02A0AQEA	1x Silver 4210 10C 85W 2.2G	1x 16GB 2Rx8 2666	530-8i	8x 2.5" SAS, Open bay	Open	Open	1x 750W	Opt	Ent	Slide
7X02A0AXEA	1x Silver 4210 10C 85W 2.2G	1x 16GB 2Rx8 2666	930-8i	4x 3.5" SAS, Open bay	Open	Open	1x 750W	Opt	Ent	Slide
7X02A0EEEA	1x Silver 4210R 10C 100W 2.4G	1x 16GB 2Rx8 2933	930-8i	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide
7X02A0EMEA	1x Silver 4210R 10C 100W 2.4G	1x 32GB 2933	530-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide
7X02A0F0EA	1x Silver 4210R 10C 100W 2.4G	1x 16GB 2Rx8 2933	530-8i	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide
7X02A0F4EA	1x Silver 4210R 10C 100W 2.4G	1x 32GB 2933	930-8i	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide
7X02A0GAEA	1x Silver 4210R 10C 100W 2.4G	1x 32GB 2933	940-8i 4GB	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide
7X02A0GQEA	1x Silver 4210R 10C 100W 2.4G	1x 32GB 2933	930-8i	8x 2.5" SAS, 2x 480GB MV SSD	Open	Open	2x 750W	Opt	Ent	Slide

Model	Intel Xeon Scalable processor†	Memory	RAID	Front drive bays and drives	LOM	Rear Slots	Power supply	Front VGA	XCC	Rail kit
7X02A0GWEA	1x Silver 4210R 10C 100W 2.4G	1x 32GB 2933	940-8i 4GB	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide
7X02A0HEEA	1x Silver 4210R 10C 100W 2.4G	1x 32GB 2933	9350-8i	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide
7X02A0HPEA	1x Silver 4210R 10C 100W 2.4G	1x 32GB 2933	9350-8i	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide
7X02A08REA	1x Silver 4214 12C 85W 2.2G	1x 16GB 2Rx8 2666	930-8i	8x 2.5" SAS, Open bay	Open	Open	1x 750W	Opt	Ent	Slide
7X02A0EFEA	1x Silver 4214R 12C 100W 2.4G	1x 16GB 2Rx8 2933	930-8i	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide
7X02A0ETEA	1x Silver 4214R 12C 100W 2.4G	1x 32GB 2933	930-8i	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide
7X02A0GBEA	1x Silver 4214R 12C 100W 2.4G	1x 32GB 2933	940-8i 4GB	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide
7X02A0H7EA	1x Silver 4214R 12C 100W 2.4G	1x 32GB 2933	940-8i 4GB	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide
7X02A0HCEA	1x Silver 4214R 12C 100W 2.4G	1x 32GB 2933	9350-8i	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide
7X02A08SEA	1x Silver 4215 8C 85W 2.5G	1x 16GB 2Rx8 2666	930-8i	8x 2.5" SAS, Open bay	Open	Open	1x 750W	Opt	Ent	Slide
7X02A0GFEA	1x Silver 4215 8C 85W 2.5G	1x 32GB 2933	940-8i 4GB	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide
7X02A0H6EA	1x Silver 4215 8C 85W 2.5G	1x 32GB 2933	940-8i 4GB	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide
7X02A0HBEA	1x Silver 4215 8C 85W 2.5G	1x 32GB 2933	9350-8i	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide
7X02A0EGEA	1x Silver 4215R 8C 130W 3.2G	1x 32GB 2933	930-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide
7X02A0EHEA	1x Silver 4215R 8C 130W 3.2G	1x 16GB 2Rx8 2933	930-8i	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide
7X02A0EJEA	1x Silver 4215R 8C 130W 3.2G	1x 32GB 2933	530-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide
7X02A0ELEA	1x Silver 4215R 8C 130W 3.2G	1x 32GB 2933	930-8i	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide
7X02A0GCEA	1x Silver 4215R 8C 130W 3.2G	1x 32GB 2933	940-8i 4GB	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide
7X02A0H4EA	1x Silver 4215R 8C 130W 3.2G	1x 32GB 2933	940-8i 4GB	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide
7X02A0HAEA	1x Silver 4215R 8C 130W 3.2G	1x 32GB 2933	9350-8i	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide
7X02A0B4EA	1x Silver 4216 16C 100W 2.1G	1x 16GB 2Rx8 2666	930-8i	8x 2.5" SAS, Open bay	Open	Open	1x 750W	Opt	Ent	Slide
7X02A0G9EA	1x Silver 4216 16C 100W 2.1G	1x 32GB 2933	940-8i 4GB	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide
7X02A0H2EA	1x Silver 4216 16C 100W 2.1G	1x 32GB 2933	940-8i 4GB	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide
7X02A0H5EA	1x Silver 4216 16C 100W 2.1G	1x 32GB 2933	9350-8i	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide
7X02A0B7EA	1x Gold 5217 8C 115W 3.0G	1x 16GB 2Rx8 2666	930-8i	8x 2.5" SAS, Open bay	Open	Open	1x 1100W	Opt	Ent	Slide
7X02A085EA	1x Gold 5218 16C 125W 2.3G	1x 16GB 2Rx8 2666	930-16i 4GB	10x 2.5" (6xSAS + 4xAny), Open bay	Open	Open	1x 750W	Opt	Ent	Slide

Model	Intel Xeon Scalable processor†	Memory	RAID	Front drive bays and drives	LOM	Rear Slots	Power supply	Front VGA	XCC	Rail kit
7X02A0AHEA	1x Gold 5218 16C 125W 2.3G	1x 16GB 2Rx8 2666	930-8i	8x 2.5" SAS, Open bay	Open	Open	1x 750W	Opt	Ent	Slide
7X02A0EAEA	1x Gold 5218 16C 125W 2.3G	1x 32GB 2933	930-8i	8x 2.5" SAS, Open bay	Open	Open	1x 750W	Opt	Ent	Slide
7X021017EA	2x Gold 5218R 20C 125W 2.1G	2x 32GB 2933	930-8i	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide
7X02A0EUEA	1x Gold 5218R 20C 125W 2.1G	1x 32GB 2933	930-16i 4GB	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide
7X02A0EZEA	1x Gold 5218R 20C 125W 2.1G	1x 32GB 2933	930-8i	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide
7X02A0GDEA	1x Gold 5218R 20C 125W 2.1G	1x 32GB 2933	940-8i 4GB	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide
7X02A0H0EA	1x Gold 5218R 20C 125W 2.1G	1x 32GB 2933	940-8i 4GB	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide
7X02A0H3EA	1x Gold 5218R 20C 125W 2.1G	1x 32GB 2933	9350-8i	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide
7X02A0EYEA	1x Gold 5220R 24C 150W 2.2G	1x 32GB 2933	930-8i	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide
7X02A0B6EA	1x Gold 5222 4C 105W 3.8G	1x 16GB 2Rx8 2933	930-8i	8x 2.5" SAS, Open bay	Open	Open	1x 750W	Opt	Ent	Slide
7X02A0E6EA	1x Gold 6226 12C 125W 2.7G	1x 32GB 2933	930-8i	8x 2.5" SAS, Open bay	Open	Open	1x 750W	Opt	Ent	Slide
7X02A0EXEA	1x Gold 6226R 16C 150W 2.9G	1x 32GB 2933	930-8i	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide
7X02A0GEEA	1x Gold 6226R 16C 150W 2.9G	1x 32GB 2933	940-8i 4GB	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide
7X02A0GZEA	1x Gold 6226R 16C 150W 2.9G	1x 32GB 2933	940-8i 4GB	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide
7X02A0H1EA	1x Gold 6226R 16C 150W 2.9G	1x 32GB 2933	9350-8i	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide
7X02A0ALEA	1x Gold 6230 20C 125W 2.1G	1x 32GB 2933	930-8i	8x 2.5" SAS, Open bay	Open	Open	1x 750W	Opt	Ent	Slide
7X02A0EWEA	1x Gold 6230R 26C 150W 2.1G	1x 32GB 2933	930-8i	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide
7X02A0E8EA	1x Gold 6234 8C 130W 3.3G	1x 32GB 2933	930-8i	8x 2.5" SAS, Open bay	Open	Open	1x 750W	Opt	Ent	Slide
7X02A0EVEA	1x Gold 6238R 28C 165W 2.2G	1x 32GB 2933	930-8i	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide
7X02A0ESEA	1x Gold 6240R 24C 165W 2.4G	1x 32GB 2933	930-8i	8x 2.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Ent	Slide
7X02A090EA	1x Gold 6242 16C 150W 2.8G	1x 16GB 2Rx8 2933	930-8i	8x 2.5" SAS, Open bay	Open	Open	1x 1100W	Opt	Ent	Slide
7X02A0EREA	1x Gold 6242R 20C 205W 3.1G	1x 32GB 2933	930-8i	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 1100W	Opt	Ent	Slide
7X02A093EA	1x Gold 6244 8C 150W 3.6G	1x 16GB 2Rx8 2933	930-8i	8x 2.5" SAS, Open bay	Open	Open	1x 1100W	Opt	Ent	Slide
7X02A0E9EA	1x Gold 6246 12C 165W 3.3G	1x 32GB 2933	930-8i	8x 2.5" SAS, Open bay	Open	Open	1x 750W	Opt	Ent	Slide
7X02A0EQEA	1x Gold 6246R 16C 205W 3.4G	1x 32GB 2933	930-8i	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 1100W	Opt	Ent	Slide
7X02A0ENEA	1x Gold 6248R 24C 205W 3.0G	1x 32GB 2933	930-8i	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 1100W	Opt	Ent	Slide

Model	Intel Xeon Scalable processor†	Memory	RAID	Front drive bays and drives	LOM	Rear Slots	Power supply	Front VGA	XCC	Rail kit
7X02A0EKEA	1x Gold 6258R 28C 205W 2.7G	1x 32GB 2933	930-8i	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 1100W	Opt	Ent	Slide

† Processor description: Processor model, number of cores, thermal design power (TDP), core frequency

Models for Hong Kong, Taiwan, Korea (HTK)

Table 8. Models for Hong Kong, Taiwan, Korea (HTK)

Model	Intel Xeon Scalable processor†	Memory	RAID	Front drive bays and drives	LOM	Rear Slots	Power supply	Front VGA	XCC	Rail kit
Standard models with a 3-year warranty (machine type 7X02)										
7X021018CN	1x Silver 4215 8C 85W 2.5G	4x 16GB 2Rx8 2933	930-8i	8x 2.5" SAS, 4x 1.2TB 10K, 2x 480GB S4510, 1x Ext DVDRW	4x1Gb	Open	2x 750W	Opt	Std	Slide
TopSeller models with a 3-year warranty (machine type 7X02)										
7X02A08VCN	1x Bronze 3204 6C 85W 1.9G	1x 16GB 2Rx8 2933	930-16i 4GB	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A0AECN	1x Bronze 3204 6C 85W 1.9G	1x 16GB 2Rx8 2933	930-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A09ECN	1x Silver 4208 8C 85W 2.1G	1x 16GB 2Rx8 2933	930-16i 4GB	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A0B3CN	1x Silver 4208 8C 85W 2.1G	1x 16GB 2Rx8 2933	930-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A089CN	1x Silver 4210 10C 85W 2.2G	1x 16GB 2Rx8 2933	930-16i 4GB	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A0ARCN	1x Silver 4210 10C 85W 2.2G	1x 16GB 2Rx8 2933	930-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A083CN	1x Silver 4214 12C 85W 2.2G	1x 16GB 2Rx8 2933	930-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A09JCN	1x Silver 4214 12C 85W 2.2G	1x 16GB 2Rx8 2933	930-16i 4GB	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A09ZCN	1x Silver 4215 8C 85W 2.5G	1x 16GB 2Rx8 2933	930-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A0APCN	1x Silver 4215 8C 85W 2.5G	1x 16GB 2Rx8 2933	930-16i 4GB	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A09FCN	1x Silver 4216 16C 100W 2.1G	1x 16GB 2Rx8 2933	930-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A0BGCN	1x Silver 4216 16C 100W 2.1G	1x 16GB 2Rx8 2933	930-16i 4GB	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt

Model	Intel Xeon Scalable processor†	Memory	RAID	Front drive bays and drives	LOM	Rear Slots	Power supply	Front VGA	XCC	Rail kit
7X02A0A6CN	1x Gold 5215 10C 85W 2.5G	1x 16GB 2Rx8 2933	930-16i 4GB	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A0BFCN	1x Gold 5215 10C 85W 2.5G	1x 16GB 2Rx8 2933	930-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A08JCN	1x Gold 5217 8C 115W 3.0G	1x 16GB 2Rx8 2933	930-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A09CCN	1x Gold 5217 8C 115W 3.0G	1x 16GB 2Rx8 2933	930-16i 4GB	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A07YCN	1x Gold 5218 16C 125W 2.3G	1x 16GB 2Rx8 2933	930-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A0ABCN	1x Gold 5218 16C 125W 2.3G	1x 16GB 2Rx8 2933	930-16i 4GB	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A098CN	1x Gold 5220 18C 125W 2.2G	1x 16GB 2Rx8 2933	930-16i 4GB	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A0AVCN	1x Gold 5220 18C 125W 2.2G	1x 16GB 2Rx8 2933	930-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A08BCN	1x Gold 6230 20C 125W 2.1G	1x 16GB 2Rx8 2933	930-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A09KCN	1x Gold 6230 20C 125W 2.1G	1x 16GB 2Rx8 2933	930-16i 4GB	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A086CN	1x Gold 6240 18C 150W 2.6G	1x 16GB 2Rx8 2933	930-16i 4GB	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A09RCN	1x Gold 6240 18C 150W 2.6G	1x 16GB 2Rx8 2933	930-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A08ZCN	1x Gold 6242 16C 150W 2.8G	1x 16GB 2Rx8 2933	930-16i 4GB	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A0BBCN	1x Gold 6242 16C 150W 2.8G	1x 16GB 2Rx8 2933	930-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A09HCN	1x Gold 6248 20C 150W 2.5G	1x 16GB 2Rx8 2933	930-16i 4GB	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A0A0CN	1x Gold 6248 20C 150W 2.5G	1x 16GB 2Rx8 2933	930-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A08DCN	1x Gold 6252 24C 150W 2.1G	1x 16GB 2Rx8 2933	930-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A0A1CN	1x Gold 6252 24C 150W 2.1G	1x 16GB 2Rx8 2933	930-16i 4GB	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt

† Processor description: Processor model, number of cores, thermal design power (TDP), core frequency

Models for India

Table 9. Models for India

Model	Intel Xeon Scalable processor†	Memory	RAID	Front drive bays and drives	LOM	Rear Slots	Power supply	Front VGA	XCC	Rail kit
TopSeller models with a 3-year warranty (machine type 7X02)										
7X02A08GSG	1x Bronze 3204 6C 85W 1.9G	1x 16GB 2Rx8 2933	930-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A0B5SG	1x Bronze 3204 6C 85W 1.9G	1x 16GB 2Rx8 2933	930-16i 4GB	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A0HNSG	1x Bronze 3204 6C 85W 1.9G	1x 8GB 2933	530-8i	4x 3.5" SAS, 1x 1TB SATA HDD	Open	x8 LP, x16 LP	1x 550W	Opt	Ent	Slide CMA
7X02A0A3SG	1x Silver 4208 8C 85W 2.1G	1x 16GB 2Rx8 2933	930-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A0A8SG	1x Silver 4208 8C 85W 2.1G	1x 16GB 2Rx8 2933	930-16i 4GB	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A0AASG	1x Silver 4210 10C 85W 2.2G	1x 16GB 2Rx8 2933	930-16i 4GB	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A0BASG	1x Silver 4210 10C 85W 2.2G	1x 16GB 2Rx8 2933	930-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A08LSG	1x Silver 4214 12C 85W 2.2G	1x 16GB 2Rx8 2933	930-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A0BCSG	1x Silver 4214 12C 85W 2.2G	1x 16GB 2Rx8 2933	930-16i 4GB	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A09USG	1x Silver 4215 8C 85W 2.5G	1x 16GB 2Rx8 2933	930-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A0ASSG	1x Silver 4215 8C 85W 2.5G	1x 16GB 2Rx8 2933	930-16i 4GB	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A07XSG	1x Silver 4216 16C 100W 2.1G	1x 16GB 2Rx8 2933	930-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A09PSG	1x Silver 4216 16C 100W 2.1G	1x 16GB 2Rx8 2933	930-16i 4GB	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A081SG	1x Gold 5215 10C 85W 2.5G	1x 16GB 2Rx8 2933	930-16i 4GB	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A097SG	1x Gold 5215 10C 85W 2.5G	1x 16GB 2Rx8 2933	930-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A094SG	1x Gold 5217 8C 115W 3.0G	1x 16GB 2Rx8 2933	930-16i 4GB	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A09ASG	1x Gold 5217 8C 115W 3.0G	1x 16GB 2Rx8 2933	930-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A092SG	1x Gold 5218 16C 125W 2.3G	1x 16GB 2Rx8 2933	930-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A096SG	1x Gold 5218 16C 125W 2.3G	1x 16GB 2Rx8 2933	930-16i 4GB	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A08XSG	1x Gold 5220 18C 125W 2.2G	1x 16GB 2Rx8 2933	930-16i 4GB	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A09YSG	1x Gold 5220 18C 125W 2.2G	1x 16GB 2Rx8 2933	930-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A09SSG	1x Gold 6230 20C 125W 2.1G	1x 16GB 2Rx8 2933	930-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A0AJSG	1x Gold 6230 20C 125W 2.1G	1x 16GB 2Rx8 2933	930-16i 4GB	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt

Model	Intel Xeon Scalable processor†	Memory	RAID	Front drive bays and drives	LOM	Rear Slots	Power supply	Front VGA	XCC	Rail kit
7X02A08MSG	1x Gold 6240 18C 150W 2.6G	1x 16GB 2Rx8 2933	930-16i 4GB	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A0BDSG	1x Gold 6240 18C 150W 2.6G	1x 16GB 2Rx8 2933	930-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A0AKSG	1x Gold 6242 16C 150W 2.8G	1x 16GB 2Rx8 2933	930-16i 4GB	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A0ANSG	1x Gold 6242 16C 150W 2.8G	1x 16GB 2Rx8 2933	930-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A08ASG	1x Gold 6248 20C 150W 2.5G	1x 16GB 2Rx8 2933	930-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A09GSG	1x Gold 6248 20C 150W 2.5G	1x 16GB 2Rx8 2933	930-16i 4GB	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A0ADSG	1x Gold 6252 24C 150W 2.1G	1x 16GB 2Rx8 2933	930-16i 4GB	10x 2.5" (6xSAS + 4xAny), Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt
7X02A0B0SG	1x Gold 6252 24C 150W 2.1G	1x 16GB 2Rx8 2933	930-8i	4x 3.5" SAS, Open bay	Open	x8 LP, x16 LP	1x 750W	Opt	Std	Opt

† Processor description: Processor model, number of cores, thermal design power (TDP), core frequency

Models for Japan

Table 10. Models for Japan

Model	Intel Xeon Scalable processor†	Memory	RAID	Front drive bays and drives	LOM	Rear Slots	Power supply	Front VGA	XCC	Rail kit
Standard models with a 3-year warranty (machine type 7X02)										
7X02101HJP	1x Bronze 3204 6C 85W 1.9G	1x 16GB 1Rx4 2933	9350-8i	10x 2.5" (6xSAS + 4xAny), Open bay	4x1Gb	x8 LP, x16 LP	1x 750W	Opt	Adv	Slide
7X02101MJP	1x Silver 4209T 8C 70W 2.2G	1x 16GB 1Rx4 2933	9350-8i	10x 2.5" (6xSAS + 4xAny), Open bay	4x1Gb	x8 LP, x16 LP	1x 750W	Opt	Adv	Slide
7X02101GJP	1x Silver 4210 10C 85W 2.2G	1x 16GB 1Rx4 2933	9350-8i	10x 2.5" (6xSAS + 4xAny), Open bay	4x1Gb	x8 LP, x16 LP	1x 750W	Opt	Adv	Slide
7X02A0FLJP	1x Silver 4215R 8C 130W 3.2G	1x 16GB 1Rx4 2666	Option	Option 2.5"/8, Open bay	4x1Gb	Open	1x 550W	Opt	Adv	Slide
7X02101JJP	1x Gold 5218 16C 125W 2.3G	1x 16GB 1Rx4 2933	9350-8i	10x 2.5" (6xSAS + 4xAny), Open bay	4x1Gb	x8 LP, x16 LP	1x 750W	Opt	Adv	Slide
7X02A0FFJP	1x Gold 5218R 20C 125W 2.1G	1x 16GB 1Rx4 2666	Option	Option 2.5"/8, Open bay	4x1Gb	Open	1x 550W	Opt	Adv	Slide
7X02A0FJJP	1x Gold 5220R 24C 150W 2.2G	1x 16GB 1Rx4 2666	Option	Option 2.5"/8, Open bay	4x1Gb	Open	1x 750W	Opt	Adv	Slide
7X02101KJP	1x Gold 5222 4C 105W 3.8G	1x 16GB 1Rx4 2933	9350-8i	10x 2.5" (6xSAS + 4xAny), Open bay	4x1Gb	x8 LP, x16 LP	1x 750W	Opt	Adv	Slide
7X02A0BKJP	1x Gold 6242 16C 150W 2.8G	1x 16GB 2Rx8 2933	930-8i	8x 2.5" SAS, Open bay	4x1Gb	x8 LP, x16 LP	1x 750W	Opt	Adv	Slide
7X02A0BEJP	1x Gold 6244 8C 150W 3.6G	1x 16GB 2Rx8 2933	930-8i	8x 2.5" SAS, Open bay	4x1Gb	x8 LP, x16 LP	1x 750W	Opt	Adv	Slide
7X02A0AFJP	1x Gold 6248 20C 150W 2.5G	1x 16GB 2Rx8 2933	930-8i	10x 2.5" (6xSAS + 4xAny), Open bay	4x1Gb	x8 LP, x16 LP	1x 750W	Opt	Adv	Slide
7X02A080JP	1x Gold 6252 24C 150W 2.1G	1x 16GB 2Rx8 2933	1610-4P NVMe	10x 2.5" NVMe, Open bay	4x1Gb	x8 LP, x16 LP	1x 1100W	Opt	Adv	Slide
7X02A0BHJP	1x Gold 6254 18C 200W 3.1G	1x 16GB 2Rx8 2933	Option	Option 2.5"/10, Open bay	4x1Gb	x8 LP, x16 LP	1x 1100W	Opt	Adv	Slide

Model	Intel Xeon Scalable processor†	Memory	RAID	Front drive bays and drives	LOM	Rear Slots	Power supply	Front VGA	XCC	Rail kit
7X02A09QJP	1x Platinum 8276 28C 165W 2.2G	1x 16GB 2Rx8 2933	930-8i	10x 2.5" (6xSAS + 4xAny), Open bay	4x1Gb	x8 LP, x16 LP	1x 1100W	Opt	Adv	Slide
TopSeller models with a 3-year warranty (machine type 7X02)										
7X02A084JP	1x Bronze 3204 6C 85W 1.9G	1x 16GB 1Rx4 2666	730-8i 2GB	8x 2.5" SAS, Open bay	4x1Gb	x8 LP, x16 LP	1x 550W	Opt	Adv	Slide
7X02A0B8JP	1x Bronze 3204 6C 85W 1.9G	1x 16GB 1Rx4 2666	530-8i	4x 3.5" SAS, Open bay	4x1Gb	x8 LP, x16 LP	1x 550W	Opt	Adv	Slide
7X02A0FEJP	1x Bronze 3206R 8C 85W 1.9G	1x 16GB 1Rx4 2666	Option	Option 2.5"/8, Open bay	4x1Gb	Open	1x 550W	Opt	Adv	Slide
7X02A09BJP	1x Silver 4208 8C 85W 2.1G	1x 16GB 1Rx4 2666	730-8i 2GB	8x 2.5" SAS, Open bay	4x1Gb	x8 LP, x16 LP	1x 550W	Opt	Adv	Slide
7X02A08WJP	1x Silver 4210 10C 85W 2.2G	1x 16GB 1Rx4 2666	730-8i 2GB	8x 2.5" SAS, Open bay	4x1Gb	x8 LP, x16 LP	1x 550W	Opt	Adv	Slide
7X02A0B9JP	1x Silver 4210 10C 85W 2.2G	1x 16GB 1Rx4 2666	530-8i	4x 3.5" SAS, Open bay	4x1Gb	x8 LP, x16 LP	1x 550W	Opt	Adv	Slide
7X02A0FNJP	1x Silver 4210R 10C 100W 2.4G	1x 16GB 1Rx4 2666	Option	Option 2.5"/8, Open bay	4x1Gb	Open	1x 550W	Opt	Adv	Slide
7X02A08YJP	1x Silver 4214 12C 85W 2.2G	1x 16GB 1Rx4 2666	730-8i 2GB	8x 2.5" SAS, Open bay	4x1Gb	x8 LP, x16 LP	1x 550W	Opt	Adv	Slide
7X02A0FMJP	1x Silver 4214R 12C 100W 2.4G	1x 16GB 1Rx4 2666	Option	Option 2.5"/8, Open bay	4x1Gb	Open	1x 550W	Opt	Adv	Slide
7X02A09LJP	1x Silver 4215 8C 85W 2.5G	1x 16GB 1Rx4 2666	730-8i 2GB	8x 2.5" SAS, Open bay	4x1Gb	x8 LP, x16 LP	1x 550W	Opt	Adv	Slide
7X02A08NJP	1x Silver 4216 16C 100W 2.1G	1x 16GB 1Rx4 2666	530-8i	4x 3.5" SAS, Open bay	4x1Gb	x8 LP, x16 LP	1x 550W	Opt	Adv	Slide
7X02A09MJP	1x Silver 4216 16C 100W 2.1G	1x 16GB 1Rx4 2666	730-8i 2GB	8x 2.5" SAS, Open bay	4x1Gb	x8 LP, x16 LP	1x 550W	Opt	Adv	Slide
7X02A09NJP	1x Gold 5215 10C 85W 2.5G	1x 16GB 1Rx4 2666	730-8i 2GB	8x 2.5" SAS, Open bay	4x1Gb	x8 LP, x16 LP	1x 750W	Opt	Adv	Slide
7X02A0B1JP	1x Gold 5217 8C 115W 3.0G	1x 16GB 1Rx4 2666	730-8i 2GB	8x 2.5" SAS, Open bay	4x1Gb	x8 LP, x16 LP	1x 750W	Opt	Adv	Slide
7X02A0AMJP	1x Gold 5218 16C 125W 2.3G	1x 16GB 1Rx4 2666	730-8i 2GB	8x 2.5" SAS, Open bay	4x1Gb	x8 LP, x16 LP	1x 750W	Opt	Adv	Slide
7X02A0A2JP	1x Gold 5220 18C 125W 2.2G	1x 16GB 1Rx4 2666	730-8i 2GB	8x 2.5" SAS, Open bay	4x1Gb	x8 LP, x16 LP	1x 750W	Opt	Adv	Slide
7X02A0A4JP	1x Gold 5222 4C 105W 3.8G	1x 16GB 1Rx4 2666	730-8i 2GB	8x 2.5" SAS, Open bay	4x1Gb	x8 LP, x16 LP	1x 750W	Opt	Adv	Slide
7X02A0A7JP	1x Gold 6230 20C 125W 2.1G	1x 16GB 2Rx8 2933	930-8i	8x 2.5" SAS, Open bay	4x1Gb	x8 LP, x16 LP	1x 750W	Opt	Adv	Slide
7X02A0AWJP	1x Gold 6240 18C 150W 2.6G	1x 16GB 2Rx8 2933	930-8i	8x 2.5" SAS, Open bay	4x1Gb	x8 LP, x16 LP	1x 750W	Opt	Adv	Slide

† Processor description: Processor model, number of cores, thermal design power (TDP), core frequency

Models for Latin American countries (except Brazil)

Table 11. Models with a 3-year warranty for Latin American countries (except Brazil)

Model	Intel Xeon Scalable processor†	Memory	RAID	Front drive bays and drives	LOM	Rear Slots	Power supply	Front VGA	XCC	Rail kit
Standard models with a 3-year warranty (machine type 7X02)										
7X02A0G6LA	1x Silver 4210R 10C 100W 2.4G	1x 16GB 2Rx8 2933	730-8i 1GB	8x 2.5" SAS, Open bay	4x1Gb	x8 LP, x16 FH	1x 750W	Yes	Std	Slide
7X02A0G4LA	1x Silver 4214R 12C 100W 2.4G	1x 16GB 2Rx8 2933	730-8i 1GB	8x 2.5" SAS, Open bay	4x1Gb	x8 LP, x16 FH	1x 750W	Yes	Std	Slide
7X02A0G5LA	1x Gold 5218R 20C 125W 2.1G	1x 32GB 2933	930-16i 4GB	8x 2.5" SAS, Open bay	2x10Gb RJ45	x8 LP, x16 FH	2x 750W	Yes	Ent	Slide
7X02A0G2LA	1x Gold 6230R 26C 150W 2.1G	2x 32GB 2933	930-16i 4GB	8x 2.5" SAS, Open bay	4x10Gb RJ45	x8 LP, x16 FH	2x 750W	Yes	Ent	Slide
7X02A0G1LA	1x Gold 6240R 24C 165W 2.4G	2x 32GB 2933	930-16i 4GB	8x 2.5" SAS, Open bay	4x10Gb RJ45	x8 LP, x16 FH	2x 750W	Yes	Ent	Slide

† Processor description: Processor model, number of cores, thermal design power (TDP), core frequency

Models for USA and Canada

Table 12. Models for USA and Canada

Model	Intel Xeon Scalable processor†	Memory	RAID	Front drive bays and drives	LOM	Rear Slots	Power supply	Front VGA	XCC	Rail kit
Standard models with a 3-year warranty (machine type 7X02)										
7X021014NA	1x Silver 4208 8C 85W 2.1G	1x 8GB 2933	Option	Option 2.5"/8, Open bay	2x1Gb	x8 LP, x16 LP	1x 1100W	Opt	Adv	Slide
7X02A0CENA	1x Silver 4208 8C 85W 2.1G	1x 16GB 2Rx8 2933	Option	Option 2.5"/8, Open bay	Open	Open	1x 750W	Yes	Ent	Slide
7X02A0FANA	1x Silver 4208 8C 85W 2.1G	1x 16GB 2Rx8 2933	530-8i	4x 3.5" SAS, Open bay	Open	Open	1x 750W	Yes	Ent	Slide
7X02A0HDNA	1x Silver 4208 8C 85W 2.1G	1x 32GB 2933	Option	Option 2.5"/8, Open bay	Open	Open	1x 750W	Yes	Ent	Slide
7X02A0CGNA	1x Silver 4214 12C 85W 2.2G	1x 16GB 2Rx8 2933	Option	Option 2.5"/8, Open bay	Open	Open	1x 750W	Yes	Ent	Slide
7X02A0H8NA	1x Silver 4214 12C 85W 2.2G	1x 32GB 2933	Option	Option 2.5"/8, Open bay	Open	Open	1x 750W	Yes	Ent	Slide
7X02A0CJNA	1x Silver 4216 16C 100W 2.1G	1x 16GB 2Rx8 2933	Option	Option 2.5"/8, Open bay	Open	Open	1x 750W	Yes	Ent	Slide
7X02A0GYNA	1x Silver 4216 16C 100W 2.1G	1x 32GB 2933	Option	Option 2.5"/8, Open bay	Open	Open	1x 750W	Yes	Ent	Slide
7X02A0CKNA	1x Gold 5218 16C 125W 2.3G	1x 32GB 2933	Option	Option 2.5"/8, Open bay	Open	Open	1x 750W	Yes	Ent	Slide
7X02A0H9NA	1x Gold 5218 16C 125W 2.3G	1x 32GB 2933	Option	Option 2.5"/8, Open bay	Open	Open	1x 750W	Yes	Ent	Slide
7X021015NA	1x Gold 5222 4C 105W 3.8G	1x 8GB 2933	Option	Option 2.5"/8, Open bay	2x1Gb	x8 LP, x16 LP	1x 1100W	Opt	Adv	Slide
7X021016NA	1x Gold 6208U 16C 150W 2.9G	1x 8GB 2933	Option	Option 2.5"/8, Open bay	2x1Gb	x8 LP, x16 LP	1x 1100W	Opt	Adv	Slide

† Processor description: Processor model, number of cores, thermal design power (TDP), core frequency

Processors

The SR630 server supports one or two Intel Xeon Bronze, Silver, Gold, or Platinum processors. The following table lists the specifications of the processors for the SR630 server.

Processor support: Both 1st Gen and 2nd Gen Intel Xeon SP processors are supported. For supported 1st Gen processors, see the Continued support for 1st Gen Intel Xeon Scalable processors section.

Abbreviations used in the table:

- UPI: Ultra Path Interconnect
- TDP: Thermal Design Power
- HT: Hyper-Threading
- TB: Turbo Boost 2.0
- VT-x: Virtualization Technology
- VT-d: Virtualization Technology for Directed I/O
- SST-PP: Speed Select Technology - Performance Profile
- FMA: Fused-Multiply Add (AVX-512)
- DCPMM: DC Persistent Memory Module
- RAS: Reliability, Availability, and Serviceability
  - Std: Standard RAS
  - Adv: Advanced RAS

Table 13. Processor specifications

CPU model	Cores / threads	Core speed (Base / TB Max)	Cache	Max DDR4 speed	Max memory capacity per socket	UPI speed	TDP	HT	TB	VT-x	VT-d	SST-PP	FMA units	DCPMM	RAS
Intel Xeon Bronze processors															
3204	6 / 6	1.9 / 1.9 GHz	8.25 MB	2133 MHz	1 TB	9.6 GT/s	85 W	N	N	Y	Y	N	1	N	Std
3206R	8 / 8	1.9 / 1.9 GHz	11 MB	2133 MHz	1 TB	9.6 GT/s	85 W	N	N	Y	Y	N	1	N	Std
Intel Xeon Silver processors															
4208	8 / 16	2.1 / 3.2 GHz	11 MB	2400 MHz	1 TB	9.6 GT/s	85 W	Y	Y	Y	Y	N	1	N	Std
4209T	8 / 16	2.2 / 3.2 GHz	11 MB	2400 MHz	1 TB	9.6 GT/s	70 W	Y	Y	Y	Y	N	1	N	Std
4210	10 / 20	2.2 / 3.2 GHz	13.75 MB	2400 MHz	1 TB	9.6 GT/s	85 W	Y	Y	Y	Y	N	1	N	Std
4210R	10 / 20	2.4 / 3.2 GHz	13.75 MB	2400 MHz	1 TB	9.6 GT/s	100 W	Y	Y	Y	Y	N	1	N	Std
4210T	10 / 20	2.3 / 3.2 GHz	13.75 MB	2400 MHz	1 TB	9.6 GT/s	95 W	Y	Y	Y	Y	N	1	N	Std
4214	12 / 24	2.2 / 3.2 GHz	16.5 MB	2400 MHz	1 TB	9.6 GT/s	85 W	Y	Y	Y	Y	N	1	N	Std
4214R	12 / 24	2.4 / 3.5 GHz	16.5 MB	2400 MHz	1 TB	9.6 GT/s	100 W	Y	Y	Y	Y	N	1	N	Std
4214Y	12 / 24	2.2 / 3.2 GHz	16.5 MB	2400 MHz	1 TB	9.6 GT/s	85 W	Y	Y	Y	Y	Y	1	N	Std
	10 / 20	2.3 / 3.2 GHz													
	8 / 16	2.4 / 3.2 GHz													
4215	8 / 16	2.5 / 3.5 GHz	11 MB	2400 MHz	1 TB	9.6 GT/s	85 W	Y	Y	Y	Y	N	1	Y	Std
4215R	8 / 16	3.2 / 4.0 GHz	11 MB	2400 MHz	1 TB	9.6 GT/s	130 W	Y	Y	Y	Y	N	1	Y	Std
4216	16 / 32	2.1 / 3.2 GHz	22 MB	2400 MHz	1 TB	9.6 GT/s	100 W	Y	Y	Y	Y	N	1	N	Std
Intel Xeon Gold processors															
5215	10 / 20	2.5 / 3.4 GHz	13.75 MB	2666 MHz	1 TB	10.4 GT/s	85 W	Y	Y	Y	Y	N	1	Y	Adv

CPU model	Cores / threads	Core speed (Base / TB Max)	Cache	Max DDR4 speed	Max memory capacity per socket	UPI speed	TDP	HT	TB	VT-x	VT-d	SST-PP	FMA units	DCPMM	RAS
5215L	10 / 20	2.5 / 3.4 GHz	13.75 MB	2666 MHz	4.5 TB	10.4 GT/s	85 W	Y	Y	Y	Y	N	1	Y	Adv
5217	8 / 16	3.0 / 3.7 GHz	11 MB	2666 MHz	1 TB	10.4 GT/s	115 W	Y	Y	Y	Y	N	1	Y	Adv
5218	16 / 32	2.3 / 3.9 GHz	22 MB	2666 MHz	1 TB	10.4 GT/s	125 W	Y	Y	Y	Y	N	1	Y	Adv
5218B	16 / 32	2.3 / 3.9 GHz	22 MB	2666 MHz	1 TB	10.4 GT/s	125 W	Y	Y	Y	Y	N	1	Y	Adv
5218R	20 / 40	2.1 / 4.0 GHz	27.5 MB	2666 MHz	1 TB	10.4 GT/s	125 W	Y	Y	Y	Y	N	1	Y	Adv
5218N	16 / 32	2.3 / 3.7 GHz	22 MB	2666 MHz	1 TB	10.4 GT/s	110 W	Y	Y	Y	Y	N	1	Y	Adv
5218T	16 / 32	2.1 / 3.8 GHz	22 MB	2666 MHz	1 TB	10.4 GT/s	105 W	Y	Y	Y	Y	N	1	Y	Adv
5220	18 / 36	2.2 / 3.9 GHz	24.75 MB	2666 MHz	1 TB	10.4 GT/s	125 W	Y	Y	Y	Y	N	1	Y	Adv
5220R	24 / 48	2.2 / 4.0 GHz	35.75 MB	2666 MHz	1 TB	10.4 GT/s	150 W	Y	Y	Y	Y	N	1	Y	Adv
5220S	18 / 36	2.7 / 3.9 GHz	24.75 MB	2666 MHz	1 TB	10.4 GT/s	125 W	Y	Y	Y	Y	N	1	Y	Adv
5220T	18 / 36	1.9 / 3.9 GHz	24.75 MB	2666 MHz	1 TB	10.4 GT/s	105 W	Y	Y	Y	Y	N	1	Y	Adv
5222	4 / 8	3.8 / 3.9 GHz	16.5 MB	2933 MHz	1 TB	10.4 GT/s	105 W	Y	Y	Y	Y	N	2	Y	Adv
6208U	16 / 32	2.9 / 3.9 GHz	22 MB	2933 MHz	1 TB	No UPI	150 W	Y	Y	Y	Y	N	2	Y	Adv
6222V	20 / 40	1.8 / 3.6 GHz	27.5 MB	2400 MHz	1 TB	10.4 GT/s	115 W	Y	Y	Y	Y	N	2	Y	Adv
6226	12 / 24	2.7 / 3.7 GHz	19.25 MB	2933 MHz	1 TB	10.4 GT/s	125 W	Y	Y	Y	Y	N	2	Y	Adv
6226R	16 / 32	2.9 / 3.9 GHz	22 MB	2933 MHz	1 TB	10.4 GT/s	150 W	Y	Y	Y	Y	N	2	Y	Adv
6230	20 / 40	2.1 / 3.9 GHz	27.5 MB	2933 MHz	1 TB	10.4 GT/s	125 W	Y	Y	Y	Y	N	2	Y	Adv
6230N	20 / 40	2.3 / 3.9 GHz	27.5 MB	2933 MHz	1 TB	10.4 GT/s	125 W	Y	Y	Y	Y	N	2	Y	Adv
6230R	26 / 52	2.1 / 4.0 GHz	35.75 MB	2933 MHz	1 TB	10.4 GT/s	150 W	Y	Y	Y	Y	N	2	Y	Adv
6230T	20 / 40	2.1 / 3.9 GHz	27.5 MB	2933 MHz	1 TB	10.4 GT/s	125 W	Y	Y	Y	Y	N	2	Y	Adv
6234	8 / 16	3.3 / 4.0 GHz	24.75 MB	2933 MHz	1 TB	10.4 GT/s	130 W	Y	Y	Y	Y	N	2	Y	Adv
6238	22 / 44	2.1 / 3.7 GHz	30.25 MB	2933 MHz	1 TB	10.4 GT/s	140 W	Y	Y	Y	Y	N	2	Y	Adv
6238L	22 / 44	2.1 / 3.7 GHz	30.25 MB	2933 MHz	4.5 TB	10.4 GT/s	140 W	Y	Y	Y	Y	N	2	Y	Adv
6238R	28 / 56	2.2 / 4.0 GHz	38.5 MB	2933 MHz	1 TB	10.4 GT/s	165 W	Y	Y	Y	Y	N	2	Y	Adv
6238T	22 / 44	1.9 / 3.7 GHz	30.25 MB	2933 MHz	1 TB	10.4 GT/s	125 W	Y	Y	Y	Y	N	2	Y	Adv
6240	18 / 36	2.6 / 3.9 GHz	24.75 MB	2933 MHz	1 TB	10.4 GT/s	150 W	Y	Y	Y	Y	N	2	Y	Adv
6240L	18 / 36	2.6 / 3.9 GHz	24.75 MB	2933 MHz	4.5 TB	10.4 GT/s	150 W	Y	Y	Y	Y	N	2	Y	Adv
6240R	24 / 48	2.4 / 4.0 GHz	35.75 MB	2933 MHz	1 TB	10.4 GT/s	165 W	Y	Y	Y	Y	N	2	Y	Adv
6240Y	18 / 36	2.6 / 3.9 GHz	24.75 MB	2933 MHz	1 TB	10.4 GT/s	150 W	Y	Y	Y	Y	Y	2	Y	Adv
	14 / 28	2.8 / 3.9 GHz													
	8 / 16	3.1 / 3.9 GHz													
6242	16 / 32	2.8 / 3.9 GHz	22 MB	2933 MHz	1 TB	10.4 GT/s	150 W	Y	Y	Y	Y	N	2	Y	Adv
6242R	20 / 40	3.1 / 4.1 GHz	35.75 MB	2933 MHz	1 TB	10.4 GT/s	205 W	Y	Y	Y	Y	N	2	Y	Adv

CPU model	Cores / threads	Core speed (Base / TB Max)	Cache	Max DDR4 speed	Max memory capacity per socket	UPI speed	TDP	HT	TB	VT-x	VT-d	SST-PP	FMA units	DCPMM	RAS
6244	8 / 16	3.6 / 4.4 GHz	24.75 MB	2933 MHz	1 TB	10.4 GT/s	150 W	Y	Y	Y	Y	N	2	Y	Adv
6246	12 / 24	3.3 / 4.2 GHz	24.75 MB	2933 MHz	1 TB	10.4 GT/s	165 W	Y	Y	Y	Y	N	2	Y	Adv
6246R	16 / 32	3.4 / 4.1 GHz	35.75 MB	2933 MHz	1 TB	10.4 GT/s	205 W	Y	Y	Y	Y	N	2	Y	Adv
6248	20 / 40	2.5 / 3.9 GHz	27.5 MB	2933 MHz	1 TB	10.4 GT/s	150 W	Y	Y	Y	Y	N	2	Y	Adv
6248R	24 / 48	3.0 / 4.0 GHz	35.75 MB	2933 MHz	1 TB	10.4 GT/s	205 W	Y	Y	Y	Y	N	2	Y	Adv
6252	24 / 48	2.1 / 3.7 GHz	35.75 MB	2933 MHz	1 TB	10.4 GT/s	150 W	Y	Y	Y	Y	N	2	Y	Adv
6252N	24 / 48	2.3 / 3.6 GHz	35.75 MB	2933 MHz	1 TB	10.4 GT/s	150 W	Y	Y	Y	Y	N	2	Y	Adv
6254	18 / 36	3.1 / 4.0 GHz	24.75 MB	2933 MHz	1 TB	10.4 GT/s	200 W	Y	Y	Y	Y	N	2	Y	Adv
6258R	28 / 56	2.7 / 4.0 GHz	38.5 MB	2933 MHz	1 TB	10.4 GT/s	205 W	Y	Y	Y	Y	N	2	Y	Adv
6262V	24 / 48	1.9 / 3.6 GHz	33 MB	2400 MHz	1 TB	10.4 GT/s	135 W	Y	Y	Y	Y	N	2	Y	Adv
Intel Xeon Platinum processors															
8253	16 / 32	2.2 / 3.0 GHz	22 MB	2933 MHz	1 TB	10.4 GT/s	125 W	Y	Y	Y	Y	N	2	Y	Adv
8256	4 / 8	3.8 / 3.9 GHz	16.5 MB	2933 MHz	1 TB	10.4 GT/s	105 W	Y	Y	Y	Y	N	2	Y	Adv
8260	24 / 48	2.4 / 3.9 GHz	35.75 MB	2933 MHz	1 TB	10.4 GT/s	165 W	Y	Y	Y	Y	N	2	Y	Adv
8260L	24 / 48	2.4 / 3.9 GHz	35.75 MB	2933 MHz	4.5 TB	10.4 GT/s	165 W	Y	Y	Y	Y	N	2	Y	Adv
8260Y	24 / 48	2.4 / 3.9 GHz	35.75 MB	2933 MHz	1 TB	10.4 GT/s	165 W	Y	Y	Y	Y	Y	2	Y	Adv
	20 / 40	2.5 / 3.9 GHz													
	16 / 32	2.7 / 3.9 GHz													
8268	24 / 48	2.9 / 3.9 GHz	35.75 MB	2933 MHz	1 TB	10.4 GT/s	205 W	Y	Y	Y	Y	N	2	Y	Adv
8270	26 / 52	2.7 / 4.0 GHz	35.75 MB	2933 MHz	1 TB	10.4 GT/s	205 W	Y	Y	Y	Y	N	2	Y	Adv
8276	28 / 56	2.2 / 4.0 GHz	38.5 MB	2933 MHz	1 TB	10.4 GT/s	165 W	Y	Y	Y	Y	N	2	Y	Adv
8276L	28 / 56	2.2 / 4.0 GHz	38.5 MB	2933 MHz	4.5 TB	10.4 GT/s	165 W	Y	Y	Y	Y	N	2	Y	Adv
8280	28 / 56	2.7 / 4.0 GHz	38.5 MB	2933 MHz	1 TB	10.4 GT/s	205 W	Y	Y	Y	Y	N	2	Y	Adv
8280L	28 / 56	2.7 / 4.0 GHz	38.5 MB	2933 MHz	4.5 TB	10.4 GT/s	205 W	Y	Y	Y	Y	N	2	Y	Adv

Configuration notes:

- The Intel Xeon Gold 5218 and 5218B processors have similar specifications; however, they use different silicon designs and cannot be mixed in the same system.
- Processors with a U suffix are only supported in single-processor configurations. This limits the PCIe slots that can be used.
- The processors that support SST-PP offer three distinct operating points that are defined by a core count with a base speed associated with that core count. The operating point is static, it is selected during the boot process and cannot be changed at runtime.

For the SR630 server models that come standard with one processor, the second processor can be ordered, if required (see the following table for ordering information). The second processor must be of the same model as the first processor. The second processor option includes a processor and a heatsink; two additional system fans are not included and need to be purchased with the second processor (see Cooling for details).

Table 14. Processor options

Part number	Feature code*	Description
Intel Xeon Bronze processors		
4XG7A37939	B4HU	SR530/SR570/SR630 Intel Xeon Bronze 3204 6C 85W 1.9GHz Processor w/o FAN
4XG7A37990	B7N3	SR530/SR570/SR630 Intel Xeon Bronze 3206R 8C 85W 1.9GHz Processor w/o FAN
Intel Xeon Silver processors		
4XG7A37936	B4HT	SR530/SR570/SR630 Intel Xeon Silver 4208 8C 85W 2.1GHz Processor w/o FAN
4XG7A37945	B4P4	SR530/SR570/SR630 Intel Xeon Silver 4209T 8C 70W 2.2GHz Processor w/o FAN
4XG7A37933	B4HS	SR530/SR570/SR630 Intel Xeon Silver 4210 10C 85W 2.2GHz Processor w/o FAN
4XG7A37988	B7N5	SR530/SR570/SR630 Intel Xeon Silver 4210R 10C 100W 2.4GHz Processor w/o FAN
4XG7A63300	BAZT	SR630 Intel Xeon Silver 4210T 10C 95W 2.3GHz Processor w/o FAN
4XG7A37930	B4HR	SR530/SR570/SR630 Intel Xeon Silver 4214 12C 85W 2.2GHz Processor w/o FAN
4XG7A37987	B7N6	SR530/SR570/SR630 Intel Xeon Silver 4214R 12C 100W 2.4GHz Processor w/o FAN
4XG7A37942	B4NW	SR530/SR570/SR630 Intel Xeon Silver 4214Y 12/10/8C 85W 2.2GHz Processor w/o FAN
4XG7A37927	B4HQ	SR530/SR570/SR630 Intel Xeon Silver 4215 8C 85W 2.5GHz Processor w/o FAN
4XG7A63298	BAZU	SR570/SR630 Intel Xeon Silver 4215R 8C 130W 3.2GHz Processor w/o FAN
4XG7A37924	B4HP	SR530/SR570/SR630 Intel Xeon Silver 4216 16C 100W 2.1GHz Processor w/o FAN
Intel Xeon Gold processors		
4XG7A37917	B4HN	SR530/SR570/SR630 Intel Xeon Gold 5215 10C 85W 2.5GHz Processor w/o FAN
4XG7A37911	B4P9	SR530/SR570/SR630 Intel Xeon Gold 5215L 10C 85W 2.5GHz Processor w/o FAN
4XG7A37920	B4HM	SR630 Intel Xeon Gold 5217 8C 115W 3.0GHz Processor w/o FAN
4XG7A37896	B4HL	SR530/SR570/SR630 Intel Xeon Gold 5218 16C 125W 2.3GHz Processor w/o FAN
4XG7A37959	B6BS	SR530/SR570/SR630 Intel Xeon Gold 5218B 16C 125W 2.3GHz Processor w/o FAN
4XG7A63296	BAZS	SR530/SR570/SR630 Intel Xeon Gold 5218R 20C 125W 2.1GHz Processor w/o FAN
4XG7A37956	B5S0	SR630 Intel Xeon Gold 5218N 16C 110W 2.3GHz Processor Option Kit w/o FAN
4XG7A37966	B4P3	SR630 Intel Xeon Gold 5218T 16C 105W 2.1GHz Processor w/o FAN
4XG7A37893	B4HK	SR530/SR570/SR630 Intel Xeon Gold 5220 18C 125W 2.2GHz Processor w/o FAN
4XG7A37984	B7N9	SR570/SR630 Intel Xeon Gold 5220R 24C 150W 2.2GHz Processor w/o FAN
4XG7A38018	B6CW	SR530/SR570/SR630 Intel Xeon Gold 5220S 18C 125W 2.7GHz Processor w/o FAN
4XG7A15809	B6CQ	SR630 Intel Xeon Gold 5220T 18C 105W 1.9GHz Processor w/o Fan
4XG7A37952	B5S1	SR630 Intel Xeon Gold 5222 4C 105W 3.8GHz Processor w/o FAN
None**	BAZV	Intel Xeon Gold 6208U 16C 150W 2.9GHz Processor
4XG7A38022	B6CV	SR530/SR570/SR630 Intel Xeon Gold 6222V 20C 115W 1.8GHz Processor w/o FAN
4XG7A38020	B6CL	SR530/SR570/SR630 Intel Xeon Gold 6226 12C 125W 2.7GHz Processor w/o FAN
4XG7A63292	BAZW	SR570/SR630 Intel Xeon Gold 6226R 16C 150W 2.9GHz Processor w/o FAN
4XG7A37890	B4HJ	SR530/SR570/SR630 Intel Xeon Gold 6230 20C 125W 2.1GHz Processor w/o FAN
4XG7A15720	B5RY	SR630 Intel Xeon Gold 6230N 20C 125W 2.3GHz Processor Option Kit w/o FAN

Part number	Feature code*	Description
4XG7A63290	BAZX	SR570/SR630 Intel Xeon Gold 6230R 26C 150W 2.1GHz Processor w/o FAN
4XG7A15808	B6CP	SR630 Intel Xeon Gold 6230T 20C 125W 2.1GHz Processor w/o Fan
4XG7A38000	B6CK	SR570/SR630 Intel Xeon Gold 6234 8C 130W 3.3GHz Processor w/o FAN
4XG7A38024	B6CJ	SR570/SR630 Intel Xeon Gold 6238 22C 140W 2.1GHz Processor w/o FAN
4XG7A38002	B6CR	SR570/SR630 Intel Xeon Gold 6238L 22C 140W 2.1GHz Processor w/o FAN
4XG7A63288	BAZL	SR630 Intel Xeon Gold 6238R 28C 165W 2.2GHz Processor w/o FAN
4XG7A37907	B4P2	SR630 Intel Xeon Gold 6238T 22C 125W 1.9GHz Processor w/o FAN
4XG7A37884	B4HH	SR570/SR630 Intel Xeon Gold 6240 18C 150W 2.6GHz Processor w/o FAN
4XG7A38014	B6CS	SR570/SR630 Intel Xeon Gold 6240L 18C 150W 2.6GHz Processor w/o FAN
4XG7A63286	BAZM	SR630 Intel Xeon Gold 6240R 24C 165W 2.4GHz Processor w/o FAN
4XG7A37904	B4NV	SR630 Intel Xeon Gold 6240Y 18/14/8C 150W 2.6GHz Processor w/o FAN
4XG7A37887	B4HG	SR630 Intel Xeon Gold 6242 16C 150W 2.8GHz Processor w/o FAN
4XG7A63284	BAZN	SR630 Intel Xeon Gold 6242R 20C 205W 3.1GHz Processor w/o FAN
4XG7A15875	B4HF	SR630 Intel Xeon Gold 6244 8C 150W 3.6GHz Processor w/o FAN
4XG7A37965	B6PD	SR630 Intel Xeon Gold 6246 12C 165W 3.3GHz Processor w/o FAN
4XG7A63282	BAZP	SR630 Intel Xeon Gold 6246R 16C 205W 3.4GHz Processor w/o FAN
4XG7A15893	B4HE	SR570/SR630 Intel Xeon Gold 6248 20C 150W 2.5GHz Processor w/o FAN
4XG7A63280	BAZQ	SR630 Intel Xeon Gold 6248R 24C 205W 3.0GHz Processor w/o FAN
4XG7A15890	B4HC	SR570/SR630 Intel Xeon Gold 6252 24C 150W 2.1GHz Processor w/o FAN
4XG7A15812	B6CT	SR630 Intel Xeon Gold 6252N 24C 150W 2.3GHz Processor w/o FAN
4XG7A15873	B4HD	SR630 Intel Xeon Gold 6254 18C 200W 3.1GHz Processor w/o FAN
4XG7A63278	BAZR	SR630 Intel Xeon Gold 6258R 28C 205W 2.7GHz Processor w/o FAN
4XG7A38009	B6CU	SR570/SR630 Intel Xeon Gold 6262V 24C 135W 1.9GHz Processor w/o FAN
Intel Xeon Platinum processors		
4XG7A37899	B5RZ	SR530/SR570/SR630 Intel Xeon Platinum 8253 16C 125W 2.2GHz Processor w/o FAN
4XG7A37948	B5S2	SR630 Intel Xeon Platinum 8256 4C 105W 3.8GHz Processor w/o FAN
4XG7A15888	B4HB	SR630 Intel Xeon Platinum 8260 24C 165W 2.4GHz Processor w/o FAN
4XG7A15884	B4P7	SR630 Intel Xeon Platinum 8260L 24C 165W 2.4GHz Processor w/o FAN
4XG7A37902	B4NU	SR630 Intel Xeon Platinum 8260Y 24/20/16C 165W 2.4GHz Processor w/o FAN
4XG7A15871	B4HA	SR630 Intel Xeon Platinum 8268 24C 205W 2.9GHz Processor w/o FAN
4XG7A15869	B4H9	SR630 Intel Xeon Platinum 8270 26C 205W 2.7GHz Processor w/o FAN
4XG7A15882	B4H8	SR630 Intel Xeon Platinum 8276 28C 165W 2.2GHz Processor w/o FAN
4XG7A15878	B4P6	SR630 Intel Xeon Platinum 8276L 28C 165W 2.2GHz Processor w/o FAN
4XG7A15867	B4H7	SR630 Intel Xeon Platinum 8280 28C 205W 2.7GHz Processor w/o FAN
4XG7A15863	B4P5	SR630 Intel Xeon Platinum 8280L 28C 205W 2.7GHz Processor w/o FAN

* For CTO configurations, the feature code represents a processor, and fans and heatsinks are derived by the configuration tool.
** Factory-installed only; no field upgrade. Supported in the single-processor configurations only.

Configuration note: If processors with 200 W or 205 W TDP are used, or if Gold 6240Y, 6244, 6246, or 6252N processors are used, the following conditions must be met:

- No rear HDD kit installed
- No U.2 NVMe 10-bay backplane installed
- No PCIe flash adapters installed
- No GPUs installed
- No persistent memory modules installed
- A maximum of one supercapacitor module installed
- ASHRAE A2 (35°C [95 °F]) only
- The server performance might be impacted in case of a fan failure
- No backplanes installed, or 2.5" 10-bay AnyBay backplane with up to 4 drives installed (For Gold 6240Y, 6242R, 6246R, 6248R, 6254, and 6258R processors only)

Continued support for 1st Gen Intel Xeon Scalable processors

The SR630 also continues to support the 1st Gen Intel Xeon Scalable processors (formerly codenamed "Skylake") listed in the following table.

Table 15. Long-life 1st Gen Intel Xeon Scalable processors

Part number	Feature code	Description
7XG7A05526	AWEH	ThinkSystem SR630 Intel Xeon Bronze 3106 8C 85W 1.7GHz Processor Option Kit
7XG7A05530	AWET	ThinkSystem SR630 Intel Xeon Silver 4109T 8C 70W 2.0GHz Processor Option Kit
7XG7A05531	AWEE	ThinkSystem SR630 Intel Xeon Silver 4110 8C 85W 2.1GHz Processor Option Kit
7XG7A05535	AWES	ThinkSystem SR630 Intel Xeon Silver 4114T 10C 85W 2.2GHz Processor Option Kit
7XG7A05532	AWER	ThinkSystem SR630 Intel Xeon Silver 4116 12C 85W 2.1GHz Processor Option Kit
7XG7A05529	AWEA	ThinkSystem SR630 Intel Xeon Silver 4116T 12C 85W 2.1GHz Processor Option Kit
7XG7A05551	No CTO*	ThinkSystem SR630 Intel Xeon Gold 5115 10C 85W 2.4GHz Processor Option Kit
7XG7A05536	AWEP	ThinkSystem SR630 Intel Xeon Gold 5118 12C 105W 2.3GHz Processor Option Kit
7XG7A05537	AWEQ	ThinkSystem SR630 Intel Xeon Gold 5119T 14C 85W 1.9GHz Processor Option Kit
7XG7A05538	AWE8	ThinkSystem SR630 Intel Xeon Gold 5120T 14C 105W 2.2GHz Processor Option Kit
7XG7A05546	AWEL	ThinkSystem SR630 Intel Xeon Gold 6126 12C 125W 2.6GHz Processor Option Kit
7XG7A05545	AWE5	ThinkSystem SR630 Intel Xeon Gold 6126T 12C 125W 2.6GHz Processor Option Kit
7XG7A05543	AWEN	ThinkSystem SR630 Intel Xeon Gold 6130 16C 125W 2.1GHz Processor Option Kit
7XG7A05542	AWE4	ThinkSystem SR630 Intel Xeon Gold 6130T 16C 125W 2.1GHz Processor Option Kit
7XG7A05560	No CTO*	ThinkSystem SR630 Intel Xeon Gold 6134 8C 130W 3.2GHz Processor Option Kit
7XG7A05541	AWDZ	ThinkSystem SR630 Intel Xeon Gold 6138 20C 125W 2.0GHz Processor Option Kit
7XG7A05540	AWEM	ThinkSystem SR630 Intel Xeon Gold 6138T 20C 125W 2.0GHz Processor Option Kit
7XG7A05555	No CTO*	ThinkSystem SR630 Intel Xeon Gold 6144 8C 155W 3.5GHz Processor Option Kit
7XG7A05554	No CTO*	Thinksystem SR630 Intel Xeon Gold 6146 12C 165W 3.2GHz Processor Option Kit
7XG7A05549	No CTO*	ThinkSystem SR630 Intel Xeon Platinum 8153 16C 125W 2.0GHz Processor Option Kit
7XG7A05569	No CTO*	ThinkSystem SR630 Intel Xeon Platinum 8158 12C 150W 3.0GHz Processor Option Kit
7XG7A05568	No CTO*	ThinkSystem SR630 Intel Xeon Platinum 8160 24C 150W 2.1GHz Processor Option Kit
7XG7A05566	AWEK	ThinkSystem SR630 Intel Xeon Platinum 8160T 24C 150W 2.1GHz Processor Option Kit
7XG7A05563	No CTO*	ThinkSystem SR630 Intel Xeon Platinum 8176 28C 165W 2.1GHz Processor Option Kit
4XG7A09411	No CTO*	ThinkSystem SR630 Intel Xeon Platinum 8180M 28C 205W 2.5GHz Processor Option Kit

* Only available as a field upgrade for existing customers. Not available in CTO (configure to order) configurations.

For specifications of these processors, see the Intel Xeon Scalable Processor Reference for Lenovo ThinkSystem Servers:
https://lenovopress.com/lp1262-intel-xeon-sp-processor-reference#term=SKL

Configuration note for 1st Gen processors: For Gold 6144, 6146, and 6154 processors and Platinum 8168, 8180, and 8180M processors, the following conditions must be met:

- No rear HDD kit installed
- No U.2 NVMe 10-bay backplane installed
- No PCIe flash adapters installed
- No GPUs installed
- A maximum of one supercap module installed
- ASHRAE A2 (35°C, 95 °F) only
- The server performance might be impacted in case of a fan failure
- For Gold 6154 processor only: No backplanes installed, or 2.5" 10-bay AnyBay backplane with up to 4 drives installed

Memory

The SR630 server supports up to 12 TruDDR4 memory when one processor is installed and up to 24 DIMMs when two processors are installed. Each processor has six memory channels (two integrated memory controllers with three memory channels per memory controller), and there are two DIMMs per channel.

With 2nd Gen Xeon SP processors, the server supports memory up to 2933 MHz. With 1st Gen Xeon SP processors, the server supports memory up to 2666 MHz.

Lenovo TruDDR4 memory uses the highest-quality components sourced from Tier 1 DRAM suppliers and only memory that meets strict requirements is selected. It is compatibility tested and tuned on every ThinkSystem server to maximize performance and reliability. TruDDR4 memory has a unique signature programmed into the DIMM, which enables Lenovo servers to verify whether the memory installed is qualified and supported. Lenovo qualified and supported TruDDR4 memory is covered by Lenovo warranty, and service and support provided worldwide.

The following maximum memory capacities are supported by the server with 2nd Gen Xeon SP processors (up to 2933 MHz):

- RDIMMs: 1.5 TB (768 GB per processor)
- 3DS RDIMMs: 3 TB (1.5 TB per processor)

The following maximum memory capacities supported by the server with 1st Gen Xeon SP processors (up to 2666 MHz):

- RDIMMs: 768 GB (384 GB per processor)
- LRDIMMs: 1.5 TB (768 GB per processor)
- 3DS RDIMMs: 3 TB (1.5 GB per processor)

The following memory protection technologies are supported by the processor's integrated memory controllers:

- ECC
- SDDC (for x4-based memory DIMMs)
- ADDDC (for x4-based memory DIMMs; Gold and Platinum processors only)
- Memory mirroring
- Memory rank sparing
- Patrol scrubbing
- Demand scrubbing

The following table lists memory options are supported by the server. The table also indicates which processor generation is supported for each memory option.

Table 16. Memory options

Part number	Feature code	Description	Maximum quantity*	Gen 1 CPU	Gen 2 CPU
RDIMMs - 2933 MHz Performance+					
4X77A12184	B5N6	ThinkSystem 16GB TruDDR4 Performance+ 2933MHz (2Rx8 1.2V) RDIMM	12 / 24	No	Yes
4X77A12185	B5N7	ThinkSystem 32GB TruDDR4 Performance+ 2933MHz (2Rx4 1.2V) RDIMM	12 / 24	No	Yes
4X77A12186	B5N8	ThinkSystem 64GB TruDDR4 Performance+ 2933MHz (2Rx4 1.2V) RDIMM	12 / 24	No	Yes
RDIMMs - 2933 MHz					
4ZC7A08706	B4H1	ThinkSystem 8GB TruDDR4 2933MHz (1Rx8 1.2V) RDIMM	12 / 24	No	Yes
4ZC7A08707	B4LY	ThinkSystem 16GB TruDDR4 2933MHz (1Rx4 1.2V) RDIMM	12 / 24	No	Yes
4ZC7A08708	B4H2	ThinkSystem 16GB TruDDR4 2933MHz (2Rx8 1.2V) RDIMM	12 / 24	No	Yes
4ZC7A08709	B4H3	ThinkSystem 32GB TruDDR4 2933MHz (2Rx4 1.2V) RDIMM	12 / 24	No	Yes
4ZC7A08710	B4H4	ThinkSystem 64GB TruDDR4 2933MHz (2Rx4 1.2V) RDIMM	12 / 24	No	Yes
3DS RDIMMs - 2933 MHz Performance+					
4X77A12187	B5N9	ThinkSystem 128GB TruDDR4 Performance+ 2933MHz (4Rx4 1.2V) 3DS RDIMM	12 / 24	No	Yes
3DS RDIMMs - 2933 MHz					
4ZC7A15113	B587	ThinkSystem 128GB TruDDR4 2933MHz (4Rx4 1.2V) 3DS RDIMM	12 / 24	No	Yes
RDIMMs - 2666 MHz					
7X77A01301	AUU1	ThinkSystem 8GB TruDDR4 2666 MHz (1Rx8 1.2V) RDIMM	12 / 24	Yes	No
7X77A01302	AUNB	ThinkSystem 16GB TruDDR4 2666 MHz (1Rx4 1.2V) RDIMM	12 / 24	Yes	Yes
7X77A01303	AUNC	ThinkSystem 16GB TruDDR4 2666 MHz (2Rx8 1.2V) RDIMM	12 / 24	Yes	Yes
7X77A01304	AUND	ThinkSystem 32GB TruDDR4 2666 MHz (2Rx4 1.2V) RDIMM	12 / 24	Yes	Yes
LRDIMMs - 2666 MHz					
7X77A01305	AUNE	ThinkSystem 64GB TruDDR4 2666 MHz (4Rx4 1.2V) LRDIMM	12 / 24	Yes	No
3DS RDIMMs - 2666 MHz					
4ZC7A08716	AUW5	ThinkSystem 64GB TruDDR4 2666MHz (4Rx4, 1.2V) 3DS RDIMM	12 / 24	No	Yes
7X77A01307	AUNF	ThinkSystem 128GB TruDDR4 2666 MHz (8Rx4 1.2V) 3DS RDIMM	12 / 24	Yes	No

* The maximum quantity shown is with one processor / two processors.

Configuration notes:

- All DIMMs in the server operate at the same speed, which is determined as the lowest value of:
    - DIMM rated speed (2666 MHz or 2933 MHz).
    - Memory speed supported by the specific processor (2133 MHz, 2400 MHz, 2666 MHz, or 2933 MHz).
    - Memory speed for the selected quantity of DIMMs per channel:
        - One DIMM per channel (1 DPC): 2933 MHz.
        - Two DIMMs per channel (2 DPC)
            - Performance+ DIMMs: 2933 MHz.
            - Other supported DIMMs: 2666 MHz.

  Note: Maximum memory speed can be achieved when Max performance mode is enabled in UEFI.
- Mixing different types of memory (RDIMMs, LRDIMMs, 3DS RDIMMs) is not supported.
- All Performance+ DIMMs in the server must be of the same type, rank, and capacity (the same part number or feature code) to operate at 2933 MHz in the configurations with two DIMMs per channel. Performance+ DIMMs cannot be mixed with other DIMMs.
- Mixing RDIMMs of different ranks (single- or dual-rank), DRAM chip types (x4 or x8), speeds (2666 MHz or 2933 MHz), and capacities (8 GB, 16 GB, 32 GB, or 64 GB) is supported in the independent channel mode (the default operational mode) (excluding Performance+ RDIMMs).
- Mixing 3DS RDIMMs of different speeds (2666 MHz or 2933 MHz) and capacities (64 GB or 128 GB) is supported in the independent channel mode (excluding Performance+ 3DS RDIMMs).
- The 128 GB Performance+ 2933 MHz 3DS RDIMMs (feature code B5N9) running at 2933 MHz with two DIMMs per channel are supported at the ambient temperature of up to 30 °C (86 °F).
- The maximum quantity of DIMMs supported is reduced by the quantity of DC Persistent Memory Modules used in the configuration.
- Server configurations with more than 1 TB of memory capacity per socket (including DCPMMs and RDIMMs or 3DS RDIMMs) require processors that support up to 4.5 TB (L-suffix) per socket.
- For server configurations with memory protection, the following rules apply:
    - Single Device Data Correction (SDDC) works only in the independent channel mode and supports only x4-based memory DIMMs.
    - Adaptive Double Device Data Correction (ADDDC) works with x4-based memory DIMMs and requires two DIMM ranks per channel, Intel Xeon Gold or Platinum processors, and the Closed Page memory access mode.
    - If memory mirroring is used, then DIMMs must be installed in quantities of 2, 4, or 8 per processor for mirroring across two memory channels, or in quantities of 3, 6, 9, or 12 per processor for mirroring across three memory channels. Mixing two- and three-channel mirroring in the server is allowed (one processor uses two-channel mirroring, and another processor uses three-channel mirroring). All DIMMs in the server must be identical in type and size.
    - If memory rank sparing is used, then a minimum of two ranks must be installed per populated channel (a least one dual-rank or quad-rank DIMM, or two single-rank DIMMs). In rank sparing mode, one rank in each populated channel is reserved as spare memory for other ranks on the same channel. All DIMMs in the server must be identical in type and size.
    - SDDC, memory mirroring, and memory rank sparing modes are mutually exclusive. Only one operational memory mode can be enabled on the server.
    - In the configurations with DCPMMs, memory mirroring is supported only in the App Direct mode (other DCPMM modes do not support memory mirroring) and applies only to the RDIMMs or 3DS RDIMMs (DCPMMs are not mirrored). Memory sparing is not supported in the configurations with DCPMMs.

Persistent memory

Intel Optane persistent memory is an innovative technology that delivers a unique combination of affordable large memory capacity and persistence (non-volatility). The persistent memory technology can help boost the performance of data-intensive applications, such as in-memory analytics, databases, content delivery networks, and high performance computing (HPC), as well as deliver consistent service levels at scale with higher virtual machine and container density.

2nd Gen processors only: Persistent Memory is only supported with 2nd Generation Intel Xeon SP processors. Not supported with 1st Generation processors.

The SR630 server supports up to six TruDDR4 DC Persistent Memory Modules (DCPMMs) when one processor is installed and up to 12 DCPMMs when two processors are installed (up to one DCPMM per processor's memory channel) for a total of up to 6 TB of persistent memory capacity. The DCPMMs are installed in the same memory DIMM slots on the system board that are used for installing RDIMMs or 3DS RDIMMs.

The DCPMMs support the following modes of operation:

- Memory Mode
  Memory Mode seamlessly brings large memory capacity at affordable cost points to legacy applications. In this mode, DCPMMs provide volatile memory that behaves much like traditional RDIMMs or 3DS RDIMMs (the data will not be saved in case of a power loss) and is transparent to the operating system and applications. DCPMMs provide memory capacity and RDIMMs or 3DS RDIMMs provide cache memory that is managed by the processor's memory controller. The total memory capacity that is seen by the operating system is the capacity of the DCPMMs; the capacity of the RDIMMs or 3DS RDIMMs is hidden and does not appear as a memory resource in the operating system. This mode is considered particularly suited for virtualized database deployments and big-data analytics applications.
- App Direct Mode
  App Direct Mode brings persistency to the data and structures (the data will be saved in case of a power loss). This mode requires operating system and application awareness of two types of system memory: Persistent (DCPMMs) and DRAM (RDIMMs or 3DS RDIMMs). The total memory capacity that is seen by the operating system includes the capacity of the DCPMMs and RDIMMs or 3DS RDIMMs. This mode is considered particularly suited for in-memory databases, in-memory analytics frameworks, and ultrafast storage applications.
- Mixed Memory Mode
  Mixed Memory Mode is a combination of Memory Mode and App Direct Mode, where a portion of the capacity of the DCPMMs is used for the Memory Mode operations, and the remaining capacity of the DCPMMs is used for the App Direct Mode operations.

The following memory protection technologies are supported by the DCPMM's onboard memory controllers:

- ECC
- SDDC
- DDDC
- Patrol scrubbing
- Demand scrubbing

The following table lists DCPMM options available for the SR630 server.

Table 17. DCPMM options

Part number	Feature code	Description	Maximum quantity*
4ZC7A15110	B4LV	ThinkSystem 128GB TruDDR4 2666MHz (1.2V) Intel Optane DC Persistent Memory	6 / 12
4ZC7A15111	B4LW	ThinkSystem 256GB TruDDR4 2666MHz (1.2V) Intel Optane DC Persistent Memory	6 / 12
4ZC7A15112	B4LX	ThinkSystem 512GB TruDDR4 2666MHz (1.2V) Intel Optane DC Persistent Memory	6 / 12

* The maximum quantity shown is with one processor / two processors.

Configuration notes:

- All DCPMMs in the server must be of the same capacity (the same part number or feature code).
- The RDIMMs or 3DS RDIMMs are required in the configurations with DCPMMs, and all RDIMMs or 3DS RDIMMs must be of the same type, rank, and capacity (the same part number or feature code).
- The DCPMMs cannot be mixed with the 8GB TruDDR4 2933 MHz RDIMM (4ZC7A08706).
- For Mixed Memory Mode, the volatile (Memory) portion of the total capacity of DCPMMs is configured in increments of 32 GB multiplied by the number of DCPMMs in the server, and the remaining capacity is allocated to the persistent (App Direct) portion. The ratio of the total capacity of RDIMMs or 3DS RDIMMs to the total capacity of the volatile portion of DCPMMs should be between 1:2 and 1:16.
- Server configurations with more than 1 TB of memory capacity per socket (including DCPMMs and RDIMMs or 3DS RDIMMs) require processors that support up to 4.5 TB (L-suffix) per socket.
- The DCPMMs are *not* supported in the following configurations:
  - Processors of 200 W or 205 W TDP, or Gold 6240Y, 6244, 6246, or 6252N processors installed
  - GPU adapters installed
  - Mellanox Innova-2 adapters installed

For more information, including supported combinations of DIMMs and Persistent Memory, refer to the Intel Optane Persistent Memory 100 Series Product Guide:
https://lenovopress.com/lp1066-intel-optane-persistent-memory-100-series

Internal storage

The SR630 server supports the following internal drive bay configurations:

1. Up to 6 SAS/SATA hot-swap drive bays: 4x 3.5" (front) + 2x 2.5" (rear)
2. 4 LFF AnyBay hot-swap drive bays: 4x 3.5" (front)
3. Up to 10 SFF SAS/SATA hot-swap drive bays: 8x 2.5" (front) + 2x 2.5" (rear)
4. Up to 12 SFF hot-swap drive bays:
   6x 2.5" SAS/SATA & 4x 2.5" AnyBay (front) + 2x 2.5" SAS/SATA (rear)
5. 10 SFF U.2 NVMe PCIe SSD hot-swap drive bays

In addition, the SR630 server models can be configured with one or two internal M.2 SATA SSDs. The server also supports configurations without drive bays.

The following figure shows the internal drive bay configurations.

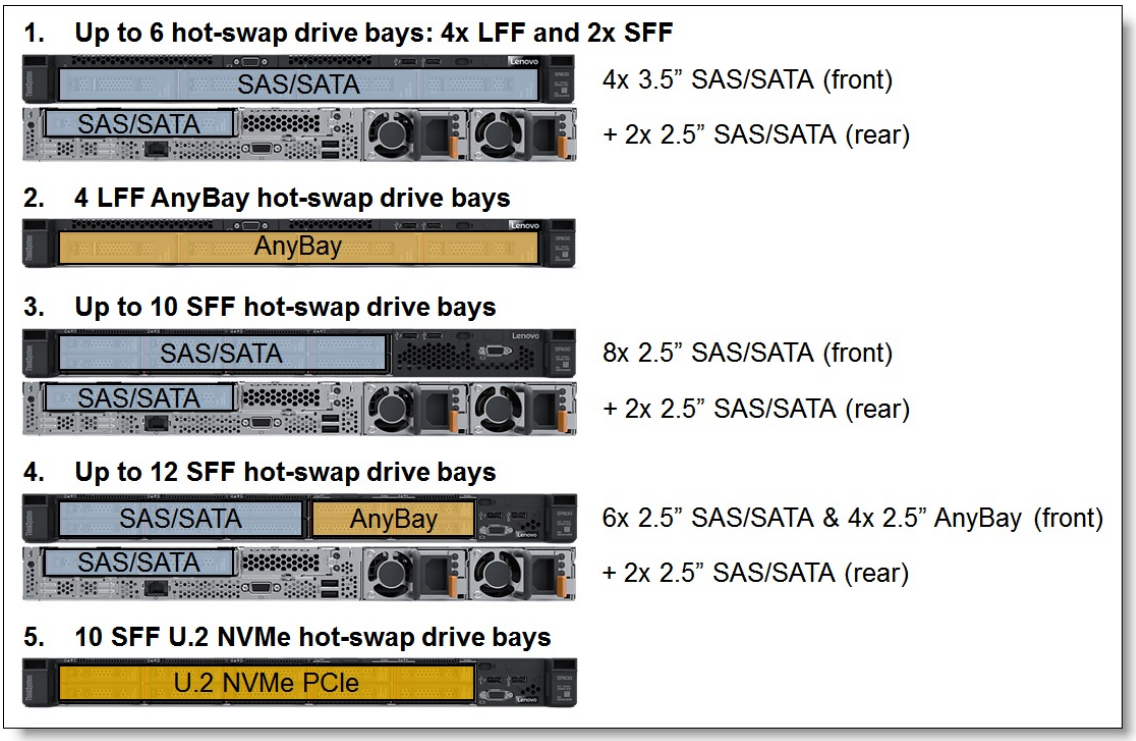


Figure 7. Internal drive bay configurations

In this section:

- Backplanes
- Supported drive bay combinations
- Field upgrades
- RAID flash power module (supercap) support
- M.2 drives
- SED encryption key management with ISKLM

Backplanes

The following table lists the internal storage options for the server.

Table 18. Internal storage options

Part number	Feature code	Description	Maximum quantity
Factory-installed backplane kits			
None*	AUW8	ThinkSystem 1U 3.5" SATA/SAS 4-Bay Backplane	1
None*	B0WJ	ThinkSystem SR630 3.5" AnyBay 4-Bay Backplane	1
None*	AUWB	ThinkSystem 1U 2.5" SATA/SAS 8-Bay Backplane	1
None*	AUW9	ThinkSystem 1U 2.5" AnyBay 10-Bay Backplane	1
None*	AVKF	ThinkSystem SR630 2.5" U.2 10-Bay Backplane Kit	1
4XH7A80454	BMXG	ThinkSystem SR630 Rear HDD/SSD Kit v2	1
7XH7A06252	AUWJ	ThinkSystem SR630 Rear HDD Kit	1

* For field upgrades, see the Field upgrades section

Configuration notes:

- Lenovo AnyBay allows a choice of drive interface types in the same drive bay: SAS drives, SATA drives, or U.2 NVMe PCIe drives.
- U.2 NVMe PCIe SSDs in the AnyBay drive bays require either the second processor (enables the onboard NVMe controller) or the 1610-4P NVMe Switch Adapter to be installed. The 1610-4P NVMe Switch Adapter connected to the AnyBay drive bays is supported only in the configurations with one processor.
- Models with 10x 2.5-inch drive bays (6x SAS/SATA + 4x AnyBay) and an 8-port SAS RAID controller or HBA support only NVMe drives in the AnyBay drive bays.

Supported drive bay combinations

The following tables list supported internal storage configurations with the SAS/SATA and AnyBay backplanes.

Table 19. Internal storage configurations: 3.5-inch front drive bays

Drive bay configuration	Backplane kit type and quantity			Storage controller type and quantity*
	4x 3.5" SATA/SAS	4x 3.5" AnyBay	2x 2.5" Rear HDD	
4x 3.5" chassis (Feature code AUW2)				
4x 3.5-in. SAS/SATA hot-swap (front)	1	0	0	1x RAID 8i or HBA 8i (4)
4x 3.5-in. SAS/SATA hot-swap (front) + 2x 2.5-in. SAS/SATA hot-swap (rear)	1	0	1	1x RAID 8i or HBA 8i (6)
4x 3.5-in. AnyBay hot-swap (front)	0	1	0	1x RAID 8i or HBA 8i (4) + NVMe (4)**

* The number in brackets (x) specifies the quantity of drive bays connected to each of the controllers.
** The 1610-4P NVMe Switch Adapter in the configurations with one processor, or the onboard NVMe controller in the configurations with two processors.

Table 20. Internal storage configurations: 2.5-inch front drive bays

Drive bay configuration	Backplane kit type and quantity				Storage controller type and quantity*
	8x 2.5" SATA/SAS	10x 2.5" AnyBay	10x U.2 NVMe	2x 2.5" Rear HDD	
8x 2.5" chassis (Feature code AUW0)					
8x 2.5-in. SAS/SATA hot-swap (front)	1	0	0	0	1x RAID 8i or HBA 8i (8)
					1x RAID 16i (8)
8x 2.5-in. SAS/SATA hot-swap (front) + 2x 2.5-in. SAS/SATA hot-swap (rear)	1	0	0	1	1x RAID 16i or HBA 16i (10)
10x 2.5" chassis (Feature code AUW1)					
4x 2.5-in. AnyBay (NVMe only) hot-swap (front)	0	1	0	0	Onboard NVMe (4)**
6x 2.5-in. SAS/SATA + 4x 2.5-in. AnyBay hot-swap (front)	0	1	0	0	1x RAID 16i or HBA 16i (10) + Onboard NVMe (4)**
6x 2.5-in. SAS/SATA + 4x 2.5-in. AnyBay (NVMe only) hot-swap (front)	0	1	0	0	1x RAID 8i or HBA 8i (6) + Onboard NVMe (4)**
					1x RAID 8i or HBA 8i (6) + 1x 1610-4P NVMe (4)^
					1x RAID 16i or HBA 16i (6) + 1x 1610-4P NVMe (4)^
6x 2.5-in. SAS/SATA + 4x 2.5-in. AnyBay hot-swap (front) + 2x 2.5-in. SAS/SATA hot-swap (rear)	0	1	0	1	1x RAID 16i or HBA 16i (12) + Onboard NVMe (4)**
10x 2.5-in. U.2 NVMe (8x NVMe drives only) hot-swap (front)	0	0	1	0	Onboard NVMe (4)** + 1x 1610-4P/810-4P (4)
					Onboard NVMe (4)** + 2x 810-4P (2+2)
10x 2.5-in. U.2 NVMe hot-swap (front)	0	0	1	0	Onboard NVMe (4)** + 1x 1610-4P (4) + 1x 810-4P (2)
					Onboard NVMe (4)** + 2x 1610-4P (4+2)

* The number in brackets (x) specifies the quantity of drive bays connected to each of the controllers.
** The onboard NVMe controller requires the second processors to be installed.
^ The 1610-4P NVMe Switch Adapter connected to the AnyBay drive bays is supported only in the configurations with one processor. CTO only, not available as a field upgrade.

Field upgrades

The following table lists the backplane options that can be installed as field upgrades.

Use with X40 adapters: These backplane kits in the table below include SAS/SATA cables for use with the onboard SATA controller or with RAID 930, 730, 530 adapters and 430 HBAs (collectively called X30 adapters). If you are adding or upgrading to RAID 940 adapters or 440 HBAs (collectively called X40 adapters), you will need to *also* order an X40 cable kit. See the Cable kits for 440 HBAs and RAID 940 adapters section for details.

Table 21. Field Upgrades

Part number	Description	Maximum quantity
Backplane kit field upgrade options		
4XH7A80455	ThinkSystem SR530/SR630 2.5" SATA/SAS 8-Bay Backplane Kit v2	1
7XH7A05896	ThinkSystem SR530/SR630 2.5" SATA/SAS 8-Bay Backplane Kit	1
4XH7A80452	ThinkSystem SR570/SR630 2.5" AnyBay 10-Bay Backplane Upgrade Kit v2	1
4XH7A08768	ThinkSystem SR570/SR630 2.5" AnyBay 10-Bay Backplane Upgrade Kit	1
4XH7A80451	ThinkSystem SR630 2.5" U.2 10-Bay Backplane Kit v2	1
4XH7A08766	ThinkSystem SR630 2.5" U.2 10-Bay Backplane Kit	1
4XH7A80454	ThinkSystem SR630 Rear HDD/SSD Kit v2	1
7XH7A06252	ThinkSystem SR630 Rear HDD/SSD Kit	1

Configuration notes:

- The backplane upgrade kits include drive backplanes and required signal cables, power cables, and drive bay fillers; storage controllers are not included.
- Models without any drive bays that are based on the 8x 2.5" chassis (feature code AUW0) support adding drive bays by using the 2.5" 8-drive backplane kit (7XH7A05896 or 4XH7A80455).
- Models without any drive bays that are based on the 10x 2.5" chassis (feature code AUW1) support adding drive bays by using the 2.5" 10-bay AnyBay backplane kit (4XH7A08768 or 4XH7A80452) or 2.5" U.2 10-bay backplane kit (4XH7A08766 or 4XH7A80451).
- Models with 10x 2.5-inch drive bays (6x SAS/SATA + 4x AnyBay) can be upgraded to all NVMe configuration by using the 2.5" U.2 10-Bay Backplane Kit (4XH7A08766 or 4XH7A80451). The kit includes a backplane, cables, and drive fillers. NVMe Switch Adapters are *not* included in the kit; two NVMe Switch Adapters are required for the 10x 2.5" U.2 NVMe configuration.
- The 2.5" Rear HDD Kit is installed in place of the PCIe Riser Card 1; PCIe slots 1 and 2 are not present.
- The 2.5" Rear HDD Kit is supported only with the processors of up to 125 W TDP.
- U.2 NVMe PCIe SSDs in the AnyBay drive bays require either the second processor (enables the onboard NVMe controller) or the 1610-4P NVMe Switch Adapter to be installed. The 1610-4P NVMe Switch Adapter connected to the AnyBay drive bays is supported only in the configurations with one processor.
- Models with 10x 2.5-inch drive bays (6x SAS/SATA + 4x AnyBay) and an 8-port SAS RAID controller or HBA support only NVMe drives in the AnyBay drive bays.

Cable kits for 440 HBAs and RAID 940 adapters

The backplane kits listed in the preceding table include cables for use with the onboard SATA controller or with RAID 930, 730, 530 adapters and 430 HBAs (collectively called X30 adapters). If you wish to use the backplane kits with RAID 940 adapters or 440 HBAs (collectively called X40 adapters), then you will also need to order an additional X40 cable kit to use instead of the cables in the backplane kit.

Tip: When adding an X40 adapter, you will order both the backplane kit and the relevant X40 cable kit, however the SAS/SATA data cable(s) in the backplane kit will not be used.

Table 22. Cable kits for 440 HBAs and RAID 940 adapters

Backplane kits with X30 cables		X40 cable kits also needed	
4XH7A80455	ThinkSystem SR530/SR630 2.5" SATA/SAS 8-Bay Backplane Kit v2	4XH7A61096	ThinkSystem SR530/SR570/SR630 2.5" SAS/SATA 8-Bay X40 RAID Cable Kit
7XH7A05896	ThinkSystem SR530/SR630 2.5" SATA/SAS 8-Bay Backplane Kit	4XH7A61096	ThinkSystem SR530/SR570/SR630 2.5" SAS/SATA 8-Bay X40 RAID Cable Kit
4XH7A80452	ThinkSystem SR570/SR630 2.5" AnyBay 10-Bay Backplane Upgrade Kit v2	4XH7A61101	ThinkSystem SR570/SR630 2.5" AnyBay 10-Bay X40 RAID Cable Kit
4XH7A08768	ThinkSystem SR570/SR630 2.5" AnyBay 10-Bay Backplane Upgrade Kit	4XH7A61101	ThinkSystem SR570/SR630 2.5" AnyBay 10-Bay X40 RAID Cable Kit
4XH7A80451	ThinkSystem SR630 2.5" U.2 10-Bay Backplane Kit v2	None	Not needed - NVMe only
4XH7A08766	ThinkSystem SR630 2.5" U.2 10-Bay Backplane Kit	None	Not needed - NVMe only
4XH7A80454	ThinkSystem SR630 Rear HDD/SSD Kit v2	4XH7A61110	ThinkSystem SR590/SR630/SR650 SAS/SATA 2-Bay Rear BP X40 RAID Cable Kit
7XH7A06252	ThinkSystem SR630 Rear HDD/SSD Kit	4XH7A61110	ThinkSystem SR590/SR630/SR650 SAS/SATA 2-Bay Rear BP X40 RAID Cable Kit

RAID flash power module (supercap) support

ThinkSystem RAID 930, 940, 9350 and 730-8i 2GB adapters all include a RAID flash power module (supercap). If you plan to install one of these adapters as a field upgrade, then may also need to order a Supercap Holder kit or a replacement server air baffle to mount the power module. For CTO orders, the components in the installation kit are automatically derived when you select the RAID adapter.

The server supports one or two RAID adapters (internal or external) with supercaps, however there are three possible locations for the supercaps to be mounted. The locations of the supercaps depends on the front drive bays (2.5-inch or 3.5-inch) and the size of the processor heatsinks (which in turn determines if the air baffle is installed). Details are summarized in the following table. Location references are shown in the figure below.

Note: Some configurations do not support a supercap-attached RAID adapter as shown in the table.

Table 23. Supercap support

Front drive configuration	Processor heatsinks	Number of adapters & supercaps	Location of supercaps
2.5-inch	Standard	2	1. Front of server behind operator panel ❸ (if present) 2. Mounted on air baffle ❶ 3. Mounted on air baffle ❷ (if no holder installed in ❸)
	High Performance (no air baffle installed)	1	1. Front of server behind operator panel ❸
3.5-inch	Standard	2	1. Mounted on air baffle ❶ 2. Mounted on air baffle ❷
	High Performance (no air baffle installed)	0	Not supported

The locations where supercaps are installed is shown in the following figure.

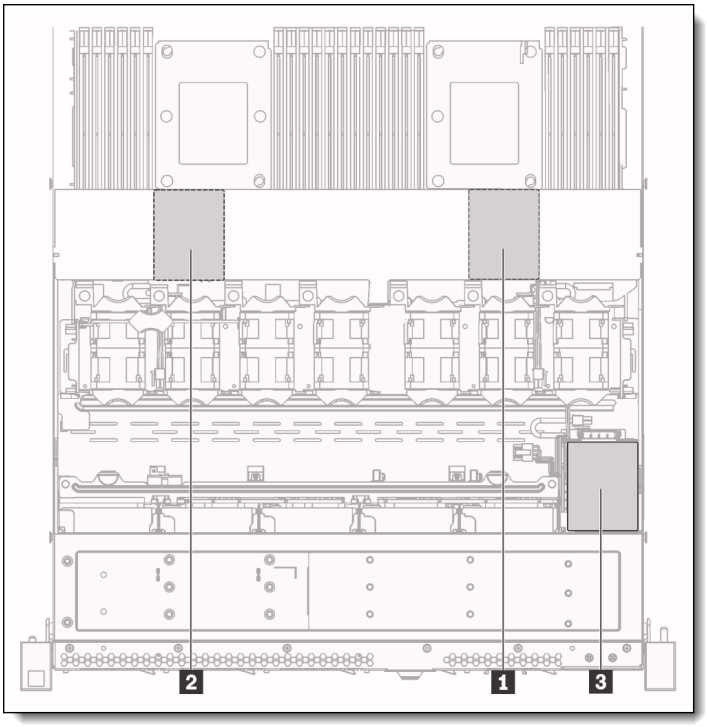


Figure 8. Location of the supercaps in the SR630

If the server does not include a supercap holder in position ❸, order the holder using part number 4M17A61348.

If one of the following RAID adapters is purchased as a field upgrade, the system air duct (air baffle) will need to be replaced with a new one to accommodate the supercap that ships with the adapter. The original air duct won't physically accomodate the supercap of these adapters.

- ThinkSystem RAID 9350-8i 2GB Flash PCIe 12Gb Adapter, 4Y37A72483
- ThinkSystem RAID 9350-16i 4GB Flash PCIe 12Gb Adapter, 4Y37A72485

Tip: If your server was built before December 2021 or was configured with base feature codes AUW0, AUW1 or AUW2, then you have original air duct, part SM27A03474, and you will need to order a replacement air duct. If you server was configured with newer "v2" base feature codes BNPN, BNPQ or BNPP, then you have the newer air duct, part SM27A99423 and no air duct replacement is necessary.

Ordering information for the Supercap Holder Kit and the replacement air duct are listed in the following table.

Table 24. Supercap installation field upgrades

Part number	Feature code	Description	Maximum supported
4M17A61304	AURN	ThinkSystem SR630/SR630 V2 Supercap Holder Kit (For use in position ❸ at the front of the server)	1
4M17A61348	BNZ0	ThinkSystem SR630 Air Duct Kit v2 (Enables ❶ and ❷ for use with the 9350-8i or 9350-16i adapters)	1

M.2 drives

The server supports one or two M.2 form-factor SATA drives for use as an operating system boot solution. With two M.2 drives configured, the drives are configured by default as a RAID-1 mirrored pair for redundancy.

The M.2 drives install into an M.2 adapter which in turn is installed in a dedicated slot on the system board. See the internal view of the server in the Components and connectors section for the location of the M.2 slot.

There are two M.2 adapters supported, as listed in the following table.

Table 25. M.2 components

Part number	Feature code	Description	Maximum supported
7Y37A01092	AUMU	ThinkSystem M.2 Enablement Kit (contains the Single M.2 Boot Adapter; supports 1 drive)	1
7Y37A01093	AUMV	ThinkSystem M.2 with Mirroring Enablement Kit (contains the Dual M.2 Boot Adapter, supports 1 or 2 drives)	1

Supported drives are listed in the Internal drive options section.

For details about M.2 components, see the *ThinkSystem M.2 Drives and M.2 Adapters* product guide: https://lenovopress.com/lp0769-thinksystem-m2-drives-adapters

SED encryption key management with ISKLM

The server supports self-encrypting drives (SEDs) as listed in the Internal drive options section. To effectively manage a large deployment of these drives in Lenovo servers, IBM Security Key Lifecycle Manager (SKLM) offers a centralized key management solution. A Lenovo Feature on Demand (FoD) upgrade is used to enable this SKLM support in the management processor of the server.

The following table lists the part numbers and feature codes for the upgrades.

Table 26. FoD upgrades for SKLM support

Part number	Feature code	Description
Security Key Lifecycle Manager - FoD (United States, Canada, Asia Pacific, and Japan)		
00D9998	A5U1	SKLM for System x/ThinkSystem w/SEDs - FoD per Install with 1 year S&S
00D9999	AS6C	SKLM for System x/ThinkSystem w/SEDs - FoD per Install with 3 year S&S
Security Key Lifecycle Manager - FoD (Latin America, Europe, Middle East, and Africa)		
00FP648	A5U1	SKLM for System x/ThinkSystem w/SEDs - FoD per Install with 1 year S&S
00FP649	AS6C	SKLM for System x/ThinkSystem w/SEDs - FoD per Install with 3 year S&S

The IBM Security Key Lifecycle Manager software is available from Lenovo using the ordering information listed in the following table.

Table 27. IBM Security Key Lifecycle Manager licenses

Part number	Description
7S0A007FWW	IBM Security Key Lifecycle Manager Basic Edition Install License + SW Subscription & Support 12 Months
7S0A007HWW	IBM Security Key Lifecycle Manager For Raw Decimal Terabyte Storage Resource Value Unit License + SW Subscription & Support 12 Months
7S0A007KWW	IBM Security Key Lifecycle Manager For Raw Decimal Petabyte Storage Resource Value Unit License + SW Subscription & Support 12 Months
7S0A007MWW	IBM Security Key Lifecycle Manager For Usable Decimal Terabyte Storage Resource Value Unit License + SW Subscription & Support 12 Months
7S0A007PWW	IBM Security Key Lifecycle Manager For Usable Decimal Petabyte Storage Resource Value Unit License + SW Subscription & Support 12 Months

Controllers for internal storage

The following table lists the storage controllers and options for internal storage of the SR630 server.

Table 28. RAID controllers and HBAs for internal storage

Part number	Feature code	Description	Maximum quantity	I/O slots supported
12 Gb SAS/SATA RAID controllers - 8-port adapters				
7Y37A01082	AUNG	ThinkSystem RAID 530-8i PCIe 12Gb Adapter	1	4, 1
4Y37A78834	BMFT	ThinkSystem RAID 540-8i PCIe Gen4 12Gb Adapter	1	1
4Y37A72482	BJHK	ThinkSystem RAID 5350-8i PCIe 12Gb Adapter	1	1
7Y37A01083	AUNH	ThinkSystem RAID 730-8i 1GB Cache PCIe 12Gb Adapter	1	4, 1
4Y37A09722	B4RQ	ThinkSystem RAID 730-8i 2GB Flash PCIe 12Gb Adapter	1	4, 1
7Y37A01084	AUNJ	ThinkSystem RAID 930-8i 2GB Flash PCIe 12Gb Adapter	1	4, 1
4Y37A72483†	BJHL†	ThinkSystem RAID 9350-8i 2GB Flash PCIe 12Gb Adapter	1	1
4Y37A09728	B8NY	ThinkSystem RAID 940-8i 4GB Flash PCIe Gen4 12Gb Adapter	1	1
12 Gb SAS/SATA RAID controllers - 16-port adapters				
4Y37A09727	B6CE	ThinkSystem RAID 530-16i PCIe 12Gb Adapter	1	4, 1
7Y37A01085	AUNK	ThinkSystem RAID 930-16i 4GB Flash PCIe 12Gb Adapter	1	4, 1
4Y37A72485†	BJHN†	ThinkSystem RAID 9350-16i 4GB Flash PCIe 12Gb Adapter	1	1
4Y37A09721	B31E	ThinkSystem RAID 930-16i 8GB Flash PCIe 12Gb Adapter	1	4, 1
4Y37A78600	BM35	ThinkSystem RAID 940-16i 4GB Flash PCIe Gen4 12Gb Adapter	1	1
4Y37A09730	B8NZ	ThinkSystem RAID 940-16i 8GB Flash PCIe Gen4 12Gb Adapter	1	1
12 Gb SAS/SATA non-RAID HBAs				
7Y37A01088	AUNL	ThinkSystem 430-8i SAS/SATA 12Gb HBA	1	4, 1
4Y37A72480	BJHH	ThinkSystem 4350-8i SAS/SATA 12Gb HBA	1	1
7Y37A01089	AUNM	ThinkSystem 430-16i SAS/SATA 12Gb HBA	1	4, 1
4Y37A72481	BJHJ	ThinkSystem 4350-16i SAS/SATA 12Gb HBA	1	1
4Y37A78601	BM51	ThinkSystem 440-8i SAS/SATA PCIe Gen4 12Gb HBA	1	1

Part number	Feature code	Description	Maximum quantity	I/O slots supported
4Y37A78602	BM50	ThinkSystem 440-16i SAS/SATA PCIe Gen4 12Gb HBA	1	1
NVMe PCIe interfaces				
None	B9X7	Intel VROC (VMD NVMe RAID) Intel SSD Only (Standard)	1	Not applicable
4L47A39164	B96G	Intel VROC (VMD NVMe RAID) Premium (license upgrade - to enable RAID support for non-Intel NVMe SSDs)	1	Not applicable
7Y37A01081	AUV2	ThinkSystem 1610-4P NVMe Switch Adapter	2	2, 3*
4Y37A09719	B22D	ThinkSystem 810-4P NVMe Switch Adapter	2	4, 1

* Requires a PCIe x16 riser card.
† Field upgrades to add this adapter may also require a replacement system air duct. See the RAID flash power module (supercap) support section for details. CTO orders that include this adapter must have base BNPS or BNPR selected. See the Models section for information.

For a comparison of the functions of the supported storage adapters, see the ThinkSystem RAID Adapter and HBA Reference:
https://lenovopress.com/lp1288-thinksystem-raid-adapter-and-hba-reference#sr630-support=SR630

Configuration notes:

- Low profile SAS RAID controllers and HBAs for internal storage are supported in the dedicated PCIe x8 slot 4 on the system board. If a full-height GPU adapter is installed, the internal slot for a storage controller cannot be used; the storage controller should be installed in the PCIe slot 1.
- The 1610-4P NVMe Switch Adapters are supported in the low-profile PCIe x16 slots supplied by the PCIe Riser Cards 1 and 2.
- The 810-4P NVMe Switch Adapters are supported in the dedicated PCIe x8 slot 4 on the system board and a low-profile PCIe x8 slot supplied by the PCIe Riser Card 1.
- The total quantity of the 1610-4P and 810-4P NVMe Switch Adapters in the server must not exceed 2.
- Configurations using onboard NVMe and NVMe switch adapters support RAID using Intel VROC NVMe RAID as described in the Intel VROC onboard SATA and NVMe RAID section.
- The onboard NVMe interface provides 4x PCIe 3.0 x4 ports for connectivity to U.2 NVMe PCIe SSDs in the AnyBay drive bays, and it requires the second processor to be installed.
- In the configurations with one processor, the 1610-4P NVMe Switch Adapter provides 4x PCIe 3.0 x4 ports for connectivity to U.2 NVMe PCIe SSDs in the AnyBay drive bays, and it is supported in the PCIe x16 slot 2 supplied by the x8/x16 Riser Card 1.
- In the configurations without GPU installed, the total quantity of the RAID 730-8i 2GB, RAID 930-8i, RAID 930-16i, and RAID 930-8e controllers in the server must not exceed 2 (up to 2 supercapacitors can be mounted in the server).
- In the configurations with GPU installed, the total quantity of the RAID 730-8i 2GB, RAID 930-8i, RAID 930-16i, and RAID 930-8e controllers in the server must not exceed 1 (the storage controller can be installed only in the PCIe slot 1; the PCIe slots 3 and 4 cannot be used if the GPU adapter is installed).
- In the configurations with the RAID 730-8i 2GB controller, the RAID 930-8e controller cannot be used.

For 10x U.2 NVMe configurations, the onboard NVMe interface and two NVMe Switch Adapters provide connectivity with PCIe 3.0 x4 links to each of the U.2 NVMe PCIe SSDs. The 1610-4P NVMe Switch Adapter has a PCIe 3.0 x16 host interface, and the 810-4P NVMe Switch Adapter has a PCIe 3.0 x8 host interface.

Intel VROC onboard SATA and NVMe RAID

Intel VROC (Virtual RAID on CPU) is a feature of the Intel processor that enables RAID support. There are two separate functions of VROC:

- Intel VROC SATA RAID, formerly known as Intel RSTe
- Intel VROC NVMe RAID

VROC SATA RAID (RSTe) is available and supported with all SATA drives , both SATA SSDs and SATA HDDs. It offers a 6 Gb/s connection to each drive and on the SR630 implements RAID levels 0, 1, 5, and 10 Hot-spare functionality is also supported.

VROC NVMe RAID offers RAID support for any NVMe drives directly connected to the ports on the server's system board or via adapters such as NVMe retimers or NVMe switch adapters. On the SR630, it implements RAID levels 0, 1, 5, and 10. Hot-spare functionality is also supported.

Performance tip: For best performance with VROC NVMe RAID, the drives in an array should all be connected to the same processor. Spanning processors is possible however performance will be unpredictable and should be evaluated based on your workload.

By default, VROC NVMe RAID support is limited to use with only Intel-branded NVMe drives (feature B9X7). If you wish to enable RAID support for non-Intel NVMe SSDs, select the VROC Premium license using the ordering information in the following table. VROC Premium is fulfilled as a Feature on Demand (FoD) license and is activated via the XCC management processor user interface.

Table 29. VROC upgrade

Part number	Feature code	Description
4L47A39164	B96G	Intel VROC (VMD NVMe RAID) Premium

VROC Premium is only needed for non-Intel NVMe drives in a RAID configuration. You do not need the VROC Premium license upgrade under any of the following conditions:

- If you have SATA drives connected to the onboard SATA ports, you do not need VROC Premium
- If you have Intel NVMe drives connected to the onboard NVMe ports, you do not need VROC Premium
- If you have non-Intel NVMe drives connected to the onboard NVMe ports, but you don’t want RAID support, you do not need VROC Premium

Virtualization support: Virtualization support for Intel VROC is as follows:

- **VROC SATA RAID (RSTe)**: VROC SATA RAID is not supported by virtualization hypervisors such as ESXi, KVM, Xen, and Hyper-V. Virtualization is only supported on the onboard SATA ports in AHCI (non-RAID) mode.
- **VROC (VMD) NVMe RAID**: VROC (VMD) NVMe RAID is supported by KVM, Xen, and Hyper-V. ESXi is currently not supported. Windows and Linux OSes support VROC RAID NVMe, both for host boot functions and for guest OS function, and RAID-0, 1, 5, and 10 are supported.

No ESXi support: VROC NVMe RAID on the SR650 and SR630 is based on VROC 7.5 firmware and is currently not generally supported with VMware ESXi. Only Windows, RHEL and SLES are supported. ESXi support is currently only available using an older VROC version and requires a factory build via a Special Bid order. Support of ESXi is planned for 1Q/2022.

Internal drive options

The following tables list the drive options for internal storage of the server.

2.5-inch hot-swap drives:

- 2.5-inch hot-swap 12 Gb SAS HDDs
- 2.5-inch hot-swap 6 Gb SATA HDDs
- 2.5-inch hot-swap 12 Gb SAS SSDs
- 2.5-inch hot-swap 6 Gb SATA SSDs
- 2.5-inch hot-swap PCIe 4.0 NVMe SSDs
- 2.5-inch hot-swap PCIe 3.0 NVMe SSDs

3.5-inch hot-swap drives:

- 3.5-inch hot-swap 12 Gb SAS HDDs
- 3.5-inch hot-swap 6 Gb SATA HDDs
- 3.5-inch hot-swap 12 Gb SAS SSDs
- 3.5-inch hot-swap 6 Gb SATA SSDs
- 3.5-inch hot-swap PCIe 4.0 NVMe SSDs
- 3.5-inch hot-swap PCIe 3.0 NVMe SSDs

M.2 drives:

- M.2 SATA drives

M.2 drive support: The use of M.2 drives requires an additional adapter as described in the M.2 drives subsection.

PCIe 4.0 NVMe drive support: When installed in this server, PCIe 4.0 NVMe drives will operate at PCIe 3.0 speeds.

Table 30. 2.5-inch hot-swap 12 Gb SAS HDDs

Part number	Feature	Description	Maximum supported
2.5-inch hot-swap HDDs - 12 Gb SAS 10K			
7XB7A00024	AULY	ThinkSystem 2.5" 300GB 10K SAS 12Gb Hot Swap 512n HDD	12
7XB7A00025	AULZ	ThinkSystem 2.5" 600GB 10K SAS 12Gb Hot Swap 512n HDD	12
7XB7A00026	AUM0	ThinkSystem 2.5" 900GB 10K SAS 12Gb Hot Swap 512n HDD	12
7XB7A00027	AUM1	ThinkSystem 2.5" 1.2TB 10K SAS 12Gb Hot Swap 512n HDD	12
7XB7A00028	AUM2	ThinkSystem 2.5" 1.8TB 10K SAS 12Gb Hot Swap 512e HDD	12
7XB7A00069	B0YS	ThinkSystem 2.5" 2.4TB 10K SAS 12Gb Hot Swap 512e HDD	12
2.5-inch hot-swap HDDs - 12 Gb SAS 15K			
7XB7A00021	AULV	ThinkSystem 2.5" 300GB 15K SAS 12Gb Hot Swap 512n HDD	12
7XB7A00022	AULW	ThinkSystem 2.5" 600GB 15K SAS 12Gb Hot Swap 512n HDD	12
7XB7A00023	AULX	ThinkSystem 2.5" 900GB 15K SAS 12Gb Hot Swap 512e HDD	12
2.5-inch hot-swap HDDs - 12 Gb NL SAS			
7XB7A00034	AUM6	ThinkSystem 2.5" 1TB 7.2K SAS 12Gb Hot Swap 512n HDD	12
7XB7A00035	AUM7	ThinkSystem 2.5" 2TB 7.2K SAS 12Gb Hot Swap 512n HDD	12
2.5-inch hot-swap SED HDDs - 12 Gb SAS 10K			
7XB7A00030	AUM4	ThinkSystem 2.5" 300GB 10K SAS 12Gb Hot Swap 512n HDD SED	12
7XB7A00031	AUM5	ThinkSystem 2.5" 600GB 10K SAS 12Gb Hot Swap 512n HDD SED	12
7XB7A00033	B0YX	ThinkSystem 2.5" 1.2TB 10K SAS 12Gb Hot Swap 512n HDD SED	12
7XB7A00070	B0YV	ThinkSystem 2.5" 2.4TB 10K SAS 12Gb Hot Swap 512e HDD FIPS	12
2.5-inch hot-swap SED HDDs - 12 Gb NL SAS			
7XB7A00064	B0YM	ThinkSystem 2.5" 2TB 7.2K SAS 12Gb Hot Swap 512e HDD FIPS	12

Table 31. 2.5-inch hot-swap 6 Gb SATA HDDs

Part number	Feature	Description	Maximum supported
2.5-inch hot-swap HDDs - 6 Gb NL SATA			
7XB7A00036	AUUE	ThinkSystem 2.5" 1TB 7.2K SATA 6Gb Hot Swap 512n HDD	12
7XB7A00037	AUUJ	ThinkSystem 2.5" 2TB 7.2K SATA 6Gb Hot Swap 512e HDD	12

Table 32. 2.5-inch hot-swap 12 Gb SAS SSDs

Part number	Feature	Description	Maximum supported
2.5-inch hot-swap SSDs - 12 Gb SAS - Write Intensive/Performance (10+ DWPD)			
4XB7A70006	BG07	ThinkSystem 2.5" Nytro 3732 400GB Performance SAS 12Gb Hot Swap SSD	12
4XB7A70005	BG06	ThinkSystem 2.5" Nytro 3732 800GB Performance SAS 12Gb Hot Swap SSD	12
4XB7A70004	BG05	ThinkSystem 2.5" Nytro 3732 1.6TB Performance SAS 12Gb Hot Swap SSD	12
4XB7A70003	BG04	ThinkSystem 2.5" Nytro 3732 3.2TB Performance SAS 12Gb Hot Swap SSD	12
4XB7A10230	B4Y5	ThinkSystem 2.5" SS530 800GB Performance SAS 12Gb Hot Swap SSD	12
2.5-inch hot-swap SSDs - 12 Gb SAS - Mixed Use/Mainstream (3-5 DWPD)			
4XB7A17062	B8HU	ThinkSystem 2.5" PM1645a 800GB Mainstream SAS 12Gb Hot Swap SSD	12
4XB7A17063	B8J4	ThinkSystem 2.5" PM1645a 1.6TB Mainstream SAS 12Gb Hot Swap SSD	12
4XB7A17064	B8JD	ThinkSystem 2.5" PM1645a 3.2TB Mainstream SAS 12Gb Hot Swap SSD	12
4XB7A17065	B8JA	ThinkSystem 2.5" PM1645a 6.4TB Mainstream SAS 12Gb Hot Swap SSD	12
4XB7A13655	B4A2	ThinkSystem 2.5" PM1645 3.2TB Mainstream SAS 12Gb Hot Swap SSD	12
2.5-inch hot-swap SSDs - 12 Gb SAS - Read Intensive/Entry/Capacity (<3 DWPD)			
4XB7A38175	B91A	ThinkSystem 2.5" PM1643a 960GB Entry SAS 12Gb Hot Swap SSD	12
4XB7A38176	B91B	ThinkSystem 2.5" PM1643a 1.92TB Entry SAS 12Gb Hot Swap SSD	12
4XB7A17054	B91C	ThinkSystem 2.5" PM1643a 3.84TB Entry SAS 12Gb Hot Swap SSD	12
4XB7A17055	B91D	ThinkSystem 2.5" PM1643a 7.68TB Entry SAS 12Gb Hot Swap SSD	12
4XB7A17056	BC4R	ThinkSystem 2.5" PM1643a 15.36TB Entry SAS 12Gb Hot Swap SSD	12
2.5-inch hot-swap SED SSDs - 12 Gb SAS - Write Intensive/Performance (10+ DWPD)			
4XB7A70007	BFZZ	ThinkSystem 2.5" Nytro 3732 800GB Performance SAS 12Gb Hot Swap SSD SED	12

Table 33. 2.5-inch hot-swap 6 Gb SATA SSDs

Part number	Feature	Description	Maximum supported
2.5-inch hot-swap SSDs - 6 Gb SATA - Mixed Use/Mainstream (3-5 DWPD)			
4XB7A17125	BA7Q	ThinkSystem 2.5" S4620 480GB Mixed Use SATA 6Gb HS SSD	12
4XB7A17126	BA4T	ThinkSystem 2.5" S4620 960GB Mixed Use SATA 6Gb HS SSD	12
4XB7A17127	BA4U	ThinkSystem 2.5" S4620 1.92TB Mixed Use SATA 6Gb HS SSD	12
4XB7A17128	BK7L	ThinkSystem 2.5" S4620 3.84TB Mixed Use SATA 6Gb HS SSD	12
4XB7A17087	B8J1	ThinkSystem 2.5" 5300 240GB Mainstream SATA 6Gb Hot Swap SSD	12
4XB7A17088	B8HY	ThinkSystem 2.5" 5300 480GB Mainstream SATA 6Gb Hot Swap SSD	12
4XB7A17089	B8J6	ThinkSystem 2.5" 5300 960GB Mainstream SATA 6Gb Hot Swap SSD	12
4XB7A17090	B8JE	ThinkSystem 2.5" 5300 1.92TB Mainstream SATA 6Gb Hot Swap SSD	12
4XB7A17091	B8J7	ThinkSystem 2.5" 5300 3.84TB Mainstream SATA 6Gb Hot Swap SSD	12
4XB7A13633	B49L	ThinkSystem 2.5" S4610 240GB Mixed Use SATA 6Gb HS SSD	12
4XB7A13634	B49M	ThinkSystem 2.5" S4610 480GB Mixed Use SATA 6Gb HS SSD	12
4XB7A13635	B49N	ThinkSystem 2.5" S4610 960GB Mixed Use SATA 6Gb HS SSD	12

Part number	Feature	Description	Maximum supported
4XB7A13636	B49P	ThinkSystem 2.5" S4610 1.92TB Mixed Use SATA 6Gb HS SSD	12
4XB7A13637	B49Q	ThinkSystem 2.5" S4610 3.84TB Mixed Use SATA 6Gb HS SSD	12
2.5-inch hot-swap SSDs - 6 Gb SATA - Read Intensive/Entry (<3 DWPD)			
4XB7A72438	BM8B	ThinkSystem 2.5" PM893 480GB Read Intensive SATA 6Gb HS SSD	12
4XB7A72439	BM8A	ThinkSystem 2.5" PM893 960GB Read Intensive SATA 6Gb HS SSD	12
4XB7A72440	BM89	ThinkSystem 2.5" PM893 1.92TB Read Intensive SATA 6Gb HS SSD	12
4XB7A72441	BM88	ThinkSystem 2.5" PM893 3.84TB Read Intensive SATA 6Gb HS SSD	12
4XB7A72442	BM87	ThinkSystem 2.5" PM893 7.68TB Read Intensive SATA 6Gb HS SSD	12
4XB7A17072	B99D	ThinkSystem 2.5" S4520 240GB Read Intensive SATA 6Gb HS SSD	12
4XB7A17101	BA7G	ThinkSystem 2.5" S4520 480GB Read Intensive SATA 6Gb HS SSD	12
4XB7A17102	BA7H	ThinkSystem 2.5" S4520 960GB Read Intensive SATA 6Gb HS SSD	12
4XB7A17103	BA7J	ThinkSystem 2.5" S4520 1.92TB Read Intensive SATA 6Gb HS SSD	12
4XB7A17104	BK77	ThinkSystem 2.5" S4520 3.84TB Read Intensive SATA 6Gb HS SSD	12
4XB7A17105	BK78	ThinkSystem 2.5" S4520 7.68TB Read Intensive SATA 6Gb HS SSD	12
4XB7A38271	BCTC	ThinkSystem 2.5" Multi Vendor 240GB Entry SATA 6Gb Hot Swap SSD	12
4XB7A38272	BCTD	ThinkSystem 2.5" Multi Vendor 480GB Entry SATA 6Gb Hot Swap SSD	12
4XB7A38273	BCTE	ThinkSystem 2.5" Multi Vendor 960GB Entry SATA 6Gb Hot Swap SSD	12
4XB7A38274	BCTF	ThinkSystem 2.5" Multi Vendor 1.92TB Entry SATA 6Gb Hot Swap SSD	12
4XB7A38275	BCTG	ThinkSystem 2.5" Multi Vendor 3.84TB Entry SATA 6Gb Hot Swap SSD	12
4XB7A17075	B8HV	ThinkSystem 2.5" 5300 240GB Entry SATA 6Gb Hot Swap SSD	12
4XB7A17076	B8JM	ThinkSystem 2.5" 5300 480GB Entry SATA 6Gb Hot Swap SSD	12
4XB7A17077	B8HP	ThinkSystem 2.5" 5300 960GB Entry SATA 6Gb Hot Swap SSD	12
4XB7A17078	B8J5	ThinkSystem 2.5" 5300 1.92TB Entry SATA 6Gb Hot Swap SSD	12
4XB7A17079	B8JP	ThinkSystem 2.5" 5300 3.84TB Entry SATA 6Gb Hot Swap SSD	12
4XB7A17080	B8J2	ThinkSystem 2.5" 5300 7.68TB Entry SATA 6Gb Hot Swap SSD	12
4XB7A38185	B9AC	ThinkSystem 2.5" 5210 960GB Entry SATA 6Gb Hot Swap QLC SSD	12
4XB7A38144	B7EW	ThinkSystem 2.5" 5210 1.92TB Entry SATA 6Gb Hot Swap QLC SSD	12
4XB7A38145	B7EX	ThinkSystem 2.5" 5210 3.84TB Entry SATA 6Gb Hot Swap QLC SSD	12
4XB7A38146	B7EY	ThinkSystem 2.5" 5210 7.68TB Entry SATA 6Gb Hot Swap QLC SSD	12
4XB7A10247	B498	ThinkSystem 2.5" S4510 240GB Read Intensive SATA 6Gb HS SSD	12
4XB7A10248	B499	ThinkSystem 2.5" S4510 480GB Read Intensive SATA 6Gb HS SSD	12
4XB7A10249	B49A	ThinkSystem 2.5" S4510 960GB Read Intensive SATA 6Gb HS SSD	12
4XB7A13622	B49B	ThinkSystem 2.5" S4510 1.92TB Read Intensive SATA 6Gb HS SSD	12
4XB7A13623	B49C	ThinkSystem 2.5" S4510 3.84TB Read Intensive SATA 6Gb HS SSD	12
4XB7A10195	B34H	ThinkSystem 2.5" PM883 240GB Entry SATA 6Gb Hot Swap SSD	12
4XB7A10196	B34J	ThinkSystem 2.5" PM883 480GB Entry SATA 6Gb Hot Swap SSD	12
4XB7A10197	B34K	ThinkSystem 2.5" PM883 960GB Entry SATA 6Gb Hot Swap SSD	12
4XB7A10198	B34L	ThinkSystem 2.5" PM883 1.92TB Entry SATA 6Gb Hot Swap SSD	12
4XB7A10199	B34M	ThinkSystem 2.5" PM883 3.84TB Entry SATA 6Gb Hot Swap SSD	12
4XB7A10200	B4D2	ThinkSystem 2.5" PM883 7.68TB Entry SATA 6Gb Hot Swap SSD	12

Part number	Feature	Description	Maximum supported
2.5-inch hot-swap SED SSDs - 6 Gb SATA - Mixed Use/Mainstream (3-5 DWPD)			
4XB7A13972	BNEW	ThinkSystem 2.5" 5300 960GB Mainstream SATA 6Gb Hot Swap SSD SED	12
4XB7A38193	B93K	ThinkSystem 2.5" 5300 1.92TB Mainstream SATA 6Gb Hot Swap SSD SED	12
4XB7A64222	BEMB	ThinkSystem 2.5" 5300 3.84TB Mainstream SATA 6Gb Hot Swap SSD SED	12
2.5-inch hot-swap SED SSDs - 6 Gb SATA - Read Intensive/Entry (<3 DWPD)			
4XB7A38141	BE29	ThinkSystem 2.5" 5300 960GB Entry SATA 6Gb Hot Swap SSD SED	12
4XB7A38191	B93L	ThinkSystem 2.5" 5300 3.84TB Entry SATA 6Gb Hot Swap SSD SED	12
4XB7A38192	B93M	ThinkSystem 2.5" 5300 7.68TB Entry SATA 6Gb Hot Swap SSD SED	12

Table 34. 2.5-inch hot-swap PCIe 4.0 NVMe SSDs (operate at PCIe 3.0 speeds in this server)

Part number	Feature	Description	Maximum supported
2.5-inch SSDs - U.2 PCIe 4.0 NVMe - Mixed Use/Mainstream (3-5 DWPD)			
4XB7A17152	BCFV	ThinkSystem 2.5" U.2 P5600 1.6TB Mixed Use NVMe PCIe 4.0 x4 HS SSD	10
4XB7A17153	BCFR	ThinkSystem 2.5" U.2 P5600 3.2TB Mixed Use NVMe PCIe 4.0 x4 HS SSD	10
4XB7A17154	BCFS	ThinkSystem 2.5" U.2 P5600 6.4TB Mixed Use NVMe PCIe 4.0 x4 HS SSD	10
2.5-inch SSDs - U.3 PCIe 4.0 NVMe - Mixed Use/Mainstream (3-5 DWPD)			
4XB7A64175	BE03	ThinkSystem U.3 Kioxia CM6-V 800GB Mainstream NVMe PCIe 4.0 x4 Hot Swap SSD	10
4XB7A17112	B96Z	ThinkSystem U.3 Kioxia CM6-V 1.6TB Mainstream NVMe PCIe4.0 x4 Hot Swap SSD	10
4XB7A17113	B96T	ThinkSystem U.3 Kioxia CM6-V 3.2TB Mainstream NVMe PCIe4.0 x4 Hot Swap SSD	10
4XB7A17114	B96P	ThinkSystem U.3 Kioxia CM6-V 6.4TB Mainstream NVMe PCIe4.0 x4 Hot Swap SSD	10
2.5-inch SSDs - U.2 PCIe 4.0 NVMe - Read Intensive/Entry (<3 DWPD)			
4XB7A17145	BCFT	ThinkSystem 2.5" U.2 P5500 1.92TB Read Intensive NVMe PCIe 4.0 x4 HS SSD	10
4XB7A17146	BCFW	ThinkSystem 2.5" U.2 P5500 3.84TB Read Intensive NVMe PCIe 4.0 x4 HS SSD	10
4XB7A17147	BCFU	ThinkSystem 2.5" U.2 P5500 7.68TB Read Intensive NVMe PCIe 4.0 x4 HS SSD	10
4XB7A38197	BC4Z	ThinkSystem U.2 PM1733 3.84TB Entry NVMe PCIe 4.0 x4 Hot Swap SSD	10
4XB7A38283	BE2E	ThinkSystem U.2 PM1733 7.68TB Entry NVMe PCIe 4.0 x4 Hot Swap SSD	10
4XB7A38284	BE2F	ThinkSystem U.2 PM1733 15.36TB Entry NVMe PCIe 4.0 x4 Hot Swap SSD	10
2.5-inch SSDs - U.3 PCIe 4.0 NVMe - Read Intensive/Entry (<3 DWPD)			
4XB7A64141	BE2G	ThinkSystem U.3 Kioxia CM6-R 3.84TB Entry NVMe PCIe 4.0 x4 Hot Swap SSD	10
4XB7A64142	BE2H	ThinkSystem U.3 Kioxia CM6-R 7.68TB Entry NVMe PCIe 4.0 x4 Hot Swap SSD	10
2.5-inch SED SSDs - U.2 PCIe 4.0 NVMe - Read Intensive/Entry (<3 DWPD)			
4XB7A38257	BE2A	ThinkSystem U.2 PM1733 3.84TB Entry NVMe PCIe 4.0 x4 Hot Swap SSD SED	10
4XB7A38258	BE2B	ThinkSystem U.2 PM1733 7.68TB Entry NVMe PCIe 4.0 x4 Hot Swap SSD SED	10
2.5-inch SED SSDs - U.3 PCIe 4.0 NVMe - Read Intensive/Entry (<3 DWPD)			
4XB7A38269	BE2C	ThinkSystem U.3 Kioxia CM6-R 3.84TB Entry NVMe PCIe 4.0 x4 Hot Swap SSD SED	10
4XB7A38270	BE2D	ThinkSystem U.3 Kioxia CM6-R 7.68TB Entry NVMe PCIe 4.0 x4 Hot Swap SSD SED	10

Table 35. 2.5-inch hot-swap PCIe 3.0 NVMe SSDs

Part number	Feature	Description	Maximum supported
2.5-inch SSDs - U.2 PCIe 3.0 NVMe - Write Intensive/Performance (10+ DWPD)			
7N47A00081	AUMJ	ThinkSystem 2.5" U.2 P4800X 375GB Write Intensive NVMe PCIe 3.0 x4 HS SSD	10
7N47A00083	B2ZJ	ThinkSystem 2.5" U.2 P4800X 750GB Write Intensive NVMe PCIe 3.0 x4 HS SSD	10
2.5-inch SSDs - U.2 PCIe 3.0 NVMe - Mixed Use/Mainstream (3-5 DWPD)			
4XB7A13936	B589	ThinkSystem U.2 Intel P4610 1.6TB Mainstream NVMe PCIe3.0 x4 Hot Swap SSD	10
4XB7A08518	B21Y	ThinkSystem U.2 Toshiba CM5-V 3.2TB Mainstream NVMe PCIe 3.0 x4 Hot Swap SSD	10
2.5-inch SSDs - U.2 PCIe 3.0 NVMe - Read Intensive/Entry (<3 DWPD)			
4XB7A10204	B58G	ThinkSystem 2.5" U.2 P4510 2.0TB Read Intensive NVMe PCIe 3.0 x4 HS SSD	10
4XB7A10205	B58H	ThinkSystem U.2 Intel P4510 4.0TB Entry NVMe PCIe3.0 x4 Hot Swap SSD	10
4XB7A08513	B58J	ThinkSystem U.2 Intel P4510 8.0TB Entry NVMe PCIe3.0 x4 Hot Swap SSD	10
4XB7A10175	B34N	ThinkSystem U.2 PM983 1.92TB Entry NVMe PCIe 3.0 x4 Hot Swap SSD	10
4XB7A10176	B34P	ThinkSystem U.2 PM983 3.84TB Entry NVMe PCIe 3.0 x4 Hot Swap SSD	10
4XB7A10177	B4D3	ThinkSystem U.2 PM983 7.68TB Entry NVMe PCIe3.0 x4 Hot Swap SSD	10

Note: NVMe PCIe SSDs support surprise hot removal and hot insertion, provided the operating system supports PCIe SSD hot-swap.

Table 36. 3.5-inch hot-swap 12 Gb SAS HDDs

Part number	Feature	Description	Maximum supported
3.5-inch hot-swap HDDs - 12 Gb SAS 10K			
7XB7A00063	B1JJ	ThinkSystem 3.5" 300GB 10K SAS 12Gb Hot Swap 512n HDD	4
3.5-inch hot-swap HDDs - 12 Gb SAS 15K			
7XB7A00038	AUU2	ThinkSystem 3.5" 300GB 15K SAS 12Gb Hot Swap 512n HDD	4
7XB7A00039	AUU3	ThinkSystem 3.5" 600GB 15K SAS 12Gb Hot Swap 512n HDD	4
7XB7A00040	AUUC	ThinkSystem 3.5" 900GB 15K SAS 12Gb Hot Swap 512e HDD	4
3.5-inch hot-swap HDDs - 12 Gb NL SAS			
7XB7A00041	AUU4	ThinkSystem 3.5" 1TB 7.2K SAS 12Gb Hot Swap 512n HDD	4
7XB7A00042	AUU5	ThinkSystem 3.5" 2TB 7.2K SAS 12Gb Hot Swap 512n HDD	4
7XB7A00043	AUU6	ThinkSystem 3.5" 4TB 7.2K SAS 12Gb Hot Swap 512n HDD	4
7XB7A00044	AUU7	ThinkSystem 3.5" 6TB 7.2K SAS 12Gb Hot Swap 512e HDD	4
7XB7A00045	B0YR	ThinkSystem 3.5" 8TB 7.2K SAS 12Gb Hot Swap 512e HDD	4
7XB7A00046	AUUG	ThinkSystem 3.5" 10TB 7.2K SAS 12Gb Hot Swap 512e HDD	4
7XB7A00067	B117	ThinkSystem 3.5" 12TB 7.2K SAS 12Gb Hot Swap 512e HDD	4
4XB7A13906	B496	ThinkSystem 3.5" 14TB 7.2K SAS 12Gb Hot Swap 512e HDD	4
4XB7A13911	B7EZ	ThinkSystem 3.5" 16TB 7.2K SAS 12Gb Hot Swap 512e HDD	4
4XB7A38266	BCFP	ThinkSystem 3.5" 18TB 7.2K SAS 12Gb Hot Swap 512e HDD	4
3.5-inch hot-swap SED HDDs - 12 Gb NL SAS			
7XB7A00065	B0YN	ThinkSystem 3.5" 2TB 7.2K SAS 12Gb Hot Swap 512e HDD FIPS	4
7XB7A00047	AUUH	ThinkSystem 3.5" 4TB 7.2K SAS 12Gb Hot Swap 512n HDD FIPS	4
7XB7A00048	B0YP	ThinkSystem 3.5" 6TB 7.2K SAS 12Gb Hot Swap 512e HDD FIPS	4
7XB7A00066	B0YQ	ThinkSystem 3.5" 8TB 7.2K SAS 12Gb Hot Swap 512e HDD FIPS	4

Table 37. 3.5-inch hot-swap 6 Gb SATA HDDs

Part number	Feature	Description	Maximum supported
3.5-inch hot-swap HDDs - 6 Gb NL SATA			
7XB7A00049	AUUF	ThinkSystem 3.5" 1TB 7.2K SATA 6Gb Hot Swap 512n HDD	4
7XB7A00050	AUUD	ThinkSystem 3.5" 2TB 7.2K SATA 6Gb Hot Swap 512n HDD	4
7XB7A00051	AUU8	ThinkSystem 3.5" 4TB 7.2K SATA 6Gb Hot Swap 512n HDD	4
7XB7A00052	AUUA	ThinkSystem 3.5" 6TB 7.2K SATA 6Gb Hot Swap 512e HDD	4
7XB7A00053	AUU9	ThinkSystem 3.5" 8TB 7.2K SATA 6Gb Hot Swap 512e HDD	4
7XB7A00054	AUUB	ThinkSystem 3.5" 10TB 7.2K SATA 6Gb Hot Swap 512e HDD	4
7XB7A00068	B118	ThinkSystem 3.5" 12TB 7.2K SATA 6Gb Hot Swap 512e HDD	4
4XB7A13907	B497	ThinkSystem 3.5" 14TB 7.2K SATA 6Gb Hot Swap 512e HDD	4
4XB7A13914	B7F0	ThinkSystem 3.5" 16TB 7.2K SATA 6Gb Hot Swap 512e HDD	4
4XB7A38130	BCFH	ThinkSystem 3.5" 18TB 7.2K SATA 6Gb Hot Swap 512e HDD	4

Table 38. 3.5-inch hot-swap 12 Gb SAS SSDs

Part number	Feature	Description	Maximum supported
3.5-inch hot-swap SSDs - 12 Gb SAS - Write Intensive/Performance (10+ DWPD)			
4XB7A70011	BG03	ThinkSystem 3.5" Nytro 3732 400GB Performance SAS 12Gb Hot Swap SSD	4
4XB7A70010	BG02	ThinkSystem 3.5" Nytro 3732 800GB Performance SAS 12Gb Hot Swap SSD	4
4XB7A70009	BG01	ThinkSystem 3.5" Nytro 3732 1.6TB Performance SAS 12Gb Hot Swap SSD	4
4XB7A70008	BG00	ThinkSystem 3.5" Nytro 3732 3.2TB Performance SAS 12Gb Hot Swap SSD	4
4XB7A10234	B4Y8	ThinkSystem 3.5" SS530 800GB Performance SAS 12Gb Hot Swap SSD	4
3.5-inch hot-swap SSDs - 12 Gb SAS - Mixed Use/Mainstream (3-5 DWPD)			
4XB7A17066	B8HT	ThinkSystem 3.5" PM1645a 800GB Mainstream SAS 12Gb Hot Swap SSD	4
4XB7A17043	B8JN	ThinkSystem 3.5" PM1645a 1.6TB Mainstream SAS 12Gb Hot Swap SSD	4
4XB7A17067	B8JK	ThinkSystem 3.5" PM1645a 3.2TB Mainstream SAS 12Gb Hot Swap SSD	4
4XB7A17068	B8JG	ThinkSystem 3.5" PM1645a 6.4TB Mainstream SAS 12Gb Hot Swap SSD	4
4XB7A13659	B4A5	ThinkSystem 3.5" PM1645 3.2TB Mainstream SAS 12Gb Hot Swap SSD	4
3.5-inch hot-swap SSDs - 12 Gb SAS - Read Intensive/Entry/Capacity (<3 DWPD)			
4XB7A17058	B91E	ThinkSystem 3.5" PM1643a 3.84TB Entry SAS 12Gb Hot Swap SSD	4
4XB7A17059	BEVK	ThinkSystem 3.5" PM1643a 7.68TB Entry SAS 12Gb Hot Swap SSD	4

Table 39. 3.5-inch hot-swap 6 Gb SATA SSDs

Part number	Feature	Description	Maximum supported
3.5-inch hot-swap SSDs - 6 Gb SATA - Mixed Use/Mainstream (3-5 DWPD)			
4XB7A17137	BA4W	ThinkSystem 3.5" S4620 480GB Mixed Use SATA 6Gb HS SSD	4
4XB7A17138	BA4X	ThinkSystem 3.5" S4620 960GB Mixed Use SATA 6Gb HS SSD	4
4XB7A17139	BA4Y	ThinkSystem 3.5" S4620 1.92TB Mixed Use SATA 6Gb HS SSD	4
4XB7A17140	BK7P	ThinkSystem 3.5" S4620 3.84TB Mixed Use SATA 6Gb HS SSD	4
4XB7A17096	B8JL	ThinkSystem 3.5" 5300 240GB Mainstream SATA 6Gb Hot Swap SSD	4
4XB7A17097	B8JF	ThinkSystem 3.5" 5300 480GB Mainstream SATA 6Gb Hot Swap SSD	4
4XB7A17098	B8J0	ThinkSystem 3.5" 5300 960GB Mainstream SATA 6Gb Hot Swap SSD	4
4XB7A17099	B8HR	ThinkSystem 3.5" 5300 1.92TB Mainstream SATA 6Gb Hot Swap SSD	4
4XB7A17100	B8HX	ThinkSystem 3.5" 5300 3.84TB Mainstream SATA 6Gb Hot Swap SSD	4
4XB7A13639	B49R	ThinkSystem 3.5" S4610 240GB Mixed Use SATA 6Gb HS SSD	4
4XB7A13640	B49S	ThinkSystem 3.5" S4610 480GB Mixed Use SATA 6Gb HS SSD	4
4XB7A13641	B49T	ThinkSystem 3.5" S4610 960GB Mixed Use SATA 6Gb HS SSD	4
4XB7A13642	B49U	ThinkSystem 3.5" S4610 1.92TB Mixed Use SATA 6Gb HS SSD	4
4XB7A13643	B49V	ThinkSystem 3.5" S4610 3.84TB Mixed Use SATA 6Gb HS SSD	4
3.5-inch hot-swap SSDs - 6 Gb SATA - Read Intensive/Entry (<3 DWPD)			
4XB7A17118	BA7K	ThinkSystem 3.5" S4520 240GB Read Intensive SATA 6Gb HS SSD	4
4XB7A17119	BA7L	ThinkSystem 3.5" S4520 480GB Read Intensive SATA 6Gb HS SSD	4

Part number	Feature	Description	Maximum supported
4XB7A17120	BA7M	ThinkSystem 3.5" S4520 960GB Read Intensive SATA 6Gb HS SSD	4
4XB7A17121	BA7N	ThinkSystem 3.5" S4520 1.92TB Read Intensive SATA 6Gb HS SSD	4
4XB7A17122	BK7F	ThinkSystem 3.5" S4520 3.84TB Read Intensive SATA 6Gb HS SSD	4
4XB7A17123	BK7G	ThinkSystem 3.5" S4520 7.68TB Read Intensive SATA 6Gb HS SSD	4
4XB7A38276	BCTH	ThinkSystem 3.5" Multi Vendor 240GB Entry SATA 6Gb Hot Swap SSD	4
4XB7A38277	BCTJ	ThinkSystem 3.5" Multi Vendor 480GB Entry SATA 6Gb Hot Swap SSD	4
4XB7A38278	BCTK	ThinkSystem 3.5" Multi Vendor 960GB Entry SATA 6Gb Hot Swap SSD	4
4XB7A38279	BCTL	ThinkSystem 3.5" Multi Vendor 1.92TB Entry SATA 6Gb Hot Swap SSD	4
4XB7A38281	BCTM	ThinkSystem 3.5" Multi Vendor 3.84TB Entry SATA 6Gb Hot Swap SSD	4
4XB7A17081	B8JB	ThinkSystem 3.5" 5300 240GB Entry SATA 6Gb Hot Swap SSD	4
4XB7A17082	B8J9	ThinkSystem 3.5" 5300 480GB Entry SATA 6Gb Hot Swap SSD	4
4XB7A17083	B8JC	ThinkSystem 3.5" 5300 960GB Entry SATA 6Gb Hot Swap SSD	4
4XB7A17084	B8HZ	ThinkSystem 3.5" 5300 1.92TB Entry SATA 6Gb Hot Swap SSD	4
4XB7A17085	B8HQ	ThinkSystem 3.5" 5300 3.84TB Entry SATA 6Gb Hot Swap SSD	4
4XB7A17086	B8J3	ThinkSystem 3.5" 5300 7.68TB Entry SATA 6Gb Hot Swap SSD	4
4XB7A13625	B49D	ThinkSystem 3.5" S4510 240GB Read Intensive SATA 6Gb HS SSD	4
4XB7A13626	B49E	ThinkSystem 3.5" S4510 480GB Read Intensive SATA 6Gb HS SSD	4
4XB7A13627	B49F	ThinkSystem 3.5" S4510 960GB Read Intensive SATA 6Gb HS SSD	4
4XB7A13628	B49G	ThinkSystem 3.5" S4510 1.92TB Read Intensive SATA 6Gb HS SSD	4
4XB7A13629	B49H	ThinkSystem 3.5" S4510 3.84TB Read Intensive SATA 6Gb HS SSD	4
4XB7A17176	B6TM	ThinkSystem 3.5" PM883 240GB Entry SATA 6Gb Hot Swap SSD	4
4XB7A17177	B6TN	ThinkSystem 3.5" PM883 480GB Entry SATA 6Gb Hot Swap SSD	4
4XB7A17178	B6TP	ThinkSystem 3.5" PM883 960GB Entry SATA 6Gb Hot Swap SSD	4
4XB7A17179	B6JY	ThinkSystem 3.5" PM883 1.92TB Entry SATA 6Gb Hot Swap SSD	4
4XB7A17180	B6JZ	ThinkSystem 3.5" PM883 3.84TB Entry SATA 6Gb Hot Swap SSD	4

Table 40. 3.5-inch hot-swap PCIe 4.0 NVMe SSDs (operate at PCIe 3.0 speeds in this server)

Part number	Feature	Description	Maximum supported
3.5-inch SSDs - U.2 PCIe 4.0 NVMe - Mixed Use/Mainstream (3-5 DWPD)			
4XB7A17155	BCFM	ThinkSystem 3.5" U.2 P5600 1.6TB Mixed Use NVMe PCIe 4.0 x4 HS SSD	4
4XB7A17156	BCFJ	ThinkSystem 3.5" U.2 P5600 3.2TB Mixed Use NVMe PCIe 4.0 x4 HS SSD	4
4XB7A17157	BCFQ	ThinkSystem 3.5" U.2 P5600 6.4TB Mixed Use NVMe PCIe 4.0 x4 HS SSD	4
3.5-inch SSDs - U.3 PCIe 4.0 NVMe - Mixed Use/Mainstream (3-5 DWPD)			
4XB7A64176	BE04	ThinkSystem 3.5" Kioxia CM6-V 800GB Mainstream NVMe PCIe 4.0 x4 Hot Swap SSD	4
4XB7A17115	B96V	ThinkSystem 3.5" Kioxia CM6-V 1.6TB Mainstream NVMe PCIe4.0 x4 Hot Swap SSD	4
4XB7A17116	B96K	ThinkSystem 3.5" Kioxia CM6-V 3.2TB Mainstream NVMe PCIe4.0 x4 Hot Swap SSD	4
4XB7A17117	B96W	ThinkSystem 3.5" Kioxia CM6-V 6.4TB Mainstream NVMe PCIe4.0 x4 Hot Swap SSD	4

Note: NVMe PCIe SSDs support surprise hot removal and hot insertion, provided the operating system supports PCIe SSD hot-swap.

Table 41. 3.5-inch hot-swap PCIe 3.0 NVMe SSDs

Part number	Feature	Description	Maximum supported
3.5-inch SSDs - U.2 PCIe 3.0 NVMe - Mixed Use/Mainstream (3-5 DWPD)			
4XB7A13944	B58C	ThinkSystem 3.5" Intel P4610 1.6TB Mainstream NVMe PCIe3.0 x4 Hot Swap SSD	4
4XB7A08534	B221	ThinkSystem 3.5" Toshiba CM5-V 3.2TB Mainstream NVMe PCIe 3.0 x4 Hot Swap SSD	4
3.5-inch SSDs - U.2 PCIe 3.0 NVMe - Read Intensive/Entry (<3 DWPD)			
4XB7A10178	B34Q	ThinkSystem 3.5" PM983 1.92TB Entry NVMe PCIe 3.0 x4 Hot Swap SSD	4
4XB7A10179	B34R	ThinkSystem 3.5" PM983 3.84TB Entry NVMe PCIe 3.0 x4 Hot Swap SSD	4
4XB7A10180	B4D4	ThinkSystem 3.5" PM983 7.68TB Entry NVMe PCIe3.0 x4 Hot Swap SSD	4

Note: NVMe PCIe SSDs support surprise hot removal and hot insertion, provided the operating system supports PCIe SSD hot-swap.

Table 42. M.2 SATA drives

Part number	Feature	Description	Maximum supported
M.2 SSDs - 6 Gb SATA - Read Intensive/Entry (<3 DWPD)			
7N47A00129	AUUL	ThinkSystem M.2 32GB SATA 6Gbps Non-Hot Swap SSD	2
7N47A00130	AUUV	ThinkSystem M.2 128GB SATA 6Gbps Non-Hot Swap SSD	2
4XB7A17071	B8HS	ThinkSystem M.2 5300 240GB SATA 6Gbps Non-Hot Swap SSD	2
4XB7A17073	B919	ThinkSystem M.2 5300 480GB SATA 6Gbps Non-Hot Swap SSD	2

USB memory key

For general portable storage needs, the server also supports the USB memory key option that is listed in the following table.

Table 43. USB memory key

Part number	Feature	Description
4X77A08621	B8NV	ThinkSystem 32GB USB Flash Drive

Optical drives

The server supports the external USB optical drive listed in the following table.

Table 44. External optical drive

Part number	Feature code	Description
7XA7A05926	AVV8	ThinkSystem External USB DVD RW Optical Disk Drive

The drive is based on the Lenovo Slim DVD Burner DB65 drive and supports the following formats: DVD-RAM, DVD-RW, DVD+RW, DVD+R, DVD-R, DVD-ROM, DVD-R DL, CD-RW, CD-R, CD-ROM.

I/O expansion

The SR630 server supports one LOM card slot and up to four PCIe slots: one slot on the system planar that is dedicated to an internal storage controller and up to three PCIe slots with different riser cards installed into two riser sockets on the system planar (one riser socket supports installation of one riser card).

The slot form factors are as follows:

- LOM card slot
- Slot 1: PCIe 3.0 x8, ML2 x8, or ML2 x16; low profile
- Slot 2: PCIe 3.0 x16 or x8; low profile or full-height, half-length
- Slot 3: PCIe 3.0 x16; low profile
- Slot 4: PCIe 3.0 x8 (dedicated to an internal storage controller)

Configuration notes:

- Slot 3 requires the second processor to be installed.
- Slots 1 and 2 are not present if the Rear HDD Kit is installed.
- The COM Port Upgrade Kit is installed in place of one of the PCIe slots 1, 2, or 3.

The locations of the PCIe slots are shown in the following figure.

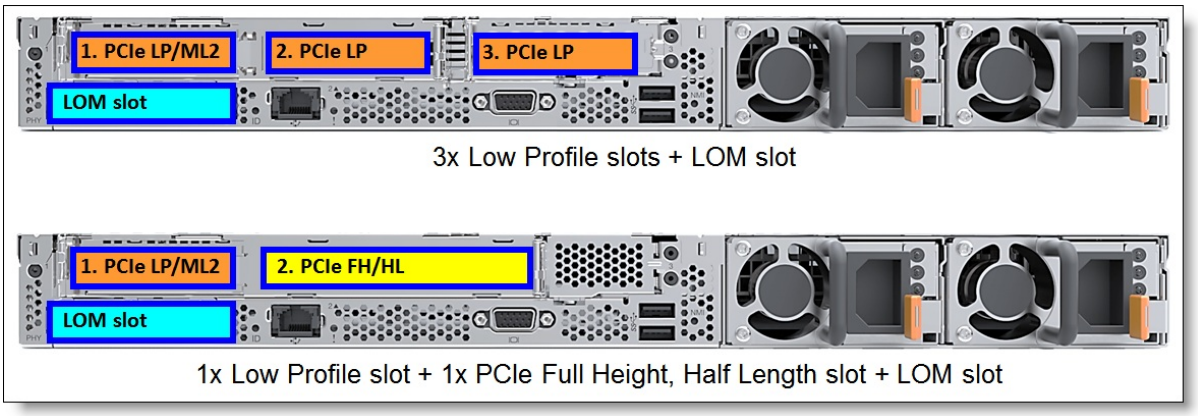


Figure 9. Slot locations

Riser 1 supplies slots 1 and 2, and Riser 2 supplies slot 3. The slots that are available for use depend on the number of riser cards that are installed and whether the second processor is installed, as shown in the following table.

Table 45. Slots available for use

Riser Card 1	Riser Card 2	Slots available for use	
		Processor 1	Processor 2
None	None	LOM, 4	-
None	PCIe x16	LOM, 4	3
PCIe x8/x16, PCIe x8ML2/x16, or PCIe x16ML2/x8	None	LOM, 1, 2, 4	-
PCIe x8/x16, PCIe x8ML2/x16, or PCIe x16ML2/x8	PCIe x16	LOM, 1, 2, 4	3

The following table lists available PCIe riser card options.

Table 46. PCIe riser cards and miscellaneous options

Part number	Feature code	Description	Maximum quantity
x8 Riser Card 1 options (Riser card 1 supplies slots 1 and 2)			
7XH7A02682	AUWC	ThinkSystem SR530/SR570/SR630 x8/x16 PCIe LP+LP Riser 1 Kit	1
7XH7A05893	None*	ThinkSystem SR530/SR570/SR630 x8/x16 PCIe LP+FH Riser 1 Kit	1
7XH7A02683	AUWE	ThinkSystem SR630 x8ML2/x16 PCIe LP+LP Riser 1 Kit	1
7XH7A05894	None*	ThinkSystem SR630 x8ML2/x16 PCIe LP+FH Riser 1 Kit	1
7XH7A02684	AUWD	ThinkSystem SR630 x16ML2/x8 PCIe LP+LP Riser 1 Kit	1
7XH7A05895	None*	ThinkSystem SR630 x16ML2/x8 PCIe LP+FH Riser 1 Kit	1
Riser Card 2 option (Riser card 2 supplies slot 3)			
7XH7A02685	AUWA	ThinkSystem SR530/SR570/SR630 x16 PCIe LP Riser 2 Kit	1
Serial port upgrade kit			
4Z17A80446	BMNJ	ThinkSystem COM Port Upgrade Kit v2	
7Z17A02577	AUSL	ThinkSystem COM Port Upgrade Kit	1

* The LP+FH Riser 1 can be factory-installed by selecting the feature codes of the LP+LP Riser 1 (AUWC, AUWD, or AUWE) and LP+FH Bracket (AUWS).

The COM Port Upgrade Kit, part number 7Z17A02577, is used for mounting the external serial port on the rear of the SR630. This option includes the bracket and the cable. The COM Port option is mounted in place of one of the PCIe slots 1, 2, or 3.

Network adapters

The SR630 server supports up to four onboard network ports with optional LOM cards that use the Intel Ethernet Connection X722 1/10 GbE technology integrated into the Intel C624 Platform Controller Hub (PCH). The server also supports ML2 adapters that are installed in the custom ML2 slot provided by an ML2 riser card. The LOM cards and ML2 network adapters support direct connectivity to the XClarity Controller via the Network Controller Sideband Interface (NSCI) for out-of-band systems management.

The integrated Intel Ethernet Connection X722 has the following features:

- Four 1/10 Gb Ethernet capable ports (no 10/100 Mb Ethernet support)
- NIC Teaming (load balancing and failover)
- Data Center Bridging
- iWARP (RDMA over IP)
- VMDq and SR-IOV virtualization (10 Gb speeds only, 4 PFs, 128 VFs per device)
- IEEE 802.1q Virtual Local Area Networks (VLANs)
- NVGRE, VXLAN, IPinGRE, and MACinUDP network virtualization
- IEEE 802.1Qbg Edge Virtual Bridging
- TCP, IP, and UDP checksum offload
- Large Send Offload (LSO) and Generic Send Offload (GSO)
- Receive Side Scaling (RSS) for TCP and UDP traffic
- Jumbo frames up to 9.5 Kbytes

The following table lists the network adapters that are supported with the SR630 server.

Table 47. Network adapters

Part number	Feature code	Description	Max qty#	I/O slots supported
LOM cards - Gigabit Ethernet				
7ZT7A00544	AUKG	ThinkSystem 1Gb 2-port RJ45 LOM	1	LOM slot
7ZT7A00545	AUKH	ThinkSystem 1Gb 4-port RJ45 LOM	1	LOM slot
LOM cards - 10 Gb Ethernet				
7ZT7A00548	AUKL	ThinkSystem 10Gb 2-port Base-T LOM	1	LOM slot
7ZT7A00546	AUKJ	ThinkSystem 10Gb 2-port SFP+ LOM	1*	LOM slot
7ZT7A00549	AUKM	ThinkSystem 10Gb 4-port Base-T LOM	1	LOM slot
7ZT7A00547	AUKK	ThinkSystem 10Gb 4-port SFP+ LOM	1*	LOM slot
ML2 adapters - Gigabit Ethernet				
7ZT7A00536	AUKW	ThinkSystem Intel I350-T4 ML2 1Gb 4-Port RJ45 Ethernet Adapter	1	1 (ML2)
ML2 adapters - 10 Gb Ethernet				
7ZT7A00497	AUKQ	Broadcom NX-E ML2 10Gb 2-Port Base-T Ethernet Adapter	1	1 (ML2)
01CV770	AU7Z	Emulex VFA5.2 ML2 2x10 GbE SFP+ Adapter and FCoE/iSCSI SW	1*	1 (ML2)
00JY940	ATRH	Intel X710-DA2 ML2 2x10GbE SFP+ Adapter	1*	1 (ML2)
ML2 adapters - 25 Gb Ethernet				
00MN990	ATZR	Mellanox ConnectX-4 Lx 10/25GbE SFP28 1-port ML2 Ethernet Adapter	1*	1 (ML2)
7ZT7A00507	AUKU	Mellanox ConnectX-4 Lx 10/25GbE SFP28 2-port ML2 Ethernet Adapter	1*	1 (ML2)

Part number	Feature code	Description	Max qty#	I/O slots supported
ML2 adapters - FDR InfiniBand				
7ZT7A00501	AUKR	Mellanox ConnectX-3 Pro ML2 FDR 2-Port QSFP VPI Adapter	1*	1 (ML2)
PCIe Low Profile adapters - 1 Gb Ethernet				
7ZT7A00482	AUZX	Broadcom 5720 1GbE RJ45 2-Port PCIe Ethernet Adapter	2 / 3	2, 3, 1
7ZT7A00484	AUZV	Broadcom 5719 1GbE RJ45 4-Port PCIe Ethernet Adapter	2 / 3	2, 3, 1
7ZT7A00533	AUZZ	ThinkSystem I350-F1 PCIe 1Gb 1-Port SFP Ethernet Adapter	2 / 3	2, 3, 1
7ZT7A00534	AUZY	ThinkSystem I350-T2 PCIe 1Gb 2-Port RJ45 Ethernet Adapter	2 / 3	2, 3, 1
7ZT7A00535	AUZW	ThinkSystem I350-T4 PCIe 1Gb 4-Port RJ45 Ethernet Adapter	2 / 3	2, 3, 1
PCIe Low Profile adapters - 10 Gb Ethernet				
7ZT7A00496	AUKP	Broadcom 57416 10GBASE-T 2-Port PCIe Ethernet Adapter	2 / 3	2, 3, 1
00AG570	AT7S	Emulex VFA5.2 2x10 GbE SFP+ PCIe Adapter	2 / 3*	2, 3, 1
00AG580	AT7T	Emulex VFA5.2 2x10 GbE SFP+ Adapter and FCoE/iSCSI SW	2 / 3*	2, 3, 1
00MM860	ATPX	Intel X550-T2 Dual Port 10GBase-T Adapter	2 / 3	2, 3, 1
7ZT7A00537	AUKX	Intel X710-DA2 PCIe 10Gb 2-Port SFP+ Ethernet Adapter	2 / 3*	2, 3, 1
4XC7A08225	B31G	QLogic QL41134 PCIe 10Gb 4-Port Base-T Ethernet Adapter	2 / 3	2, 3, 1
PCIe Full Height adapters - 10 Gb Ethernet				
7ZT7A00493	AUKN	Emulex OCe14104B-NX PCIe 10Gb 4-Port SFP+ Ethernet Adapter	1 / 1*	2
PCIe Low Profile adapters - 25 Gb Ethernet				
7ZT7A00505	AUKS	Broadcom 57412 10/25GbE SFP28 1-Port PCIe Ethernet Adapter	2 / 3*	2, 3, 1
4XC7A08238	B5T0	Broadcom 57414 10/25GbE SFP28 2-port PCIe Ethernet Adapter	2 / 3*	2, 3, 1
4XC7A08316	BD49	Broadcom 57454 10/25GbE SFP28 4-port PCIe Ethernet Adapter V2	1 / 1	2§
4XC7A08295	BCD6	Intel E810-DA2 10/25GbE SFP28 2-Port PCIe Ethernet Adapter	2 / 3*	2, 3, 1
7XC7A05523	B0WY	Intel XXV710-DA2 10/25GbE SFP28 2-Port PCIe Ethernet Adapter	2 / 3*	2, 3, 1
01GR250	AUAJ	Mellanox ConnectX-4 Lx 10/25GbE SFP28 2-Port PCIe Ethernet Adapter	2 / 3*	2, 3, 1
4XC7A62574	BEAP	Mellanox ConnectX-5 EN 10/25GbE SFP28 2-port PCIe Ethernet Adapter	2 / 3*	2, 3, 1
4XC7A62580	BE4U	Mellanox ConnectX-6 Lx 10/25GbE SFP28 2-Port PCIe Ethernet Adapter	2 / 3*	2, 3, 1
4XC7A16683	B5XZ	Mellanox Innova-2 ConnectX-5 FPGA 25GbE 2-port Adapter	1 / 2*	2, 3†
4XC7A08228	B21R	QLogic QL41262 10/25GbE SFP28 2-Port PCIe Ethernet Adapter	2 / 3*	2, 3, 1
4XC7A62581	BHE2	Solarflare X2522-Plus 10/25GbE SFP28 2-Port PCIe Ethernet Adapter	2 / 3*	2, 3, 1
PCIe Low Profile adapters - 40 Gb Ethernet				
00MM950	ATRN	Mellanox ConnectX-4 Lx 1x40GbE QSFP+ Adapter	2 / 3*	2, 3, 1
4XC7A08229	B31C	Mellanox ConnectX-5 Ex 25/40GbE 2-port Low-Latency Adapter	1 / 2*^	2, 3†
PCIe Low Profiles adapters - FDR InfiniBand				
7XC7A05524	B0WX	Mellanox ConnectX-4 PCIe FDR 1-Port QSFP VPI Adapter	2 / 3*	2, 3, 1
7ZT7A00500	AUVG	Mellanox ConnectX-4 PCIe FDR 2-Port QSFP VPI Adapter	2 / 3*	2, 3, 1
PCIe x16 Low Profile adapters - 100 Gb Ethernet / EDR InfiniBand				
00MM960	ATRP	Mellanox ConnectX-4 2x100GbE/EDR IB QSFP28 VPI Adapter	1 / 2*	2, 3†

Part number	Feature code	Description	Max qty#	I/O slots supported
4XC7A08248	B8PP	Mellanox ConnectX-6 Dx 100GbE QSFP56 2-port PCIe Ethernet Adapter	1 / 2*	2, 3†
4XC7A76757	BLC2	ThinkSystem Xilinx Alveo U50 Data Center Accelerator Adapter	1 / 2*	2, 3†
PCIe x16 Low Profile adapters - HDR InfiniBand				
4C57A15326	B4RC	Mellanox ConnectX-6 HDR/200GbE QSFP56 1-port PCIe 4 VPI Adapter	1 / 1*	2†‡
4C57A14179	B4RB	Mellanox HDR/200GbE 2x PCIe Aux Kit	1 / 1*	3†‡
4C57A14177	B4R9	Mellanox ConnectX-6 HDR100/100GbE QSFP56 1-port PCIe VPI Adapter	1 / 2*	2, 3†
4C57A14178	B4RA	Mellanox ConnectX-6 HDR100/100GbE QSFP56 2-port PCIe VPI Adapter	1 / 2*	2, 3†
PCIe Low Profile adapters - Omni-Path				
00WE023	AU0A	Intel OPA 100 Series Single-port PCIe 3.0 x8 HFA	1 / 3*	2, 3, 1
00WE027	AU0B	Intel OPA 100 Series Single-port PCIe 3.0 x16 HFA	1 / 2*	2, 3†

The maximum quantity shown is with one processor / two processors (this does not apply to LOM cards and ML2 adapters).
* The adapter comes without transceivers or cables; for ordering transceivers or cables, see the product guide for the adapter.
† The adapter is supported in the PCIe x16 slots supplied by the riser cards 1 and 2.
^ The 25 GbE connectivity requires the optional Mellanox QSA 100G to 25G Cable Adapter (4G17A10853) (one per port); the supported cables include 25 GbE passive DAC and active optical cables (25 GbE transceivers not supported)
‡ The ConnectX-6 HDR adapter is supported only in the configurations with two processors, and it requires the HDR/200GbE 2x PCIe Aux Kit.
§ In the SR630, this adapter requires a full-height bracket and must be installed in a full-height slot. The use of a low-profile bracket and slot is not supported.

Configuration notes:

- ML2 network adapters are supported in the ML2 x8 slot 1 supplied by the x8ML2/x16 Riser Card 1 (7XH7A02683 or 7XH7A05894).
- PCIe full-height network adapters are supported in the full-height PCIe x8 or x16 slot 2 supplied by the riser card 1.
- PCIe x16 Low Profile network adapters and Omni-Path adapters are supported in the full-height and low profile PCIe x16 slots supplied by the riser cards 1 and 2.
- PCIe Low Profile network adapters (except PCIe x16 adapters and Omni-Path adapters) are supported in the full-height and low profile PCIe x8 and x16 slots supplied by the riser cards 1 and 2.
- Supported transceivers or DAC cables should be purchased for the SFP+, SFP28, QSFP+, and QSFP28 adapters, and UTP Category 6 or Category 5e cables should be purchased for the 10 GbE (Cat6) or 1 GbE (Cat5e or Cat6) RJ-45 adapters. The maximum number of transceivers or cables that are supported per adapter equals the quantity of the adapter ports, and all adapter ports must have the same type of the transceiver or cable selected.

For more information, see the list of Product Guides in the Ethernet Adapters category:
http://lenovopress.com/servers/options/ethernet#rt=product-guide

SAS adapters for external storage

The following table lists SAS RAID controllers and HBAs for external storage attachments that are supported by the SR630 server.

Table 48. SAS RAID adapters and HBAs for external storage

Part number	Feature code	Description	Maximum quantity*	I/O slots supported
12 Gbps SAS RAID adapters				
7Y37A01087	AUNQ	ThinkSystem RAID 930-8e 4GB Flash PCIe 12Gb Adapter	2 / 2	2, 3, 1
12 Gbps SAS HBAs				
7Y37A01090	AUNR	ThinkSystem 430-8e SAS/SATA 12Gb HBA	2 / 2	2, 3, 1
7Y37A01091	AUNN	ThinkSystem 430-16e SAS/SATA 12Gb HBA	2 / 2	2, 3, 1
4Y37A09724	B8P7	ThinkSystem 440-16e SAS/SATA PCIe Gen4 12Gb HBA	2 / 2	2, 3, 1

* The maximum quantity shown is with one processor / two processors.

For a comparison of the functions of the supported external storage adapters, see the ThinkSystem RAID Adapter and HBA Reference:
https://lenovopress.com/lp1288#sr630-support=SR630&internal-or-external-ports=External

Configuration notes:

- Low profile SAS RAID controllers and HBAs for external storage are supported in the low profile and full-high PCIe x8 and x16 slots supplied by the riser cards 1 and 2.
- In the configurations without GPU installed, the total quantity of the RAID 730-8i 2GB, RAID 930-8i, RAID 930-16i, and RAID 930-8e controllers in the server must not exceed 2 (up to 2 supercapacitors can be mounted in the server).
- In the configurations with GPU installed, the total quantity of the RAID 730-8i 2GB, RAID 930-8i, RAID 930-16i, and RAID 930-8e controllers in the server must not exceed 1 (the storage controller can be installed only in the PCIe slot 1; the PCIe slots 3 and 4 cannot be used if the GPU adapter is installed).
- The RAID 930-8e controller cannot be used in the configurations with the RAID 730-8i 2GB controller.

For more information, see the list of Product Guides in the following categories:

- RAID adapters
  http://lenovopress.com/servers/options/raid#rt=product-guide
- Host bus adapters
  http://lenovopress.com/servers/options/hba#rt=product-guide

Fibre Channel host bus adapters

The following table lists Fibre Channel HBAs supported by the SR630 server.

Table 49. Fibre Channel HBAs

Part number	Feature code	Description	Maximum quantity*	I/O slots supported
64 Gb Fibre Channel - PCIe				
4XC7A77485	BLC1	ThinkSystem Emulex LPe36002 64Gb 2-port PCIe Fibre Channel Adapter	2 / 3	2, 3, 1
32 Gb Fibre Channel - PCIe				
4XC7A08250	B5SX	Emulex LPe35000 32Gb 1-port PCIe Fibre Channel Adapter	2 / 3	2, 3, 1
4XC7A08251	B5SY	Emulex LPe35002 32Gb 2-port PCIe Fibre Channel Adapter	2 / 3	2, 3, 1
7ZT7A00519	AUNV	Emulex LPe32002-M2-L PCIe 32Gb 2-Port SFP+ FC HBA	2 / 3	2, 3, 1
7ZT7A00516	AUNS	QLogic QLE2740 PCIe 32Gb 1-Port SFP+ FC HBA	2 / 3	2, 3, 1
7ZT7A00518	AUNU	QLogic QLE2742 PCIe 32Gb 2-Port SFP+ FC HBA	2 / 3	2, 3, 1
16 Gb Fibre Channel - PCIe				
01CV830	ATZU	Emulex 16Gb Gen6 FC Single-port HBA	2 / 3	2, 3, 1
01CV840	ATZV	Emulex 16Gb Gen6 FC Dual-port HBA	2 / 3	2, 3, 1
01CV750	ATZB	QLogic 16Gb Enhanced Gen5 FC Single-port HBA	2 / 3	2, 3, 1
01CV760	ATZC	QLogic 16Gb Enhanced Gen5 FC Dual-port HBA	2 / 3	2, 3, 1
8 Gb Fibre Channel - PCIe (available only in AP and PRC)				
4XC7A08221	B0X0	Emulex LPe12002-M8-L PCIe 8Gb 2-Port SFP+ FC HBA	2 / 3	2, 3, 1

* The maximum quantity shown is with one processor / two processors.

Configuration note: FC HBAs are supported in the low profile and full-high PCIe x8 and x16 slots supplied by the riser cards 1 and 2.

For more information, see the list of Product Guides in the Host bus adapters category:
http://lenovopress.com/servers/options/hba#rt=product-guide

Flash storage adapters

The SR630 server supports the flash storage adapters listed in the following table.

Table 50. Flash storage adapters

Part number	Feature code	Description	Maximum quantity*	I/O slots supported
Mainstream Flash Adapters - CM5-V				
4XB7A38234	BCGJ	CM5-V 1.6TB Mainstream NVMe PCIe 3.0 x4 Flash Adapter	2 / 3	2, 3, 1
4XB7A38237	BCGK	CM5-V 3.2TB Mainstream NVMe PCIe 3.0 x4 Flash Adapter	2 / 3	2, 3, 1
4XB7A38240	BCGL	CM5-V 6.4TB Mainstream NVMe PCIe 3.0 x4 Flash Adapter	2 / 3	2, 3, 1

* The maximum quantity shown is with one processor / two processors.

Configuration notes:

- Flash storage adapters are supported in the low profile and full-high PCIe x8 and x16 slots supplied by the riser cards 1 and 2.
- The Flash storage adapters are supported only in the environments with the air temperature of up to 35 °C (95 °F).

For more information, see the list of Product Guides in the Flash storage adapters category:
http://lenovopress.com/servers/options/ssdadapter#rt=product-guide

GPU adapters

The SR630 server supports graphics processing units (GPUs) listed in the following table.

Table 51. GPU adapters

Part number	Feature code	Description	Maximum quantity*	I/O slots supported
Full-height, full-length PCIe 3.0 x16 single-wide GPU adapters				
4X67A14934	B6CG	ThinkSystem NVIDIA Quadro RTX 4000 8GB PCIe Active GPU	1 / 1	2
Low profile PCIe 3.0 x16 single-wide GPU adapters				
4X67A11584	B31D	ThinkSystem NVIDIA Quadro P620 2GB PCIe Active GPU	2 / 3	3, 1, 2
4X67A14926	B4YB	ThinkSystem NVIDIA Tesla T4 16GB PCIe Passive GPU	1 / 2	3, 1
4X67A81547	BP05	ThinkSystem NVIDIA A2 16GB PCIe Gen4 Passive GPU	1 / 2	3, 1
GPU upgrade kits				
4XH7A08765	None**	ThinkSystem SR630 GPU Upgrade Kit	1	-

* The maximum quantity shown is with one processor / two processors.
** Field upgrade only.

Configuration notes:

- The full-height GPU adapters require the PCIe x16 FH riser card for the PCIe slot 2, and the PCIe slot 3 cannot be used. The low profile GPU adapters are supported in the PCIe x8 and x16 low-profile slots supplied by the riser cards 1 and 2.
- The full-height GPU adapters require the GPU Upgrade Kit (4XH7A08765).
- If the full-height GPU adapter is installed, the internal slot for a storage controller cannot be used.

- The GPU adapters are *not* supported in the following configurations:
  - Rear HDD kit installed.
  - 10x U.2 PCIe NVMe drive bays installed.
  - DCPMMs installed.
- Configurations with any one GPU adapter or two P620 GPU adapters are supported only with the 750 W or 1100 W power supplies. Configurations with two P4 or three P620 GPU adapters are supported only with the 1100 W power supplies.
- The P2000 or P2200 GPU adapters are supported with the processors of up to 140 W TDP. The GPU adapters other than P2000 or P2200 are supported with the processors of up to 165 W TDP (excluding Gold 6240Y, 6244, 6246, and 6252N processors).
- The RTX4000 GPU adapter is supported with the processors of up to 165 W TDP in the configurations with 8x 2.5-inch drive bays, or with the processors of up to 150 W TDP in the configurations with 4x 3.5-inch SAS/SATA or AnyBay, or 6x 2.5-inch SAS/SATA & 4x 2.5-inch AnyBay drive bays.
- The GPU adapters are supported only in the ASHRAE A2 environments (up to 35 °C [95 °F]) with the following exceptions:
  - Two P4 or T4 or A2 GPU adapters are supported at the ambient temperature of up to 30 °C [86 °F]), and the server performance might be impacted in case of a system fan failure.
  - Three P620 adapters are supported at the ambient temperature of up to 30 °C [86 °F]).

Cooling

The SR630 server supports up to seven hot-swap system fans that provide N+1 cooling redundancy. SR630 server models with one processor include five system fans, and server models with two processors include seven system fans. The fans are dual-rotor counter-rotating units, which means that the fans have two separate propellors, one in front of the other, and that the propellors rotate in opposite directions.

In the case of a system fan failure, performance might be impacted if any of the following processors are installed:

- 2nd Gen processors with 200 W or 205 W TDP, or Xeon 6240Y, 6244, 6246, or 6252N processors
- 1st Gen processors Xeon 6144, 6146, 6154, 8168, 8180, and 8180M

The installation of a 2nd processor requires an extra cooling fan be installed. For CTO orders, fans are derived by the configurator. For field upgrades, 1st Gen processor option part numbers include this fan however 2nd Gen processor options do not included the fan and it must be ordered separately using SR630 Fan Option Kit listed in the following table.

Table 52. Cooling options

Part number	Feature code	Description	Maximum quantity
4F17A12350	AUW7	ThinkSystem SR630 Fan Option Kit (2nd Gen processors only)	1

Power supplies and cables

The SR630 server supports up to two redundant power supplies and is capable of N+N redundancy depending on the configuration. A second power supply can be added to the models that come with one power supply.

The following table lists the power supply options.

Table 53. Power supplies

Part number	Feature code	Description	Maximum quantity
7N67A00882	AVW8	ThinkSystem 550W (230V/115V) Platinum Hot-Swap Power Supply	2
7N67A00883	AVWA	ThinkSystem 750W (230/115V) Platinum Hot-Swap Power Supply	2
7N67A00884	AVW9	ThinkSystem 750W (230V) Titanium Hot-Swap Power Supply	2
7N67A00885	AVWB	ThinkSystem 1100W (230V/115V) Platinum Hot-Swap Power Supply	2
4P57A15363*	B4Z5*	ThinkSystem 1100W -48V DC Power Supply	2

* -48V DC power supply is only available via Special Bid

Power supply options do not include a line cord. For server configurations, the inclusion of a power cord is model dependent. Configure-to-order models can be configured without power cords if desired.

Configuration notes:

- Minimum of 1 and maximum of 2 power supplies per system.
- If 2 are installed, power supplies must be identical.
- AC power supplies support AC (Worldwide) and HVDC (PRC only) power sources
- AC power supplies have a C14 connector. The -48V DC power supply has a Weidmuller TOP 4GS/3 7.6 terminal.

Important: The Standalone Solution Configuration Tool (SSCT) and Lenovo Data Center Solution Configurator (DCSC) power supply selection rules allow a subset of possible configurations due to power restrictions. Configurations that cannot be built in SSCT or DCSC due to power restrictions may still be supported. To verify support and ensure that the right power supply is chosen for optimal performance, you should always validate your server configuration using the latest version of the Lenovo Capacity Planner:
http://datacentersupport.lenovo.com/us/en/solutions/lnvo-lcp

Power cords

Line cords and rack power cables with C13 connectors can be ordered as listed in the following table.

110V customers: If you plan to use the ThinkSystem 1100W power supply with a 110V power source, select a power cable that is rated above 10A. Power cables that are rated at 10A or below are not supported with 110V power.

Table 54. Power cords

Part number	Feature code	Description
Rack cables		
00Y3043	A4VP	1.0m, 10A/100-250V, C13 to C14 Jumper Cord
39Y7937	6201	1.5m, 10A/100-250V, C13 to C14 Jumper Cord
4L67A08369	6570	2.0m, 13A/100-250V, C13 to C14 Jumper Cord
4L67A08366	6311	2.8m, 10A/100-250V, C13 to C14 Jumper Cord
4L67A08370	6400	2.8m, 13A/100-250V, C13 to C14 Jumper Cord
39Y7932	6263	4.3m, 10A/100-250V, C13 to C14 Jumper Cord
4L67A08371	6583	4.3m, 13A/100-250V, C13 to C14 Jumper Cord
Line cords		

Part number	Feature code	Description
39Y7930	6222	2.8m, 10A/250V, C13 to IRAM 2073 (Argentina) Line Cord
81Y2384	6492	4.3m, 10A/250V, C13 to IRAM 2073 (Argentina) Line Cord
39Y7924	6211	2.8m, 10A/250V, C13 to AS/NZS 3112 (Australia/NZ) Line Cord
81Y2383	6574	4.3m, 10A/250V, C13 to AS/NZS 3112 (Australia/NZ) Line Cord
69Y1988	6532	2.8m, 10A/250V, C13 to NBR 14136 (Brazil) Line Cord
81Y2387	6404	4.3m, 10A/250V, C13 to NBR 14136 (Brazil) Line Cord
39Y7928	6210	2.8m, 10A/220V, C13 to GB 2099.1 (China) Line Cord
81Y2378	6580	4.3m, 10A/250V, C13 to GB 2099.1 (China) Line Cord
39Y7918	6213	2.8m, 10A/250V, C13 to DK2-5a (Denmark) Line Cord
81Y2382	6575	4.3m, 10A/250V, C13 to DK2-5a (Denmark) Line Cord
39Y7917	6212	2.8m, 10A/250V, C13 to CEE 7/7 (Europe) Line Cord
81Y2376	6572	4.3m, 10A/250V, C13 to CEE 7/7 (Europe) Line Cord
39Y7927	6269	2.8m, 10A/250V, C13 to IS 6538 (India) Line Cord
81Y2386	6567	4.3m, 10A/250V, C13 to IS 6538 (India) Line Cord
39Y7920	6218	2.8m, 10A/250V, C13 to SI 32 (Israel) Line Cord
81Y2381	6579	4.3m, 10A/250V, C13 to SI 32 (Israel) Line Cord
39Y7921	6217	2.8m, 10A/250V, C13 to CEI 23-16 (Italy) Line Cord
81Y2380	6493	4.3m, 10A/250V, C13 to CEI 23-16 (Italy) Line Cord
4L67A08362	6495	4.3m, 12A/200V, C13 to JIS C-8303 (Japan) Line Cord
39Y7922	6214	2.8m, 10A/250V, C13 to SABS 164-1 (South Africa) Line Cord
81Y2379	6576	4.3m, 10A/250V, C13 to SANS 164-1 (South Africa) Line Cord
39Y7926	6335	4.3m, 12A/100V, C13 to JIS C-8303 (Japan) Line Cord
39Y7925	6219	2.8m, 12A/220V, C13 to KSC 8305 (S. Korea) Line Cord
81Y2385	6494	4.3m, 12A/250V, C13 to KSC 8305 (S. Korea) Line Cord
39Y7919	6216	2.8m, 10A/250V, C13 to SEV 1011-S24507 (Swiss) Line Cord
81Y2390	6578	4.3m, 10A/250V, C13 to SEV 1011-S24507 (Swiss) Line Cord
23R7158	6386	2.8m, 10A/125V, C13 to CNS 10917 (Taiwan) Line Cord
81Y2375	6317	2.8m, 10A/250V, C13 to CNS 10917 (Taiwan) Line Cord
81Y2374	6402	2.8m, 13A/125V, C13 to CNS 10917 (Taiwan) Line Cord
4L67A08363	AX8B	4.3m, 10A/125V, C13 to CNS 10917 (Taiwan) Line Cord
81Y2389	6531	4.3m, 10A/250V, C13 to CNS 10917 (Taiwan) Line Cord
81Y2388	6530	4.3m, 13A/125V, C13 to CNS 10917 (Taiwan) Line Cord
39Y7923	6215	2.8m, 10A/250V, C13 to BS 1363/A (UK) Line Cord
81Y2377	6577	4.3m, 10A/250V, C13 to BS 1363/A (UK) Line Cord
90Y3016	6313	2.8M, 10A/125V, C13 to NEMA 5-15P (US) Line Cord
46M2592	A1RF	2.8m, 10A/250V, C13 to NEMA 6-15P (US) Line Cord
00WH545	6401	2.8M, 13A/125V, C13 to NEMA 5-15P (US) Line Cord
4L67A08359	6370	4.3m, 10A/125V, C13 to NEMA 5-15P (US) Line Cord
4L67A08361	6373	4.3m, 10A/250V, C13 to NEMA 6-15P (US) Line Cord
4L67A08360	AX8A	4.3m, 13A/125V, C13 to NEMA 5-15P (US) Line Cord

For the -48V DC Power Supply, the following power cable is supported.

Table 55. -48V DC power cable

Part number	Feature code	Description
CTO only	B93F	ThinkSystem 2.5m,DC Cable

Systems management

The SR630 supports the following systems management tools:

- Lenovo XClarity Controller
- Light path diagnostics
- Lenovo XClarity Provisioning Manager
- Lenovo XClarity Essentials
- Lenovo XClarity Administrator
- Lenovo XClarity Integrators
- Lenovo XClarity Energy Manager
- Lenovo Capacity Planner

Lenovo XClarity Controller

The SR630 server contains Lenovo XClarity Controller (XCC), which provides advanced service-processor control, monitoring, and alerting functions. XClarity Controller offers three functional levels: Standard, Advanced, and Enterprise.

By default, the SR630 server includes XClarity Controller Standard features, and it can be upgraded to Advanced or Enterprise functionality by using the Features on Demand (FoD) upgrades.

XClarity Controller Standard offers the following capabilities:

- Gathering and viewing system information and inventory
- Monitoring system status and health
- Alerting and notifications
- Event logging
- Configuring network connectivity
- Configuring security
- Updating system firmware
- Configuring server settings and devices
- Real-time power usage monitoring
- Remotely controlling server power (Power on, Power off, Restart)
- Managing FoD activation keys
- Redirecting serial console via IPMI
- Capturing the video display contents when an operating system hang condition is detected

XClarity Controller Advanced Upgrade adds the following functionality to the Standard features:

- Remotely viewing video with the following graphics resolutions:
  - Up to 1600x1200 with up to 23 bits per pixel; or
  - Up to 1920x1200 with up to 15 bits per pixel
- Remotely accessing the server using the keyboard and mouse from a remote client
- Remotely deploying an operating system
- Syslog alerting
- Redirecting serial console via SSH
- Displaying graphics for real-time and historical power usage data and temperature

XClarity Controller Enterprise Upgrade adds the following functionality to the Advanced features:

- Capping power usage
- Mapping the ISO and image files located on the local client as virtual drives for use by the server
- Mounting the remote ISO and image files via HTTPS, SFTP, CIFS, and NFS
- Collaborating across up to six users of the virtual console
- Controlling quality and bandwidth usage

The XClarity Controller provides remote server management through industry-standard interfaces:

- Intelligent Platform Management Interface (IPMI) Version 2.0
- Simple Network Management Protocol (SNMP) Version 3
- Common Information Model (CIM)
- Data Center Manageability Interface (DCMI) Version 1.5
- Redfish REpresentational State Transfer (REST) API
- Web browser with HTML5 support
- Command-line interface
- Virtual Operator Panel with XClarity Mobile App via the front USB port with XClarity Controller access

Virtual Operator Panel provides quick access to system status, firmware, network, health, and alerts information. With proper authentication, it also allows to configure systems management and network settings and to control system power (Power on, Power off, Restart). The Virtual Operator Panel can be accessed from the XClarity Mobile App running on the Android or iOS mobile device that is connected to the front USB port with XClarity Controller access (See Components and connectors).

Note: Depending on the system settings, the front USB port can be assigned to XClarity Controller for management functions, or to the system as a regular USB 2.0 port, or switched between two functions by using the system ID button.

The following table lists the XClarity Controller FoD upgrades.

Table 56. XClarity Controller FoD upgrades

Description	Part number	Feature code	Maximum quantity
ThinkSystem XClarity Controller Standard to Advanced Upgrade	4L47A09132	AVUT	1
ThinkSystem XClarity Controller Standard to Enterprise Upgrade	None*	AUPW	1
ThinkSystem XClarity Controller Advanced to Enterprise Upgrade	4L47A09133	None**	1

* Factory-installed only.
** Field-upgrade only.

Configuration notes:

- For factory-installed upgrades, either Standard to Advanced Upgrade (feature AVUT) or Standard to Enterprise Upgrade (feature AUPW) can be selected, but not both.
- For field upgrades, the Advanced to Enterprise Upgrade (4L47A09133) requires the Standard to Advanced Upgrade to be activated on the server previously with either the factory-installed feature AVUT or field upgrade 4L47A09132.

Light path diagnostics

All SR630 server models include basic light path diagnostics, which provides the system LEDs on the front of the server (see Components and connectors) and the LEDs near the monitored components (for example, the DIMM error LEDs on the system board).

Lenovo XClarity Provisioning Manager

Lenovo XClarity Provisioning Manager is a UEFI-embedded GUI application that combines the functions of configuring system setup settings, configuring RAID, and updating applications and firmware. It also enables you to install the supported operating systems and associated device drivers, run diagnostics, and collect service data.

Lenovo XClarity Provisioning Manager has the following features:

- Automatic hardware detection
- Collecting and viewing system inventory information
- Configuring UEFI system setup settings
- Updating the system firmware
- Configuring RAID by using the RAID Setup Wizard or Advanced mode
- Installing an operating system and device drivers automatically or manually
- Running diagnostics and collecting service data

Lenovo XClarity Essentials

Lenovo offers the following XClarity Essentials software tools that can help you set up, use, and maintain the server at no additional cost:

- Lenovo XClarity Essentials OneCLI
  OneCLI is a collection of server management tools that utilize a command line interface program to manage firmware, hardware, and operating systems. It provides functions to collect full system health information (including health status), configure system setting, and update system firmware and drivers.
- Lenovo XClarity Essentials UpdateXpress
  The UpdateXpress tool is a standalone GUI application for firmware and device driver updates that enables you to maintain your server firmware and device drivers up-to-date and help you avoid unnecessary server outages. The tool acquires and deploys individual updates and UpdateXpress System Packs (UXSPs) which are integration-tested bundles.
- Lenovo XClarity Essentials Bootable Media Creator
  The Bootable Media Creator (BOMC) tool is used to create bootable media for offline firmware update.

For more information and downloads, visit the Lenovo XClarity Essentials web page:
http://support.lenovo.com/us/en/documents/LNVO-center

Lenovo XClarity Administrator

Lenovo XClarity is a centralized systems management solution that helps administrators deliver infrastructure faster. This solution integrates easily with Lenovo x86 servers, RackSwitch switches, and DS Series storage, providing automated agent-less discovery, monitoring, firmware updates, configuration management, and bare metal deployment of operating systems and hypervisors across multiple servers.

Lenovo XClarity Administrator is an optional software component for the SR630 server which can be downloaded and used at no charge to discover and monitor the SR630 and manage firmware upgrades for them.

If software support is required for Lenovo XClarity Administrator, or Lenovo XClarity Administrator premium features (such as configuration management and operating system deployment) are required, or both, Lenovo XClarity Pro software subscription should be ordered. Lenovo XClarity Pro is licensed on a per managed system basis, that is, each managed Lenovo system requires a license.

The following table lists the geo-specific Lenovo XClarity software license options.

Table 57. Lenovo XClarity software options

Description	Part number (NA, AP, Japan)*	Part number (EMEA, LA)**	Quantity
Lenovo XClarity Pro, per Managed Endpoint w/1 Yr SW S&S	00MT201	00MT207	1
Lenovo XClarity Pro, per Managed Endpoint w/3 Yr SW S&S	00MT202	00MT208	1
Lenovo XClarity Pro, per Managed Endpoint w/5 Yr SW S&S	00MT203	00MT209	1

* NA = North America; AP = Asia Pacific
** EMEA = Europe, Middle East, Africa; LA = Latin America

Lenovo XClarity Administrator offers the following standard features that are available at no charge:

- Auto-discovery and monitoring of Lenovo x86 servers, RackSwitch switches, Flex System chassis, and DS Series storage systems
- Firmware updates and compliance enforcement
- External alerts and notifications via SNMP traps, syslog remote logging, and e-mail
- Secure connections to managed endpoints
- NIST 800-131A or FIPS 140-2 compliant cryptographic standards between the management solution and managed endpoints
- Integration into existing higher-level management systems such as cloud automation and orchestration tools through REST APIs, providing extensive external visibility and control over hardware resources
- An intuitive, easy-to-use GUI
- Scripting with Windows PowerShell, providing command-line visibility and control over hardware resources

Lenovo XClarity Administrator offers the following premium features that require an optional Pro license:

- Pattern-based configuration management that allows to define configurations once and apply repeatedly without errors when deploying new servers or redeploying existing servers without disrupting the fabric
- Bare-metal deployment of operating systems and hypervisors to streamline infrastructure provisioning

For more information, refer to the Lenovo XClarity Administrator Product Guide:
http://lenovopress.com/tips1200

Lenovo XClarity Integrators

Lenovo offers at no charge (if software support is required, a Lenovo XClarity Pro software subscription license should be ordered) two software plug-in modules, Lenovo XClarity Integrators, to manage physical infrastructure from leading external virtualization management software tools from Microsoft and VMware:

- Lenovo XClarity Integrator for Microsoft System Center
- Lenovo XClarity Integrator for VMware vCenter

Lenovo XClarity Integrators offer the following additional features:

- Ability to discover, manage, and monitor Lenovo server hardware from VMware vCenter or Microsoft System Center
- Deployment of firmware updates and configuration patterns to Lenovo x86 rack servers and Flex System from the virtualization management tool

- Non-disruptive server maintenance in clustered environments that reduces workload downtime by dynamically migrating workloads from affected hosts during rolling server updates or reboots
- Greater service level uptime and assurance in clustered environments during unplanned hardware events by dynamically triggering workload migration from impacted hosts when impending hardware failures are predicted

For more information about XClarity Integrators, refer to the Lenovo Systems Management web page: https://www.lenovo.com/us/en/data-center/software/management/

Lenovo XClarity Energy Manager

Lenovo XClarity Energy Manager provides a stand-alone, web-based agent-less power management console that provides real time data and enables you to observe, plan and manage power and cooling for Lenovo servers. Using built-in intelligence, it identifies server power consumption trends and ideal power settings and performs cooling analysis so that you can define and optimize power-saving policies.

Lenovo XClarity Energy Manager offers the following capabilities:

- Monitors room, row, rack, and device levels in the data center
- Reports vital server information, such as power, temperature and resource utilization
- Monitors inlet temperature to locate hot spots, reducing the risk of data or device damage
- Provides finely-grained controls to limit platform power in compliance with IT policy
- Generates alerts when a user-defined threshold is reached

Lenovo XClarity Energy Manager is an optional software component for the SR630 server that is licensed on a per managed node basis, that is, each managed server requires a license. The 1-node Energy Manager license is included in the XClarity Controller Enterprise upgrade.

To manage systems without XClarity Controller Enterprise licenses, a node license pack should be purchased. The following table lists the geo-specific Lenovo XClarity Energy Manager software license options.

Table 58. Lenovo XClarity Energy Manager software options

Description	Part number (NA, AP, Japan)*	Part number (EMEA, LA)**	Quantity
Lenovo XClarity Energy Manager, 1 Node w/ 1 Yr S&S	01DA225	01DA228	1

* NA = North America; AP = Asia Pacific.
** EMEA = Europe, Middle East, Africa; LA = Latin America.

For more information, refer to the Lenovo XClarity Energy Manager web page: http://datacentersupport.lenovo.com/us/en/solutions/lnvo-lxem

Lenovo Capacity Planner

Lenovo Capacity Planner is a power consumption evaluation tool that enhances data center planning by enabling IT administrators and pre-sales professionals to understand various power characteristics of racks, servers, and other devices. Capacity Planner can dynamically calculate the power consumption, current, British Thermal Unit (BTU), and volt-ampere (VA) rating at the rack level, improving the planning efficiency for large scale deployments.

For more information, refer to the Capacity Planner web page: http://datacentersupport.lenovo.com/us/en/solutions/lnvo-lcp

Security

The SR630 server offers the following security features:

- Power-on password
- Administrator's password
- Secure firmware updates
- Onboard Trusted Platform Module (TPM) version 1.2 or 2.0 (configurable UEFI system setting)
- Trusted Cryptographic Module (TCM) (optional; PRC only)
- Nationz Trusted Platform Module v2.0 (optional; PRC only)
- Lockable front bezel (optional)
- Self-encrypting drives (SEDs) with support for enterprise key managers - see the SED encryption key management section
- Lenovo Business Vantage security software (optional; PRC only)

The server is NIST SP 800-147B compliant.

The following table lists the security options that are available for the SR630 server.

Table 59. Security options

Part number	Feature code	Description	Maximum quantity
Lockable front bezel			
7Z17A02581	AUWR	ThinkSystem 1U Security Bezel	1
Trusted Cryptographic Module (PRC only)			
None*	AVKE	ThinkSystem Trusted Cryptographic Module	1
Trusted Platform Module (PRC only)			
None*	B22N	ThinkSystem Nationz Trusted Platform Module v2.0	1

* Factory-installed only; no field upgrade.

Lenovo Business Vantage is a security software tool suite (available only in PRC) designed to work with the TCM or Nationz TPM for enhanced security, to keep user data safe, and to erase confidential data completely from a drive.

Lenovo Business Vantage provides the following features:

- Encrypts files to ensure data safety by using the TCM or Nationz TPM.
- Erases confidential data from a hard disk.
- Prohibits unauthorized access to the USB port of devices.
- Encrypts files to ensure data security on a USB storage device.

For more information, refer to the Lenovo Business Vantage web page:
http://support.lenovo.com.cn/lenovo/wsi/es/es.html

Intel Transparent Supply Chain

Add a layer of protection in your data center and have peace of mind that the server hardware you bring into it is safe authentic and with documented, testable, and provable origin.

Lenovo has one of the world's best supply chains, as ranked by Gartner Group, backed by extensive and mature supply chain security programs that exceed industry norms and US Government standards. Now we are the first Tier 1 manufacturer to offer Intel® Transparent Supply Chain in partnership with Intel, offering you an unprecedented degree of supply chain transparency and assurance.

To enable Intel Transparent Supply Chain for the Intel-based servers in your order, add the following feature code in the DCSC configurator, under the Security tab.

Table 60. Intel Transparent Supply Chain ordering information

Feature code	Description
BB0P	Intel Transparent Supply Chain

For more information on this offering, see the paper *Introduction to Intel Transparent Supply Chain on Lenovo ThinkSystem Servers*, available from https://lenovopress.com/lp1434-introduction-to-intel-transparent-supply-chain-on-thinksystem-servers.

Rack installation

The following table lists the rack installation options that are available for the SR630 server.

Table 61. Rack installation options

Part number	Feature code	Description	Maximum quantity
4-post rail kits			
7M27A05702	AXCA	ThinkSystem Tool-less Slide Rail	1
7M27A05701	AXCB	ThinkSystem Tool-less Slide Rail Kit with 1U CMA	1
4M17A07274	AXFN	ThinkSystem Screw-in Slide Rail	1
4M17A07281	B0TE	ThinkSystem Screw-in Slide Rail Kit with 1U CMA	1
4M17A07273	BK7W	ThinkSystem Toolless Friction Rail v2	1
Cable management arm (CMA) upgrade			
7M27A05699	None^	ThinkSystem 1U CMA Upgrade Kit for Tool-less Slide Rail	1*
4M17A07276	AXFP	ThinkSystem 1U CMA Upgrade Kit for Screw-in Slide Rail	1**
Front VGA port			
CTO only	BMNW	Front VGA Connector Upgrade Kit for 1U v2 (for 3.5" models)	
CTO only	AUWU	ThinkSystem SR530/SR630 Front VGA Connector (for 3.5" models)	1
4Z17A80447	BMNV	ThinkSystem SR530/SR570/SR630 Front VGA Connector Upgrade Kit v2 (for 2.5" models)	
7Z17A02579	AUWW	ThinkSystem SR530/SR570/SR630 Front VGA Connector Upgrade Kit (for 2.5" models)	1

^ Field upgrade only.
* The CMA Upgrade Kit for Tool-less Slide Rail is supported with the Tool-less Slide Rail (7M27A05702) only.
** The CMA Upgrade Kit for Screw-in Slide Rail is supported with the Screw-in Slide Rail (4M17A07274) only.

The following table summarizes the rail kit features and specifications.

Table 62. Rail kit features and specifications summary

Feature	Tool-less Slide Rail		Screw-in Slide Rail		Tool-less Friction Rail
	Without CMA	With CMA	Without CMA	With CMA	
Part number	7M27A05702	7M27A05701	4M17A07274	4M17A07281	4M17A07273
CMA	7M27A05699	Included	4M17A07276	Included	No support
Rail length	730 mm (28.74 in.)	807 mm (31.8 in.)	836.8 mm (32.9 in.)	836.8 mm (32.9 in.)	728.1 mm (28.7 in.)

Feature	Tool-less Slide Rail		Screw-in Slide Rail		Tool-less Friction Rail
	Without CMA	With CMA	Without CMA	With CMA	
Rail type	Full-out slide (ball bearing)		Full-out slide (ball bearing)		Half-out slide (friction)
Tool-less installation	Yes		No		Yes
In-rack server maintenance	Yes		Yes		No
1U PDU support	Yes		Yes		Yes
0U PDU support	Limited*		Limited*		Limited**
Rack type	IBM and Lenovo 4-post, IEC standard-compliant		IBM and Lenovo 4-post, IEC standard-compliant		IBM and Lenovo 4-post, IEC standard-compliant
Mounting holes	Square or round		Square, round, or threaded		Square or round
Mounting flange thickness	2 mm (0.08 in.) – 3.3 mm (0.13 in.)		2 mm (0.08 in.) – 3.3 mm (0.13 in.)		2 mm (0.08 in.) – 3.3 mm (0.13 in.)
Distance between front and rear mounting flanges^	609.6 mm (24 in.) – 863.6 mm (34 in.)		609.6 mm (24 in.) – 812.8 mm (32 in.)		609.6 mm (24 in.) – 863.6 mm (34 in.)

* If a 0U PDU is used, the rack cabinet must be at least 1100 mm (43.31 in.) deep if no CMA is used, or at least 1200 mm (47.24 in.) deep if a CMA is used.
** If a 0U PDU used, the rack must be at least 1000 mm (39.37 in.) deep.
^ Measured when mounted on the rack, from the front surface of the front mounting flange to the rear most point of the rail.

Operating systems

The SR630 server with 2nd Gen Intel Xeon SP processors supports the following operating systems:

- Microsoft Windows Server 2016
- Microsoft Windows Server 2019
- Microsoft Windows Server 2022
- Red Hat Enterprise Linux 7.6
- Red Hat Enterprise Linux 7.7
- Red Hat Enterprise Linux 7.8
- Red Hat Enterprise Linux 7.9
- Red Hat Enterprise Linux 8.0
- Red Hat Enterprise Linux 8.1
- Red Hat Enterprise Linux 8.2
- Red Hat Enterprise Linux 8.3
- Red Hat Enterprise Linux 8.4
- Red Hat Enterprise Linux 8.5
- SUSE Linux Enterprise Server 12 SP4
- SUSE Linux Enterprise Server 12 SP5
- SUSE Linux Enterprise Server 12 Xen SP4
- SUSE Linux Enterprise Server 12 Xen SP5
- SUSE Linux Enterprise Server 15
- SUSE Linux Enterprise Server 15 SP1
- SUSE Linux Enterprise Server 15 SP2
- SUSE Linux Enterprise Server 15 SP3
- SUSE Linux Enterprise Server 15 Xen
- SUSE Linux Enterprise Server 15 Xen SP1
- SUSE Linux Enterprise Server 15 Xen SP2
- SUSE Linux Enterprise Server 15 Xen SP3
- VMware ESXi 6.5 U2
- VMware ESXi 6.5 U3
- VMware ESXi 6.7 U1
- VMware ESXi 6.7 U2
- VMware ESXi 6.7 U3
- VMware ESXi 7.0

- VMware ESXi 7.0 U1
- VMware ESXi 7.0 U2
- VMware ESXi 7.0 U3

The SR630 server with 1st Gen Intel Xeon SP processors supports the following operating systems:

- Microsoft Windows Server 2012 R2
- Microsoft Windows Server 2016
- Microsoft Windows Server 2019
- Microsoft Windows Server 2022
- Microsoft Windows Server, version 1709
- Microsoft Windows Server, version 1803
- Red Hat Enterprise Linux 6.10 x64
- Red Hat Enterprise Linux 6.9 x64
- Red Hat Enterprise Linux 7.3
- Red Hat Enterprise Linux 7.4
- Red Hat Enterprise Linux 7.5
- Red Hat Enterprise Linux 7.6
- Red Hat Enterprise Linux 7.7
- Red Hat Enterprise Linux 7.8
- Red Hat Enterprise Linux 7.9
- Red Hat Enterprise Linux 8.0
- Red Hat Enterprise Linux 8.1
- Red Hat Enterprise Linux 8.2
- Red Hat Enterprise Linux 8.3
- Red Hat Enterprise Linux 8.4
- Red Hat Enterprise Linux 8.5
- SUSE Linux Enterprise Server 11 Xen x64 SP4
- SUSE Linux Enterprise Server 11 x64 SP4
- SUSE Linux Enterprise Server 12 SP2
- SUSE Linux Enterprise Server 12 SP3
- SUSE Linux Enterprise Server 12 SP4
- SUSE Linux Enterprise Server 12 SP5
- SUSE Linux Enterprise Server 12 Xen SP2
- SUSE Linux Enterprise Server 12 Xen SP3
- SUSE Linux Enterprise Server 12 Xen SP4
- SUSE Linux Enterprise Server 12 Xen SP5
- SUSE Linux Enterprise Server 15
- SUSE Linux Enterprise Server 15 SP1
- SUSE Linux Enterprise Server 15 SP2
- SUSE Linux Enterprise Server 15 SP3
- SUSE Linux Enterprise Server 15 Xen
- SUSE Linux Enterprise Server 15 Xen SP1
- SUSE Linux Enterprise Server 15 Xen SP2
- SUSE Linux Enterprise Server 15 Xen SP3
- VMware ESXi 6.0 U3
- VMware ESXi 6.5
- VMware ESXi 6.5 U1
- VMware ESXi 6.5 U2
- VMware ESXi 6.5 U3
- VMware ESXi 6.7
- VMware ESXi 6.7 U1
- VMware ESXi 6.7 U2
- VMware ESXi 6.7 U3
- VMware ESXi 7.0
- VMware ESXi 7.0 U1
- VMware ESXi 7.0 U2

- VMware ESXi 7.0 U3

For a complete list of supported, certified and tested operating systems, plus additional details and links to relevant web sites, see the Operating System Interoperability Guide:
https://lenovopress.com/osig#servers=sr630-7x01-7x02-sp-gen-2

For configure-to-order configurations, the server can be preloaded with VMware ESXi installed on M.2 cards. Ordering information is listed in the following table.

Table 63. VMware ESXi preload

Part number	Feature code	Description
CTO only	B3VW	VMware ESXi 6.5 U2 (Factory Installed)
CTO only	B6U0	VMware ESXi 6.5 U3 (factory installed)
CTO only	B3VX	VMware ESXi 6.7 (Factory Installed)
CTO only	B4XA	VMware ESXi 6.7 U1 (Factory Installed)
CTO only	B6U1	VMware ESXi 6.7 U2 (factory installed)
CTO only	B88T	VMware ESXi 6.7 U3 (factory installed)
CTO only	BBZG	VMware ESXi 7.0 (Factory Installed)
CTO only	BE5E	VMware ESXi 7.0 U1 (Factory Installed)
CTO only	BHSR	VMware ESXi 7.0 U2 (Factory Installed)

Physical and electrical specifications

The SR630 has the following overall physical dimensions, excluding components that extend outside the standard chassis, such as EIA flanges, front security bezel (if any), and power supply handles:

- Width: 435 mm (17.1 inches)
- Height: 43 mm (1.7 inches)
- Depth: 750 mm (29.5 inches)

The following table lists the detailed dimensions. See the figure below for the definition of each dimension.

Table 64. Detailed dimensions

Dimension	Description
482 mm	X_a = Width, to the outsides of the front EIA flanges
435 mm	X_b = Width, to the rack rail mating surfaces
435 mm	X_c = Width, to the outer most chassis body feature
43 mm	Y_a = Height, from the bottom of chassis to the top of the chassis
715 mm	Z_a = Depth, from the rack flange mating surface to the rearmost I/O port surface
716 mm	Z_b = Depth, from the rack flange mating surface to the rearmost feature of the chassis body
744 mm	Z_c = Depth, from the rack flange mating surface to the rearmost feature such as power supply handle
35 mm	Z_d = Depth, from the forwardmost feature on front of EIA flange to the rack flange mating surface
47 mm	Z_e = Depth, from the front of security bezel (if applicable) or forwardmost feature to the rack flange mating surface

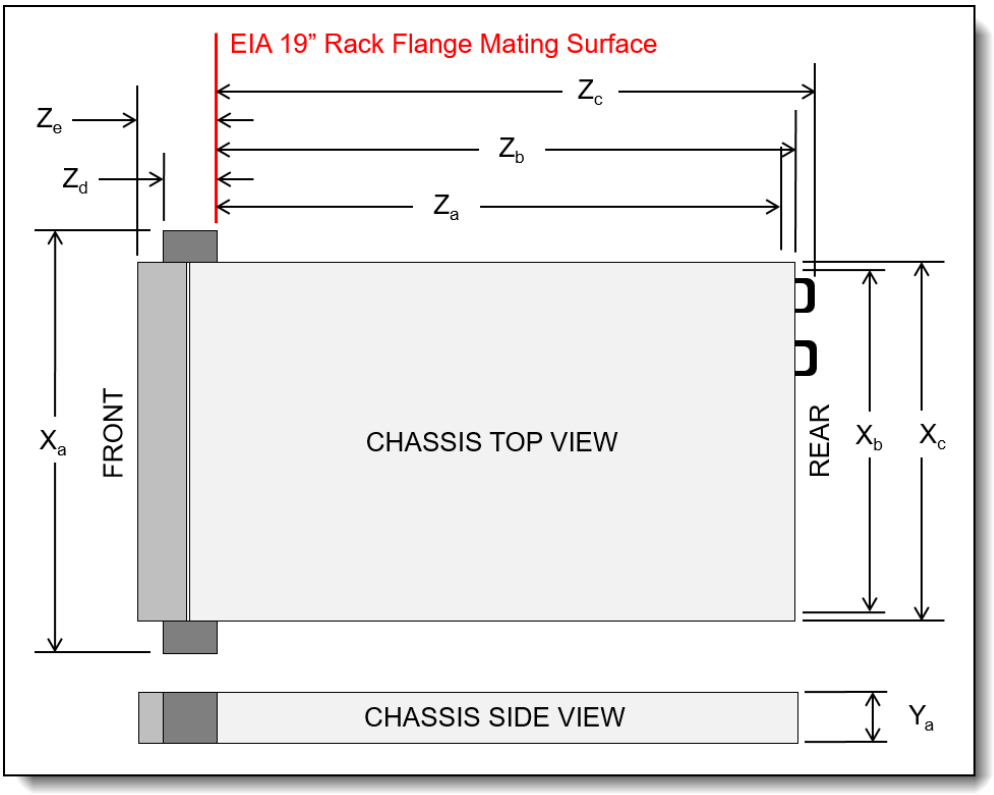


Figure 10. Server dimensions

The shipping dimensions (cardboard packaging) of the SR630 are as follows:

- Width: 587 mm (23.1 inches)
- Height: 225 mm (8.9 inches)
- Depth: 998 mm (39.3 inches)

The server has the following weight:

- Minimum configuration: 11.9 kg (26.2 lb)
- Maximum configuration: 18.8 kg (41.4 lb)

Electrical specifications for AC power supplies:

- 100 - 127 (nominal) V AC; 50 Hz / 60 Hz
- 200 - 240 (nominal) V AC; 50 Hz / 60 Hz
- 180 - 300 V DC (HVDC; supported in PRC only)

Power load and inlet current

The following table lists the maximum system power load, rated inlet current, and system heat output based on the power supply and source voltage.

Table 65. Rated system power, inlet current, and system heat output

Power supply	Source voltage	Maximum power load per system (two power supplies)	Rated current per inlet	System heat output
550 W Platinum	100 - 127 V AC	722 W	6.2 A	2463 BTU/hour
	200 - 240 V AC	704 W	3 A	2402 BTU/hour
	180 - 300 V DC	702 W	2.5 A	2395 BTU/hour
750 W Platinum	100 - 127 V AC	984 W	8.4 A	3357 BTU/hour
	200 - 240 V AC	958 W	4.1 A	3269 BTU/hour
	180 - 300 V DC	958 W	3.5 A	3269 BTU/hour
750 W Titanium	200 - 240 V AC	949 W	4.1 A	3238 BTU/hour
	180 - 300 V DC	948 W	3.5 A	3235 BTU/hour
1100 W Platinum	100 - 127 V AC	1382 W	12 A	4715 BTU/hour
	200 - 240 V AC	1408 W	6 A	4804 BTU/hour
	180 - 300 V DC	1408 W	5.1 A	4804 BTU/hour

Operating environment

The SR630 server complies with ASHRAE class A2 specifications. The server performance might be impacted when the operating temperature is outside the ASHRAE A2 specifications. Depending on the hardware configuration, some server models comply with ASHRAE class A3 and class A4 specifications. To comply with ASHRAE class A3 and class A4 specifications, the server models must meet the following hardware configuration requirements at the same time:

- Two power supplies installed
- Persistent memory modules not installed
- NVMe drives not installed
- NVMe PCIe flash adapters not installed
- Graphic processing units (GPUs) not installed
- Processors with TDP more than or equal to 150 W not installed
- For 8x 2.5" and 10x 2.5" chassis only: A maximum of one supercapacitor module installed

Temperature and humidity

The SR630 server is supported in the following environment:

- Air temperature:
  - Operating:
    - ASHRAE Class A4: 5 °C - 45 °C (41 °F - 113 °F); for altitudes above 900 m (2,953 ft), decrease the maximum ambient temperature by 1 °C for every 125-m (410-ft) increase in altitude
    - ASHRAE Class A3: 5 °C - 40 °C (41 °F - 104 °F); for altitudes above 900 m (2,953 ft), decrease the maximum ambient temperature by 1 °C for every 175-m (574-ft) increase in altitude
    - ASHRAE Class A2: 10 °C - 35 °C (50 °F - 95 °F); for altitudes above 900 m (2,953 ft), decrease the maximum ambient temperature by 1 °C for every 300-m (984-ft) increase in altitude
  - Non-operating: 5 °C - 45 °C (41 °F - 113 °F)

    - Storage: -40 °C - +60 °C (-40 °F - 140 °F)
- Maximum altitude: 3,050 m (10,000 ft)
- Humidity:
    - Operating:
        - ASHRAE Class A4: 8% - 90% (non-condensing); maximum dew point: 24 °C (75 °F)
        - ASHRAE Class A3: 8% - 85% (non-condensing); maximum dew point: 24 °C (75 °F)
        - ASHRAE Class A2: 8% - 80% (non-condensing); maximum dew point: 21 °C (70 °F)
    - Storage: 8% - 90% (non-condensing)

Acoustical noise emissions

The server has the following acoustic noise emissions declaration:

- Minimum configuration:
    - Operating: 5.3 bels
    - Idle: 4.9 bels
- Maximum configuration:
    - Operating: 6.0 bels
    - Idle: 5.8 bels

Shock and vibration

The server has the following vibration and shock limits:

- Vibration:
    - Operating: 0.21 G rms at 5 Hz to 500 Hz for 15 minutes across 3 axes
    - Non-operating: 1.04 G rms at 2 Hz to 200 Hz for 15 minutes across 6 surfaces
- Shock:
    - Operating: 15 G for 3 milliseconds in each direction (positive and negative X, Y, and Z axes)
    - Non-operating:
        - 12 kg - 22 kg: 50 G for 152 in./sec velocity change across 6 surfaces
        - 23 kg - 31 kg: 35 G for 152 in./sec velocity change across 6 surfaces

Particulate contamination

Airborne particulates (including metal flakes or particles) and reactive gases acting alone or in combination with other environmental factors such as humidity or temperature might damage the system that might cause the system to malfunction or stop working altogether.

The following specifications indicate the limits of particulates that the system can tolerate:

- Reactive gases:
    - The reactivity rate of copper coupons shall be less than 200 Angstroms per month (Å/month)
    - The reactivity rate of silver coupons shall be less than 200 Å/month
- Airborne particulates:
    - The room air should be continuously filtered with MERV 8 filters.
    - Air entering a data center should be filtered with MERV 11 or preferably MERV 13 filters.
    - The deliquescent relative humidity of the particulate contamination should be more than 60% RH
    - Data centers must be free of zinc whiskers

For additional information, see the Specifications section of the Setup Guide for the server, available from the Lenovo ThinkSystem Information Center, https://thinksystem.lenovofiles.com/help/index.jsp

Warranty and support

The SR630 server has a one-year (7X01) or three-year (Machine Type 7X02) warranty.

The standard warranty terms are customer-replaceable unit (CRU) and onsite (for field-replaceable units FRUs only) with standard call center support during normal business hours and 9x5 Next Business Day Parts Delivered.

Lenovo's additional support services provide a sophisticated, unified support structure for your data center, with an experience consistently ranked number one in customer satisfaction worldwide. Available offerings include:

- **Premier Support**

  Premier Support provides a Lenovo-owned customer experience and delivers direct access to technicians skilled in hardware, software, and advanced troubleshooting, in addition to the following:

  - Direct technician-to-technician access through a dedicated phone line
  - 24x7x365 remote support
  - Single point of contact service
  - End to end case management
  - Third-party collaborative software support
  - Online case tools and live chat support
  - On-demand remote system analysis

- **Warranty Upgrade (Preconfigured Support)**

  Services are available to meet the on-site response time targets that match the criticality of your systems.

  - 3, 4, or 5 years of service coverage
  - 1-year or 2-year post-warranty extensions
  - **Foundation Service**: 9x5 service coverage with next business day onsite response. YourDrive YourData is an optional extra (see below).
  - **Essential Service:** 24x7 service coverage with 4-hour onsite response or 24-hour committed repair (available only in select markets). Bundled with YourDrive YourData.
  - **Advanced Service:** 24x7 service coverage with 2-hour onsite response or 6-hour committed repair (available only in select markets). Bundled with YourDrive YourData.

- **Managed Services**

  Lenovo Managed Services provides continuous 24x7 remote monitoring (plus 24x7 call center availability) and proactive management of your data center using state-of-the-art tools, systems, and practices by a team of highly skilled and experienced Lenovo services professionals.

  Quarterly reviews check error logs, verify firmware & OS device driver levels, and software as needed. We'll also maintain records of latest patches, critical updates, and firmware levels, to ensure you systems are providing business value through optimized performance.

- **Technical Account Management (TAM)**

  A Lenovo Technical Account Manager helps you optimize the operation of your data center based on a deep understanding of your business. You gain direct access to your Lenovo TAM, who serves as your single point of contact to expedite service requests, provide status updates, and furnish reports to track incidents over time. In addition, your TAM will help proactively make service recommendations and manage your service relationship with Lenovo to make certain your needs are met.

- **Enterprise Server Software Support**

  Enterprise Software Support is an additional support service providing customers with software support on Microsoft, Red Hat, SUSE, and VMware applications and systems. Around the clock availability for critical problems plus unlimited calls and incidents helps customers address challenges fast, without incremental costs. Support staff can answer troubleshooting and diagnostic questions, address product comparability and interoperability issues, isolate causes of problems, report defects to software vendors, and more.

- **YourDrive YourData**

  Lenovo's YourDrive YourData is a multi-drive retention offering that ensures your data is always under your control, regardless of the number of drives that are installed in your Lenovo server. In the unlikely event of a drive failure, you retain possession of your drive while Lenovo replaces the failed drive part. Your data stays safely on your premises, in your hands. The YourDrive YourData service can be purchased in convenient bundles and is optional with Foundation Service. It is bundled with Essential Service and Advanced Service.

- **Health Check**

  Having a trusted partner who can perform regular and detailed health checks is central to maintaining efficiency and ensuring that your systems and business are always running at their best. Health Check supports Lenovo-branded server, storage, and networking devices, as well as select Lenovo-supported products from other vendors that are sold by Lenovo or a Lenovo-Authorized Reseller.

Examples of region-specific warranty terms are second or longer business day parts delivery or parts-only base warranty.

If warranty terms and conditions include onsite labor for repair or replacement of parts, Lenovo will dispatch a service technician to the customer site to perform the replacement. Onsite labor under base warranty is limited to labor for replacement of parts that have been determined to be field-replaceable units (FRUs). Parts that are determined to be customer-replaceable units (CRUs) do not include onsite labor under base warranty.

If warranty terms include parts-only base warranty, Lenovo is responsible for delivering only replacement parts that are under base warranty (including FRUs) that will be sent to a requested location for self-service. Parts-only service does not include a service technician being dispatched onsite. Parts must be changed at customer's own cost and labor and defective parts must be returned following the instructions supplied with the spare parts.

Lenovo Service offerings are region-specific. Not all preconfigured support and upgrade options are available in every region. For information about Lenovo service upgrade offerings that are available in your region, refer to the following resources:

- Service part numbers in Lenovo Data Center Solution Configurator (DCSC):
  http://dcsc.lenovo.com/#/services
- Lenovo Services Availability Locator
  http://lenovolocator.com/

For service definitions, region-specific details, and service limitations, please refer to the following documents:

- Lenovo Statement of Limited Warranty for Infrastructure Solutions Group (ISG) Servers and System Storage
  http://pcsupport.lenovo.com/us/en/solutions/ht503310
- Lenovo Data Center Services Agreement
  http://support.lenovo.com/us/en/solutions/ht116628

Services

Lenovo Services is a dedicated partner to your success. Our goal is to reduce your capital outlays, mitigate your IT risks, and accelerate your time to productivity.

Note: Some service options may not be available in all countries. For more information, go to https://www.lenovo.com/services. For information about Lenovo service upgrade offerings that are available in your region, contact your local Lenovo sales representative or business partner.

Here's a more in-depth look at what we can do for you:

- **Asset Recovery Services**

  Asset Recovery Services (ARS) helps customers recover the maximum value from their end-of-life equipment in a cost-effective and secure way. On top of simplifying the transition from old to new equipment, ARS mitigates environmental and data security risks associated with data center equipment disposal. Lenovo ARS is a cash-back solution for equipment based on its remaining market value, yielding maximum value from aging assets and lowering total cost of ownership for your customers. For more information, see the ARS page, https://lenovopress.com/lp1266-reduce-e-waste-and-grow-your-bottom-line-with-lenovo-ars.

- **Assessment Services**

  An Assessment helps solve your IT challenges through an onsite, multi-day session with a Lenovo technology expert. We perform a tools-based assessment which provides a comprehensive and thorough review of a company's environment and technology systems. In addition to the technology based functional requirements, the consultant also discusses and records the non-functional business requirements, challenges, and constraints. Assessments help organizations like yours, no matter how large or small, get a better return on your IT investment and overcome challenges in the ever-changing technology landscape.

- **Design Services**

  Professional Services consultants perform infrastructure design and implementation planning to support your strategy. The high-level architectures provided by the assessment service are turned into low level designs and wiring diagrams, which are reviewed and approved prior to implementation. The implementation plan will demonstrate an outcome-based proposal to provide business capabilities through infrastructure with a risk-mitigated project plan.

- **Basic Hardware Installation**

  Lenovo experts can seamlessly manage the physical installation of your server, storage, or networking hardware. Working at a time convenient for you (business hours or off shift), the technician will unpack and inspect the systems on your site, install options, mount in a rack cabinet, connect to power and network, check and update firmware to the latest levels, verify operation, and dispose of the packaging, allowing your team to focus on other priorities.

- **Deployment Services**

  When investing in new IT infrastructures, you need to ensure your business will see quick time to value with little to no disruption. Lenovo deployments are designed by development and engineering teams who know our Products & Solutions better than anyone else, and our technicians own the process from delivery to completion. Lenovo will conduct remote preparation and planning, configure & integrate systems, validate systems, verify and update appliance firmware, train on administrative tasks, and provide post-deployment documentation. Customer's IT teams leverage our skills to enable IT staff to transform with higher level roles and tasks.

- **Integration, Migration, and Expansion Services**

  Move existing physical & virtual workloads easily, or determine technical requirements to support increased workloads while maximizing performance. Includes tuning, validation, and documenting ongoing run processes. Leverage migration assessment planning documents to perform necessary migrations.

Regulatory compliance

The ThinkSystem SR630 server conforms to the following regulations:

- United States: FCC Part 15, Class A; UL 60950-1
- Canada: ICES-003/NMB-03, Class A; CAN/CSA-C22.2 60950-1
- Mexico: NOM-19
- Argentina: IEC60950-1
- European Union: CE Mark (EN55022 Class A, IEC/EN60950-1, EN55024, EN61000-3-2, EN61000-3-3)
- Germany: TUV-GS (IEC/EN 60950-1, EK1-ITB2000)
- Russia, Kazakhstan, Belarus: EAC (TR CU 004/2011, TR CU 020/2011)
- China: CCC GB4943.1, GB9254 Class A, GB17625.1
- India: BIS
- Japan: VCCI, Class A
- Taiwan: BSMI CNS13438, Class A; CNS14336-1
- Korea: KN22, Class A; KN24
- Australia/New Zealand: AS/NZS CISPR 22 Class A
- Reduction of Hazardous Substances (ROHS)
- Energy Star 3.0 (excluding configurations with Bronze 3204, Gold 5222, or Platinum 8256 processors)

Note: For more information on the Energy Star 3.0 certification, refer to the *Energy Star 3.0 Certifications for ThinkSystem Servers* publication:
http://lenovopress.com/lp1230

External drive enclosures

The server supports attachment to external drive enclosures using a RAID controller with external ports or a SAS host bus adapter. Adapters supported by the server are listed in the SAS adapters for external storage section.

Note: Information provided in this section is for ordering reference purposes only. For the operating system and adapter support details, refer to the interoperability matrix for a particular storage enclosure that can be found on the Lenovo Data Center Support web site:
http://datacentersupport.lenovo.com

Table 66. External drive enclosures

Description	Part number		
	Worldwide	Japan	PRC
Lenovo Storage D1212 LFF Disk Expansion with Dual SAS IO Modules	4587A11	4587A1J	4587A1C
Lenovo Storage D1224 SFF Disk Expansion with Dual SAS IO Modules	4587A31	4587A3J	4587A3C
Lenovo Storage D3284 4TB x 84 HD Expansion Enclosure	641311F		
Lenovo Storage D3284 6TB x 84 HD Expansion Enclosure	641312F		
Lenovo Storage D3284 8TB x 84 HD Expansion Enclosure	641313F		
Lenovo Storage D3284 10TB x 84 HD Expansion Enclosure	641314F		

For details about supported drives, adapters, and cables, see the following Lenovo Press Product Guides:

- Lenovo Storage D1212 and D1224
  http://lenovopress.com/lp0512
- Lenovo Storage D3284
  http://lenovopress.com/lp0513

External storage systems

Lenovo offers the ThinkSystem DE Series and ThinkSystem DM Series external storage systems for high-performance storage. See the DE Series and DM Series product guides for specific controller models, expansion enclosures and configuration options:

- ThinkSystem DE Series Storage
  https://lenovopress.com/storage/thinksystem/de-series#rt=product-guide
- ThinkSystem DM Series Storage
  https://lenovopress.com/storage/thinksystem/dm-series#rt=product-guide

External backup units

The following table lists the external backup options that are offered by Lenovo.

Table 67. External backup options

Part number	Description
External RDX USB drives	
4T27A10725	ThinkSystem RDX External USB 3.0 Dock
External SAS tape backup drives	
6160S6E	IBM TS2260 Tape Drive Model H6S
6160S7E	IBM TS2270 Tape Drive Model H7S
6160S8E	IBM TS2280 Tape Drive Model H8S
External SAS tape backup autoloaders	
6171S5R	IBM TS2900 Tape Autoloader w/LTO5 HH SAS
6171S6R	IBM TS2900 Tape Autoloader w/LTO6 HH SAS
6171S7R	IBM TS2900 Tape Autoloader w/LTO7 HH SAS
External tape backup libraries	
6741A1F	IBM TS4300 3U Tape Library-Base Unit
6741A3F	IBM TS4300 3U Tape Library-Expansion Unit
Full High 8 Gb Fibre Channel for TS4300	
01KP954	LTO 8 FH Fibre Channel Drive
01KP938	LTO 7 FH Fibre Channel Drive
01KP935	LTO 6 FH Fibre Channel Drive
Half High 8 Gb Fibre Channel for TS4300	
01KP952	LTO 8 HH Fibre Channel Drive
01KP936	LTO 7 HH Fibre Channel Drive
01KP933	LTO 6 HH Fibre Channel Drive
Half High 6 Gb SAS for TS4300	
01KP953	LTO 8 HH SAS Drive
01KP937	LTO 7 HH SAS Drive
01KP934	LTO 6 HH SAS Drive

For more information, see the list of Product Guides in the Backup units category:
https://lenovopress.com/servers/options/backup

Fibre Channel SAN switches

Lenovo offers the ThinkSystem DB Series of Fibre Channel SAN switches for high-performance storage expansion. See the DB Series product guides for models and configuration options:

- ThinkSystem DB Series SAN Switches:
  https://lenovopress.com/storage/switches/rack#rt=product-guide

Rack cabinets

The following table lists the supported rack cabinets.

Table 68. Rack cabinets

Part number	Description
7D2B0001WW / 7D2N0001WW	12U 1200mm Deep Micro Datacenter Rack
7D2C0001WW / 7D2P0001WW	18U 1200mm Deep Micro Datacenter Rack
93072RX	25U Standard Rack
93072PX	25U Static S2 Standard Rack
93634PX	42U 1100mm Dynamic Rack
93634EX	42U 1100mm Dynamic Expansion Rack
93604PX	42U 1200mm Deep Dynamic Rack
93614PX	42U 1200mm Deep Static Rack
93084EX	42U Enterprise Expansion Rack
93084PX	42U Enterprise Rack
93074RX	42U Standard Rack

For specifications about these racks, see the Lenovo Rack Cabinet Reference, available from:
https://lenovopress.com/lp1287-lenovo-rack-cabinet-reference

For more information, see the list of Product Guides in the Rack cabinets category:
https://lenovopress.com/servers/options/racks

KVM switches and consoles

The following table lists the supported KVM consoles.

Table 69. KVM console

Part number	Description
4XF7A73009	ThinkSystem 18.5" LCD Console (with English keyboard)

The following table lists the available KVM switches and the options that are supported with them.

Table 71. KVM switches and options

Part number	Description
KVM Console switches	
1754D1T	ThinkSystem Digital 2x1x16 KVM Switch (DVI video output port)
1754A1T	ThinkSystem Analog 1x8 KVM Switch (DVI video output port)
1754D2X	Global 4x2x32 Console Manager (GCM32)
1754D1X	Global 2x2x16 Console Manager (GCM16)
1754A2X	Local 2x16 Console Manager (LCM16)
1754A1X	Local 1x8 Console Manager (LCM8)
Cables for ThinkSystem Digital and Analog KVM Console switches	
4X97A11108	ThinkSystem VGA to DVI Conversion Cable
4X97A11109	ThinkSystem Single-USB Conversion Cable for Digital KVM
4X97A11107	ThinkSystem Dual-USB Conversion Cable for Digital KVM
4X97A11106	ThinkSystem USB Conversion Cable for Analog KVM
Cables for GCM and LCM Console switches	
46M5383	Virtual Media Conversion Option Gen2 (VCO2)
46M5382	Serial Conversion Option (SCO)

For more information, see the list of Product Guides in the KVM Switches and Consoles category: http://lenovopress.com/servers/options/kvm

Power distribution units

The following table lists the power distribution units (PDUs) that are offered by Lenovo.

Table 72. Power distribution units

Part number	Feature code	Description	ANZ	ASEAN	Brazil	EET	MEA	RUCIS	WE	HTK	INDIA	JAPAN	LA	NA	PRC
0U Basic PDUs															
00YJ776	ATZY	0U 36 C13/6 C19 24A 1 Phase PDU	N	Y	Y	N	N	N	N	N	N	Y	Y	Y	N
00YJ777	ATZZ	0U 36 C13/6 C19 32A 1 Phase PDU	Y	Y	N	Y	Y	Y	Y	Y	Y	N	N	Y	Y
00YJ778	AU00	0U 21 C13/12 C19 32A 3 Phase PDU	Y	Y	N	Y	Y	Y	Y	Y	Y	N	N	Y	Y
00YJ779	ATZX	0U 21 C13/12 C19 48A 3 Phase PDU	N	N	Y	N	N	N	Y	N	N	Y	Y	Y	N
0U Switched and Monitored PDUs															
00YJ780	AU01	0U 20 C13/4 C19 Switched and Monitored 32A 1 Phase PDU	Y	Y	Y	Y	Y	Y	Y	Y	Y	N	Y	N	Y
00YJ781	AU03	0U 20 C13/4 C19 Switched and Monitored 24A 1 Phase PDU	N	N	Y	N	Y	N	Y	N	N	Y	Y	Y	N
00YJ782	AU02	0U 18 C13/6 C19 Switched and Monitored 32A 3 Phase PDU	Y	Y	Y	Y	Y	Y	Y	Y	Y	N	Y	N	Y
00YJ783	AU04	0U 12 C13/12 C19 Switched and Monitored 48A 3 Phase PDU	N	N	Y	N	N	N	Y	N	N	Y	Y	Y	N
1U Ultra Density Enterprise PDUs (9x IEC 320 C13 + 3x IEC 320 C19 outlets)															

Part number	Feature code	Description	ANZ	ASEAN	Brazil	EET	MEA	RUCIS	WE	HTK	INDIA	JAPAN	LA	NA	PRC
71763NU	6051	Ultra Density Enterprise C19/C13 PDU 60A/208V/3PH	N	N	Y	N	N	N	N	N	N	Y	Y	Y	N
71762NX	6091	Ultra Density Enterprise C19/C13 PDU Module	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
1U C13 Enterprise PDUs (12x IEC 320 C13 outlets)															
39M2816	6030	DPI C13 Enterprise PDU Plus Module (WW)	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
39Y8941	6010	DPI C13 Enterprise PDU Module (WW)	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
1U C19 Enterprise PDUs (6x IEC 320 C19 outlets)															
39Y8948	6060	DPI C19 Enterprise PDU Module (WW)	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
39Y8923	6061	DPI Three-phase 60A/208V C19 Enterprise PDU (US)	N	N	Y	N	N	N	Y	N	N	N	Y	Y	N
1U Front-end PDUs (3x IEC 320 C19 outlets)															
39Y8938	6002	DPI Single-phase 30A/120V Front-end PDU (US)	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
39Y8939	6003	DPI Single-phase 30A/208V Front-end PDU (US)	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
39Y8934	6005	DPI Single-phase 32A/230V Front-end PDU (International)	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
39Y8940	6004	DPI Single-phase 60A/208V Front-end PDU (US)	Y	N	Y	Y	Y	Y	Y	N	N	Y	Y	Y	N
39Y8935	6006	DPI Single-phase 63A/230V Front-end PDU (International)	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
1U NEMA PDUs (6x NEMA 5-15R outlets)															
39Y8905	5900	DPI 100-127V NEMA PDU	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Line cords for 1U PDUs that ship without a line cord															
40K9611	6504	4.3m, 32A/380-415V, EPDU/IEC 309 3P+N+G 3ph wye (non-US) Line Cord	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
40K9612	6502	4.3m, 32A/230V, EPDU to IEC 309 P+N+G (non-US) Line Cord	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
40K9613	6503	4.3m, 63A/230V, EPDU to IEC 309 P+N+G (non-US) Line Cord	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
40K9614	6500	4.3m, 30A/208V, EPDU to NEMA L6-30P (US) Line Cord	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
40K9615	6501	4.3m, 60A/208V, EPDU to IEC 309 2P+G (US) Line Cord	N	N	Y	N	N	N	Y	N	N	Y	Y	Y	N
40K9617	6505	4.3m, 32A/230V, Souriau UTG Female to AS/NZ 3112 (Aus/NZ) Line Cord	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
40K9618	6506	4.3m, 32A/250V, Souriau UTG Female to KSC 8305 (S. Korea) Line Cord	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y

For more information, see the Lenovo Press documents in the PDU category:
https://lenovopress.com/servers/options/pdu

Uninterruptible power supply units

The following table lists the uninterruptible power supply (UPS) units that are offered by Lenovo.

Table 73. Uninterruptible power supply units

Part number	Description
55941AX	RT1.5kVA 2U Rack or Tower UPS (100-125VAC)
55941KX	RT1.5kVA 2U Rack or Tower UPS (200-240VAC)
55942AX	RT2.2kVA 2U Rack or Tower UPS (100-125VAC)
55942KX	RT2.2kVA 2U Rack or Tower UPS (200-240VAC)
55943AX	RT3kVA 2U Rack or Tower UPS (100-125VAC)
55943KX	RT3kVA 2U Rack or Tower UPS (200-240VAC)
55945KX	RT5kVA 3U Rack or Tower UPS (200-240VAC)
55946KX	RT6kVA 3U Rack or Tower UPS (200-240VAC)
55948KX	RT8kVA 6U Rack or Tower UPS (200-240VAC)
55949KX	RT11kVA 6U Rack or Tower UPS (200-240VAC)
55948PX	RT8kVA 6U 3:1 Phase Rack or Tower UPS (380-415VAC)
55949PX	RT11kVA 6U 3:1 Phase Rack or Tower UPS (380-415VAC)
55943KT†	ThinkSystem RT3kVA 2U Standard UPS (200-230VAC) (2x C13 10A, 2x GB 10A, 1x C19 16A outlets)
55943LT†	ThinkSystem RT3kVA 2U Long Backup UPS (200-230VAC) (2x C13 10A, 2x GB 10A, 1x C19 16A outlets)
55946KT†	ThinkSystem RT6kVA 5U UPS (200-230VAC) (2x C13 10A outlets, 1x Terminal Block output)
5594XKT†	ThinkSystem RT10kVA 5U UPS (200-230VAC) (2x C13 10A outlets, 1x Terminal Block output)

† Only available in China and countries in the Asia Pacific region.

For more information, see the list of Product Guides in the UPS category:
https://lenovopress.com/servers/options/ups

Lenovo Financial Services

Lenovo Financial Services reinforces Lenovo's commitment to deliver pioneering products and services that are recognized for their quality, excellence, and trustworthiness. Lenovo Financial Services offers financing solutions and services that complement your technology solution anywhere in the world.

We are dedicated to delivering a positive finance experience for customers like you who want to maximize your purchase power by obtaining the technology you need today, protect against technology obsolescence, and preserve your capital for other uses.

We work with businesses, non-profit organizations, governments and educational institutions to finance their entire technology solution. We focus on making it easy to do business with us. Our highly experienced team of finance professionals operates in a work culture that emphasizes the importance of providing outstanding customer service. Our systems, processes and flexible policies support our goal of providing customers with a positive experience.

We finance your entire solution. Unlike others, we allow you to bundle everything you need from hardware and software to service contracts, installation costs, training fees, and sales tax. If you decide weeks or months later to add to your solution, we can consolidate everything into a single invoice.

Our Premier Client services provide large accounts with special handling services to ensure these complex transactions are serviced properly. As a premier client, you have a dedicated finance specialist who manages your account through its life, from first invoice through asset return or purchase. This specialist develops an in-depth understanding of your invoice and payment requirements. For you, this dedication provides a high-quality, easy, and positive financing experience.

For your region-specific offers, please ask your Lenovo sales representative or your technology provider about the use of Lenovo Financial Services. For more information, see the following Lenovo website:

https://www.lenovo.com/us/en/landingpage/lenovo-financial-services/

Related publications and links

For more information, see these resources:

- ThinkSystem SR630 product page
  https://www.lenovo.com/us/en/data-center/servers/racks/ThinkSystem-SR630/p/77XX7SRSR63
- Datasheet for the ThinkSystem SR630:
  https://lenovopress.com/ds0031-lenovo-thinksystem-sr630
- 3D Interactive Tour of the ThinkSystem SR630:
  https://lenovopress.com/lp0672-3d-tour-thinksystem-sr630
- Walkthrough Video for the ThinkSystem SR630:
  https://lenovopress.com/lp0699-thinksystem-sr630-server-video-walkthrough
- User Manuals for the ThinkSystem SR630:
  https://thinksystem.lenovofiles.com/help/topic/7X01/introduction.html?cp=4_6
  - Quick Start Guide
  - Setup Guide
  - Rack Installation Guides
  - Maintenance Manual
  - Messages and Codes Reference
  - UEFI Manual
- Lenovo Data Center Support Downloads - ThinkSystem SR630:
  http://datacentersupport.lenovo.com/products/servers/thinksystem/sr630/7x01/downloads
  http://datacentersupport.lenovo.com/products/servers/thinksystem/sr630/7x02/downloads
- Lenovo Hardware Installation & Removal Videos on the ThinkSystem SR630:
  - YouTube: https://www.youtube.com/playlist?list=PLYV5R7hVcs-C_GGjNudxYy8-c69INMcmE
  - Youku: https://list.youku.com/albumlist/show/id_50462008
- Lenovo Data Center Solution Configurator (DCSC):
  http://dcsc.lenovo.com

Related product families

Product families related to this document are the following:

- 2-Socket Rack Servers
- ThinkSystem SR630 Server

Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area. Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service. Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

Lenovo (United States), Inc.
8001 Development Drive
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION ”AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary. Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk. Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

© Copyright Lenovo 2022. All rights reserved.

This document, LP1049, was created or updated on March 22, 2022.

Send us your comments in one of the following ways:

- Use the online Contact us review form found at:
  https://lenovopress.com/LP1049
- Send your comments in an e-mail to:
  comments@lenovopress.com

This document is available online at https://lenovopress.com/LP1049.

Trademarks

Lenovo and the Lenovo logo are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. A current list of Lenovo trademarks is available on the Web at https://www.lenovo.com/us/en/legal/copytrade/.

The following terms are trademarks of Lenovo in the United States, other countries, or both:
Lenovo®
AnyBay®
Bootable Media Creator
Flex System
Lenovo Services
RackSwitch
System x®
ThinkServer®
ThinkSystem
TopSeller
TruDDR4
UpdateXpress System Packs
XClarity®

The following terms are trademarks of other companies:

Intel®, Intel Optane™, and Xeon® are trademarks of Intel Corporation or its subsidiaries.

Linux® is the trademark of Linus Torvalds in the U.S. and other countries.

Hyper-V®, Microsoft®, PowerShell, Windows PowerShell®, Windows Server®, and Windows® are trademarks of Microsoft Corporation in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Veeam Backup & Replication

Version 11

User Guide for VMware vSphere

April, 2022

© 2022 Veeam Software.

All rights reserved. All trademarks are the property of their respective owners.

No part of this publication may be reproduced, transmitted, transcribed, stored in a retrieval system, or translated into any language in any form by any means, without written permission from Veeam Software (Veeam). The information contained in this document represents the current view of Veeam on the issue discussed as of the date of publication and is subject to change without notice. Veeam shall not be liable for technical or editorial errors or omissions contained herein. Veeam makes no warranties, express or implied, in this document. Veeam may have patents, patent applications, trademark, copyright, or other intellectual property rights covering the subject matter of this document. All other trademarks mentioned herein are the property of their respective owners. Except as expressly provided in any written license agreement from Veeam, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

NOTE

Read the End User Software License Agreement before using the accompanying software programs. Using any part of the software indicates that you accept the terms of the End User Software License Agreement.

Contents

CONTACTING VEEAM SOFTWARE 17
ABOUT THIS DOCUMENT 18
ABOUT VEEAM BACKUP & REPLICATION 19
PLANNING AND PREPARATION 22
Platform Support 23
System Requirements 29
Required Permissions 57
Used Ports 74
Trusted Root Certificates 115
Infrastructure Planning 116
Naming Conventions 119
Security Considerations 120
Kerberos Authentication for Guest OS Processing 122
LICENSING 126
Types of Licenses 128
Instance Consumption for NAS Backup Jobs 129
Obtaining and Renewing License 132
Installing License 134
Merging Licenses 137
Viewing License Information 139
Revoking License 142
Removing License 143
Exceeding License Limit 146
License Expiration 148
Updating License 150
Updating License Manually 151
Updating License Automatically 153
Automatic License Usage Reporting 156
GETTING TO KNOW VEEAM BACKUP & REPLICATION 157
Veeam Backup & Replication UI 158
Main Menu 159
Navigation Pane 160
Ribbon and Tabs 161
Views 162
Working Area 164
Job Filter 165

Changing Color Theme 168
Restore Points Icons 169
VMware vSphere Icons 171
Job Priorities 173
DEPLOYMENT 174
Installing Veeam Backup & Replication 175
Before You Begin 176
Step 1. Start Setup Wizard 178
Step 2. Read and Accept License Agreement 179
Step 3. Provide License File 180
Step 4. Review Components and Select Installation Folder 181
Step 5. Install Missing Software 183
Step 6. Specify Installation Settings 184
Step 7. Specify Service Account Settings 186
Step 8. Select Microsoft SQL Server Instance 187
Step 9. Specify Service Ports 188
Step 10. Specify Data Locations 189
Step 11. Verify Installation Settings 190
Step 12. Begin Installation 192
Step 13. Install Available Updates 193
Upgrading to Veeam Backup & Replication 11 or 11a 194
Updating Veeam Backup & Replication 11 or 11a 198
Uninstalling Veeam Backup & Replication 199
Installing Veeam Backup & Replication Console 200
Before You Begin 201
Step 1. Start Setup Wizard 202
Step 2. Read and Accept License Agreement 203
Step 3. Install Missing Software 204
Step 4. Specify Installation Settings 205
Step 5. Specify Installation Path 206
Step 6. Begin Installation 207
Installing Veeam Backup & Replication in Unattended Mode 208
Before You Begin 210
Installation Command-Line Syntax 211
Installing Updates in Unattended Mode 242
BACKUP INFRASTRUCTURE 244
Backup Infrastructure Components 245
Backup Server 246
Backup & Replication Console 279

Virtualization Servers and Hosts 282
Backup Proxy 314
VMware CDP Proxy 340
Cache Repository 350
Backup Repository 351
External Repository 484
Scale-Out Backup Repository 516
Guest Interaction Proxy 588
Gateway Server 591
Mount Server 594
Veeam Data Movers 596
Veeam vPower NFS Service 598
WAN Accelerators 600
Log Shipping Servers 601
Tape Servers 602
NDMP Servers 603
Veeam Backup Enterprise Manager 604
Deployment Scenarios 605
Simple Deployment 606
Advanced Deployment 607
Distributed Deployment 609
Resource Scheduling 610
Limitation of Concurrent Tasks 611
Limitation of Read and Write Data Rates for Backup Repositories 614
Network Traffic Management 615
Performance Bottlenecks 624
Locations 626
Creating and Assigning Locations to Infrastructure Objects 629
Editing Locations 631
Deleting Locations 632
Exporting and Importing Locations List 633
Veeam Backup & Replication Settings 634
Specifying I/O Settings 635
Configuring Security Settings 638
Specifying Email Notification Settings 652
Specifying SNMP Settings 656
Specifying Other Notification Settings 660
Specifying Session History Settings 663
Roles and Users 664

Update Notification 666
Server Components Upgrade 668
Logging 670
Exporting Logs 671
Configuration Backup and Restore 676
Configuration Backup 677
Restoring Configuration Database 684
Migrating Veeam Backup & Replication to Another Server 700
Migrating Configuration Database to Another SQL Server 705
BACKUP 714
About Backup 715
How Backup Works 716
Backup Architecture 718
Backup Chain 721
Changed Block Tracking 755
Data Compression and Deduplication 757
Data Exclusion 760
VMware Tools Quiescence 770
Guest Processing 773
Microsoft SQL Server Log Backup 788
Oracle Log Backup 798
Backup Job Scheduling 807
Health Check for Backup Files 815
Compact of Full Backup File 820
Resume on Disconnect 822
Snapshot Hunter 823
Creating Backup Jobs 825
Before You Begin 826
Step 1. Launch New Backup Job Wizard 828
Step 2. Specify Job Name and Description 829
Step 3. Select VMs to Back Up 830
Step 4. Exclude Objects from Backup Job 831
Step 5. Define VM Backup Order 834
Step 6. Specify Backup Storage Settings 835
Step 7. Configure Long-Term Retention 837
Step 8. Specify Advanced Backup Settings 839
Step 9. Specify Secondary Target 850
Step 10. Specify Guest Processing Settings 851
Step 11. Define Job Schedule 862

Step 12. Finish Working with Wizard 864
Performing Active Full Backup 865
Quick Backup 866
Retention Policy for Quick Backups 867
Performing Quick Backup 868
Importing Backups Manually 869
Importing Encrypted Backups 871
Importing Transaction Logs 872
Importing Backup Files from Scale-Out Backup Repositories 873
Exporting Backups 874
Performing Export 875
Viewing Session Statistics 882
Managing Backups 884
Viewing Backup Properties 885
Removing Backups from Configuration 887
Deleting Backups from Disk 888
Deleting Backups from Object Storage 890
Removing Missing Restore Points 891
Managing Jobs 894
Editing Job Settings 895
Cloning Jobs 896
Disabling and Deleting Jobs 897
Starting and Stopping Jobs 898
Starting and Stopping Transaction Log Backup Jobs 900
Reconfiguring Jobs with Microsoft SQL Server VMs 902
Reporting 903
Viewing Real-Time Statistics 904
Viewing Job Session Results 907
Viewing Job and Job Session Reports 908
REPLICATION 910
Requirements and Limitations 911
Backup Infrastructure for Replication 912
Replication Scenarios 914
How Replication Works 917
Replication Chain 918
Replica Seeding and Mapping 919
Replica from Backup 922
Creating Replication Jobs 925
Before You Begin 926

Step 1. Launch New Replication Job Wizard 927
Step 2. Specify Job Name and Description 928
Step 3. Select VMs to Replicate 930
Step 4. Specify Data Source 931
Step 5. Exclude Objects from Replication Job 932
Step 6. Define VM Processing Order 935
Step 7. Specify Replica Destination 936
Step 8. Create Network Map Table 939
Step 9. Configure Re-IP Rules 940
Step 10. Specify Replication Job Settings 942
Step 11. Specify Advanced Replica Settings 943
Step 12. Specify Data Transfer Settings 949
Step 13. Define Seeding and Mapping Settings 951
Step 14. Specify Guest Processing Settings 953
Step 15. Define Job Schedule 963
Step 16. Finish Working with Wizard 965
Creating Replica Seeds 966
Managing Replicas 967
Viewing Replica Properties 968
Rescanning Replicas 969
Removing Replicas from Configuration 970
Deleting Replicas from Disk 971
Failover and Failback for Replication 972
Failover Plans 975
Failover 985
Permanent Failover 992
Planned Failover 994
Failover Undo 1000
Failback 1002
Failback Commit 1020
Failback Undo 1022
CONTINUOUS DATA PROTECTION (CDP) 1023
Requirements and Limitations 1024
Backup Infrastructure for CDP 1026
How CDP Works 1028
CDP Replication Chain 1032
Retention Policies 1034
Guaranteed Delivery 1036
Replica Seeding and Mapping 1037

Installing I/O Filter 1040
Step 1. Launch I/O Filter Management Wizard 1041
Step 2. Select Clusters 1042
Step 3. Apply Filter Settings 1043
Step 4. Finish Working with Wizard 1044
Upgrading and Uninstalling I/O Filter 1045
Creating CDP Policies 1047
Before You Begin 1048
Step 1. Launch New CDP Policy Wizard 1049
Step 2. Specify Policy Name and Advanced Settings 1050
Step 3. Select VMs to Replicate 1051
Step 4. Exclude Objects 1053
Step 5. Define VM Processing Order 1055
Step 6. Select Replica Destination 1056
Step 7. Configure Network Mapping 1058
Step 8. Configure Re-IP Rules 1059
Step 9. Configure Seeding and Mapping 1061
Step 10. Specify Data Transfer and Replica Settings 1063
Step 11. Specify Notification Settings 1065
Step 12. Specify Replica Schedule 1067
Step 13. Specify Guest Processing Settings 1069
Step 14. Finish Working with Wizard 1077
Creating Replica Seeds for CDP 1078
Managing CDP Policies 1079
Viewing Session Statistics and Results 1080
Editing Policies 1083
Disabling and Deleting Policies 1084
Managing CDP Replicas 1085
Viewing Replica Properties 1086
Removing Replicas from Configuration 1087
Deleting Replicas from Disk 1088
Failover and Failback for CDP 1089
Failover 1092
Failover Plans 1098
Permanent Failover 1099
Failover Undo 1101
Failback 1102
Failback Commit 1119
Failback Undo 1121

VEEAMZIP 1122
Creating VeeamZIP Backups 1123
BACKUP COPY 1126
About Backup Copy 1127
How Backup Copy Works 1128
Backup Copy Architecture 1130
Backup Copy Modes 1132
Backup Copy Intervals 1134
Restore Point Selection 1137
Transformation Processes 1138
Backup Copy Window 1139
Retention Policy for Backup Copy Jobs 1141
Health Check for Backup Files 1155
Compact of Full Backup File 1157
Backup Copy Job Mapping 1158
Active Full Backup Copies 1160
Automatic Job Retries 1161
Creating Backup Copy Jobs for VMs and Physical Machines 1163
Before You Begin 1164
Step 1. Launch New Backup Copy Job Wizard 1165
Step 2. Specify Job Name and Copy Mode 1167
Step 3. Select Workloads to Process 1168
Step 4. Exclude Objects from Backup Copy Job 1171
Step 5. Select Source Backup Repositories 1172
Step 6. Define Processing Order 1173
Step 7. Specify Target Repository and Retention Settings 1174
Step 8. Map Backup File 1176
Step 9. Specify Advanced Settings 1178
Step 10. Specify Data Path Settings 1186
Step 11. Define Backup Copy Window 1188
Step 12. Finish Working with Wizard 1189
Creating Backup Copy Jobs for HPE StoreOnce Repositories 1190
Before You Begin 1191
Step 1. Launch New Backup Copy Job Wizard 1192
Step 2. Specify Job Name and Description 1193
Step 3. Select Source and Target Repositories 1194
Step 4. Specify Advanced Settings 1196
Step 5. Define Backup Copy Window 1200
Step 6. Finish Working with Wizard 1201

Creating Backup Copy Jobs for Oracle and SAP HANA Databases 1202
Linking Backup Jobs to Backup Copy Jobs 1203
Managing Backups 1205
Viewing Backup Properties 1206
Removing Backups from Target Repositories 1208
Removing Missing Restore Points 1210
Reporting 1213
Managing Jobs 1214
Editing Backup Copy Jobs 1215
Starting Backup Copy Jobs Manually 1217
Creating Active Full Backups 1219
Disabling and Deleting Jobs 1220
VM COPY 1222
Copying VMs 1223
Before You Begin 1224
Step 1. Launch VM Copy Job Wizard 1225
Step 2. Specify Job Name and Description 1226
Step 3. Select VMs to Copy 1227
Step 4. Exclude Objects from VM Copy Job 1228
Step 5. Specify Copy Destination 1231
Step 6. Specify Guest Processing Settings 1233
Step 7. Define Job Schedule 1240
Step 8. Finish Working with Wizard 1242
FILE COPY 1243
Creating File Copy Jobs 1244
Before You Begin 1245
Step 1. Launch New File Copy Job Wizard 1246
Step 2. Specify Job Name and Description 1247
Step 3. Select Files and Folders to Be Copied 1248
Step 4. Select Destination for Copying 1249
Step 5. Define Job Schedule 1250
Step 6. Finish Working with Wizard 1252
Copying Files and Folders Manually 1253
Managing Folders 1255
Editing and Deleting Files 1256
QUICK MIGRATION 1257
Quick Migration Architecture 1259
Migrating VMs 1260
Before You Begin 1261

Step 1. Launch Quick Migration Wizard 1262
Step 2. Select VMs to Relocate 1263
Step 3. Specify VM Destination 1264
Step 4. Select Infrastructure Components for Data Transfer 1266
Step 5. Finish Working with Wizard 1268
Migrating First Class Disks (FCDs) 1270
Before You Begin 1271
Step 1. Launch FCD Quick Migration Wizard 1272
Step 2. Select FCDs to Migrate 1273
Step 3. Specify FCD Destination 1274
Step 4. Finish Working with Wizard 1275
RECOVERY VERIFICATION 1276
SureBackup 1277
How SureBackup Works 1278
Backup Recovery Verification Tests 1279
Application Group 1285
Virtual Lab 1296
SureBackup Job 1320
XML Files with VM Roles Description 1344
Manual Recovery Verification 1347
SureReplica 1348
How SureReplica Works 1349
Replica Recovery Verification Tests 1351
Application Group 1352
Virtual Lab Configuration 1353
SureBackup Job for VM Replicas 1358
ON-DEMAND SANDBOX 1361
On-Demand Sandbox for Storage Snapshots 1362
Mixed Scenarios 1364
Configuring On-Demand Sandbox 1365
DATA RECOVERY 1367
VM Recovery 1368
Instant Recovery to VMware vSphere 1369
Instant Recovery to Microsoft Hyper-V 1391
Entire VM Restore 1412
Staged Restore 1432
Restore to Microsoft Azure 1434
Restore to Amazon EC2 1474
Restore to Google Compute Engine 1493

Disk Recovery 1512
Instant Disk Recovery 1513
Instant First Class Disk (FCD) Recovery 1525
Virtual Disk Restore 1538
Disk Export 1547
Item Recovery 1557
VM Files Restore 1558
Guest OS File Recovery 1565
Application Item Restore 1599
Secure Restore 1607
How Secure Restore Works 1608
Antivirus XML Configuration File 1609
Viewing Malware Scan Results 1616
VCLOUD DIRECTOR SUPPORT 1617
Viewing vCloud Director VMs 1618
Backup and Restore of vApps 1619
Backup of vCloud Director VMs 1620
Data to Back Up 1621
vCD Backup Jobs 1623
Performing Backup of vCloud Director VMs 1624
Creating VeeamZIP Files for vCloud Director VMs 1626
vCloud Director Replication 1627
Creating vCD Replication Job 1628
Managing vCD Replication Jobs 1662
Managing vCD Replicas 1664
Failover and Failback for vCD 1668
Restore of vCloud Director VMs 1698
Restoring Regular and Standalone VMs to vCloud Director 1699
Restoring Linked Clone VMs to vCloud Director 1700
Performing Instant Recovery for VMs 1702
Restoring vCloud vApps 1715
Restoring VMs into vCloud vApp 1728
Restoring Entire VMs into vSphere Infrastructure 1741
Restoring VM Files 1742
Restoring VM Hard Disks 1743
Restoring VM Guest OS Files 1744
VMWARE CLOUD ON AWS SUPPORT 1745
WAN ACCELERATION 1749
Global Data Deduplication 1750

WAN Accelerators 1751
How WAN Acceleration Works 1753
WAN Global Cache 1755
WAN Accelerator Sizing 1761
Adding WAN Accelerators 1766
Removing WAN Accelerators 1774
Data Block Verification 1775
Data Transport on WAN Disconnect 1776
DATA ENCRYPTION 1777
Encryption Standards 1779
Encryption Algorithms 1781
Encryption Keys 1782
How Data Encryption Works 1787
How Data Decryption Works 1789
How Decryption Without Password Works 1791
Encrypted Objects 1793
Backup Job Encryption 1794
Backup Copy Job Encryption 1796
VeeamZIP Encryption 1799
Tape Encryption 1800
Encryption Best Practices 1803
Restoring Data from Encrypted Backups 1805
Decrypting Data with Password 1806
Decrypting Data Without Password 1807
Restoring Encrypted Data from Tapes 1813
Decrypting Tapes with Password 1814
Decrypting Tapes Without Password 1816
TAPE DEVICES SUPPORT 1822
NAS BACKUP SUPPORT 1823
How File Share Backup Works 1830
NAS Backup Retention Scenarios 1831
NAS Backup Integration with Storage Systems 1836
Scale-Out Repository as NAS Backup Repository 1838
Scale-Out Repository with Extents in Metadata and Data Roles 1839
Adding File Share 1841
Adding Managed Server File Share 1842
Adding Enterprise Storage System as NAS Filer 1849
Adding NFS File Share 1856
Adding SMB File Share 1865

Adding Backup Proxy 1875
Step 1. Launch New Backup Proxy Wizard 1876
Step 2. Choose Microsoft Windows Server 1877
Step 3. Configure Traffic Rules 1878
Step 4. Review Components to Install 1879
Step 5. Apply Backup Proxy Settings 1880
Step 6. Finish Working with Wizard 1881
Creating File Share Backup Jobs 1882
Before You Begin 1883
Step 1. Launch New File Backup Job Wizard 1884
Step 2. Specify Job Name and Description 1885
Step 3. Select Files and Folders to Back Up 1886
Step 4. Define Target Backup Storage Settings 1889
Step 5. Specify Advanced Backup Settings 1891
Step 6. Specify Secondary Repository 1898
Step 7. Define Job Schedule 1900
Step 8. Finish Working with Wizard 1902
Starting New Backup Chain 1903
Performing Health Check and Repair for File Share Backup Files 1904
Viewing File Share Backup Properties 1906
Converting Backups from Non-Root to Root Shared Folders 1907
Converting Backups from SMB or NFS Shares to NAS Filer Shares 1908
File Share Data Recovery 1910
Performing Instant File Share Recovery 1911
Restoring Entire File Share 1919
Rolling Back to a Point in Time 1925
Restoring Specific Files and Folders 1929
Restoring Backup Files from Archive Repository 1939
INTEGRATION WITH VEEAM BACKUP FOR AWS 1941
INTEGRATION WITH VEEAM BACKUP FOR MICROSOFT AZURE 1942
INTEGRATION WITH GOOGLE CLOUD PLATFORM 1943
INTEGRATION WITH VEEAM BACKUP FOR NUTANIX AHV 1944
INTEGRATION WITH VEEAM BACKUP FOR RED HAT VIRTUALIZATION 1945
INTEGRATION WITH KASTEN K10 1946
VEEAM AGENT MANAGEMENT 1947
VEEAM CLOUD CONNECT 1948
ADVANCED VMWARE VSPHERE FEATURES 1949
VM Tags 1950
Encrypted VMs 1952

Storage Profiles 1954
VEEAM BACKUP & REPLICATION UTILITIES 1955
Extract Utility 1956
Using Extract Utility in GUI 1957
Using Extract Utility in Interactive Mode 1958
Using Extract Utility from Command Line 1959
Veeam Configuration Database Connection Utility 1963
Using Veeam.Backup.DBConfig.exe Utility 1964
Veeam Backup Validator 1971
Using Veeam Backup Validator 1972
VEEAM BACKUP & REPLICATION EVENTS 1978
About this Document 1979
General 1980
License 1981
Servers and Hosts 1982
Repositories 1983
Proxy Server 1985
WAN Accelerator 1986
Jobs 1987
Nas Backup 1991
Restore 1992
SureBackup 1993
Tape 1995

Contacting Veeam Software

At Veeam Software we value feedback from our customers. It is important not only to help you quickly with your technical issues, but it is our mission to listen to your input and build products that incorporate your suggestions.

Customer Support

Should you have a technical concern, suggestion or question, visit the Veeam Customer Support Portal to open a case, search our knowledge base, reference documentation, manage your license or obtain the latest product release.

Company Contacts

For the most up-to-date information about company contacts and office locations, visit the Veeam Contacts Webpage.

Online Support

If you have any questions about Veeam products, you can use the following resources:

- Full documentation set: veeam.com/documentation-guides-datasheets.html
- Veeam R&D Forums: forums.veeam.com

About This Document

This user guide provides information about main features, installation and use of Veeam Backup & Replication in VMware vSphere environments. The document applies to version 11 and all subsequent versions until it is replaced with a new edition.

Intended Audience

The user guide is intended for anyone who wants to use Veeam Backup & Replication. It is primarily aimed at VMware administrators, consultants, analysts and any other IT professionals using the product.

About Veeam Backup & Replication

Veeam Backup & Replication is a comprehensive data protection and disaster recovery solution. With Veeam Backup & Replication, you can create image-level backups of virtual, physical, cloud machines and restore from them. Technology used in the product optimizes data transfer and resource consumption, which helps to minimize storage costs and the recovery time in case of a disaster.

Veeam Backup & Replication provides a centralized console for administering backup/restore/replication operations in all supported platforms (virtual, physical, cloud). Also, the console allows you to automate and schedule routine data protection operations and integrate with solutions for alerting and generating compliance reports.

This section contains an overview of Veeam Backup & Replication and solutions integrated with it.

Main Features

Main functionality of Veeam Backup & Replication includes:

- Backup: creating image-level backups of virtual, physical, cloud machines and backups of NAS share files.
- Restore: performing restore from backup files to the original or a new location.
  Veeam Backup & Replication offers a number of recovery options for various disaster recovery scenarios, including Instant Recovery, image-level restore, file-level restore, restore of application items and so on.
- Replication: creating an exact copy of a VM and maintaining the copy in sync with the original VM.
- Continuous Data Protection (CDP): replication technology that helps you protect mission-critical VMs and reach recovery point objective (RPO) up to seconds.
- Backup Copy: copying backup files to a secondary repository.
- Storage Systems Support: backing up and restoring VMs using capabilities of native snapshots created on storage systems. For more information, see the Integration with Storage Systems section in the Veeam Backup & Replication User Guide.
- Tape Devices Support: storing copies of backups in tape devices. For more information, see the Tape Devices Support section in the Veeam Backup & Replication User Guide.
- Recovery Verification: testing VM backups and replicas before recovery.

Protected Objects

With Veeam Backup & Replication, you can back up and restore the following objects:

- **Virtual machines**:
  - VMware vSphere VMs
  - Microsoft Hyper-V VMs
  - Nutanix AHV VMs (requires an additional plug-in)
- **Cloud VMs**:
  - AWS EC2 instances
  - Microsoft Azure VMs

    - Google Cloud VMs

- **NAS file shares**. For details, see NAS Backup Support.
- **Physical machines**. To back up machines running Windows, Linux or macOS operating systems, Veeam Backup & Replication uses backup agents installed on each computer. Veeam Backup & Replication operates as a centralized control center for deploying and managing Veeam Agent for Microsoft Windows, Veeam Agent for Linux and Veeam Agent for Mac. For details, see the Veeam Agent Management Guide.

  Also, you can use standalone backup agents and integrate them with Veeam Backup & Replication: Veeam Agent for Oracle Solaris and Veeam Agent for IBM AIX.

Protected Applications

Native functionality of Veeam Backup & Replication allows you to create application-consistent backups for:

- Microsoft SQL Server
- Oracle Database
- Active Directory
- Microsoft Exchange
- Microsoft OneDrive
- Microsoft SharePoint

Also, you can install the following additional tools:

- Veeam Backup for Microsoft Office 365: for full protection of Microsoft Office applications.
- Veeam Plug-ins for Enterprise Applications: for integration of Veeam Backup & Replication with Oracle RMAN, SAP HANA Backint, and BR*Tools.

Management and Reporting

Veeam Backup & Replication integrates with a set of solutions that provide reporting and management capabilities for enterprise environments:

- Veeam ONE: a solution that enables real-time monitoring, business documentation and management reporting for Veeam Backup & Replication, VMware vSphere and Microsoft Hyper-V.
- Veeam Backup Enterprise Manager: a management and reporting component that allows you to manage multiple Veeam Backup & Replication installations from a single web console.
- Management Pack for Veeam Backup & Replication: a component that integrates Veeam Backup & Replication infrastructure, services and jobs into Microsoft System Center Operations Manager.
- Veeam Disaster Recovery Orchestrator: a solution that orchestrates disaster recovery processes in VMware vSphere environments, supports one-click recovery for critical applications and sites, and provides features for documentation and testing.

Service Providers

If you are a service provider, you can use Veeam Service Provider Console to deliver Veeam-powered Backup-as-a-Service (BaaS) and Disaster Recovery-as-a-Service (DRaaS) services to your customers.

You can also use Veeam Backup & Replication to offer cloud repository as a service and disaster recovery as a service. For details, see Veeam Cloud Connect.

Planning and Preparation

Before you install Veeam Backup & Replication, you must make sure that the virtual environment and machines that you plan to use as backup infrastructure components meet product hardware recommendations and system requirements.

Platform Support

Make sure platforms that you plan to use for various operations are supported.

Veeam Backup & Replication provides support for the following versions of the VMware vSphere platform.

Virtual Infrastructure

Specification	Requirement
Platform	• vSphere 7.0 (up to 7.0 U3[1]) • vSphere 6.x • vSphere 5.5 • VMware Cloud Foundation (VCF) This platform is supported as individual VMware software components. VMware components listed on this page can be part of VCF. For the information on the correspondence of VMware components to the VCF version, see this VMware KB article. • VMware Cloud on AWS • VMware Cloud on Dell EMC • Azure VMware Solution For more information on Azure VMware Solution support, see this Veeam KB article.
Hypervisor	• ESXi 7.0 (up to 7.0 U3[1]) • ESXi 6.x • ESXi 5.5 Free ESXi is not supported. Veeam Backup & Replication leverages vSphere and vStorage APIs that are disabled by VMware in free ESXi.
Management Server (optional)	• vCenter Server 7.0 (optional) (up to 7.0 U3[1]) • vCenter Server 6.x (optional) • vCenter Server 5.5 (optional)

[1] This version is supported starting from Veeam Backup & Replication 11a (build 11.0.1.1261).

Veeam Continuous Data Protection (CDP)

The following infrastructure requirements apply only when you protect VMs with Continuous Data Protection (CDP):

- Minimum ESXi version required is 6.5 U2.
- Minimum 16GB RAM is required for source and target ESXi hosts.

- vCenter Server is required. Standalone ESXi hosts are not supported.
- Backup server, VMware CDP proxies, vCenter Server and ESXi hosts must be able to resolve each other DNS names.
- VMware Cloud on AWS is not supported.

For more information on Veeam CDP, its requirements and limitations, see Continuous Data Protection (CDP).

VMs

Specification	Requirement
Virtual Hardware	• All types and versions of virtual hardware are supported, including 62 TB VMDK. • Virtual machines with virtual NVDIMM devices, with virtual disks engaged in SCSI bus sharing or residing on PMem datastores are not supported for host-based backup, because VMware does not support snapshotting such VMs. To protect such VMs, use agent-based backup. • RDM virtual disks in physical mode, independent disks, and disks connected through in-guest iSCSI initiator are not supported for host-based backup, and are skipped from processing automatically. If backup of these disks is required, use agent-based backup. Network shares and mount points targeted to 3rd party storage devices are also skipped as these volumes/disks are not visible in the VM configuration file. RDM virtual disks in virtual mode are supported to create backups based on VMware Changed Block Tracking technology, although there are some restrictions on the virtual disk restore operation. To learn more about them, see Restoring Virtual Disks.

Specification	Requirement
OS	• All operating systems supported by VMware. • Guest processing (which includes application-aware processing and indexing) is supported for Microsoft Windows 2008/Windows Vista or later except Nano Server, due to the absence of VSS framework. • Guest processing (which includes application-aware processing and indexing) is supported for 64-bit versions of the following Linux operating systems: o CentOS 7 to 8.3, and 8.4[1], and 8.5[2] o CentOS Stream o Debian 9.0 to 10.8, and 11.0[1] o Fedora 30 to 33, and 34[1], and 35[2] o openSUSE Leap 15.2 and 15.3[1], Tumbleweed o Oracle Linux 6 (UEK3) to 8.3 (UEK R6 and UEK R6 U2[1]) o Oracle Linux 6 to 8.3, and 8.4[1], and 8.5[2] (RHCK) o RHEL 6.0 to 8.3, and 8.4[1], and 8.5[2] o SLES 11 SP4, 12 SP1-SP5, 15 SP0-SP2, and SP3[1] o Ubuntu: 14.04 LTS, 16.04 LTS, 18.04 LTS, 19.10, 20.04 LTS, and 21.04[1], and 21.10[2]
Software	• VMware Tools (optional, recommended). VMware Tools are required for the following operations: application-aware processing, file-level restore from Microsoft Windows guest OS and SureBackup testing functions. • Open VM Tools (OVT, optional). Open VM Tools are a set of services and modules used by VMware for interaction with VMs running Linux or other VMware supported Unix-like guest operating systems. • All latest OS service packs and patches (required for application-aware processing).

[1] This version is supported starting from Veeam Backup & Replication 11a (build 11.0.1.1261).

[2] This version is supported starting from Veeam Backup & Replication 11a (build 11.0.1.1261) with Cumulative Patch P20211211.

vCloud Director

Specification	Requirement
vCloud Director	vCloud Director 9.5 to 10.2, and 10.3[1] (9.7 or later for vCD to vCD replication functionality)

[1] This version is supported starting from Veeam Backup & Replication 11a (build 11.0.1.1261).

OS	Supported File Systems
Microsoft Windows	• FAT, FAT32 • NTFS • ReFS (ReFS is supported only if Veeam Backup & Replication is installed on Microsoft Windows Server 2012 or later). Windows file-level restore to original location is supported for Microsoft Windows 2008/Windows Vista or later except Nano Server.
Linux	• ext2, ext3, ext4 • ReiserFS • JFS • XFS • Btrfs DRBD (Distributed Replicated Block Devices) are not supported.
BSD	UFS, UFS2
Mac	HFS, HFS+ (volumes up to 2 TB)
Micro Focus OES	NSS AD-enabled NSS volumes on Open Enterprise Server 2015 are supported. Restore of NSS file/folder permissions is not supported. **Note:** Starting from Veeam Backup & Replication 9.5 Update 4a, restore of NSS trustee rights on files/folders is supported. File-level restore is supported for the following OSes: • Open Enterprise Server (OES) versions 2015(SP1), 2018(SP1-SP2). • SUSE Linux Enterprise Server 11 SP1-SP4 with OES11(SP1-SP3) • NetWare 6.5 (only restore to a new location is supported)
Solaris	• UFS • ZFS (except any pool versions of Oracle Solaris) The helper appliance uses module ZFSonLinux version 0.7.0. For this reason, Veeam Backup & Replication supports only those versions of pools and features that are available in ZFSonLinux version 0.7.0.

For other requirements and limitations of guest OS file restore, see Requirements and Limitations.

File Servers

Veeam Backup & Replication supports backup of files and folders from the file servers managed by the following operating systems:

- 64-bit versions of the Microsoft Windows operating systems:
  - Microsoft Windows Server Semi-Annual Channel (from version 1803 to version 20H2)
  - Microsoft Windows Server 2022[1]
  - Microsoft Windows Server 2019
  - Microsoft Windows Server 2016
  - Microsoft Windows Server 2012 R2
  - Microsoft Windows Server 2012
  - Microsoft Windows Server 2008 R2 SP1
  - Microsoft Windows 11 (version 21H2[1])
  - Microsoft Windows 10 (from version 1803 to version 21H1[1], and 21H2[2])
  - Microsoft Windows 8.1
  - Microsoft Windows 7 SP1
- 64-bit versions of the Linux operating systems:
  - CentOS 7 to 8.3, and 8.4[1], and 8.5[2]
  - CentOS Stream
  - Debian 9.0 to 10.8, and 11.0[1]
  - Fedora 30 to 33, and 34[1], and 35[2]
  - openSUSE Leap 15.2 and 15.3[1], Tumbleweed
  - Oracle Linux 6 (UEK3) to 8.3 (UEK R6 and UEK R6 U2[1])
  - Oracle Linux 6 to 8.3, and 8.4[1], and 8.5[2] (RHCK)
  - RHEL 6.0 to 8.3, and 8.4[1], and 8.5[2]
  - SLES 11 SP4, 12 SP1-SP5, 15 SP0-SP2, and SP3[1]
  - Ubuntu: 14.04 LTS, 16.04 LTS, 18.04 LTS, 19.10, 20.04 LTS, and 21.04[1], and 21.10[2]

[1] This version is supported starting from Veeam Backup & Replication 11a (build 11.0.1.1261).

[2] This version is supported starting from Veeam Backup & Replication 11a (build 11.0.1.1261) with Cumulative Patch P20211211.

NAS Backup Support

Veeam Backup & Replication supports backup of file shares from the following sources:

- Microsoft Windows- or Linux-managed server
- Enterprise NAS systems: NetApp Data ONTAP, Lenovo ThinkSystem DM Series, Dell EMC Isilon

- NFS file shares
- SMB file shares

Consider the following requirements and limitations:

- Only 64-bit versions of operating systems are supported for Microsoft Windows- or Linux-managed server file share.
- Backup of file shares on Linux hosts added with single use credentials is not supported.
- NFS file share must run NFS protocol version 3 or 4.1.
- Anonymous or AD/Kerberos authentication is not supported for access to file shares through NFS.

  In NFS settings of the source file share, you must explicitly specify what servers will have access to the file share.
- SMB file share must run on SMB version 1.x, 2.x or 3.x.
- To support the **VSS for SMB File Shares** feature, make sure that requirements listed in this Veeam KB article are met.
- To correctly back up SACL (Ownership) files and folders from the SMB file share and restore them:
  a. When you are specifying access settings for the SMB file share, select the **This share requires access credentials** check box.
  b. Make sure that the account you use to access the file share is either added to the **Backup Operators** group or has the **SeBackupPrivilege** and **SeRestorePrivilege** privileges in Windows Server on the file share.

For more information on NAS backup, see NAS Backup Support.

Network

Consider the following requirements and limitations:

- Domain names of all managed servers added to the Veeam backup infrastructure and machines you want to back up must be resolvable into IPv4 addresses.
- VMware NSX-t 2.3 or later is supported with N-VDS for VMware vSphere and VMware Cloud on AWS/Dell EMC except Virtual Lab. Virtual Lab does not support NSX-t for all environments.
- VMware NSX-t 3.0 or later is supported with VDS for VMware vSphere and VMware Cloud on AWS/Dell EMC except Virtual Lab. Virtual Lab does not support NSX-t for all environments.
- VMware NSX-v is supported (see details on the support of vSphere and VMware Cloud version above).

System Requirements

Make sure that servers that you plan to use as backup infrastructure components meet the system requirements listed below.

Limitations and Recommendations

Coexistence with Mission-Critical Production Servers

We do not recommend to install Veeam Backup & Replication and its components on mission-critical machines in the production environment such as vCenter Server, Domain Controller, Microsoft Exchange Server, Small Business Server/ Windows Server Essentials and so on. If possible, install Veeam Backup & Replication and its components on dedicated machines. Backup infrastructure component roles can be co-installed.

Microsoft Windows Server Core

You can assign roles of a backup proxy, backup repository, WAN accelerator, Veeam Cloud Connect infrastructure components and tape infrastructure components to machines running Microsoft Windows Server Core.
Mind that you cannot install Veeam Backup & Replication and Veeam Backup Enterprise Manager on a machine running Microsoft Windows Server Core.

Domain Member

The machine on which you plan to install Veeam Backup & Replication does not necessarily need to be a domain member. However, if you plan to restore Microsoft Exchange items from the Veeam Backup Enterprise Manager UI, you must install Veeam Backup Enterprise Manager on the domain member server from the Microsoft Active Directory forest in which Microsoft Exchange mailboxes are located.

Linux Server

To avoid negotiation problems between Veeam Backup & Replication and a Linux server, the latter should use ciphers, Key Exchange (KEX) algorithms, and MAC algorithms compatible with SSH libraries supported by Veeam Backup & Replication:

Algorithms	Supported values
Ciphers	Recommended: *aes128-ctr*, *aes192-ctr*, *aes256-ctr*, *aes128-gcm@openssh.com*, *aes256-gcm@openssh.com*, *chacha20-poly1305@openssh.com* Supported for backward compatibility: *3des-cbc*, *3des-ctr*, *aes128-cbc*, *aes192-cbc*, *aes256-cbc*, *arcfour*, *arcfour128*, *arcfour256*, *blowfish-cbc*, *blowfish-ctr*, *cast128-cbc*, *twofish-cbc*, *twofish128-cbc*, *twofish128-ctr*, *twofish192-cbc*, *twofish192-ctr*, *twofish256-cbc*, *twofish256-ctr*

Algorithms	Supported values
KEX algorithms	Recommended: *diffie-hellman-group-exchange-sha256*, *diffie-hellman-group14-sha256*, *diffie-hellman-group15-sha512*, *diffie-hellman-group16-sha512*, *ecdh-sha2-nistp256*, *ecdh-sha2-nistp384*, *ecdh-sha2-nistp521*, *curve25519-sha256*, *curve25519-sha256@libssh.org* Supported for backward compatibility: *diffie-hellman-group1-sha1*, *diffie-hellman-group14-sha1*, *diffie-hellman-group-exchange-sha1*
MAC algorithms	Recommended: *hmac-sha2-256*, *hmac-sha2-512* Supported for backward compatibility: *hmac-md5*, *hmac-md5-96*, *hmac-sha-256-96*, *hmac-sha1-96*, *hmac-sha2-512-96*, *hmac-ripemd160*, *hmac-ripemd160@openssh.com*, *hmac-sha1*

Ensure that your SSH configuration on the Linux server allows to use at least one cipher, KEX algorithm, and MAC algorithm from the table above. You can run the following command to verify the list of allowed algorithms:

```
sudo sshd -T | grep "\(ciphers\|macs\|kexalgorithms\)"
```

All-in-One Installations

For all-in-one installations, you can subtract 2 GB of memory resources from each but one role. These 2 GB are allotted to the OS itself, assuming each component is installed on the dedicated server.

NAS Backup

Each of the following NAS backup components may consume up to 4 GB RAM per task (in case of deduplicating storage appliances, up to 8 GB RAM): backup repository, backup proxy, cache repository. Make sure you allocate enough memory resources for your installation. For all-in-one installations, the server where all the components are installed must have enough memory resources for all components.

Sharing Backup Infrastructure Components Across Veeam Installations

We do not recommend to use shared backup infrastructure components across different Veeam installations due to several reasons:

- Veeam installations compete for resources.
- Backup components cannot simultaneously interact with Veeam Backup & Replication of different versions.
- Adding the same repository to different Veeam Backup & Replication installations may lead to corrupted backup and data in the database.

Backup Server

Specification	Requirement
Hardware	*CPU*: x86-64 processor (minimum 4 cores recommended). *Memory*: 4 GB RAM plus 500 MB RAM for each enabled job. Memory consumption varies according to the number of VMs in the job, size of VM metadata, size of production infrastructure, and so on. Additionally, for users with tape installations (for file to tape jobs processing more than 1,000,000 files): • 1.5 GB RAM for file to tape backup for each 1,000,000 files • 2.6 GB RAM for file restore for each 1,000,000 files • 1.3 GB RAM for catalog jobs for each 1,000,000 files *Disk Space*: 5 GB[1] for product installation and 4.5 GB for Microsoft .NET Framework 4.7.2 installation. 10 GB per 100 VM for guest file system catalog folder (persistent data). Additional free disk space for Instant Recovery cache folder (non-persistent data, at least 100 GB recommended). *Network*: 1 Gbps or faster for on-site backup and replication, and 1 Mbps or faster for off-site backup and replication. High latency and reasonably unstable WAN links are supported. [1] Here and throughout this document GB is considered as 2^30 bytes, TB as 2^40 bytes.
OS	Only 64-bit version of the following operating systems are supported[1]: • Microsoft Windows Server 2022[2] • Microsoft Windows Server 2019 • Microsoft Windows Server 2016 • Microsoft Windows Server 2012 R2 • Microsoft Windows Server 2012 • Microsoft Windows Server 2008 R2 SP1 • Microsoft Windows 11 (version 21H2[2]) • Microsoft Windows 10 (from version 1803 to version 21H1[2], and 21H2[3]) • Microsoft Windows 8.1 • Microsoft Windows 7 SP1 [1] Running Veeam backup server or any of Veeam backup infrastructure components on Insider versions of Microsoft Windows OS (both Client and Server) is not supported. [2] This version is supported starting from Veeam Backup & Replication 11a (build 11.0.1.1261). [3] This version is supported starting from Veeam Backup & Replication 11a (build 11.0.1.1261) with Cumulative Patch P20211211.

Specification	Requirement
Software	During setup, the system configuration check is performed to determine if all prerequisite software is available on the machine where you plan to install Veeam Backup & Replication. If some of the required software components are missing, the setup wizard will offer you to install missing software automatically. This refers to: • Microsoft .NET Framework 4.7.2 • Microsoft Windows Installer 4.5 • Microsoft PowerShell 5.1 • Microsoft SQL Server Management Objects • Microsoft SQL Server System CLR Types • Microsoft Report Viewer Redistributable 2015 • Microsoft Universal C Runtime The backup server installation also requires the automatic installation of the prerequisite software for Veeam Cloud Plug-Ins. The following software must be installed manually: • Firefox, Google Chrome, Microsoft Edge, or Microsoft Internet Explorer 11.0 or later.

Specification	Requirement
SQL Database	Local or remote installation of the following versions of Microsoft SQL Server: • Microsoft SQL Server 2019 • Microsoft SQL Server 2017 • Microsoft SQL Server 2016 (Microsoft SQL Server 2016 SP2 Express Edition is included in the setup)[1] • Microsoft SQL Server 2014 • Microsoft SQL Server 2012 (Microsoft SQL Server 2012 SP4 Express Edition is included in the setup)[2] • Microsoft SQL Server 2008 R2 • Microsoft SQL Server 2008 All editions of Microsoft SQL Server are supported. The usage of Microsoft SQL Server Express Edition is limited by the database size up to 10 GB. If you plan to have larger databases, use other editions of Microsoft SQL Server. Veeam Backup & Replication and Veeam Backup Enterprise Manager configuration databases can be deployed in Microsoft SQL AlwaysOn Availability Groups. For more information, see this Veeam KB article. [1] For machines running Microsoft Windows Server 2012 or later. [2] For machines running Microsoft Windows 7 or Microsoft Windows Server 2008 R2.

Mind the following:

- If you plan to back up VMs running Microsoft Windows Server 2012 R2 or later, and Data Deduplication is enabled for some VM volumes, it is recommended that you deploy the Veeam Backup & Replication console and mount server on a machine running same or later version of Microsoft Windows Server with Data Deduplication feature enabled. Otherwise, some types of restore operations for these VMs (such as Microsoft Windows File-Level Recovery) may fail.
- Due to its limitations, Microsoft SQL Server Express Edition can only be used for evaluation purposes or in case of a small-scale production environment. For environments with a lot of VMs, it is necessary to install a fully functional commercial version of Microsoft SQL Server.

For more information, see Backup Server.

Veeam Backup & Replication Console

Specification	Requirement
Hardware	*CPU*: x86-64 processor. *Memory*: 2 GB RAM *Disk Space*: 500 MB for product installation and 4.5 GB for Microsoft .NET Framework 4.7.2 installation. *Network*: 1 Mbps connection to the backup server. High latency and low bandwidth impact user interface responsiveness.
OS	Only 64-bit version of the following operating systems are supported: • Microsoft Windows Server 2022[1] • Microsoft Windows Server 2019 • Microsoft Windows Server 2016 • Microsoft Windows Server 2012 R2 • Microsoft Windows Server 2012 • Microsoft Windows Server 2008 R2 SP1 • Microsoft Windows 11 (version 21H2[1]) • Microsoft Windows 10 (from version 1803 to version 21H1[1], and 21H2[2]) • Microsoft Windows 8.1 • Microsoft Windows 7 SP1 [1] This version is supported starting from Veeam Backup & Replication 11a (build 11.0.1.1261). [2] This version is supported starting from Veeam Backup & Replication 11a (build 11.0.1.1261) with Cumulative Patch P20211211.
Software	• Microsoft .NET Framework 4.7.2 (included in the setup) • Windows Installer 4.5 (included in the setup) • Microsoft PowerShell 5.1 (included in the setup) • Firefox, Google Chrome, Microsoft Edge or Microsoft Internet Explorer 11.0 or later

For more information, see Backup & Replication Console.

Backup Proxy Server

Specification	Requirement
Hardware	*CPU*: x86-64 processor with minimum of 2 cores (vCPUs), plus 1 core (vCPU) for each additional concurrent task. Using faster processors improves data processing performance. For more information, see Limitation of Concurrent Tasks. *Memory*: 2 GB RAM plus 200 MB for each concurrent task. Using faster memory (DDR3/DDR4) improves data processing performance. *Disk Space*: 750 MB for Microsoft Windows-based proxies; 400 MB for Linux-based proxies. *Network*: 1 Gbps or faster for on-site backup and replication, and 1 Mbps or faster for off-site backup and replication. High latency and reasonably unstable WAN links are supported.

Specification	Requirement
OS	64-bit versions of the following Microsoft Windows operating systems are supported, including Core edition: • Microsoft Windows Server Semi-Annual Channel (from version 1803 to version 20H2) • Microsoft Windows Server 2022[1] • Microsoft Windows Server 2019 • Microsoft Windows Server 2016 • Microsoft Windows Server 2012 R2 • Microsoft Windows Server 2012 • Microsoft Windows Server 2008 R2 SP1 • Microsoft Windows 11 (version 21H2[1]) • Microsoft Windows 10 (from version 1803 to version 21H1[1], and 21H2[2]) • Microsoft Windows 8.1 • Microsoft Windows 7 SP1 64-bit versions of the following Linux distributions are supported[3]. Note that bash shell and SSH are required. Perl is required only for non-persistent Veeam Data Movers. Check the full list of required Perl modules in this Veeam KB article. • CentOS 7 to 8.3, and 8.4[1], and 8.5[2] • CentOS Stream • Debian 9.0 to 10.8, and 11.0[1] • Fedora 30 to 33, and 34[1], and 35[2] • openSUSE Leap 15.2 and 15.3[1], Tumbleweed • Oracle Linux 6 (UEK3) to 8.3 (UEK R6 and UEK R6 U2[1]) • Oracle Linux 6 to 8.3, and 8.4[1], and 8.5[2] (RHCK) • RHEL 6.0 to 8.3, and 8.4[1], and 8.5[2] • SLES 11 SP4, 12 SP1-SP5, 15 SP0-SP2, and SP3[1] • Ubuntu: 14.04 LTS, 16.04 LTS, 18.04 LTS, 19.10, 20.04 LTS, and 21.04[1], and 21.10[2] [1] This version is supported starting from Veeam Backup & Replication 11a (build 11.0.1.1261). [2] This version is supported starting from Veeam Backup & Replication 11a (build 11.0.1.1261) with Cumulative Patch P20211211. [3] Linux distributions are not supported as a backup proxy in NAS backup. To check compatibility for storage integration, see the system requirements provided by the storage system vendor.

For more information on the backup proxy, its requirements and limitations, see Backup Proxy.

IMPORTANT

To protect VMs running on ESXi 5.5 and newer with Microsoft Windows, you must deploy backup proxies on machines running a 64-bit version of Microsoft Windows. VDDK 5.5 and newer does not support 32-bit versions of Microsoft Windows.

The following table shows the minimum system requirements for a backup proxy used in NAS backup.

Specification	Requirement
Hardware	*CPU*: x86-64 processor (minimum 2 cores or vCPUs). Using multi-core processors improves data processing performance and allows for more tasks to be processed concurrently. *Memory*: 4 GB RAM for each concurrent task. Using faster memory (DDR3/DDR4) improves data processing performance. *Disk Space*: 300 MB. *Network*: High latency and reasonably unstable WAN links are supported.
OS	64-bit versions of the following Microsoft Windows operating systems are supported: • Microsoft Windows Server Semi-Annual Channel (from version 1803 to version 20H2) • Microsoft Windows Server 2022[1] • Microsoft Windows Server 2019 • Microsoft Windows Server 2016 • Microsoft Windows Server 2012 R2 • Microsoft Windows Server 2012 • Microsoft Windows Server 2008 R2 SP1 • Microsoft Windows 11 (version 21H2[1]) • Microsoft Windows 10 (from version 1803 to version 21H1[1], and 21H2[2]) • Microsoft Windows 8.1 • Microsoft Windows 7 SP1 Backup proxies running Microsoft Windows Server 2012 R2 or later support the **VSS for SMB File Shares** feature. To use it, make sure that all requirements listed in this Veeam KB article are met. [1] This version is supported starting from Veeam Backup & Replication 11a (build 11.0.1.1261). [2] This version is supported starting from Veeam Backup & Replication 11a (build 11.0.1.1261) with Cumulative Patch P20211211.

For more information on using the backup proxy for NAS backup, see **Backup Proxy** in NAS Backup Support.

VMware CDP Proxy Server

The following table shows the minimum system requirements for a VMware CDP proxy.

Specification	Requirement
Hardware	*CPU*: x86-64 processor with minimum of 8 cores (vCPUs). Using multi-core processors improves data processing performance and allows for more tasks to be processed concurrently. *Memory*: 16 GB RAM. Using more memory allows for longer peak write I/O periods before a CDP policy switches to the disk-based write I/O cache. Using faster memory improves data processing performance. *Disk Space*: 300 MB plus disk-based write I/O cache (non-persistent data, at least 50 GB recommended). Larger cache allows for longer network downtime periods before a CDP policy switches to the CBT mode. *Network*: 1 Gbps or faster.
OS	64-bit versions of the following Microsoft Windows operating systems are supported, including Core edition: • Microsoft Windows Server Semi-Annual Channel (from version 1803 to version 20H2) • Microsoft Windows Server 2022[1] • Microsoft Windows Server 2019 • Microsoft Windows Server 2016 • Microsoft Windows Server 2012 R2 • Microsoft Windows Server 2012 • Microsoft Windows Server 2008 R2 SP1 • Microsoft Windows 11 (version 21H2[1]) • Microsoft Windows 10 (from version 1803 to version 21H1[1], and 21H2[2]) • Microsoft Windows 8.1 • Microsoft Windows 7 SP1 [1] This version is supported starting from Veeam Backup & Replication 11a (build 11.0.1.1261). [2] This version is supported starting from Veeam Backup & Replication 11a (build 11.0.1.1261) with Cumulative Patch P20211211.

For more information on the VMware CDP proxy, see VMware CDP Proxy.

Backup Repository Server

These requirements also apply to mount servers (if separate from the repository server), gateway servers for file share and deduplicating appliance-based repositories, and cache repository servers (64-bit OS only).

Specification	Requirement
Hardware	*CPU*: x86 processor (x86-64 recommended). *Memory*: 4 GB RAM, plus up to 2 GB RAM (32-bit OS) or up to 4 GB RAM (64-bit OS) for each concurrently processed machine or file share. For more information, see Limitation of Concurrent Tasks. For RAM allocation recommendations for NAS backup, see Limitations and Recommendations for NAS backup. *Network*: 1 Gbps or faster for on-site backup and replication, and 1 Mbps or faster for off-site backup and replication. High latency and reasonably unstable WAN links are supported.

OS	Both 32-bit and 64-bit (recommended) versions of the following operating systems are supported, including Core edition: • Microsoft Windows Server Semi-Annual Channel (from version 1803 to version 20H2) • Microsoft Windows Server 2022[1] • Microsoft Windows Server 2019 • Microsoft Windows Server 2016 • Microsoft Windows Server 2012 R2 • Microsoft Windows Server 2012 • Microsoft Windows Server 2008 R2 SP1 • Microsoft Windows 11 (version 21H2[1]) • Microsoft Windows 10 (from version 1803 to version 21H1[1], and 21H2[2]) • Microsoft Windows 8.1 • Microsoft Windows 7 SP1 In addition, both 32-bit and 64-bit (recommended) versions of the following Linux distributions are supported: • CentOS 7 to 8.3, and 8.4[1], and 8.5[2] • CentOS Stream • Debian 9.0 to 10.8, and 11.0[1] • Fedora 30 to 33, and 34[1], and 35[2] • RHEL 6.0 to 8.3, and 8.4[1], and 8.5[2] • openSUSE Leap 15.2 and 15.3[1], Tumbleweed • Oracle Linux 6 (UEK3) to 8.3 (UEK R6 and UEK R6 U2[1]) • Oracle Linux 6 to 8.3, and 8.4[1], and 8.5[2] (RHCK) • SLES 11 SP4, 12 SP1-SP5, 15 SP0-SP2, and SP3[1] • Ubuntu 14.04 LTS, 16.04 LTS, 18.04 LTS, 19.10, 20.04 LTS, and 21.04[1], and 21.10[2] Bash shell and SSH are required. Perl is required only for non-persistent Veeam Data Movers. Check the full list of required Perl modules in this Veeam KB article. 64-bit edition of Linux must be able to run 32-bit programs. Persistent Veeam Data Movers (pure 64-bit Linux editions) are not supported, Perl installation must support 32-bit variables. For advanced XFS integration, only the following 64-bit Linux distributions are supported: • CentOS 8.2, 8.3, and 8.4[1], and 8.5[2] • Debian 10.x, and 11[1] • RHEL 8.2, 8.3, and 8.4[1], and 8.5[2] • SLES 15 SP2, and SP3[1]

Specification	Requirement
	• Ubuntu 18.04 LTS, 20.04 LTS, and 21.04[1], and 21.10[2] For other distributions, XFS integration support is experimental, with kernel version 5.4 or later recommended. For more information, see this Veeam KB article. [1] This version is supported starting from Veeam Backup & Replication 11a (build 11.0.1.1261). [2] This version is supported starting from Veeam Backup & Replication 11a (build 11.0.1.1261) with Cumulative Patch P20211211.

For more information, see Backup Repository.

> **NOTE**
>
> If you plan to use a Microsoft Windows backup repository with Data Deduplication, make sure that you set up the Microsoft Windows server correctly. For more information, see this Veeam KB article.

Cache Repository

The following storage types can be used as a cache repository for NAS backup:

- Direct attached storage. You can add virtual and physical servers as cache repositories:
  - Microsoft Windows server (only 64-bit versions are supported).
  - Linux server (only 64-bit versions are supported).
- Network attached storage. You can add SMB (CIFS) Share or NFS Share as a cache repository.

For system requirements of backup repository servers that can be used as a cache repository, see Backup Repository Server.

For more information on the cache repository, see **Cache Repository** in NAS Backup Support.

Tape Server

Specification	Requirement
Hardware	*CPU*: x86 processor (x86-64 recommended). *Memory*: 2 GB RAM plus 200MB for each concurrent task. Depending on the source of tape jobs, different entities are considered tasks: for machine backup to tape, a task covers a source job or a source chain if tape paralleling is enabled; for file backup to tape, a task covers an entire server or a file share. Restoring VMs directly from tape requires 400MB of RAM per 1TB of virtual disk size. Additionally (for file to tape jobs processing more than 1,000,000 files): o 800 MB RAM for file to tape backup for each 1,000,000 files o 800 MB RAM catalog jobs for each 1,000,000 files *Disk Space*: 300 MB, plus 10 GB for temporary data storage for backup and restore operations. *Network*: 1 Gbps or faster.

Specification	Requirement
OS	Both 32-bit and 64-bit (recommended) versions of the following operating systems are supported, including Core edition: ○ Microsoft Windows Server Semi-Annual Channel (from version 1803 to version 20H2) ○ Microsoft Windows Server 2022[1] ○ Microsoft Windows Server 2019 ○ Microsoft Windows Server 2016 ○ Microsoft Windows Server 2012 R2 ○ Microsoft Windows Server 2012 ○ Microsoft Windows Server 2008 R2 SP1 ○ Microsoft Windows 11 (version 21H2[1]) ○ Microsoft Windows 10 (from version 1803 to version 21H1[1], and 21H2[2]) ○ Microsoft Windows 8.1 ○ Microsoft Windows 7 SP1 [1] This version is supported starting from Veeam Backup & Replication 11a (build 11.0.1.1261). [2] This version is supported starting from Veeam Backup & Replication 11a (build 11.0.1.1261) with Cumulative Patch P20211211.

WAN Accelerator

Specification	Requirement
Hardware	*CPU*: x86-64 processor. Using multi-core processors improves data processing performance, and is highly recommended on WAN links faster than 10 Mbps. *Memory*: 8 GB RAM. Using faster memory (DDR3/DDR4) improves data processing performance. *Disk Space*: Disk space requirements depend on the WAN Accelerator role. Source WAN Accelerator requires 20 GB per 1 TB of source data to store digests of data blocks of source VM disks. Disk space consumption is dynamic and changes as unique VMs are added to (or removed from) jobs with WAN Acceleration enabled. Target WAN Accelerator requires global cache size as defined by user (fixed amount). Disk space is reserved immediately upon selecting the WAN Accelerator as a target one in any job. For more information, see WAN Accelerator Sizing. *Network*: 1 Gbps or faster for on-site backup and replication, and 1 Mbps or faster for off-site backup and replication. High latency and reasonably unstable WAN links are supported.

Specification	Requirement
OS	Only 64-bit versions of the following operating systems are supported, including Core edition: ○ Microsoft Windows Server Semi-Annual Channel (from version 1803 to version 20H2) ○ Microsoft Windows Server 2022[1] ○ Microsoft Windows Server 2019 ○ Microsoft Windows Server 2016 ○ Microsoft Windows Server 2012 R2 ○ Microsoft Windows Server 2012 ○ Microsoft Windows Server 2008 R2 SP1 ○ Microsoft Windows 11 (version 21H2[1]) ○ Microsoft Windows 10 (from version 1803 to version 21H1[1], and 21H2[2]) ○ Microsoft Windows 8.1 ○ Microsoft Windows 7 SP1 [1] This version is supported starting from Veeam Backup & Replication 11a (build 11.0.1.1261). [2] This version is supported starting from Veeam Backup & Replication 11a (build 11.0.1.1261) with Cumulative Patch P20211211.

For more information, see WAN Accelerators.

> **NOTE**
>
> Global cache is not leveraged by source WAN accelerators, or by WAN accelerators operating in the **High bandwidth mode**, so it does not need to be allocated and populated in such cases.

Backup Target

Backups can be performed to the following disk-based storage:

- Local (internal) storage of the backup repository server.
- Direct Attached Storage (DAS) connected to the backup repository server, including external USB/eSATA drives, USB pass through and raw device mapping (RDM) volumes.
- Storage Area Network (SAN). Backup repository server must be connected into the SAN fabric through hardware or virtual HBA, or software iSCSI initiator.
- Network Attached Storage (NAS) able to present its capacity as NFS share (protocol version 3.0 and 4.1 only) or SMB/CIFS share (any protocol version). Using SMB protocol for non-continuously available (CA) file shares is not recommended for reliability reasons. NFS shares are supported both for direct operation and when mounted to a Linux repository server.
- Dell EMC Data Domain (DD OS version 6.2 to 7.6, and 7.7[1]) with DDBoost license. Both Ethernet and Fibre Channel (FC) connectivity is supported.
- ExaGrid (firmware version 5.0.0 or later).

- HPE StoreOnce (firmware version 3.15.1 or later) with Catalyst license. Both Ethernet and Fibre Channel (FC) connectivity is supported. HPE Cloud Volumes Backup is supported.
- Quantum and OEM partners (DXi software 3.4.0 or later). Supported Quantum DXi systems include DXi4700 (NAS configuration), DXi4700 (multi-protocol configuration), DXi 4800, DXi 6900, DXi 6900-S, DXi 9000. FIPS-compliant operations mode requires DXi software 4.0 or later.

[1] This version is supported starting from Veeam Backup & Replication 11a (build 11.0.1.1261).

Once backups are created, they can be copied (for redundancy) or offloaded (for long-term retention) to one of the following hot object storage types using the scale-out backup repository Capacity Tier:

- Amazon S3 (including AWS Snowball Edge)
- Google Cloud Storage
- IBM Cloud Object Storage
- Microsoft Azure Blob storage (including Microsoft Azure Data Box)
- Any S3-compatible object storage (on-premises appliance, or cloud storage provider)

Once backups are copied or offloaded to Amazon S3 or Microsoft Azure Blob storage, they can be further d to one of the following respective cold object storage classes using the scale-out backup repository Archive Tier:

- Amazon S3 Glacier
- Amazon S3 Glacier Deep Archive
- Microsoft Azure Archive Tier

For the full list of partner-tested solutions including primary backup storage solutions, S3-compatible object storage solutions and offline storage solutions, see this Veeam page.

For information on NAS backup target, see **Backup, Archive, Secondary Repositories** in NAS Backup Support.

Veeam CDP Source and Target Datastores

The following source and target datastores are supported for Veeam CDP:

- NFS on file storage
- VMFS on block storage
- VMFS on internal ESXi storage
- VSAN

  VSAN is supported for all hyper-converged infrastructure (HCI) appliances.
- VVOL

  VVOL is supported for the following vendors: NetApp, HPE Nimble, Pure Storage and HPE 3PAR. For the list of tested vendor product lines, see this Veeam KB article.

Support for HCI appliances is pending validation by Veeam. The documentation will be updated based on the testing results. For the updates, see CDP Requirements and Limitations.

NOTE

[For Veeam Backup & Replication starting from 11a (build 11.0.1.1261)] VVOL and VSAN of each vendor have limits for the number of storage objects. Once the limit is reached, CDP starts to fail because creation of new objects cannot be completed. To restore the CDP process, remove some objects from VSAN/VVOL.

[For Veeam Backup & Replication prior 11a (build 11.0.1.1261)] I/O journal size is limited to 254 GB per VM for VSAN; to 4 GB per VM for VVOL.

Tape

Specification	Requirement
Hardware	LTO3 or later tape libraries (including VTL) and standalone drives are supported. Tape device must be directly attached to the backup server, to a tape server via SAS, FC or iSCSI interface. Note that VMware does not support connecting tape libraries to ESXi for VM pass-through.
Software	○ Tape devices without device-specific, vendor-supplied OEM drivers for Windows installed will appear in Windows Device Manager as Unknown or Generic and require enabling native SCSI commands mode. ○ If multiple driver installation modes are available for your tape device, use the one that allows for multiple open handles from a host to a drive to exist at the same time. Usually, such drivers are referred to as "non-exclusive". ○ No other backup server must be interacting with the tape device.

Storage Systems

Veeam Backup & Replication offers integration with the the following storage systems:

Cisco HyperFlex HX-Series

- NFS connectivity only
- HyperFlex 4.0(2x) or later (Backup from Storage Snapshots, Full Integration mode)
- Basic authentication is not supported for SSO users in HyperFlex

Dell EMC Isilon/PowerScale (NAS Backup only)

- NFS or CIFS connectivity
- OneFS 8.1.2 to 9.1

Dell EMC VNX, VNX2, VNXe and Unity XT/Unity

- NFS, Fibre Channel (FC) or iSCSI connectivity
- Dell EMC VNX/VNX2 all OE versions are supported

- Dell EMC Unity XT/Unity, VNXe OE versions 3.x up to 5.1

HPE 3PAR StoreServ

- Fibre Channel (FC) or iSCSI connectivity
- 3PAR OS versions 3.2.2 up to 3.3.1 MU5
- WSAPI 1.5 and later
- iSCSI VLAN tags are supported
- Virtual Domains are supported

HPE Nimble Storage AF-Series, HF-Series and CS-Series

- Fibre Channel (FC) or iSCSI connectivity
- Nimble OS 2.3 or later
- HPE Nimble synchronous replication is not supported

HPE Alletra 6000

- Fibre Channel (FC) or iSCSI connectivity
- OS version 6.0
- Synchronous replication is not supported

HPE Primera

- Fibre Channel (FC) or iSCSI (starting from OS versions 4.3 or later) connectivity
- OS versions 4.x
- Virtual Domains are supported

HPE Alletra 9000

- Fibre Channel (FC) or iSCSI connectivity
- OS version 9.3
- Virtual Domains are supported

HPE StoreVirtual/LeftHand/P4000 series and StoreVirtual VSA

- iSCSI connectivity only
- LeftHand OS versions 9.5 up to 12.8
- HPE SV3200 (LeftHand OS version 13) is not supported

IBM FlashSystem (StorWize), IBM SVC, Lenovo Storage V Series

- Fibre Channel (FC) or iSCSI connectivity

- IBM Spectrum Virtualize OS 7.6 or later

NetApp FAS/AFF, FlexArray (V-Series), ONTAP Edge/Select/Cloud VSA and FAS OEM (Fujitsu ETERNUS HX/AX, IBM N series and Lenovo DM series)

- NFS, Fibre Channel or iSCSI connectivity
- ONTAP 7-mode versions 8.2 up to 8.2.5
- ONTAP cluster-mode versions 8.3 to 9.9
- MetroCluster is supported
- ONTAP features application-aware data management and SVM-DR are not supported
- NetApp Synchronous SnapMirror is not supported

> **NOTE**
>
> To start working with the following storage systems, you must install storage system plug-in.

DataCore SANsymphony

- Fibre Channel (FC) or iSCSI connectivity
- DataCore SANsymphony 10.0 PSP12 or later

Dell EMC PowerMax

- Fibre Channel (FC) or iSCSI connectivity
- Dell EMC PowerMax/VMAX All Flash (PowerMax OS microcode family 5978 or later)
- Unisphere for PowerMax 9.2.1.6 or later

Dell EMC PowerStore

- Fibre Channel (FC) or iSCSI connectivity
- Dell EMC PowerStore T and PowerStore X series (PowerStore OS 2.0 or later)

Dell EMC SC Series/Compellent

- Fibre Channel (FC) or iSCSI connectivity
- Storage Center OS 7.4.2 or later
- FluidFS volumes and Live Volumes are not supported

Fujitsu ETERNUS AF/DX series

- Fibre Channel (FC) or iSCSI connectivity
- ETERNUS AF series: AF250 S2, AF650 S2, AF150 S3, AF250 S3, AF650 S3
- ETERNUS DX series: DX60 S4, DX100 S4, DX200 S4, DX500 S4, DX600 S4, DX8900 S4, DX60 S5, DX100 S5, DX200 S5, DX500 S5, DX600 S5, DX900 S5

- Storage firmware version:
  - ETERNUS AF S2 and DX S4 series (except DX8900 S4): V10L88-1000 or later
  - ETERNUS AF S3 and DX S5 series, DX8900 S4: V11L30-5000 or later

Hitachi VSP

- Fibre Channel (FC) or iSCSI connectivity
- VSP E590, E790, E990 (93-03-01-60/00 or later)
- VSP F350, F370, F700, F900 (88-07-01-x0/00 or later)
- VSP G350, G370, G700, G900 (88-07-01-x0/00 or later)
- VSP 5000 series (90-05-01-00/00 or later)

HPE XP

- Fibre Channel (FC) or iSCSI connectivity
- HPE XP8 (90-05-01-00/00 or later)

Huawei OceanStor

- NFS, Fibre Channel (FC) or iSCSI connectivity
- Huawei OceanStor Dorado V3 Series with software version V300R001 or later
- Huawei OceanStor Dorado V6 Series with software version 6.0 or later
- Huawei OceanStor V3 and F V3 Series with software version V300R006 or later
- Huawei OceanStor V5 and F V5 Series with software version V500R007 or later

INFINIDAT Infinibox F-series

- NFS, Fibre Channel (FC) or iSCSI connectivity
- InfiniBox 3.0 or later

> **NOTE**
>
> You must add to the backup infrastructure only one of the two InfiniBox storage arrays for which Active/Active Replication is configured, or exclude the replicating volumes on one of these arrays from rescan. For details on how to exclude volumes from rescan, see the Rescanning Storage Systems section in the Veeam Backup & Replication User Guide.

NEC Storage M Series

- M120, M320, M320F, M520, M720, M720F (Storage Control Software revision 1234 or later)
- Veeam Backup & Replication 11 P20210525 or later

NetApp SolidFire/HCI

- iSCSI connectivity
- NetApp SolidFire support requires Element OS 9.0 or later
- NetApp HCI support requires Element OS 10.0 or later

Pure Storage FlashArray

- Fibre Channel (FC) or iSCSI connectivity
- Purity 4.10 or later
- Purity ActiveCluster is supported. Direct access to a single proxy server from both storage arrays in the cluster is required.
- Replicated volume snapshots on the target array are supported. For details, see this Veeam KB article.

Tintri IntelliFlash/Western Digital IntelliFlash/Tegile

- NFS, Fibre Channel (FC) or iSCSI connectivity
- Tintri IntelliFlash 3.9.2, 3.10.1 or later

Gateway Server

Specification	Requirement
Hardware	*CPU*: x86-64 processor, 2 or more cores. *Memory*: 4 GB RAM, plus up to 4 GB RAM for each concurrently processed machine or file share. For more information, see Limitation of Concurrent Tasks. For RAM allocation recommendations for NAS backup, see Limitations and Recommendations for NAS backup. *Network*: 1 Gbps or faster for on-site backup and replication, and 1 Mbps or faster for off-site backup and replication. High latency and reasonably unstable WAN links are supported.

Specification	Requirement
OS	Both 32-bit and 64-bit versions of the following Microsoft Windows operating systems are supported: ○ Microsoft Windows Server Semi-Annual Channel (from version 1803 to version 20H2) ○ Microsoft Windows Server 2022[1] ○ Microsoft Windows Server 2019 ○ Microsoft Windows Server 2016 ○ Microsoft Windows Server 2012 R2 ○ Microsoft Windows Server 2012 ○ Microsoft Windows Server 2008 R2 SP1 ○ Microsoft Windows 11 (version 21H2[1]) ○ Microsoft Windows 10 (from version 1803 to version 21H1[1], and 21H2[2]) ○ Microsoft Windows 8.1 ○ Microsoft Windows 7 SP1 For the communication with object storage repositories, external repositories and NFS backup repositories, you can use machines running 64-bit versions of the following Linux distributions: ○ CentOS 7 to 8.3, and 8.4[1], and 8.5[2] ○ CentOS Stream ○ Debian 9.0 to 10.8, and 11.0[1] ○ Fedora 30 to 33, and 34[1], and 35[2] ○ openSUSE Leap 15.2 and 15.3[1], Tumbleweed ○ Oracle Linux 6 (UEK3) to 8.3 (UEK R6 and UEK R6 U2[1]) ○ Oracle Linux 6 to 8.3, and 8.4[1], and 8.5[2] (RHCK) ○ RHEL 6.0 to 8.3, and 8.4[1], and 8.5[2] ○ SLES 11 SP4, 12 SP1-SP5, 15 SP0-SP2, and SP3[1] ○ Ubuntu: 14.04 LTS, 16.04 LTS, 18.04 LTS, 19.10, 20.04 LTS, and 21.04[1], and 21.10[2] [1] This version is supported starting from Veeam Backup & Replication 11a (build 11.0.1.1261). [2] This version is supported starting from Veeam Backup & Replication 11a (build 11.0.1.1261) with Cumulative Patch P20211211.

For more information, see Gateway Server.

Mount Server

The mount server role must be assigned to a 64-bit Microsoft Windows machine added to the backup infrastructure. For the list of supported OSes, see Veeam Backup & Replication Console.

On the mount server machine, Veeam Backup & Replication installs the Veeam Mount Service. The Veeam Mount Service requires Microsoft .NET Framework 4.7.2. If Microsoft .NET Framework 4.7.2 is not installed on the machine, Veeam Backup & Replication will install it automatically.

If you plan to restore VM guest OS files from VMs running Microsoft Windows ReFS or from VMs with data deduplication enabled for some volumes, you must assign the mount server role to machines running specific OS versions. For more information, see ReFS and Data Deduplication subsections in Restoring VM Guest OS Files (FAT, NTFS or ReFS).

[For Microsoft Windows 2008R2/7] Make sure that you have SHA-2 code signing support installed. Normally, this component is included in Microsoft Windows updates. For more information, see Microsoft Docs.

Starting from Veeam Backup & Replication 10, the following jobs consume 512 MB per guest VM disk on mount server + 100 MB per VM:

- SureBackup
- Instant Recovery
- Instant Disk Recovery

For more information on the mount server, see Mount Server.

Veeam Backup Enterprise Manager Server

The machine where you plan to install Veeam Backup Enterprise Manager must meet the requirements listed in the System Requirements section of the Enterprise Manager User Guide.

Veeam Cloud Plug-Ins

The machine where you plan to install cloud plug-ins must meet the following requirements:

Cloud Plug-In	Requirement
AWS Plug-in for Veeam Backup & Replication	Microsoft .NET Core Runtime 3.1.10 or later Microsoft ASP.NET Core Shared Framework 3.1.10 or later
Microsoft Azure Plug-in for Veeam Backup & Replication	Microsoft .NET Core Runtime 3.1.10 or later Microsoft ASP.NET Core Shared Framework 3.1.10 or later
Google Cloud Platform Plug-in for Veeam Backup & Replication	Microsoft .NET Core Runtime 3.1.10 or later Microsoft ASP.NET Core Shared Framework 3.1.10 or later

Veeam Backup for Nutanix AHV

The machine where you plan to install Veeam Backup for Nutanix AHV must have Microsoft .NET Core Runtime 2.1.23 or any later minor updates for version 2.1 installed.

Supported Applications

You can create transactionally-consistent backups or replicas of VMs that run the following applications.

Application	Requirement
Microsoft Active Directory	Veeam Backup & Replication supports domain controller backup for the following operating systems: o Microsoft Windows Server Semi-Annual Channel (from version 1803 to version 20H2) o Microsoft Windows Server 2022[1] o Microsoft Windows Server 2019 o Microsoft Windows Server 2016 o Microsoft Windows Server 2012 R2 o Microsoft Windows Server 2012 o Microsoft Windows Server 2008 R2 o Microsoft Windows Server 2008 Minimum supported domain and forest functional level is Windows 2008. [1] This version is supported starting from Veeam Backup & Replication 11a (build 11.0.1.1261).
Microsoft Exchange	The following versions of Microsoft Exchange are supported: o Microsoft Exchange 2019 o Microsoft Exchange 2016 o Microsoft Exchange 2013 SP1 o Microsoft Exchange 2013 o Microsoft Exchange 2010 SP1, SP2, or SP3
Microsoft SharePoint	The following versions of Microsoft SharePoint Server (virtualized either on VMware or Hyper-V platform) are supported: o Microsoft SharePoint 2019 o Microsoft SharePoint 2016 o Microsoft SharePoint 2013 o Microsoft SharePoint 2010 All editions are supported (Foundation, Standard, Enterprise).

Application	Requirement
Microsoft SQL Server	The following versions of Microsoft SQL Server are supported (for application-aware processing and transaction log backup): ○ Microsoft SQL Server 2019 (only for Windows) ○ Microsoft SQL Server 2017 (only for Windows) ○ Microsoft SQL Server 2016 SP2 ○ Microsoft SQL Server 2014 SP3 ○ Microsoft SQL Server 2012 SP4 ○ Microsoft SQL Server 2008 R2 SP3 ○ Microsoft SQL Server 2008 SP4 ○ Microsoft SQL Server 2005 SP4 All editions of Microsoft SQL Server are supported. The database whose logs you want to back up must use the *Full* or *Bulk-logged* recovery model. In this case, all changes of the Microsoft SQL Server state will be written to transaction logs, and you will be able to replay transaction logs to restore the Microsoft SQL Server. You can use the Microsoft SQL Server Management Studio to switch to one of these models. For more information, see Microsoft Docs.
Oracle on Windows OS	Veeam Backup & Replication supports backup of the following Oracle Database versions: ○ **Oracle Database 11g Release 2** For supported operating systems, see Oracle documentation. ○ **Oracle Database 12c Release 1** For supported operating systems, see Oracle documentation. ○ **Oracle Database 12c Release 2** For supported operating systems, see Oracle documentation. ○ **Oracle Database 18c** For supported operating systems, see Oracle documentation. ○ **Oracle Database 19c** For supported operating systems, see Oracle documentation. ○ **Oracle Database 21c**[1] For supported operating systems, see Oracle documentation. [1] This version is supported starting from Veeam Backup & Replication 11a (build 11.0.1.1261) with Cumulative Patch P20211211.

Application	Requirement
Oracle on Linux OS	Veeam Backup & Replication supports backup of the following Oracle Database versions: ○ **Oracle Database 11g Release 2** For supported operating systems, see Oracle documentation. ○ **Oracle Database 12c Release 1** For supported operating systems, see Oracle documentation. ○ **Oracle Database 12c Release 2** For supported operating systems, see Oracle documentation. ○ **Oracle Database 18c** For supported operating systems, see Oracle documentation. ○ **Oracle Database 19c** For supported operating systems, see Oracle documentation. ○ **Oracle Database 21c**[1] For supported operating systems, see Oracle documentation. [1] This version is supported starting from Veeam Backup & Replication 11a (build 11.0.1.1261) with Cumulative Patch P20211211.
Oracle Database configuration	Consider the following: ○ Automatic Storage Management (ASM) is supported for Oracle 11g and later; requires *ASMlib* present. ○ Oracle Real Application Clusters (RAC) are not supported within the image-level backup functionality. Use Veeam Plug-in for Oracle RMAN. For details, see the Veeam Plug-in for Oracle RMAN section of the Veeam Plug-ins for Enterprise Applications User Guide. ○ Oracle Database Express Edition (XE) is supported for Windows-based machines only. ○ Configurations with different versions of Oracle Database deployed on the same server are not supported. Use Veeam Plug-in for Oracle RMAN. For details, see the Veeam Plug-in for Oracle RMAN section of the Veeam Plug-ins for Enterprise Applications User Guide. ○ To create Oracle database backups, all Oracle servers that use Data Guard must be added to the backup job. ○ You can use Veeam Plug-in for Oracle RMAN to integrate RMAN with Veeam Backup & Replication repositories. For details, see the Veeam Plug-in for Oracle RMAN section of the Veeam Plug-ins for Enterprise Applications User Guide.

> **NOTE**
>
> Consider that 32-bit Oracle running on 64-bit operating systems and Oracle Express Edition (XE) on Linux are not supported.

Veeam Explorers

- Veeam Explorer for Microsoft Active Directory
- Veeam Explorer for Microsoft Exchange
- Veeam Explorer for Microsoft SharePoint
- Veeam Explorer for Microsoft SQL
- Veeam Explorer for Microsoft Teams
- Veeam Explorer for Oracle

Required Permissions

Make sure the user accounts that you plan to use have permissions described in the following sections.

Installing and Using Veeam Backup & Replication

The accounts used for installing and using Veeam Backup & Replication must have the following permissions.

Account	Required Permission
Setup Account	The account used for product installation must have the local Administrator permissions on the target machine.
Veeam Backup & Replication Console Permissions	When you open the Veeam Backup & Replication console for the first time or after a cumulative patch is installed on the backup server, you must run the console under an account with the local Administrator permissions on the machine where the console is installed. In other cases (except file-level restore), you can run it under an account that is a member of the Users group on the machine where the console is installed. To perform file-level restore for Microsoft Windows VMs, the account must have the following permissions and privileges: • Local Administrator permissions to start the Veeam Backup & Replication console • *SeBackupPrivilege* and *SeRestorePrivilege* to connect to the Veeam backup server and start the restore process In most environments, *SeBackupPrivilege* and *SeRestorePrivilege* are assigned to user accounts added to the Administrators group. For more information, see Microsoft Docs. Accounts that are members of the Protected Users Active Directory group cannot be used to access the backup server remotely over the Veeam Backup & Replication console. For more information, see Microsoft Docs.
Veeam Backup Service Account	The account used to run the Veeam Backup Service must be a LocalSystem account or must have the local Administrator permissions on the backup server.

Account	Required Permission
Microsoft SQL Server (where the configuration database is stored)	You require different sets of Microsoft SQL permissions in the following cases: • Installation (remote or local): current account needs CREATE ANY DATABASE permission on the SQL server level. After database creation this account automatically gets a *db_owner* role and can perform all operations with the database. If the current account does not have this permission, a Database Administrator may create an empty database in advance and grant the *db_owner* role to the account that will be used for installing Veeam Backup & Replication. • Upgrade: current account should have sufficient permissions for that database. To grant these permissions through role assignment, it is recommended that you use the account with *db_owner* role. • Operation: the account used to run Veeam Backup Service requires *db_datareader* and *db_datawriter* roles as well as permissions to execute stored procedures for the configuration database on the Microsoft SQL Server. Alternatively, you can assign *db_owner* role for this database to the service account. For more information, see Microsoft Docs.

Using Virtualization Servers and Hosts

The following are required permissions to work with virtualization servers and hosts during data protection tasks.

Role	Required Permission
Source/Target Host	Root permissions on the ESXi host. If the vCenter Server is added to the backup infrastructure, an account that has administrative permissions is required. You can either grant the Administrator role to the account or configure granular vCenter Server permissions for certain Veeam Backup & Replication operations in the VMware vSphere environment. For more information, see the Required Permissions Reference.
Windows Server added to the backup infrastructure	Administrator permissions

Role	Required Permission
Linux Server added to the backup infrastructure	Permissions for the account that you specify when adding a Linux server differ depending on the role that you plan to assign to this server: • Roles for which Veeam Data Movers must be persistent (backup proxy, hardened/immutable repository) require root or equivalent permissions. For the full list of roles, see Veeam Data Movers. • Gateway server that communicates with NFS share requires root or equivalent permissions. • Backup repository requires read and write permissions on the folder where backups will be stored. You will configure this folder at the Configure Backup Repository Settings step of the backup repository wizard. • Other roles require read and write permissions on files and folders with which the server will work.
SMB Backup Repository	Write permission on the target folder and share.

Performing Guest Processing

To use guest OS processing (application-aware processing, pre-freeze and post-thaw scripts, transaction log processing, guest file indexing and file exclusions), make sure to configure your accounts according to the requirements listed in this section. For more information on guest processing, see Guest Processing.

All user accounts used for guest processing must have the following permissions:

- *Logon as a batch job* granted
- *Deny logon as a batch job* not set

Other permissions depend on applications that you back up. You can find permissions for backup operations in the following table. For restore operation permissions, see **Required Permissions** sections in the Veeam Explorers User Guide.

Application	Required Permission
Microsoft SQL Server	To back up Microsoft SQL Server data, the following roles must be assigned: • *Administrator* role on the target VM. • *Sysadmin* role on the target Microsoft SQL Server. If you need to provide minimal permissions, the account must be assigned the following roles and permissions: • SQL Server instance-level role: *public* and *dbcreator*. • Database-level roles and roles for the model system database: *db_backupoperator*, *db_denydatareader*, *public*; for the master system database — *db_backupoperator*, *db_datareader*, *public*; for the msdb system database — *db_backupoperator*, *db_datareader*, *public, db_datawriter*. • Securables: *view any definition*, *view server state, connect SQL*.
Microsoft Active Directory	To back up Microsoft Active Directory data, the account must be a member of the built-in *Administrators* group.
Microsoft Exchange	To back up Microsoft Exchange data, the account must have the local Administrator permissions on the machine where Microsoft Exchange is installed.
Oracle	The account specified at the **Guest Processing** step must be configured as follows: • For a Windows-based VM, the account must be a member of both the *Local Administrator* group and the *ORA_DBA* group (if OS authentication is used). In addition, if *ASM* is used, then such an account must be a member of the *ORA_ASMADMIN* group (for Oracle 12 and higher). • For a Linux-based VM, the account must be a Linux user elevated to *root.* To back up Oracle databases, make sure the account specified on the **Oracle** tab has been granted *SYSDBA* privileges. You can use either the same account that was specified at the **Guest Processing** step if such an account is a member of the *ORA_DBA* group for a Windows-based VM and *OSASM, OSDBA and OINSTALL* groups for a Linux-based VM, or you can use, for example, the *SYS* Oracle account or any other Oracle account that has been granted *SYSDBA* privileges. To perform guest processing for Oracle databases on Linux servers, make sure that the `/tmp` directory is mounted with the `exec` option. Otherwise, you will get an error with the permission denial.

Application	Required Permission
Microsoft SharePoint	To back up Microsoft SharePoint server, the account must be assigned the *Farm Administrator* role. To back up Microsoft SQL databases of the Microsoft SharePoint Server, the account must have the same privileges as that of Veeam Explorer for Microsoft SQL Server.

Consider the following general requirements when choosing a user account:

- [For guest OS file indexing] For Windows-based workloads, choose an account that has administrator privileges. For Linux-based workloads, choose an account of a root user or user elevated to root.
- To use networkless guest processing over VMware VIX/vSphere Web Services, you must specify one of the following accounts at the **Guest Processing** step of the backup wizard. Check that the account also has permissions listed in the table.
  - If Windows User Account Control (UAC) is enabled, specify Local Administrator (MACHINE\Administrator) or Domain Administrator (DOMAIN\Administrator) account.
  - If UAC is disabled, specify an account that is a member of the built-in Administrators group.
  - For Linux-based VMs, specify a root account.
- [For networkless guest processing over VMware VIX] To be able to perform more than 1000 guest processing operations, the user that you specify for guest processing must be logged into the VM at least once.
- When using Active Directory accounts, make sure to provide an account in the *DOMAIN|Username* format.
- When using local user accounts, make sure to provide an account in the *Username or HOST|Username* format.
- To process a Domain Controller server, make sure that you are using an account that is a member of the *DOMAIN|Administrators* group.
- To back up a Read-Only Domain controller, a delegated RODC administrator account is sufficient. For more information, see this Microsoft article.

Using Object Storage Repositories

Amazon S3 Object Storage Permissions

Consider the following:

- Make sure the account you are using has access to Amazon buckets and folders.
- The ListAllMyBuckets permission is not required if you specify the bucket name explicitly at the **Bucket** step of the New Object Repository wizard.

The following are required permissions to use Amazon S3 object storage with immutability disabled.

```
{
 "s3:ListBucket",
 "s3:GetBucketLocation",
 "s3:GetObject",
```

```
 "s3:PutObject",
 "s3:DeleteObject",
 "s3:ListAllMyBuckets",
 "s3:GetBucketVersioning"
}
```

The following are required permissions to use Amazon S3 object storage with immutability enabled. For more information on immutability, see Immutability.

```
{
   "s3:ListBucket",
   "s3:GetBucketLocation",
   "s3:GetObject",
   "s3:PutObject",
   "s3:DeleteObject",
   "s3:ListAllMyBuckets",
   "s3:GetBucketVersioning",
   "s3:GetBucketObjectLockConfiguration",
   "s3:ListBucketVersions",
   "s3:GetObjectVersion",
   "s3:GetObjectRetention",
   "s3:GetObjectLegalHold",
   "s3:PutObjectRetention",
   "s3:PutObjectLegalHold",
   "s3:DeleteObjectVersion"
  }
```

For examples, see this Veeam KB article. For more information on permissions, see this Amazon article.

Google Cloud Object Storage Permissions

Consider the following: the storage.buckets.list permission is not required if you specify the bucket name explicitly at the **Bucket** step of the New Object Repository wizard.

> **NOTE**
>
> The Owner IAM role does not necessarily grant the permissions required for working with Google Cloud Storage.

The following are required permissions to use Google Cloud object storage.

```
{
 "storage.buckets.get",
 "storage.buckets.list",
 "storage.objects.create",
 "storage.objects.delete",
 "storage.objects.get",
 "storage.objects.list"
}
```

Amazon S3 Glacier Storage Permissions

Permissions for Amazon S3 Glacier depend on whether you use the storage with immutability enabled or disabled.

The following are required permissions to use Amazon S3 Glacier object storage with disabled immutability.

- With VPC, subnet and security group set as Create new

```
{
 "Version": "2012-10-17",
 "Statement": [
   {
     "Sid": "VisualEditor0",
     "Effect": "Allow",
     "Action": [
       "s3:DeleteObject",
       "s3:PutObject",
       "s3:GetObject",
       "s3:RestoreObject",
       "s3:ListBucket",
       "s3:AbortMultipartUpload",
       "s3:GetBucketVersioning",
       "s3:ListAllMyBuckets",
       "s3:GetBucketLocation",
       "s3:GetBucketObjectLockConfiguration",
       "ec2:DescribeInstances",
       "ec2:CreateKeyPair",
       "ec2:DescribeKeyPairs",
       "ec2:RunInstances",
       "ec2:DeleteKeyPair",
       "ec2:DescribeVpcAttribute",
       "ec2:CreateTags",
       "ec2:DescribeSubnets",
       "ec2:TerminateInstances",
       "ec2:DescribeSecurityGroups",
       "ec2:DescribeImages",
       "ec2:DescribeVpcs",
       "ec2:CreateVpc",
       "ec2:CreateSubnet",
       "ec2:DescribeAvailabilityZones",
       "ec2:CreateRoute",
       "ec2:CreateInternetGateway",
       "ec2:AttachInternetGateway",
       "ec2:ModifyVpcAttribute",
       "ec2:CreateSecurityGroup",
       "ec2:DeleteSecurityGroup",
       "ec2:AuthorizeSecurityGroupIngress",
       "ec2:AuthorizeSecurityGroupEgress",
       "ec2:DescribeRouteTables",
       "ec2:DescribeInstanceTypes"
     ],
     "Resource": "*"
   }
 ]
}
```

- With preconfigured VPC, subnet and security group

```
{
 "Version": "2012-10-17",
 "Statement": [
   {
     "Sid": "VisualEditor0",
     "Effect": "Allow",
     "Action": [
       "s3:DeleteObject",
       "s3:PutObject",
       "s3:GetObject",
       "s3:RestoreObject",
       "s3:ListBucket",
       "s3:AbortMultipartUpload",
       "s3:GetBucketVersioning",
       "s3:ListAllMyBuckets",
       "s3:GetBucketLocation",
       "s3:GetBucketObjectLockConfiguration",
       "ec2:DescribeInstances",
       "ec2:CreateKeyPair",
       "ec2:DescribeKeyPairs",
       "ec2:RunInstances",
       "ec2:DeleteKeyPair",
       "ec2:DescribeVpcAttribute",
       "ec2:CreateTags",
       "ec2:DescribeSubnets",
       "ec2:TerminateInstances",
       "ec2:DescribeSecurityGroups",
       "ec2:DescribeImages",
       "ec2:DescribeVpcs"
     ],
     "Resource": "*"
   }
 ]
}
```

The following are required permissions to use Amazon S3 object storage with immutability enabled.

- With VPC, subnet and security group set as Create new

```
{
 "Version": "2012-10-17",
 "Statement": [
   {
     "Sid": "VisualEditor0",
     "Effect": "Allow",
     "Action": [
       "s3:DeleteObject",
       "s3:PutObject",
       "s3:GetObject",
       "s3:RestoreObject",
       "s3:ListBucket",
       "s3:AbortMultipartUpload",
       "s3:GetBucketVersioning",
       "s3:ListAllMyBuckets",
       "s3:GetBucketLocation",
       "s3:GetBucketObjectLockConfiguration",
       "s3:PutObjectRetention",
       "s3:GetObjectVersion",
       "s3:PutObjectLegalHold",
       "s3:GetObjectRetention",
       "s3:DeleteObjectVersion",
       "s3:ListBucketVersions",
       "ec2:DescribeInstances",
       "ec2:CreateKeyPair",
       "ec2:DescribeKeyPairs",
       "ec2:RunInstances",
       "ec2:DeleteKeyPair",
       "ec2:DescribeVpcAttribute",
       "ec2:CreateTags",
       "ec2:DescribeSubnets",
       "ec2:TerminateInstances",
       "ec2:DescribeSecurityGroups",
       "ec2:DescribeImages",
       "ec2:DescribeVpcs",
       "ec2:CreateVpc",
       "ec2:CreateSubnet",
       "ec2:DescribeAvailabilityZones",
       "ec2:CreateRoute",
       "ec2:CreateInternetGateway",
       "ec2:AttachInternetGateway",
       "ec2:ModifyVpcAttribute",
       "ec2:CreateSecurityGroup",
       "ec2:DeleteSecurityGroup",
       "ec2:AuthorizeSecurityGroupIngress",
       "ec2:AuthorizeSecurityGroupEgress",
       "ec2:DescribeRouteTables",
       "ec2:DescribeInstanceTypes"
     ],
     "Resource": "*"
   }
 ]
}
```

- With preconfigured VPC, subnet and security group

```
{
 "Version": "2012-10-17",
 "Statement": [
   {
     "Sid": "VisualEditor0",
     "Effect": "Allow",
     "Action": [
       "s3:DeleteObject",
       "s3:PutObject",
       "s3:GetObject",
       "s3:RestoreObject",
       "s3:ListBucket",
       "s3:AbortMultipartUpload",
       "s3:GetBucketVersioning",
       "s3:ListAllMyBuckets",
       "s3:GetBucketLocation",
       "s3:GetBucketObjectLockConfiguration",
       "s3:PutObjectRetention",
       "s3:GetObjectVersion",
       "s3:PutObjectLegalHold",
       "s3:GetObjectRetention",
       "s3:DeleteObjectVersion",
       "s3:ListBucketVersions",
       "ec2:DescribeInstances",
       "ec2:CreateKeyPair",
       "ec2:DescribeKeyPairs",
       "ec2:RunInstances",
       "ec2:DeleteKeyPair",
       "ec2:DescribeVpcAttribute",
       "ec2:CreateTags",
       "ec2:DescribeSubnets",
       "ec2:TerminateInstances",
       "ec2:DescribeSecurityGroups",
       "ec2:DescribeImages",
       "ec2:DescribeVpcs"
     ],
     "Resource": "*"
   }
 ]
}
```

Azure Archive Object Storage Permissions

The following are required permissions to use Azure Archive object storage.

```
{
 "properties": {
   "roleName": "CUSTOM_ROLE_MINIMAL_PERMISSIONS",
   "description": "CUSTOM_ROLE_MINIMAL_PERMISSIONS",
   "assignableScopes": [
     "/subscriptions/111111-1111-1111-0000-00000000000"
   ],
   "permissions": [
     {
       "actions": [
         "Microsoft.Authorization/*/read",
         "Microsoft.Compute/locations/*",
         "Microsoft.Compute/virtualMachines/*",
         "Microsoft.Network/locations/*",
         "Microsoft.Network/networkInterfaces/*",
         "Microsoft.Network/networkSecurityGroups/join/action",
         "Microsoft.Network/networkSecurityGroups/read",
         "Microsoft.Network/networkSecurityGroups/write",
         "Microsoft.Network/networkSecurityGroups/delete",
         "Microsoft.Network/publicIPAddresses/join/action",
         "Microsoft.Network/publicIPAddresses/read",
         "Microsoft.Network/publicIPAddresses/write",
         "Microsoft.Network/publicIPAddresses/delete",
         "Microsoft.Network/virtualNetworks/read",
         "Microsoft.Network/virtualNetworks/write",
         "Microsoft.Network/virtualNetworks/subnets/join/action",
         "Microsoft.Storage/storageAccounts/listKeys/action",
         "Microsoft.Storage/storageAccounts/read",
         "Microsoft.Resources/deployments/*",
         "Microsoft.Resources/subscriptions/resourceGroups/read",
         "Microsoft.Resources/checkResourceName/action",
         "Microsoft.Resources/subscriptions/resourceGroups/write",
         "Microsoft.Resources/subscriptions/locations/read"
       ],
       "notActions": [],
       "dataActions": [],
       "notDataActions": []
     }
   ]
 }
}
```

Integration with Storage Systems

To perform data protection and disaster recovery operations with storage snapshots, the account used to connect to a storage system must have necessary permissions.

NetApp Data ONTAP/Lenovo Thinksystem DM Permissions

The account used to connect to a NetApp Data ONTAP/Lenovo Thinksystem DM storage system must have the following permissions:

7-Mode

- login-http-admin
- api-system-*
- api-license-* (api-license-list-info)
- api-volume-*
- api-net-*
- api-options-*
- api-vfiler-*
- api-qtree-*
- api-nfs-*
- api-snapshot-*
- api-lun-*
- api-iscsi-*
- api-feature-*
- api-registry-*
- api-fcp-*
- api-file-*
- api-igroup-*
- api-clone-*
- api-snapvault-*
- api-snapmirror-*
- api-cf-*
- cli-options
- security-api-vfiler

CDOT (VMware Integration)

Command/Directory	Access/Query Level
DEFAULT	readonly
cluster	readonly
metrocluster	readonly

Command/Directory	Access/Query Level
fcp	readonly
file	readonly
igroup	all
iscsi	all
network	readonly
node	readonly
security	readonly
security login	readonly
set	readonly
snapmirror	all
system	readonly
version	readonly
qtree	readonly
lun	all
nfs	all
snapshot	all
volume	all
vserver	all

Only as SVM (VMware Integration)

Command/Directory	Access/Query Level
DEFAULT	none
lun	all
lun igroup	all
network	readonly
security	readonly
security login	readonly
snapmirror	all
system	readonly
version	readonly
volume	all
volume file	readonly
volume qtree	all
volume snapshot	all
vserver	all
vserver fcp	all
vserver iscsi	all
vserver nfs	all

CDOT (NAS Backup Integration)

Command/Directory	Access/Query Level
DEFAULT	readonly
security	readonly
security login	readonly
volume snapshot	all
vserver	all
vserver nfs	all

Only as SVM (NAS Backup Integration)

Command/Directory	Access/Query Level
DEFAULT	none
lun	readonly
network	readonly
security	readonly
security login	readonly
snapmirror	readonly
version	readonly
volume	readonly
volume snapshot	all
vserver	all

CDOT (Veeam Agent Integration)

Command/Directory	Access/Query Level
cluster	readonly
lun	all
metrocluster	readonly
network	readonly
system license	readonly
system node	readonly
version	readonly
volume	all
volume snapshot	all
vserver	all

Only as SVM (Veeam Agent Integration)

Command/Directory	Access/Query Level
lun	all
network	readonly
version	readonly
volume	all
volume snapshot	all
vserver	all

Universal Storage API Integrated Systems Permissions

The account used to connect to a Universal Storage API integrated system must be assigned a necessary role in the storage system console and/or have a set of necessary permissions.

- For Dell EMC PowerMax, the account must be assigned the Storage Administrator role.
- For Fujitsu ETERNUS, the account must be assigned the Software role.
- For NetApp SolidFire/HCI, the account must have the following permissions:
  - Volumes
  - Cluster Admins
- For Western Digital IntelliFlash, the account must be assigned the Veeam Admin Role.
- For DataCore, the account must have the following permissions:
  - General
  - Port
  - Host
  - Virtual disk
  - Snapshot
  - Physical disk
- For Hitachi VSP, the account must be assigned the following roles:
  - Storage Administrator (View Only)
  - Storage Administrator (Provisioning)
  - Storage Administrator (Local Copy)
- For HPE XP, the account must be assigned the following roles:
  - Storage Administrator (View Only)
  - Storage Administrator (Provisioning)
  - Storage Administrator (Local Copy)
- For Dell EMC PowerStore, the account must be assigned the following roles:
  - Administrator
  - Storage Administrator
  - Storage Operator
- For NEC Storage M Series, the account must be assigned the Administrator role

For privileges required to integrate the NAS backup feature with Dell EMC Isilon/PowerScale, see Integration with Dell EMC Isilon/PowerScale in the NAS Backup Support section.

For storage systems not mentioned above, the account must have Administrator role.

Used Ports

On backup infrastructure components, Veeam Backup & Replication automatically creates firewall rules for the required ports. These rules allow communication between the components.

> **IMPORTANT**
>
> Some Linux distributions require firewall and/or security rules to be created manually. For details, see this Veeam KB article.

You can find the full list of the ports below.

Microsoft Windows Server

The following table describes network ports that must be opened to ensure proper communication with Microsoft Windows servers.

Each Microsoft Windows server that is a backup infrastructure component or a machine for which you enable application-aware processing must have these ports opened. If you want to use the server as a backup infrastructure component, you must also open ports that the component role requires.

For example, if you assign the role of a backup proxy to your Microsoft Windows server, you must open ports listed below and also ports listed in the Backup Proxy section.

The Microsoft Windows server that acts as an NFS file share requires network ports listed below and also ports listed in the NFS Backup Repository. The Microsoft Windows server that acts as an SMB file share requires network ports listed below and also ports listed in the SMB Backup Repository.

From	To	Protocol	Port	Notes
Backup server	Microsoft Windows server	TCP	445	Port required for deploying Veeam Backup & Replication components.
Backup proxy		TCP	6160	Default port used by the Veeam Installer Service.
Backup repository		TCP	2500 to 3300[1]	Default range of ports used as data transmission channels and for collecting log files. For every TCP connection that a job uses, one port from this range is assigned.
Gateway server		TCP	6161	[For Microsoft Windows servers running the vPower NFS Service] Default port used by the Veeam vPower NFS Service.
Mount server		TCP	6162	Default port used by the Veeam Data Mover.

From	To	Protocol	Port	Notes
WAN accelerator		TCP	49152 to 65535 (for Microsoft Windows 2008 and newer)	Dynamic port range. For more information, see this Microsoft KB article.
Tape server				

[1] This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components.

Linux Server

The following table describes network ports that must be opened to ensure proper communication with Linux servers.

Each Linux server that is a backup infrastructure component or a machine for which you enable application-aware processing must have these ports opened. If you want to use the server as a backup infrastructure component, you must also open ports that the component role requires.

For example, if you assign the role of a backup repository to your Linux server, you must open ports listed below and also ports listed in the Microsoft Windows/Linux-based Backup Repository section.

The Linux server that acts as an NFS file share requires network ports listed below and also ports listed in the NFS Backup Repository. The Linux server that acts as an SMB file share requires network ports listed below and also ports listed in the SMB Backup Repository.

From	To	Protocol	Port	Notes
Backup server	Linux server	TCP	22	Port used as a control channel from the console to the target Linux host.
		TCP	6162	Default port used by the Veeam Data Mover. You can specify a different port while adding the Linux server to the Veeam Backup & Replication infrastructure. Note that you can specify a different port only if there is no previously installed Veeam Data Mover on this Linux server. For more information, see Specify Credentials and SSH Settings.
		TCP	2500 to 3300[1]	Default range of ports used as data transmission channels. For every TCP connection that a job uses, one port from this range is assigned.

From	To	Protocol	Port	Notes
Linux server	Backup server	TCP	2500 to 3300[1]	Default range of ports used as data transmission channels. For every TCP connection that a job uses, one port from this range is assigned.

[1] This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components.

Backup Server

The following table describes network ports that must be opened to ensure proper communication of the backup server with backup infrastructure components.

From	To	Protocol	Port	Notes
Virtualization Servers				
Backup server	vCenter Server	HTTPS TCP	443	Default port used for connections to vCenter Server. If you use vCloud Director, make sure you open port 443 on underlying vCenter Servers.
		HTTPS TCP	10443	Port used to access vCenter Inventory Service (HTTP or HTTPS) and collect vCenter Server tags. This port is used for communication with vCenter Server 5.x only. This port is not required for VMware Cloud on AWS.
	ESXi server	HTTPS TCP	443	Default port used for connections to ESXi host. This port is not required for VMware Cloud on AWS.
		TCP	902	Port used for data transfer to ESXi host. This port is not required for VMware Cloud on AWS.

From	To	Protocol	Port	Notes
	vCloud Director	HTTPS TCP	443	Default port used for connections to vCloud Director.
Other Servers				
Backup server	Microsoft SQL Server hosting the Veeam Backup & Replication configuration database	TCP	1433	Port used for communication with Microsoft SQL Server on which the Veeam Backup & Replication configuration database is deployed (if you use a Microsoft SQL Server default instance). Additional ports may need to be open depending on your configuration. For more information, see Microsoft Docs.
	DNS server with forward/reverse name resolution of all backup servers	UDP	53	Port used for communication with the DNS Server.
	Veeam Update Notification Server (dev.veeam.com)	HTTPS TCP	443	Default port used to download information about available updates from the Veeam Update Notification Server over the Internet.
	Veeam License Update Server (vbr.butler.veeam.com, autolk.veeam.com)	TCP	443	Default port used for license auto-update.
Backup Server				
Backup server	Backup server	TCP	9501	Port used locally on the backup server for communication between Veeam Broker Service and Veeam services and components.
Backup server	Backup server	TCP	6172	Port used to provide REST access to the Veeam Backup & Replication database.
Remote Access				

From	To	Protocol	Port	Notes
Management client PC (remote access)	Backup server	TCP	3389	Default port used by the Remote Desktop Services. If you use third-party solutions to connect to the backup server, other ports may need to be open.
REST API				
REST client	Backup server	TCP	9419	Default port for communication with REST API service.

The following table describes network ports that must be opened to ensure proper communication with the Veeam Backup & Replication console installed remotely.

From	To	Protocol	Port	Notes
Veeam Backup & Replication Console	Backup server	TCP	9392	Port used by the Veeam Backup & Replication console to connect to the backup server.
		TCP	10003	Port used by the Veeam Backup & Replication console to connect to the backup server only when managing the Veeam Cloud Connect infrastructure.
		TCP	9396	Port used by the Veeam.Backup.UIService process for managing database connections.
Veeam Backup & Replication Console	Mount server (if the mount server is not located on the console)	TCP	2500 to 3300[1]	Default range of ports used as data transmission channels. For every TCP connection that a job uses, one port from this range is assigned.

[1] This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components.

Backup Proxy

The following table describes network ports that must be opened to ensure proper communication of backup proxies with other backup components.

From	To	Protocol	Port	Notes
Backup server	Backup proxy	Backup proxy can be a Microsoft Windows or Linux server. Depending on which server you use, the ports listed in Microsoft Windows Server or Linux Server must be opened.		
Communication with Backup Server				
Backup server	Backup proxy	TCP	6210	Default port used by the Veeam Backup VSS Integration Service for taking a VSS snapshot during the SMB file share backup.
Communication with VMware Servers				
Backup proxy	vCenter Server	HTTPS	443	Default VMware web service port that can be customized in vCenter settings.
	ESXi server	TCP	902	Default VMware port used for data transfer. This port is not required for VMware Cloud on AWS.
		HTTPS	443	Default VMware web service port that can be customized in ESXi host settings. Not required if vCenter connection is used. This port is not required for VMware Cloud on AWS.
Communication with Backup Repositories				
Backup proxy	Microsoft Windows server	TCP	49152 to 65535 (for Microsoft Windows 2008 and newer)	Dynamic port range. For more information, see this Microsoft KB article.

From	To	Protocol	Port	Notes
	SMB (CIFS) share	TCP	445 135[1]	Ports used as a transmission channel from the backup proxy to the target SMB (CIFS) share. Traffic goes between the backup proxy and the SMB (CIFS) share only if a gateway server is not specified explicitly in SMB (CIFS) backup repository settings (the **Automatic selection** option is used). If a gateway server is specified explicitly, traffic goes between the gateway server and the SMB (CIFS) share. For more information about required ports, see the Gateway server > SMB (CIFS) share line below in this table.
	NFS share	TCP, UDP	111, 2049	Ports used as a transmission channel from the backup proxy to the target NFS share. Traffic goes between the backup proxy and the NFS share only if a gateway server is not specified explicitly in NFS backup repository settings (the **Automatic selection** option is used). If a gateway server is specified explicitly, traffic goes between the gateway server and the NFS share. For more information about required ports, see the Gateway server > NFS share line below in this table.
	Gateway server	TCP	49152 to 65535 (for Microsoft Windows 2008 and newer)	Dynamic port range. For more information, see this Microsoft KB article.
Gateway server (if a gateway server is specified explicitly in SMB (CIFS) backup repository settings)	SMB (CIFS) share	TCP	445 135[1]	Ports used as a transmission channel from the gateway server to the target SMB (CIFS) share.

From	To	Protocol	Port	Notes
Gateway server (if a gateway server is specified explicitly in NFS backup repository settings)	NFS share	TCP, UDP	111, 2049	Ports used as a transmission channel from the gateway server to the target NFS share.
Communication with Backup Proxies				
Backup proxy	Backup proxy	TCP	2500 to 3300[2]	Default range of ports used as transmission channels for replication jobs. For every TCP connection that a job uses, one port from this range is assigned.

[1] Port 135 is optional to provide faster deployment.

[2] This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components.

Backup Repositories

- Microsoft Windows/Linux-based Backup Repository
- NFS Backup Repository
- SMB Backup Repository
- Dell EMC Data Domain System
- ExaGrid
- HPE StoreOnce
- Quantum DXi
- Object Storage Repository
- External Repository
- Archive Object Storage Repository

Microsoft Windows/Linux-based Backup Repository

The following table describes network ports that must be opened to ensure proper communication with backup repositories. Cache repositories in NAS backup use the same network ports as backup repositories.

From	To	Protocol	Port	Notes
Backup proxy	Microsoft Windows server performing the role of the backup repository/file server	Ports listed in Microsoft Windows Server must be opened.		
Backup proxy	Linux server performing the role of the backup repository/file server	Ports listed in Linux Server must be opened.		
Backup proxy	Backup repository	TCP	2500 to 3300[1]	Default range of ports used as transmission channels for replication jobs. For every TCP connection that a job uses, one port from this range is assigned.
Source backup repository	Target backup repository	TCP	2500 to 3300[1]	Default range of ports used as transmission channels for backup copy jobs. For every TCP connection that a job uses, one port from this range is assigned. Ports 2500 to 3300 are used for backup copy jobs that do not utilize WAN accelerators. If the backup copy job utilizes WAN accelerators, make sure that ports specific for WAN accelerators are open.
Source backup repository	Object storage repository gateway server	TCP	2500 to 3300[1]	Default range of ports used as transmission channels for replication jobs. For every TCP connection that a job uses, one port from this range is assigned.
Backup repository/ secondary backup repository	Cache repository in NAS backup	TCP	2500 to 3300[1]	Default range of ports used as transmission channels for file share backup restore jobs. For every TCP connection that a job uses, one port from this range is assigned.

From	To	Protocol	Port	Notes
Microsoft Windows server running vPower NFS Service	Backup repository gateway server working with backup repository	TCP	2500 to 3300[1]	Default range of ports used as transmission channels during Instant Recovery, SureBackup or Linux file-level recovery. For every TCP connection that a job uses, one port from this range is assigned.

[1] This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components.

NFS Backup Repository

The following table describes network ports that must be opened to ensure proper communication with NFS shares added as backup repositories.

From	To	Protocol	Port	Notes
Microsoft Windows server performing the role of the gateway server/backup proxy	NFS backup repository/file share	Ports listed in Microsoft Windows Server must be opened.		
Linux server performing the role of the gateway server/backup proxy	NFS backup repository/file share	Ports listed in Linux Server must be opened.		
Gateway server/backup proxy (Microsoft Windows/Linux)	NFS backup repository/file share	TCP UDP	2049	Default NFS port.
		TCP UDP	111	Port used for rpcbind service.
Gateway server/backup proxy (Microsoft Windows/Linux)	NFS backup repository/file share (for repositories supporting	TCP UDP	mountd_port	Dynamic port used for mountd service. Can be assigned statically.
		TCP UDP	statd_port	Dynamic port used for statd service. Can be assigned statically.

From	To	Protocol	Port	Notes
	NFS protocol version 3)	TCP	lockd_port	Dynamic TCP port used for lockd service. Can be assigned statically.
		UDP	lockd_port	Dynamic UDP port used for lockd service. Can be assigned statically.
Gateway server/backup proxy (specified in the NFS repository settings)	NFS backup repository/file share	TCP UDP	111, 2049	Standard NFS ports used as a transmission channel from the gateway server to the target NFS share.

SMB Backup Repository

The following table describes network ports that must be opened to ensure proper communication with SMB (CIFS) shares added as backup repositories.

From	To	Protocol	Port	Notes
Microsoft Windows server performing the role of the gateway server/backup proxy	SMB (CIFS) backup repository/file share	Ports listed in Microsoft Windows Server must be opened.		
Gateway server/backup proxy (Microsoft Windows)	SMB (CIFS) backup repository	TCP	445 135[1]	Ports used as a transmission channel from the gateway server to the target SMB (CIFS) share.

[1] Port 135 is optional to provide faster deployment.

Dell EMC Data Domain System

For more information, see Dell EMC Documents.

From	To	Protocol	Port	Notes
Backup server or Gateway server	Dell EMC Data Domain	TCP	111	Port used to assign a random port for the mountd service used by NFS and DDBOOST. Mountd service port can be statically assigned.
		TCP	2049	Main port used by NFS. Can be modified using the 'nfs set server-port' command. Command requires SE mode.
		TCP	2052	Main port used by NFS MOUNTD. Can be modified using the 'nfs set mountd-port' command in SE mode.
Backup server	Gateway server	Ports listed in Gateway Server must be opened.		

ExaGrid

From	To	Protocol	Port	Notes
Backup server or Gateway server	ExaGrid	TCP	22	Default command port used for communication with ExaGrid.
Backup server	Gateway server	Ports listed in Gateway Server must be opened.		

HPE StoreOnce

From	To	Protocol	Port	Notes
Backup server or Gateway server	HPE StoreOnce	TCP	9387	Default command port used for communication with HPE StoreOnce.
			9388	Default data port used for communication with HPE StoreOnce.

From	To	Protocol	Port	Notes
Backup server	Gateway server	Ports listed in Gateway Server must be opened.		

Quantum DXi

From	To	Protocol	Port	Notes
Backup server or Gateway server	Quantum DXi	TCP	22	Default command port used for communication with Quantum DXi.
Backup server	Gateway server	Ports listed in Gateway Server must be opened.		

Object Storage Repository

The following table describes network ports and endpoints that must be opened to ensure proper communication with object storage repositories. For more information, see Object Storage Repository.

From	To	Protocol	Port/Endpoint	Notes
Gateway server	Amazon S3 Object Storage	TCP	443	Used to communicate with Amazon S3 Object Storage.
		HTTPS	**Cloud endpoints:** • *.amazonaws.com (for both *Global* and Government regions) • *.amazonaws.com.cn (for *China* region) A complete list of connection endpoints can be found in this Amazon article.	
		TCP	80	

From	To	Protocol	Port/Endpoint	Notes
		HTTP	**Certificate verification endpoints:** • *.amazontrust.com	Used to verify the certificate status. Consider that certificate verification endpoints (CRL URLs and OCSP servers) are subject to change. The actual list of addresses can be found in the certificate itself.
	Microsoft Azure Object Storage	TCP	443	Used to communicate with Microsoft Azure Object Storage. Consider that the <xxx> part of the address must be replaced with your actual storage account URL, which can be found in the Azure management portal.
		HTTPS	**Cloud endpoints:** • xxx.blob.core.windows.net (for *Global* region) • xxx.blob.core.chinacloudapi.cn (for *China* region) • xxx.blob.core.cloudapi.de (for *Germany* region) • xxx.blob.core.usgovcloudapi.net (for *Government* region)	
		TCP	80	

From	To	Protocol	Port/Endpoint	Notes
		HTTP	**Certificate verification endpoints:** • ocsp.digicert.com • ocsp.msocsp.com • *.d-trust.net	Used to verify the certificate status. Consider the following: • Certificate verification endpoints (CRL URLs and OCSP servers) are subject to change. The actual list of addresses can be found in the certificate itself. • The **.d-trust.net* endpoint is used for the *Germany* region only.
	Google Cloud Storage	TCP	443	Used to communicate with Google Cloud Storage.
		HTTPS	**Cloud endpoints:** • storage.googleapis.com A complete list of connection endpoints can be found in this Google article.	
		TCP	80	Used to verify the certificate status. Consider that certificate verification endpoints (CRL URLs and OCSP servers) are subject to change. The actual list of addresses can be found in the certificate itself.
		HTTP	**Certificate verification endpoints:** • ocsp.pki.goog • pki.goog • crl.pki.goog	

From	To	Protocol	Port/Endpoint	Notes
	IBM Cloud Object Storage	TCP/HTTPS	Customizable and depends on device configuration	Used to communicate with IBM Cloud Object Storage.
	S3 Compatible Object Storage	TCP/HTTPS	Customizable and depends on device configuration	Used to communicate with S3 Compatible Object Storage.

External Repository

The following table describes network ports and endpoints that must be opened to ensure proper communication with external repositories. For more information, see External Repository.

From	To	Protocol	Port/Endpoint	Notes
Gateway server	Amazon S3 Object Storage	TCP	443	Used to communicate with Amazon S3 Object Storage.
		HTTPS	**Cloud endpoints:** • *.amazonaws.com (for both *Global* and Government regions) • *.amazonaws.com.cn (for *China* region) A complete list of connection endpoints can be found in this Amazon article.	
		TCP	80	Used to verify the certificate status. Consider that certificate verification endpoints (CRL URLs and OCSP servers) are subject to change. The actual list of addresses can be found in the certificate itself.
		HTTP	**Certificate verification endpoints:** • *.amazontrust.com	
		TCP	443	

From	To	Protocol	Port/Endpoint	Notes
	Microsoft Azure Object Storage	HTTPS	**Cloud endpoints:** • xxx.blob.core.windows.net (for *Global* region) • xxx.blob.core.chinacloudapi.cn (for *China* region) • xxx.blob.core.cloudapi.de (for *Germany* region) • xxx.blob.core.usgovcloudapi.net (for *Government* region)	Used to communicate with Microsoft Azure Object Storage. Consider that the <xxx> part of the address must be replaced with your actual storage account URL, which can be found in the Azure management portal.
		TCP	80	Used to verify the certificate status. Consider the following: • Certificate verification endpoints (CRL URLs and OCSP servers) are subject to change. The actual list of addresses can be found in the certificate itself. • The **.d-trust.net* endpoint is used for the *Germany* region only.
		HTTP	**Certificate verification endpoints:** • ocsp.digicert.com • ocsp.msocsp.com • *.d-trust.net	
	Google Cloud Storage	TCP	443	Used to communicate with Google Cloud Storage.
		HTTPS	**Cloud endpoints:** • storage.googleapis.com A complete list of connection endpoints can be found in this Google article.	
		TCP	80	

From	To	Protocol	Port/Endpoint	Notes
		HTTP	**Certificate verification endpoints:** • ocsp.pki.goog • pki.goog • crl.pki.goog	Used to verify the certificate status. Consider that certificate verification endpoints (CRL URLs and OCSP servers) are subject to change. The actual list of addresses can be found in the certificate itself.

Archive Object Storage Repository

The following table describes network ports and endpoints that must be opened to ensure proper communication with object storage repositories used as a part of Archive Tier. For more information, see Archive Tier.

From	To	Protocol	Port/Endpoint	Notes
Gateway server	Amazon EC2 proxy appliance	TCP	443 (default, adjustable via Amazon S3 Glacier wizard)	If there is no gateway server selected, VBR server will be used as a gateway server
		SSH	22	
		HTTPS	**Cloud endpoints:** • Public/private IPv4 addresses of EC2 appliances.	
	Microsoft Azure proxy appliance	TCP	443 (default, adjustable via Azure Archive wizard)	
		SSH	22	
		HTTPS	**Cloud endpoints:** • Public/private IPv4 addresses of Azure appliances.	
		TCP	443	

From	To	Protocol	Port/Endpoint	Notes
Amazon EC2 proxy appliance	Amazon S3 Object Storage	HTTPS	**Cloud endpoints:** • *.amazonaws.com (for both Global and Government regions) • *.amazonaws.com.cn (for China region) A complete list of connection endpoints can be found in this Amazon article.	Used to communicate with Amazon S3 Object Storage.
		TCP	80	Used to verify the certificate status. Consider that certificate verification endpoints (CRL URLs and OCSP servers) are subject to change. The actual list of addresses can be found in the certificate itself.
		HTTP	**Certificate verification endpoints:** • *.amazontrust.com	
Microsoft Azure proxy appliance	Microsoft Azure Object Storage	TCP	443	Used to communicate with Microsoft Azure Object Storage. The <xxx> part of the address must be replaced with your actual storage account URL, which can be found in the Microsoft Azure management portal.
		HTTPS	**Cloud endpoints:** • xxx.blob.core.windows.net (for Global region) • xxx.blob.core.chinacloudapi.cn (for China region) • xxx.blob.core.cloudapi.de (for Germany region) • xxx.blob.core.usgovcloudapi.net (for Government region)	
		TCP	80	

From	To	Protocol	Port/Endpoint	Notes
		HTTP	**Certificate verification endpoints:** • ocsp.digicert.com • ocsp.msocsp.com • *.d-trust.net	Used to verify the certificate status. Certificate verification endpoints (CRL URLs and OCSP servers) are subject to change. The actual list of addresses can be found in the certificate itself. The **.d-trust.net* endpoint is used for the Germany region only.

Storage Systems

- Dell EMC VNX(e) Storage
- HPE 3PAR StoreServ Storage
- HPE Lefthand Storage
- HPE Nimble Storage
- IBM Spectrum Virtualize Storage
- NetApp Data ONTAP Storage
- Universal Storage API Integrated System

Dell EMC VNX(e) Storage

From	To	Protocol	Port	Notes
Backup server	VNX File	SSH	22	Default command port used for communication with VNX File over SSH.
	VNX Block	HTTPS	443	Default port used for communication with Dell EMC VNX Block.
	VNXe	HTTPS	443	Default port used for communication with Dell EMC VNXe and sending REST API calls.
Backup proxy	VNX Block VNXe	TCP	3260	Default iSCSI target port.

From	To	Protocol	Port	Notes
	VNX File VNXe	TCP, UDP	2049, 111	Standard NFS ports. Port 111 is used by the port mapper service.

HPE 3PAR StoreServ Storage

From	To	Protocol	Port	Notes
Backup server	HPE 3PAR StoreServ storage system	HTTP	8008	Default port used for communication with HPE 3PAR StoreServ over HTTP.
		HTTPS	8080	Default port used for communication with HPE 3PAR StoreServ over HTTPS.
		SSH	22	Default command port used for communication with HPE 3PAR StoreServ over SSH.
Backup proxy	HPE 3PAR StoreServ storage system	TCP	3260	Default iSCSI target port.

HPE Lefthand Storage

From	To	Protocol	Port	Notes
Backup server	HPE Lefthand storage system	SSH	16022	Default command port used for communication with HPE Lefthand.
Backup proxy	HPE Lefthand storage system	TCP	3260	Default iSCSI target port.

HPE Nimble Storage

From	To	Protocol	Port	Notes
Backup server	HPE Nimble storage system	TCP	5392	Default command port used for communication with HPE Nimble (used for Nimble OS 2.3 and later).
Backup proxy	HPE Nimble storage system	TCP	3260	Default iSCSI target port.

IBM Spectrum Virtualize Storage

From	To	Protocol	Port	Notes
Backup server	IBM Spectrum Virtualize storage system	SSH	22	Default command port used for communication with IBM Spectrum Virtualize over SSH.
Backup proxy	IBM Spectrum Virtualize storage system	TCP	3260	Default iSCSI target port.

NetApp Data ONTAP Storage

From	To	Protocol	Port	Notes
Backup server	NetApp Data ONTAP storage system	HTTP	80	Default command port used for communication with NetApp Data ONTAP over HTTP.
		HTTPS	443	Default command port used for communication with NetApp Data ONTAP over HTTPS.
Backup proxy	NetApp Data ONTAP storage system	TCP, UDP	2049, 111	Standard NFS ports. Port 111 is used by the port mapper service.
		TCP	3260	Default iSCSI target port.

Universal Storage API Integrated System

The following tables describe network ports that must be opened to ensure proper communication with Universal Storage API integrated systems:

- DataCore SANsymphony
- Dell EMC SC Series/Compellent
- Fujitsu ETERNUS DX/AF
- INFINIDAT InfiniBox
- Huawei OceanStor
- NetApp SolidFire/HCI
- Pure Storage FlashArray
- Tintri IntelliFlash/Western Digital/Tegile

DataCore SANsymphony

From	To	Protocol	Port	Notes
Backup server	DataCore SANsymphony storage system	HTTPS	443	Default command port used for communication with DataCore SANsymphony over HTTPS.
Backup proxy	DataCore SANsymphony storage system	TCP	3260	Default iSCSI target port.

Dell EMC SC Series/Compellent

From	To	Protocol	Port	Notes
Backup server	Dell EMC SC Series storage system	HTTPS	3033	Default command port used for communication with Dell EMC SC Series over HTTPS.
Backup proxy	Dell EMC SC Series storage system	TCP	3260	Default iSCSI target port.

Fujitsu ETERNUS DX/AF

From	To	Protocol	Port	Notes
Backup server	Fujitsu ETERNUS DX/AF storage system	SSH	22	Default command port used for communication with Fujitsu ETERNUS DX/AF over SSH.
Backup proxy	Fujitsu ETERNUS DX/AF storage system	TCP	3260	Default iSCSI target port.

INFINIDAT InfiniBox

From	To	Protocol	Port	Notes
Backup server	INFINIDAT InfiniBox storage system	HTTPS	443	Default command port used for communication with INFINIDAT InfiniBox over HTTPS.
Backup proxy	INFINIDAT InfiniBox storage system	TCP	3260	Default iSCSI target port.

Huawei OceanStor

From	To	Protocol	Port	Notes
Backup server	Huawei OceanStor storage system	HTTPS	8080	Default port used for communication with Huawei OceanStor over HTTPS.
Backup proxy	Huawei OceanStor storage system	TCP	3260	Default iSCSI target port.

NetApp SolidFire/HCI

From	To	Protocol	Port	Notes
Backup server	NetApp SolidFire/HCI storage system	HTTPS	443	Default command port used for communication with NetApp SolidFire/HCI over HTTPS.
Backup proxy	NetApp SolidFire/HCI storage system	TCP	3260	Default iSCSI target port.

Pure Storage FlashArray

From	To	Protocol	Port	Notes
Backup server	Pure Storage FlashArray system	HTTPS	443	Default command port used for communication with Pure Storage FlashArray over HTTPS.
Backup proxy	Pure Storage FlashArray system	TCP	3260	Default iSCSI target port.

Tintri IntelliFlash/Western Digital/Tegile

From	To	Protocol	Port	Notes
Backup server	Tintri IntelliFlash system	HTTPS	443	Default command port used for communication with Tintri IntelliFlash over HTTPS.
Backup proxy	Tintri IntelliFlash system	TCP	3260	Default iSCSI target port.
	Tintri IntelliFlash system	TCP, UDP	2049, 111	Standard NFS ports. Port 111 is used by the port mapper service.

Gateway Server

The following table describes network ports that must be opened to ensure proper communication with gateway servers.

From	To	Protocol	Port	Notes
Backup server	Microsoft Windows server performing the role of the gateway server	Ports listed in Microsoft Windows Server must be opened.		
Backup server	Linux server performing the role of the gateway server (if a gateway server is specified explicitly in NFS backup repository settings)	Ports listed in Linux Server must be opened.		
Gateway server (if a gateway server is specified explicitly in SMB (CIFS) backup repository settings)	SMB (CIFS) share	TCP	445 135[1]	Ports used as a transmission channel from the gateway server to the target SMB (CIFS) share.
Gateway server (if a gateway server is specified explicitly in NFS backup repository settings)	NFS share	TCP, UDP	111, 2049	Ports used as a transmission channel from the gateway server to the target NFS share.

[1] Port 135 is optional to provide faster deployment.

Tape Server

The following table describes network ports that must be opened to ensure proper communication with tape servers.

From	To	Protocol	Port	Notes
Backup server	Tape server	Tape server is a Microsoft Windows server, and it requires the ports listed in Microsoft Windows Server to be opened.		
		TCP	6166	Controlling port for RPC calls.
		TCP	2500 to 3300[1]	Default range of ports used as data transmission channels. For every TCP connection that a job uses, one port from this range is assigned.
Tape server	Backup server	TCP	2500 to 3300[1]	Default range of ports used as data transmission channels. For every TCP connection that a job uses, one port from this range is assigned.
	Backup repository, gateway server or proxy server	Tape server is a Microsoft Windows server, and it requires the ports listed in Microsoft Windows Server to be opened.		

[1] This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components.

WAN Accelerator

The following table describes network ports that must be opened to ensure proper communication between WAN accelerators used in backup copy jobs and replication jobs.

From	To	Protocol	Port	Notes
Backup server	WAN accelerator (source and target)	WAN accelerator is a Microsoft Windows server, and it requires the ports listed in Microsoft Windows Server to be opened.		
		TCP	6160	Default port used by the Veeam Installer Service.
		TCP	6162	Default port used by the Veeam Data Mover.

From	To	Protocol	Port	Notes
		TCP	6164	Controlling port for RPC calls.
WAN accelerator (source and target)	Backup repository (source and target)	TCP	2500 to 3300[1]	Default range of ports used as data transmission channels. For every TCP connection that a job uses, one port from this range is selected dynamically.
WAN accelerator	WAN accelerator	TCP	6164	Controlling port for RPC calls.
		TCP	6165	Default port used for data transfer between WAN accelerators. Ensure this port is open in firewall between sites where WAN accelerators are deployed.

[1] This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components.

VM Guest OS

Connections with Non-Persistent Runtime Components

The following tables describe network ports that must be opened to ensure proper communication of the backup server and backup infrastructure components with the non-persistent runtime components deployed inside the VM guest OS for application-aware processing and indexing.

From	To	Protocol	Port	Notes
Backup server	Linux VM guest OS	TCP	22	Default SSH port used as a control channel.
	Guest interaction proxy	TCP	6190	Port used for communication with the guest interaction proxy.
		TCP	6290	Port used as a control channel for communication with the guest interaction proxy.
		TCP	445	Port used as a transmission channel.

From	To	Protocol	Port	Notes
Guest interaction proxy	ESXi server	TCP	443	Default port used for connections to ESXi host. [For VMware vSphere earlier than 6.5] Not required if vCenter connection is used. In VMware vSphere versions 6.5 and later, port 443 is required by VMware web services.

Network ports described in the table below are NOT required when working in networkless mode over VMware VIX/vSphere Web Services.

From	To	Protocol	Port	Notes
Guest interaction proxy or Mount server	Microsoft Windows VM guest OS	TCP	445 135[1]	Ports required to deploy the runtime coordination process on the VM guest OS.
		TCP	49152 to 65535 (for Microsoft Windows 2008 and newer)	Dynamic port range used by the runtime process deployed inside the VM for guest OS interaction (when working over the network, not over VIX API).[2] For more information, see this Microsoft KB article.
		TCP	6167	[For Microsoft SQL logs shipping] Port used for preparing the database and taking logs.
	Linux VM guest OS	TCP	22	Default SSH port used as a control channel.
VM guest OS	Guest interaction proxy or Mount server	TCP	2500 to 3300[3]	Default range of ports used as transmission channels for log shipping. For every TCP connection that a job uses, one port from this range is assigned.

[1] Port 135 is optional to provide faster deployment.

[2] If you use default Microsoft Windows firewall settings, you do not need to configure dynamic RPC ports: during setup, Veeam Backup & Replication automatically creates a firewall rule for the runtime process. If you use firewall settings other than default ones or application-aware processing fails with the "*RPC function call failed*" error, you need to configure dynamic RPC ports. For more information on how to configure RPC dynamic port allocation to work with firewalls, see this Microsoft KB article.

[3] This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components.

Connections with Persistent Agent Components

The following table describes network ports that must be opened to ensure proper communication of the backup server with the persistent agent components deployed inside the VM guest OS for application-aware processing and indexing.

From	To	Protocol	Port	Notes
Guest interaction proxy or Mount server	VM guest OS	TCP	6160, 11731	Default port and failover port used by the Veeam Installer Service.
		TCP	6167	Port used by the Veeam Log Shipping Service to collect and transfer logs.
		TCP	6173 2500	Ports used by the Veeam Guest Helper for guest OS processing and file-level restore.

CDP Components

The following table describes network ports that must be opened to ensure proper communication of Veeam CDP components with other backup components.

From	To	Protocol	Port	Notes
ESXi host (source)	CDP proxy (source)	TCP	33032	Default port used as a transmission channel to the source CDP proxy.
	ESXi host (source)	TCP	33033	Port used locally on the source ESXi host for data transfer between I/O filter components.
	ESXi host (source)	TCP	33038	Port used locally on the source ESXi host for communication between CDP components.
CDP proxy (source)	CDP proxy (target)	TCP	33033	Default port used as a transmission channel to the target CDP proxy.
	ESXi host (source and target)	TCP	902	Default VMware port used for data transfer. Used during the initial synchronization.
	vCenter Server (source and target)	TCP	443	Default VMware web service port that can be customized in vCenter settings. Used during the initial synchronization.

From	To	Protocol	Port	Notes
CDP proxy (target)	ESXi host (target)	TCP	33032	Default port used as a transmission channel to the target ESXi host.
	ESXi host (source and target)	TCP	902	Default VMware port used for data transfer. Used during the initial synchronization.
	vCenter Server (source and target)	TCP	443	Default VMware web service port that can be customized in vCenter settings. Used during the initial synchronization.
ESXi host (target)	ESXi host (target)	TCP	33034	Port used locally on the target ESXi host for communication between the I/O filter components during failover.
	ESXi host (target)	TCP	33038	Port used locally on the source ESXi host for communication between CDP components.
Backup server	ESXi host (source and target)	TCP	443	Port used as a control channel.
	CDP proxy (source and target)	TCP	6182	Port used as a control channel.
	Backup server	TCP	9509	Port used locally on the backup server for communication between Veeam Backup Service and Veeam CDP Coordinator Service.
ESXi host (source and target)	Backup server	TCP	33034	Port used for communication with Veeam CDP Coordinator Service.
CDP proxy (source and target)	Backup server	TCP	33034	Port used for communication with Veeam CDP Coordinator Service.

Recovery Components

- Mount Server

- Multi-OS FLR
- Veeam vPower NFS Service
- SureReplica Recovery Verification
- Veeam U-AIR
- Microsoft Active Directory Domain Controller Connections During Application Item Restore
- Microsoft Exchange Server Connections During Application Item Restore
- Microsoft SQL Server Connections During Application Item Restore
- Proxy Appliance (Restore to Amazon EC2, Google Cloud)
- Azure Proxy
- Azure Helper Appliance
- Azure Stack

Mount Server

The following table describes network ports that must be opened to ensure proper communication with mount servers.

From	To	Protocol	Port	Notes
Backup server	Mount server	Mount server is a Microsoft Windows server, and it requires the ports listed in Microsoft Windows Server to be opened.		
		TCP	6170	Port used for communication with a local or remote Mount Service.
Mount server (or machine running the Veeam Backup & Replication console)	Backup server	TCP	9401	Port used for communication with the Veeam Backup Service.
Mount server (or machine running the Veeam Backup & Replication console)	Backup repository	TCP	2500 to 3300[1]	Default range of ports used for communication with a backup repository.
Mount server	Helper appliance	TCP	22	Default SSH port used as a control channel.
		TCP	2500 to 3300[1]	Default range of ports used for communicating with the appliance.
Mount server	VM guest OS	Ports listed in VM Guest OS must be opened.		

[1] This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components.

Multi-OS FLR

From	To	Protocol	Port	Notes
Backup server	Helper appliance	TCP	22	Port used as a communication channel from the backup server to the helper appliance in the multi-OS file-level recovery process.
		TCP	2500 to 3300[1]	Default range of ports used as data transmission channels. For every TCP connection that a job uses, one port from this range is assigned.
	VM guest OS	TCP	22	Port used as a communication channel from the backup server to the VM guest OS during multi-OS file-level recovery process.
		TCP	2500 to 3300[1]	Default range of ports used as data transmission channels. For every TCP connection that a job uses, one port from this range is assigned.
Helper appliance	VM guest OS	TCP	22	Port used as a communication channel from the helper appliance to the VM guest OS during multi-OS file-level recovery process.
		TCP	20	[If FTP option is used] Default port used for data transfer.
		TCP	2500 to 3300[1]	Default range of ports used as data transmission channels. For every TCP connection that a job uses, one port from this range is assigned.
VM guest OS	Helper appliance	TCP	22	Port used as a communication channel from the VM guest OS to the helper appliance during multi-OS file-level recovery process.

From	To	Protocol	Port	Notes
		TCP	21	[If FTP option is used] Default port used for protocol control messages.
		TCP	2500 to 3300[1]	Default range of ports used as data transmission channels. For every TCP connection that a job uses, one port from this range is assigned.
Mount server	Helper appliance	TCP	22	Port used as a communication channel from the mount server to the helper appliance in the multi-OS file-level recovery process.
		TCP	2500 to 3300[1]	Default range of ports used as data transmission channels. For every TCP connection that a job uses, one port from this range is assigned.
	Backup repository	TCP	2500 to 3300[1]	Default range of ports used for communication with a backup repository.
Helper host	Backup repository	TCP	2500 to 3300[1]	Default range of ports used as data transmission channels. For every TCP connection that a job uses, one port from this range is assigned.

[1] This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components.

Veeam vPower NFS Service

From	To	Protocol	Port	Notes
Backup server	Microsoft Windows server running vPower NFS Service	TCP	6160	Default port used by the Veeam Installer Service.
		TCP	6161	Default port used by the Veeam vPower NFS Service.

From	To	Protocol	Port	Notes
ESXi host	Microsoft Windows server running vPower NFS Service	TCP UDP	111	Standard port used by the port mapper service.
		TCP UDP	1058+ or 1063+	Default mount port. The number of port depends on where the vPower NFS Service is located: • 1058+: If the vPower NFS Service is located on the backup server. • 1063+: If the vPower NFS Service is located on a separate Microsoft Windows machine. If port 1058/1063 is occupied, the succeeding port numbers will be used.
		TCP UDP	2049+	Standard NFS port. If port 2049 is occupied, the succeeding port numbers will be used.
Backup repository or Gateway server working with backup repository	Microsoft Windows server running vPower NFS Service	TCP	2500 to 3300[1]	Default range of ports used as transmission channels during Instant Recovery, SureBackup or Linux file-level recovery. For every TCP connection that a job uses, one port from this range is assigned.

[1] This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components.

SureReplica Recovery Verification

From	To	Protocol	Port	Notes
Backup server	vCenter Server	HTTPS TCP	443	Default port used for connections to vCenter Server.
	ESXi server	HTTPS TCP	443	Default port used for connections to ESXi host. Not required if vCenter connection is used.
	Proxy appliance	TCP	443	Port used for communication with the proxy appliance in the virtual lab.

From	To	Protocol	Port	Notes
			22	Port used for communication with the proxy appliance in the virtual lab.
	Applications on VMs in the virtual lab	—	—	Application-specific ports to perform port probing test. For example, to verify a DC, Veeam Backup & Replication probes port 389 for a response.
Internet-facing proxy server	VMs in the virtual lab	HTTP	8080	Port used to let VMs in the virtual lab access the Internet.

Veeam U-AIR

The following table describes network ports that must be opened to ensure proper communication of U-AIR wizards with other components.

From	To	Protocol	Port	Notes
U-AIR wizards	Veeam Backup Enterprise Manager	TCP	9394	Default port used for communication with Veeam Backup Enterprise Manager. Can be customized during Veeam Backup Enterprise Manager installation.

Microsoft Active Directory Domain Controller Connections During Application Item Restore

The following table describes network ports that must be opened to ensure proper communication of the backup server with the Microsoft Active Directory VM during application-item restore.

From	To	Protocol	Port	Notes
Backup server	Microsoft Active Directory VM guest OS	TCP	135	Port required for communication between the domain controller and backup server.
		TCP, UDP	389	LDAP connections.
		TCP	636, 3268, 3269	LDAP connections.

From	To	Protocol	Port	Notes
		TCP	49152 to 65535 (for Microsoft Windows 2008 and newer)	Dynamic port range used by the runtime coordination process deployed inside the VM guest OS for application-aware processing (when working over the network, not over VIX API).[1] For more information, see this Microsoft KB article.

[1] If you use default Microsoft Windows firewall settings, you do not need to configure dynamic RPC ports: during setup, Veeam Backup & Replication automatically creates a firewall rule for the runtime process. If you use firewall settings other than default ones or application-aware processing fails with the *"RPC function call failed"* error, you need to configure dynamic RPC ports. For more information on how to configure RPC dynamic port allocation to work with firewalls, see this Microsoft KB article.

Microsoft Exchange Server Connections During Application Item Restore

The following table describes network ports that must be opened to ensure proper communication of the Veeam backup server with the Microsoft Exchange Server system during application-item restore.

From	To	Protocol	Port	Notes
Backup server	Microsoft Exchange 2003/2007 CAS Server	TCP	80, 443	WebDAV connections.
	Microsoft Exchange 2010/2013/2016/2019 CAS Server	TCP	443	Microsoft Exchange Web Services Connections.

Microsoft SQL Server Connections During Application Item Restore

The following table describes network ports that must be opened to ensure proper communication of the backup server with the VM guest OS system during application-item restore.

From	To	Protocol	Port	Notes
Backup server	Microsoft SQL VM guest OS	TCP	1433, 1434 and other	Port used for communication with the Microsoft SQL Server installed inside the VM. Port numbers depends on configuration of your Microsoft SQL server. For more information, see Microsoft Docs.

Proxy Appliance (Restore to Amazon EC2, Google Cloud)

From	To	Protocol	Port	Notes
Backup server/Backup Repository	Proxy appliance	TCP	22	Port used as a communication channel to the proxy appliance in the restore to Amazon EC2 or Google Cloud process.
		TCP	443	Default redirector port. You can change the port in proxy appliance settings. For details, see **Specify Proxy Appliance** in Restore to Amazon EC2 and Restore to Google Cloud.

Azure Proxy

From	To	Protocol	Port	Notes
Backup server/ Backup repository	Azure proxy	TCP	443	Default management and data transport port required for communication with the Azure proxy. The port must be opened on the backup server and backup repository storing VM backups. The default port is 443, but you can change it in the settings of the Azure Proxy. For details, see Specify Credentials and Transport Port

Azure Helper Appliance

From	To	Protocol	Port	Notes
Backup server	Azure helper appliance	TCP	22	Port used as a communication channel to the proxy appliance in the Restore to Azure process The default port is 22, but you can change it during helper appliance deployment. For details, see Configuring Helper Appliances.

Azure Stack

From	To	Protocol	Port	Notes
Backup server	Azure Stack	HTTPS	443, 30024	Default management and data transport port required for communication with the Azure Stack.

Veeam Backup Enterprise Manager

Veeam Backup Enterprise Manager Connections

Veeam Explorers

- Veeam Explorer for Microsoft Active Directory Connections
- Veeam Explorer for Microsoft Exchange Connections
- Veeam Explorer for Microsoft SharePoint and Veeam Explorer for Microsoft OneDrive for Business Connections
- Veeam Explorer for Microsoft SQL Server Connections
- Veeam Explorer for Microsoft Teams Connections
- Veeam Explorer for Oracle Connections

Veeam Cloud Connect

Veeam Cloud Connect Connections

Veeam Agents

Veeam Agent for Microsoft Windows

- Connections for Veeam Agent for Microsoft Windows Operating in Managed Mode
- Connections for Veeam Agent for Microsoft Windows Operating in Standalone Mode

Veeam Agent for Linux

- Connections for Veeam Agent for Linux Operating in Managed Mode
- Connections for Veeam Agent for Linux Operating in Standalone Mode

Veeam Agent for Mac

- Connections for Veeam Agent for Mac Operating in Managed Mode
- Connections for Veeam Agent for Mac Operating in Standalone Mode

Veeam Plug-ins for Enterprise Applications

- Veeam Plug-in for SAP HANA Connections
- Veeam Plug-in for Oracle RMAN Connections

Veeam Plug-ins for Cloud Solutions

- AWS Plug-in for Veeam Backup & Replication
- Microsoft Azure Plug-in for Veeam Backup & Replication
- Google Cloud Platform Plug-in for Veeam Backup & Replication

Kasten K10

Kasten K10 Connections

Other Connections

NDMP Server

The following table describes network ports that must be opened to ensure proper communication with NDMP servers.

From	To	Protocol	Port	Notes
Gateway server	NDMP server	NDMP	10000	Port used for data transfer between the components.

SMTP Server

The following table describes network ports that must be opened to ensure proper communication of the backup server with the SMTP server.

From	To	Protocol	Port	Notes
Backup server	SMTP server	TCP	25	Port used by the SMTP server.

Internet Connections

If you use an HTTP(S) proxy server to access the Internet, make sure that WinHTTP settings are properly configured on Microsoft Windows machines with Veeam backup infrastructure components. For information on how to configure WinHTTP settings, see Microsoft Docs.

NOTE

Tenants cannot access Veeam Cloud Connect infrastructure components through HTTP(S) proxy servers. For information on supported protocols for Veeam Cloud Connect, see the Used Ports section in the Veeam Cloud Connect Guide.

Trusted Root Certificates

There are several root certificates necessary for the Veeam Backup & Replication to operate correctly.

In most cases, these certificates are already installed on Microsoft Windows machines. Some Microsoft Windows installations do not contain needed root certificates authority as trusted root certificates, or have non-current certificates. This may happen on servers with locked down security settings, or servers with no internet access or if the latest updates are not installed.

Removal of these certificates from the backup server may limit the functionality of Veeam Backup & Replication or may cause it to fail.

Make sure the following root certificates are installed on the backup server:

- https://www.digicert.com/CACerts/DigiCertAssuredIDRootCA.crt (DigiCert Assured ID Root CA)
- http://secure.globalsign.com/cacert/gscodesigng3ocsp.crt (GlobalSign CodeSigning CA - G3)
- https://www.digicert.com/CACerts/DigiCertHighAssuranceEVRootCA.crt (DigiCert High Assurance EV Root CA)
- https://www.digicert.com/CACerts/DigiCertEVCodeSigningCA-SHA2.crt (DigiCert EV Code Signing CA - SHA2)
- https://support.globalsign.com/customer/portal/articles/1426602-globalsign-root-certificates (install R1, R2, and R3 certificates)

If your backup server does not have internet access, download certificate files from another computer.

Infrastructure Planning

Infrastructure of Veeam Backup & Replication depends on business needs and resources of your company. There are a number of things to consider before you start deploying Veeam Backup & Replication. In this section you will find tips and recommendations that may help you design your backup infrastructure.

Before you deploy Veeam Backup & Replication, consider the following steps:

1. Define protection scope.
2. Define RTO and RPO goals.
3. Select Veeam Backup & Replication features that you will need.
4. Plan how many copies of your data you need to store (3-2-1 rule).
5. Design Veeam Backup & Replication infrastructure.

Step 1. Define Protection Scope

Define how many machines do you need to protect and the amount of disk space in use by those machines.

After defining the protection scope, calculate how much of the total amount of data is actually changing on a daily basis. This information is required because of the mechanism of how Veeam Backup & Replication creates a backup chain. At the first run, Veeam Backup & Replication creates a full backup file; at the second and further runs, Veeam Backup & Replication creates an incremental backup file that contains only the blocks that has been changes since the last backup. As a result, the daily change rate has a significant impact on the backup window and the storage capacity needed to store the backups. As Veeam Backup & Replication creates image-level or block level backups, you need to know the daily change rate on the block level. For VMware vSphere or Microsoft Hyper-V, you can use Veeam ONE to measure and generate a report on the daily change rate of VMs.

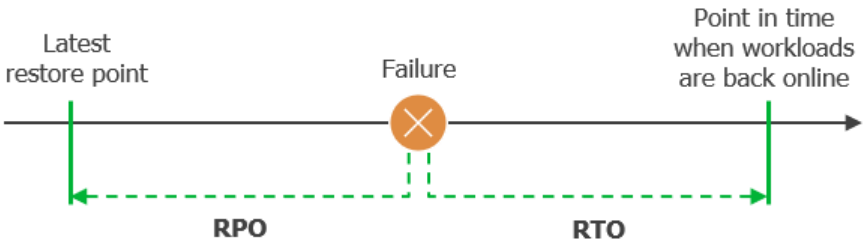
As a result of this step, you can make a list of machines that need to be protected, including the data on which of the machines contain databases, which of the machines host business critical applications, and how much of the total amount of data is changing on these machines on a daily basis. This information will help you in further steps of deployment planning.

Step 2. Define RPO and RTO

When making a business continuity and disaster recovery plan, you must define two important parameters: Recovery Point Objective (RPO) and Recovery Time Objective (RTO).

- **RPO** defines a period during which you may accept to lose data. Basically, it is the age of the the latest backup that will be used for recovery in case of a failure. It means that your company accepts that in case of a failure you may lose the data that has been accumulated since the latest restore point. RPO set by the policy of your company defines how often do you need to create a recovery point. This will help you estimate how much storage you will need to store backups, how many copies of your data do you need, and which Veeam Backup & Replication features are the most suitable for business needs of your company.

- **RTO** is related to downtime. RTO represents the amount of time from the beginning of an incident until all services are back online and available to users.



Define a list of your workloads grouped and organized by how fast they must be recoverable. Divide the list into categories. The higher the recovery priority, the lower the RTO will be required relative to the rest of your workloads.

Step 3. Select Veeam Backup & Replication Features

Based on the analysis of your RTO and RPO, you can define your protection plan and select which features are the most suitable for your business needs. It is a common practice to divide servers and applications into categories and use different protection functionality for each category based on SLA (service level agreement). You can take the following table as a reference.

	RPO: Seconds	RPO: Minutes	RPO: Hours (<24h)	RPO: Hours (24-48)
RTO: Seconds	Continuous Data Protection (for VMware vSphere)	Replication		
RTO: Minutes		Snapshot Orchestration (for VMware vSphere)	Backup	Backup Copy
RTO: Hours				Tape Devices Support

Apart from backup and replication options, the RTO also depends on the method of recovery and recovery verification. Veeam Backup & Replication offers a number of recovery options for various disaster recovery scenarios, including Instant Recovery, image-level restore, file-level restore, restore of application items and so on. For details, see the following sections:

- Restore: performing restore from backup files to the original or a new location.
- Recovery Verification: testing VM backups and replicas before recovery.

Step 4. Plan How Many Copies of Data You Need (3-2-1 rule)

To build a successful data protection and disaster recovery plan, we recommend to follow the 3-2-1 rule:

- 3: You must have at least three copies of your data: the original production data and two backups. The first copy of your data is the original production data. The second copy of your data is a backup created by a backup job. To create the third copy of data you can use Backup Copy or Backup to Tape jobs. Also, if you use cloud repositories (S3, Azure Blob, Google Cloud, IBM Cloud), you can copy backups to a capacity tier.
- 2: You must use at least two different types of media to store the copies of your data.
- 1: You must keep at least one backup off-site. For example, in the cloud or in a remote site. One of the repositories must be offline, air-gaped or immutable.

Veeam Backup & Replication provides integration with various types of repositories. Select where you want to store your backup files. For the full list of supported backup repositories, see Backup Repository, Tape Devices Support, Storage Systems Support.

To plan the required space on repositories, you might also need to analyze for how long will you store the backups. Veeam Backup & Replication provides short-term and long-term (GFS) retention policies to effectively store the backup files.

Step 5. Design Veeam Backup & Replication Infrastructure

Veeam Backup & Replication can be used in virtual environments of any size and complexity. The architecture of the solution supports on-site and off-site data protection, operations across remote sites and geographically dispersed locations. Veeam Backup & Replication provides flexible scalability and easily adapts to the needs of your virtual environment.

Before you install Veeam Backup & Replication, familiarize yourself with common deployment scenarios and plan your backup infrastructure layout. For details, see Deployment Scenarios.

The easiest way to start is to deploy a Veeam Backup & Replication server, one dedicated server for a backup proxy and one repository. While you keep adding backup jobs, add more proxies and repositories. Each backup infrastructure component has its own specifics and requirements that are described in corresponding sections of this guide:

- Veeam Backup & Replication Server
- Backup Proxy
- Backup Repository

Also, keep in mind that in most cases, it is not recommended to deploy Veeam Backup & Replication, Veeam Backup Enterprise Manager and Veeam ONE on the same server. It is recommended to deploy them on separate servers.

Naming Conventions

Do not use Microsoft Windows reserved names for names of the backup server, managed servers, backup repositories, jobs, tenants and other objects created in Veeam Backup & Replication: CON, PRN, AUX, NUL, COM1, COM2, COM3, COM4, COM5, COM6, COM7, COM8, COM9, LPT1, LPT2, LPT3, LPT4, LPT5, LPT6, LPT7, LPT8 and LPT9.

If you plan to store backups on a repository operating in the per-machine mode, do not use Microsoft Windows reserved names for names of the virtual machines to back up.

If you use a reserved name, Veeam Backup & Replication may not work as expected. For more information on naming conventions in Microsoft Windows, see Microsoft Docs.

Security Considerations

When you set up the backup infrastructure, one thing that you must not overlook is security. The backup infrastructure can be potentially used as a backdoor to gain access to your systems and data.

This section includes a number of recommendations that will help you prevent potential security issues and reduce the risk of compromising sensitive data.

General Considerations

Ensure timely guest OS updates on backup infrastructure servers. Install the latest updates and patches on backup infrastructure servers to minimize the risk of exploiting guest OS vulnerabilities by attackers.

Backups and Replicas

A potential source of vulnerability is the backup or replica itself. To secure data stored in backups and replicas, consider the following recommendations:

- **Ensure physical security of target servers**. Check that only authorized personnel have access to the room where your target servers (backup repositories and hosts) reside.
- **Restrict user access to backups and replicas**. Check that only authorized users have permissions to access backups and replicas on target servers.
- **Encrypt data in backups**. Use Veeam Backup & Replication built-in encryption to protect data in backups. To guarantee security of data in backups, follow Encryption Best Practices.

Data Communication Channel

Backup data can be intercepted in-transit, when it is communicated from source to target over a network. To secure the communication channel for backup traffic, consider the following recommendations:

- **Isolate backup traffic**. Use an isolated network to transport data between backup infrastructure components — backup server, backup proxies, repositories and so on.
- **Encrypt network traffic**. By default, Veeam Backup & Replication encrypts network traffic travelling between public networks. To ensure secure communication of sensitive data within the boundaries of the same network, you can also encrypt backup traffic in private networks. For details, see Enabling Network Data Encryption.

Internet Access for Backup Servers

Some Veeam Backup & Replication functionality requires that backup servers have outbound Internet access. For example, to enable product update check, automatic license update and license usage reporting, a backup server must be connected to the Internet and be able to send requests to servers on the Internet.

However, inbound connectivity to backup servers from the Internet must not be allowed. If you want to manage backup servers remotely over the Internet, you can deploy the Veeam Backup & Replication console on a jump server. Service providers who want to manage backup servers remotely can use the Veeam Backup Remote Access functionality. For more information, see the Using Remote Access Console section in the Veeam Cloud Connect Guide.

The account used for RDP access must not have local Administrator privileges on the jump server, and you must never use the saved credentials functionality for RDP access or any other remote console connections. To restrict users from saving RDP credentials, you can use Group Policies. For more information, see Experts Exchange.

Credentials

An attacker who gained high-privilege access to backup infrastructure servers can get credentials of user accounts and compromise other systems in your environment.

Particularly, backup proxies must be considered the target for compromise. During backup, proxies obtain from the backup server credentials required to access virtual infrastructure servers. A person having administrator privileges on a backup proxy can intercept the credentials and use them to access the virtual infrastructure.

One of the most possible causes of a credential theft are missing guest OS updates and use of outdated authentication protocols. To mitigate risks, consider the following recommendations:

- **Choose strong encryption algorithms for SSH**. To communicate with Linux servers deployed as part of the backup infrastructure, Veeam Backup & Replication uses SSH. Make sure that for the SSH tunnel you use a strong and proven encryption algorithm, with sufficient key length. Ensure that private keys are kept in a highly secure place, and cannot be uncovered by a 3rd party.

  For Linux hardened repository, instead of SSH Veeam Backup & Replication uses self-signed certification (SHA256RSA self-signed certificates with 2048-bit RSA key).

- **Avoid using password authentication to connect to remote servers over SSH**. Using key-based SSH authentication is generally considered more secure than using password authentication and helps averting man-in-the middle (MITM) attacks.

Veeam Backup & Replication Database

Another security concern you must consider is protecting the Veeam Backup & Replication configuration database. The database stores credentials of user accounts required to connect to virtual servers and other systems in the backup infrastructure. All passwords stored in the database are encrypted. However, a user with administrator privileges on the backup server can decrypt the passwords, which presents a potential threat.

To secure the Veeam Backup & Replication configuration database, consider the following recommendations:

- **Restrict user access to the database**. Check that only authorized users can access the backup server and the server that hosts the Veeam Backup & Replication configuration database (if the database runs on a remote server).

- **Encrypt data in configuration backups**. Enable data encryption for configuration backup to secure sensitive data stored in the configuration database. For details, see Creating Encrypted Configuration Backups.

Veeam Cloud Connect

Veeam Cloud Connect secures communication between the provider side and tenant side with TLS. If an attacker obtains a provider's private key, backup traffic can be eavesdropped and decrypted. The attacker can also use the certificate to impersonate the provider (MITM attack).

Veeam Cloud Connect providers must consider the following recommendations:

Keep the certificate in a secure place. Make sure that the TLS certificate is kept in a highly secure place and cannot be uncovered by a 3rd party.

Kerberos Authentication for Guest OS Processing

Starting from version 9.5 Update 4, Veeam Backup & Replication supports Kerberos authentication for guest OS processing of VMware vSphere VMs. However NTLM authentication is still required for communication between Veeam backup infrastructure servers (backup server, backup proxies, backup repositories, guest interaction proxies, log shipping servers, mount servers).

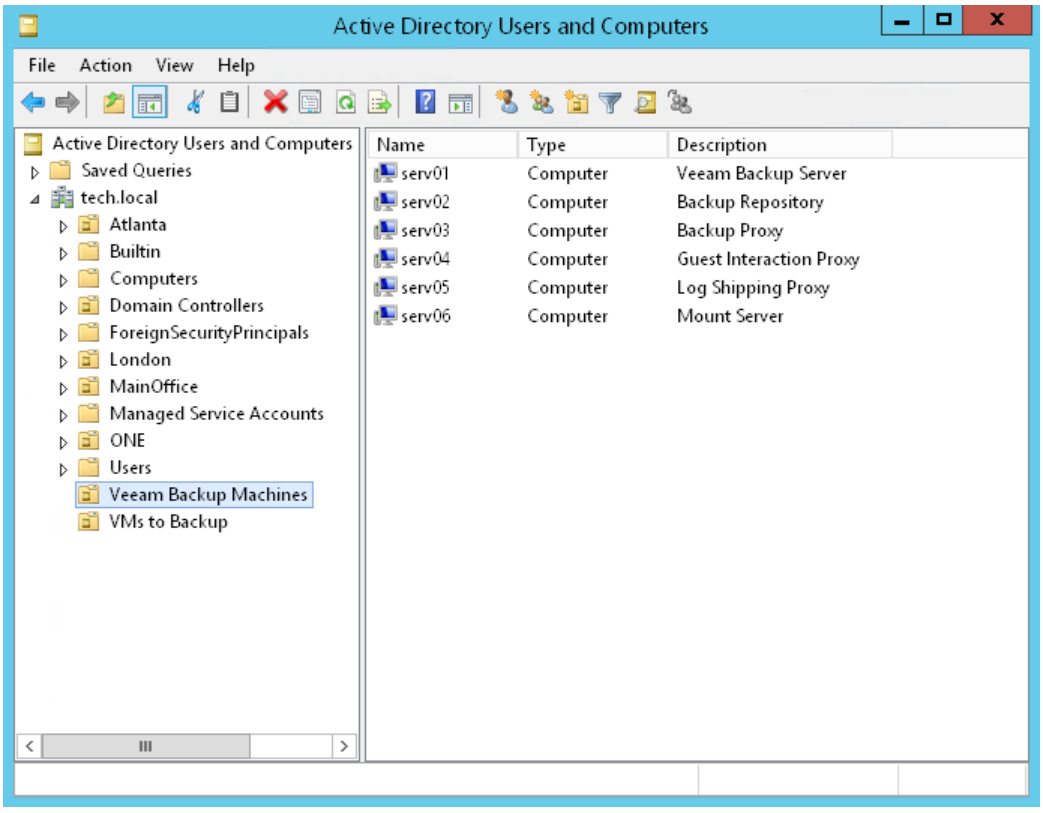
To back up or replicate VMware vSphere VMs where Kerberos is used, you must make sure that NTLM traffic is allowed in Veeam backup infrastructure machines. To do this, you must configure Active Directory group policies as shown below or in a similar way.

Configuring Active Directory Group Policies

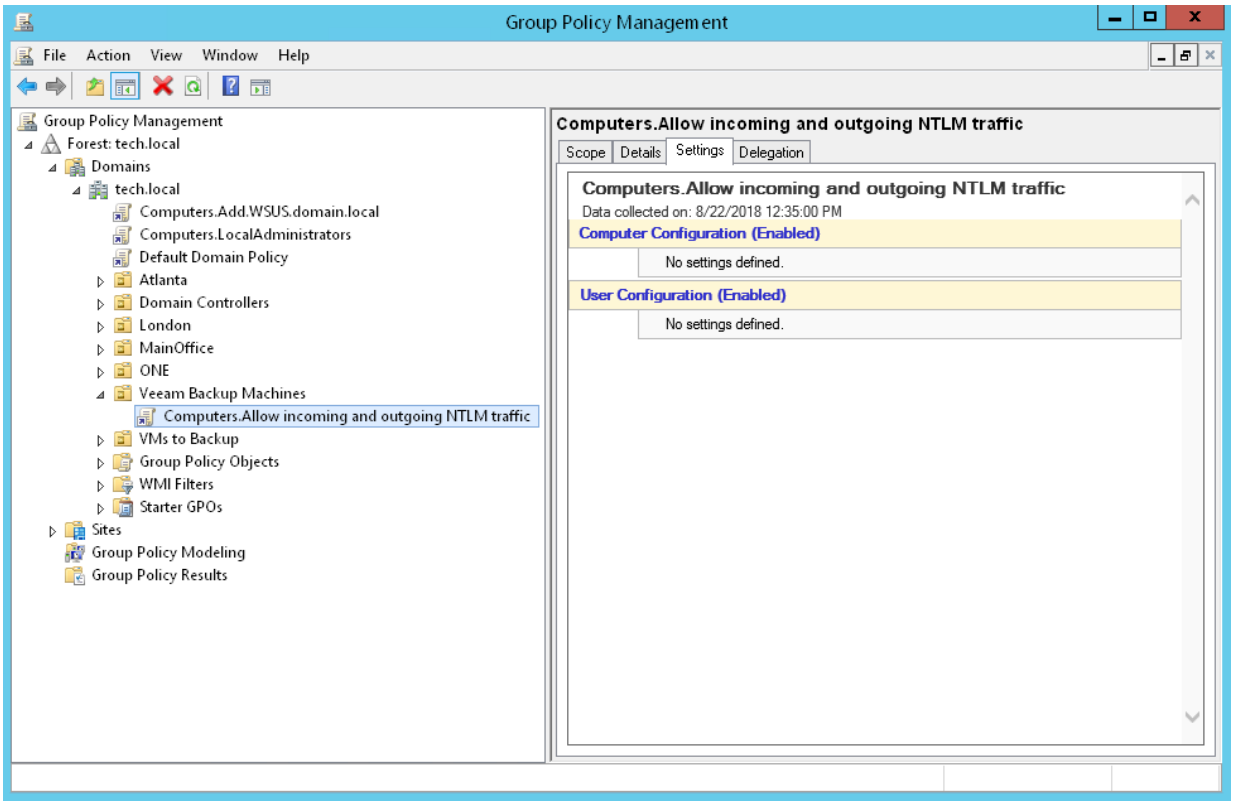
If you want to back up or replicate VMs where Kerberos protocol is used, you must make sure that NTLM traffic is allowed in the Veeam backup infrastructure machines. You can add all Veeam infrastructure servers to a separate Active Directory organizational unit and create a GPO that allows NTLM traffic for this unit.

To allow NTLM traffic in Veeam infrastructure servers, do the following:

1. On the domain controller server or management workstation, open the **Active Directory Users and Computers MMC snap-in**.
2. Create a new Active Directory organizational unit and move all Veeam infrastructure servers to the organizational unit.

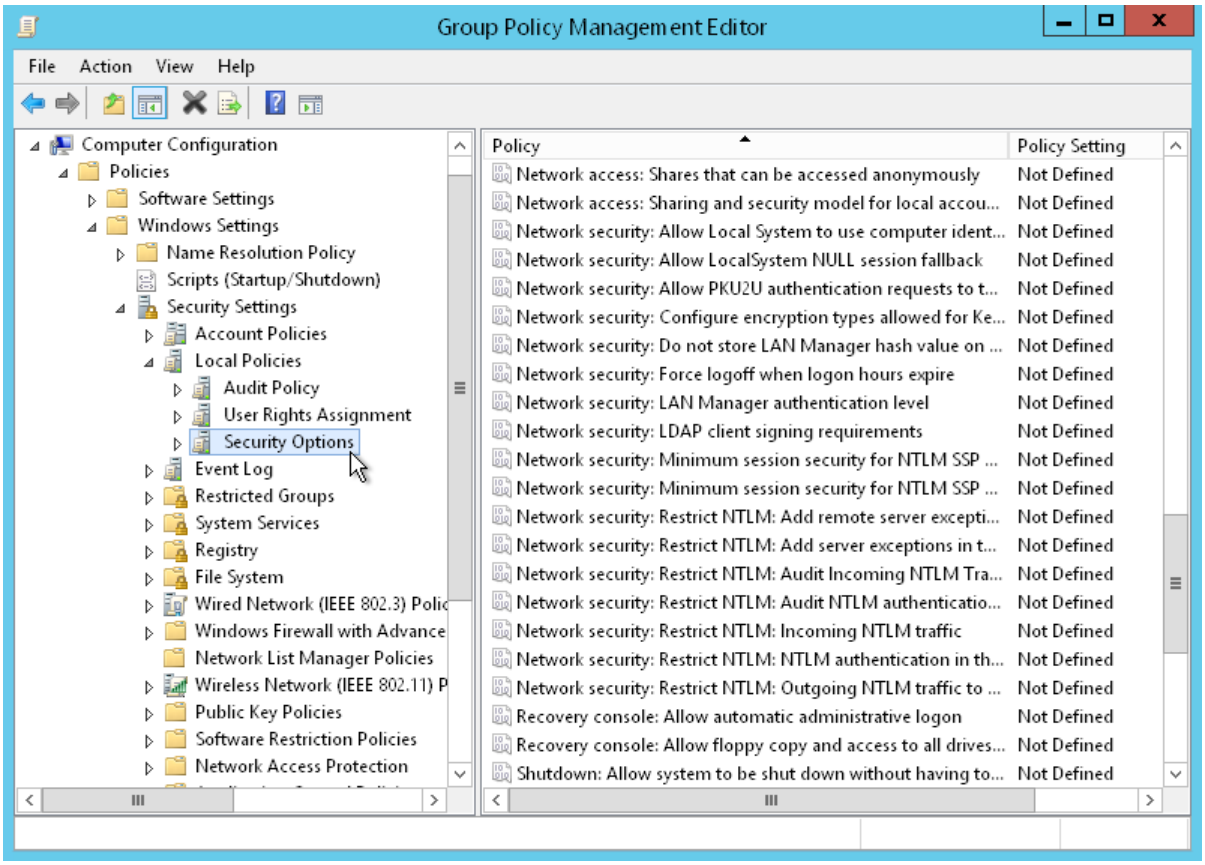


3. Open **Group Policy Management** and create a new GPO for the organizational unit with Veeam infrastructure servers.

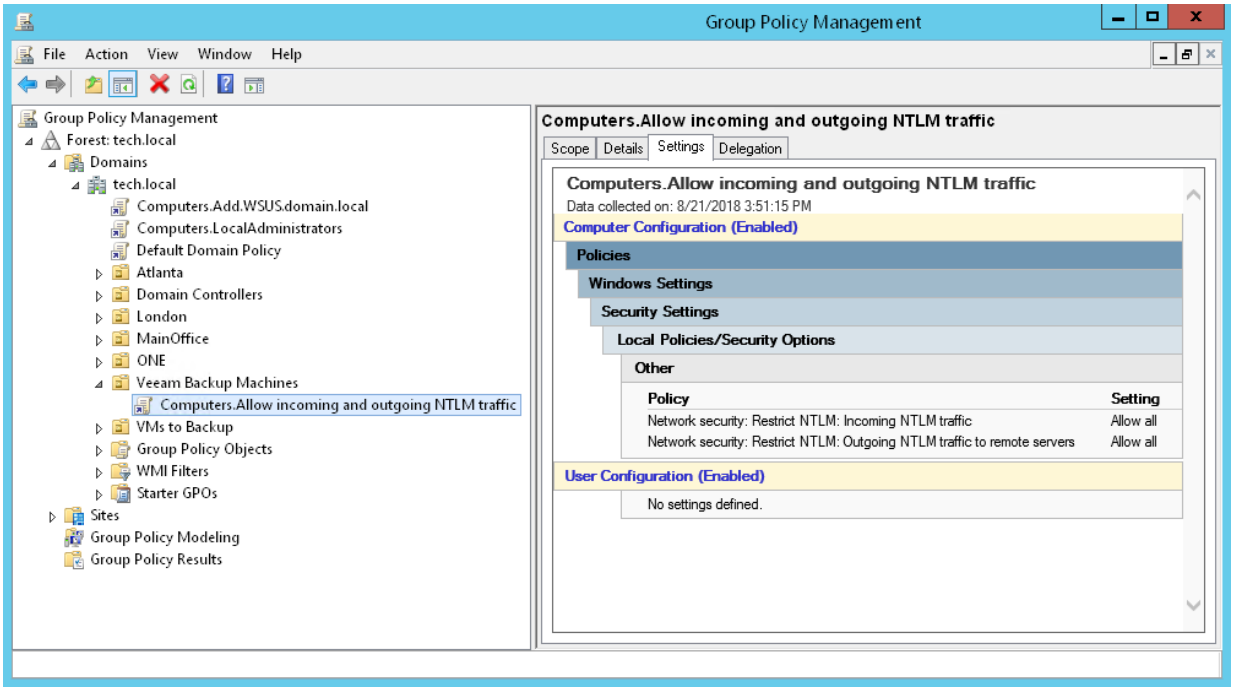


4. Right-click the created GPO and select **Edit**.

5. In the infrastructure tree of the **Group Policy Management Editor** interface, go to *Policies/Windows Settings/Security Settings/Local Policies/Security Options*.



6. In the **Security Options** folder, go to properties of the following two policies and change the policy setting to *Allow all*:
   - Network Security: Restrict NTLM: Incoming NTLM traffic
   - Network Security: Restrict NTLM: Outgoing traffic to remote servers



After you configure group policies for NTLM traffic, Veeam backup infrastructure servers will be able to authenticate to each other using NTLM, while the servers will use Kerberos to authenticate to guest OS of VMs.

Licensing

To work with Veeam Backup & Replication, you must obtain a license key and install it on the backup server. If you do not install the license key, the product will operate in the Veeam Backup & Replication Community (free) Edition. For more information, see Veeam Backup & Replication Community Edition.

Veeam licenses Veeam Backup & Replication in two ways: per instance and per socket.

You can use both instance and socket licenses. For more information, see Merging Licenses.

Instance Licensing

Veeam Backup & Replication can be licensed by the number of instances. Instances are units (or tokens) that you can use to protect your virtual, physical or cloud-based workloads. For more information, see Veeam Licensing Policy.

You must obtain a license with the total number of instances for workloads that you plan to protect in Veeam Backup & Replication.

Workloads that have been processed in the past 31 days are considered protected. Every protected workload consumes instances from the license scope. The number of instances that a workload requires depends on the workload type and product edition.

This licensing model allows you to obtain a license with a certain number of instances without knowing in advance what types of workloads you plan to protect. When a need arises, you can revoke instances from a protected workload, and reuse them to protect other workloads regardless of the workload type.

Veeam Backup & Replication keeps track of instances consumed by protected workloads. If the number of consumed instances exceeds the license limit, Veeam Backup & Replication displays a warning when you open the Veeam Backup & Replication console. For more information, see Exceeding License Limit.

Mind the following:

- VM templates are regarded as protected VMs and consume license instances.
- VMs and NAS shares processed with backup copy and tape jobs are not regarded as protected VMs and NAS shares and do not consume license instances. These types of jobs provide an additional protection level for VMs and NAS shares that are already protected with backup jobs.
- VMs processed by snapshot-only jobs are regarded as protected VMs and consume license instances. Veeam Backup & Replication will revoke instances from these VMs if you re-add a storage array to the backup infrastructure.
- For more information on how Veeam Backup & Replication calculates license instances to consume when protecting NAS file shares, see License Instance Calculation for NAS Backup.

Socket Licensing

With the socket licensing model, Veeam Backup & Replication is licensed by the number of CPU sockets on protected hosts. For more information, see Veeam Licensing Policy.

A license is required for every occupied motherboard socket as reported by the hypervisor API.

License is required only for source hosts — hosts on which VMs that you back up or replicate reside. Target hosts (for replication and migration jobs) do not need to be licensed.

NOTE

If you use a socket license that was obtained for an earlier version of Veeam Backup & Replication, Veeam Software adds up to 6 instances free of charge to your license scope. You can use these instances to protect any type of supported workloads except VMware and Hyper-V VMs — they are covered by the licensed CPU sockets on virtualization hosts.

If the number of licensed sockets is less than 6, you can use the number of instances that equals the number of licensed sockets. For example, if the number of licensed sockets is 5, you can use 5 instances. If the number of licensed sockets is 100, you can use 6 instances.

Types of Licenses

Veeam Software offers the following types of licenses for Veeam Backup & Replication:

Paid Licenses

- **Subscription license** — license that expires at the end of the subscription term. The Subscription license term is normally 1-3 years from the date of license issue.
- **Perpetual license** — permanent license. The support and maintenance period included with the license is specified in months or years. Typically, one year of basic support and maintenance is included with the Perpetual license.
- **Rental license** — license with the license expiration date set according to the chosen rental program (normally 1-12 months from the date of license issue). The Rental license can be automatically updated upon expiration.

  Rental licenses are provided to Veeam Cloud & Service Provides (VCSPs) only. For more information, see the Rental License section in the Veeam Cloud Connect Guide.

Free Licenses

- **Evaluation license** — license used for product evaluation. The Evaluation license is valid for 30 days from the moment of product download.
- **NFR license** — license used for product demonstration, training and education. The person to whom the license is provided agrees that the license is not for resell or commercial use.
- **Promo license** — license that grants additional instances. You can install it only on top of an existing Perpetual or Subscription license (primary license). The primary license can have any units (only sockets, only instances, or both). Number of additional instances and duration of promo period are decided by a sales representative.

Instance Consumption for NAS Backup Jobs

For file share backup, there are the following peculiarities in calculating the number of license instances to consume:

- Veeam Backup & Replication rounds the protected amount of data for each file share down to 500 GB.
- One license instance covers 500 GB of the protected amount of data.
- If different file shares are protected by different backup jobs, Veeam Backup & Replication rounds the protected amount to 500 GB separately for each file share and calculates the number of instances required to protect each of them. After that, it sums up the total number of license instances to consume for NAS backup support.
- Veeam Backup & Replication calculates the protected amount of data for each NAS during every run of the file share backup job that protects file shares on this NAS and keeps the result for 30 days. To calculate the number of license instances to consume for the file share protection, Veeam Backup & Replication takes the largest protected amount of data on file shares within the last 30 days.

  If the size of the file share reduces and does not increase or the file share is removed from the file share backup job, after 30 days Veeam Backup & Replication recalculates the protected amount of data and automatically revokes the excessively consumed license instances.You can manually revoke the licenses without waiting for 30 days, as described in Revoking License. During the next NAS backup job run, Veeam Backup & Replication will recalculate the license instance consumption as of the current date.
- If the same file share is protected by more than one backup job, to calculate the size of the protected amount of data Veeam Backup & Replication first sums the size of the file share data protected by all the jobs. After that, it rounds the overall protected amount of data down to 500 GB and calculates the number of license instances to consume for NAS backup support.
- If a file share backup job protects several file shares residing on the same NAS device (the same share root or NAS filer), the approach to calculating license consumption depends on how the file shares are added to the infrastructure:
  - If you have added the whole share root (`\\root\`) to the infrastructure and the file share backup job protects shares `\\root\share1` and `\\root\share2`, Veeam Backup & Replication first sums the protected amount of both file shares, rounds it down to 500 GB, and then calculates the number of license instances to consume for NAS backup support.
  - If you have separately added shares `\\root\share1` and `\\root\share2` to the infrastructure and the file share backup job protects both of them, Veeam Backup & Replication first rounds the protected amount for each file share down to 500 GB, separately calculates the number of license instances to consume for each file share, and then sums the license instances to calculate the total number of license instances to consume for NAS backup support.

Examples

Case 1

If the protected amount of data for the file share is 499 GB or less, Veeam Backup & Replication rounds it down to 0 GB. In this case, protection of this file share will not consume license instances.

If the protected amount of data for the file share is 590 GB or 990 GB, Veeam Backup & Replication rounds it down to 500 GB. In this case, protection of this file share will consume 1 license instance.

Case 2

If the protected amount of data for the file share is 1000 GB after rounding, Veeam Backup & Replication divides this amount by 500 GB to calculate the number of instances to consume:

1000 GB / 500 GB = 2 license instances

Case 3

You have 2 file shares File Share 1 (990 GB) and File Share 2 (890 GB) each protected by a separate file share backup job. In this case, Veeam Backup & Replication rounds the protected amount of each file share down to 500 GB, calculates the number of license instances required to protect each of the file shares. After that, it sums the calculated number of license instances required to protect the shares:

990 GB ~ 500 GB = 1 license instance

890 GB ~ 500 GB = 1 license instance

1 + 1 = 2 — protection of 2 file shares with separate file share backup jobs consumes 2 license instances.

Case 4

You have a file share with the protected amount of data 1490 GB. Veeam Backup & Replication runs the file share backup job and after rounding down the amount, it calculates that it will consume 2 license instances for protecting this file share:

1490 GB ~ 1000 GB = 2 license instances

Two days later, the size of the file share increases to 1510 GB. Veeam Backup & Replication runs the file share backup job and recalculates the number of license instances to consume based on the increased size of the NAS share:

1510 GB ~ 1500 GB = 3 license instances

Two days later, the size of the file share decreases back to 1490 GB and does not increase any more. Although the protected amount of data decreases, for the next 30 days Veeam Backup & Replication uses value 1510 GB as a basis to calculate the consumption of license instances.

30 days later, Veeam Backup & Replication runs the file share backup job and recalculates the number of instances to consume taking into account that the largest protected amount of data within the last 30 days is 1490 GB. After that, protection of the file share starts consuming 2 instances again.

Case 5

You have a file share (100 GB) protected by 2 file share backup jobs (1 and 2). During the 1st run of job 1, Veeam Backup & Replication rounds the protected amount down to 500 GB (0 GB) and calculates that the protection of this file share with job 1 does not consume license instances. After that the file share size increases to 270 GB. During the 1st run of job 2 which is scheduled after the increase of the file share, Veeam Backup & Replication calculates amount of data protected by all file share backup jobs protecting this file share:

100 GB (run 1 of job 1) + 270 GB (run 1 of job 2) = 370 GB

Veeam Backup & Replication rounds the protected amount down to 500 GB (0 GB) and calculates that the protection of this file share with backup job 1 and 2 does not consume license instances. After that the file share size remains 270 GB. During the 2nd run of job 1, Veeam Backup & Replication calculates amount of data protected by all file share backup jobs protecting this file share:

270 GB (run 2 of job 1) + 270 GB (run 1 of job 2) = 540 GB

Veeam Backup & Replication rounds the protected amount down to 500 GB (500 GB) and calculates that the protection of this file share with backup job 1 and 2 now consumes 1 license instance.

Case 6

You have a single NAS device with 2 NFS file shares residing on it: `\\root\share1` (490 GB) and `\\root\share2` (600 GB). You have added the root server folder (`\\root\`) of this NAS device as an NFS file share to the infrastructure. The file shares `\\root\share1` and `\\root\share2` are added to one file share backup job. In this case, Veeam Backup & Replication sums the protected amount of both file shares, rounds it down to 500 GB, and then calculates the number of license instances:

490 GB + 600 GB = 1090 GB ~ 1000 GB = 2 — protection of 2 file shares in this case consumes 2 license instances.

Case 7

You have a single NAS device, but you have added 2 of its shared folders `\\root\share1` (490 GB) and `\\root\share2` (600 GB) as separate file shares to the inventory. The file shares `\\root\share1` and `\\root\share2` are added to one file share backup job. In this case, Veeam Backup & Replication first rounds the protected amount for each file share down to 500 GB, separately calculates the number of license instances to consume for each file share, and then sums the license instances:

490 GB ~ 0 GB = 0 license instances

600 GB ~ 500 GB = 1 license instance

0 + 1 = 1 — protection of 2 file shares in this case consumes 1 license instance.

Obtaining and Renewing License

You can obtain an Evaluation or paid license for the product when you download the product from the Veeam website.

Obtaining Paid License

To obtain a paid license, refer to the Veeam Backup & Replication Pricing page.

To obtain a Perpetual Instance license, find a reseller.

Obtaining Evaluation License

To obtain an Evaluation license:

1. Sign in to veeam.com.
2. On the Download Veeam products page, click the product link.
3. In the **Get trial key** section, click the **Request Trial Key** link to download the Evaluation license.

Renewing License

To renew your maintenance plan, contact Veeam Renewals Team at renewals@veeam.com.

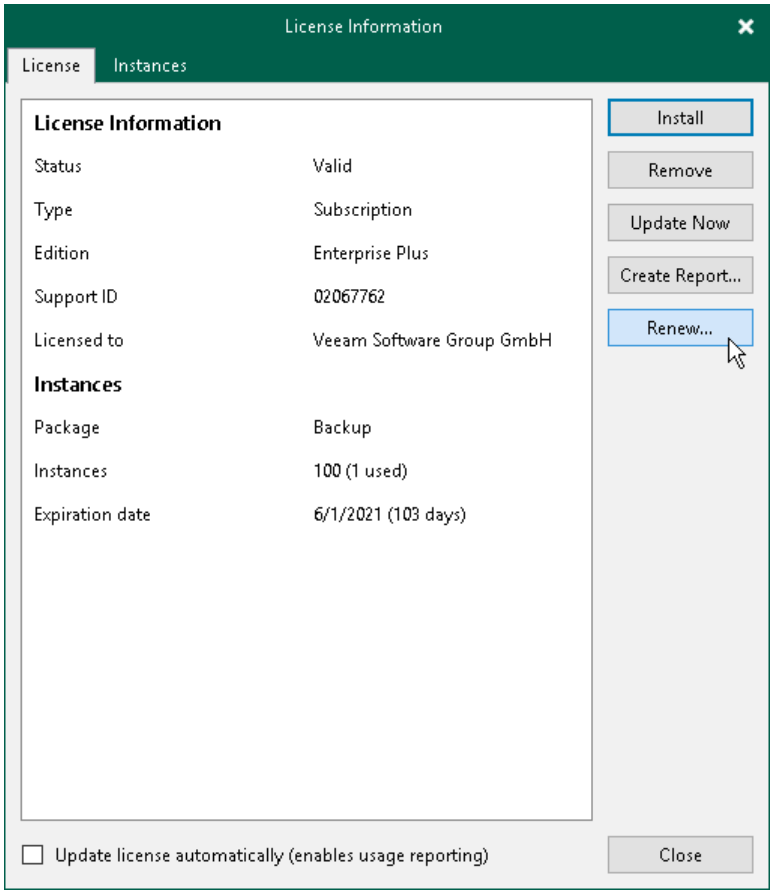
If you have a Perpetual or Subscription license, you can also renew your license contract online.

To renew the license:

1. From the main menu, select **License**.
2. In the **License Information** window, click **Renew**.

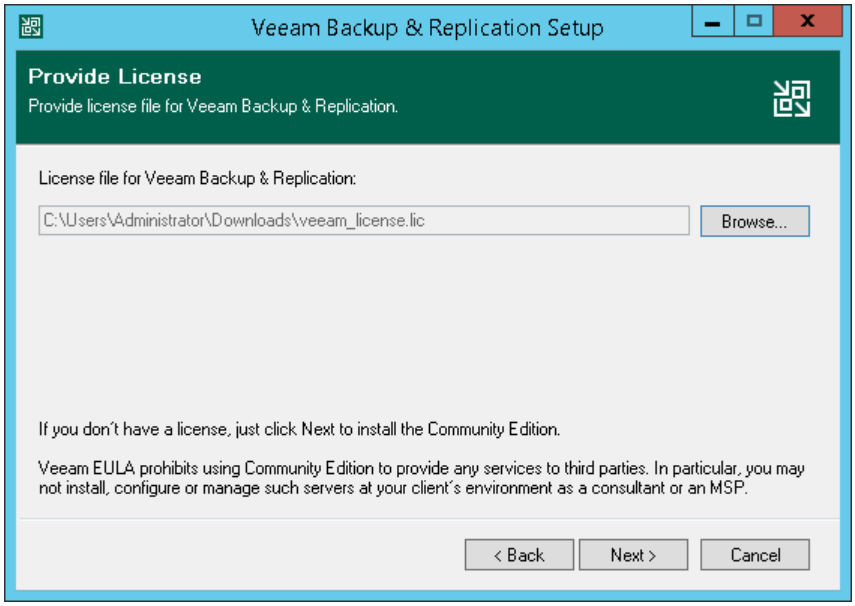
Veeam Backup & Replication will forward you to the Renewals page of Veeam website, where you can select your new maintenance plan. When your contract is renewed, you have to update your license.

Keep in mind that the Renew option is subject to restrictions. If online renewal is not possible for your account, you will be redirected to the Renewal Request page. There you will be able to submit a request for Veeam Renewals Team.



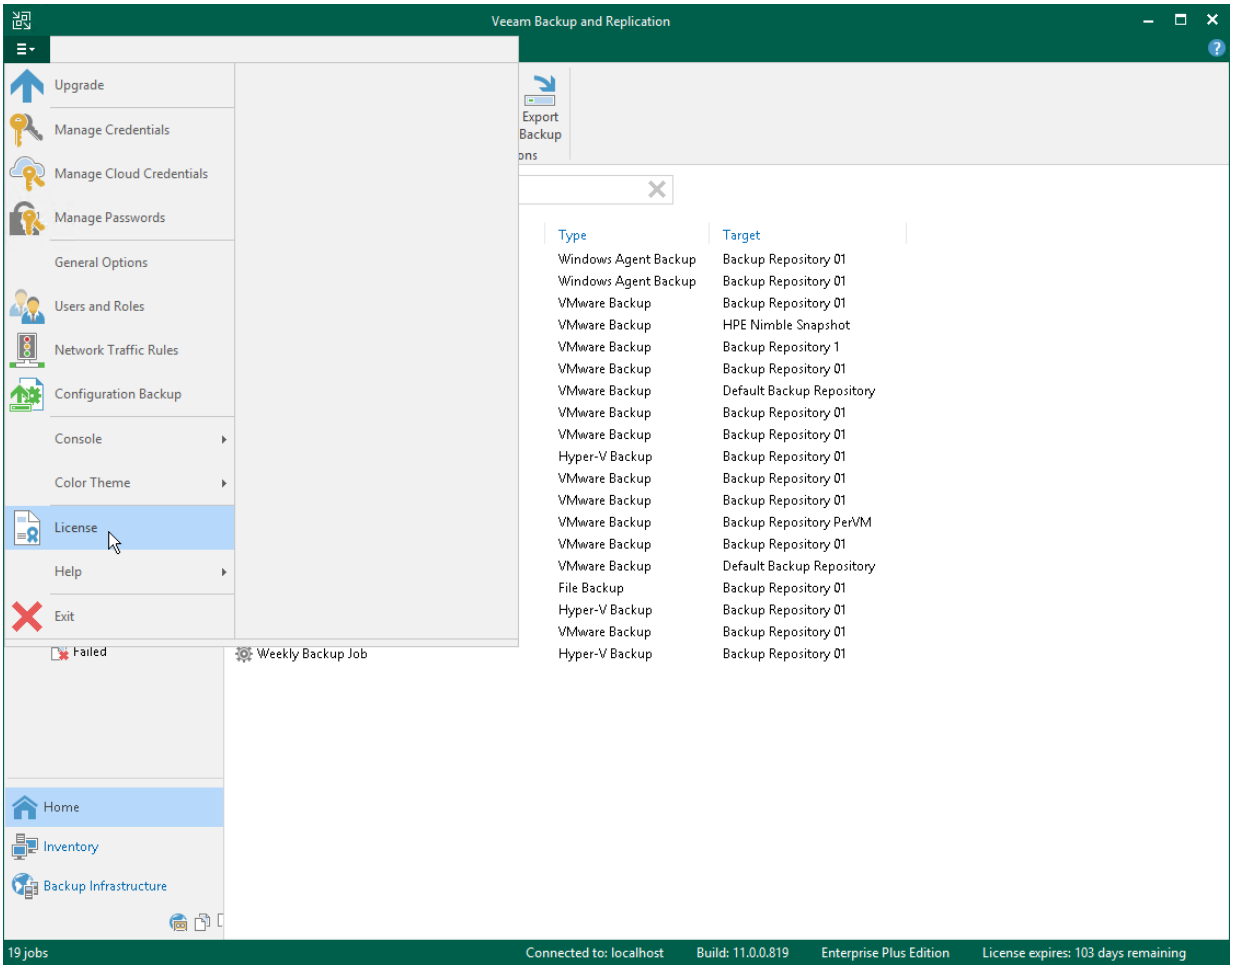
Installing License

When you install Veeam Backup & Replication, you are asked to specify a path to the license file. If you do not specify a path to the license file, Veeam Backup & Replication will run in the Veeam Backup & Replication Community (free) Edition. For more information, see Veeam Backup & Replication Community Edition.

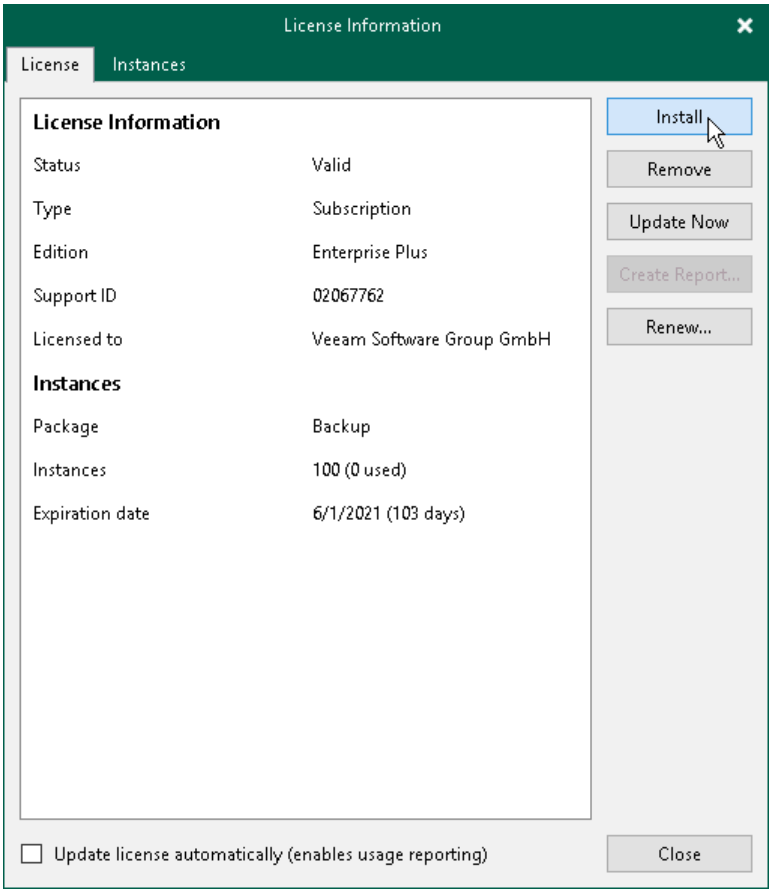


You can install or change the license after product installation:

1. From the main menu, select **License**.



2. In the **License Information** window, click **Install**.



3. Browse to the LIC file and click **Open**.

Licenses in Veeam Backup Enterprise Manager

If backup servers are connected to Veeam Backup Enterprise Manager, Veeam Backup Enterprise Manager collects information about all licenses installed on backup servers. When Veeam Backup Enterprise Manager replicates databases from backup servers, it also synchronizes license data: checks if the license installed on the backup server coincides with the license installed on the Veeam Backup Enterprise Manager server. If the licenses do not coincide, the license on the backup server is automatically replaced with the license installed on the Veeam Backup Enterprise Manager server.

For information on Veeam Backup Enterprise Manager license, see the Licensing section of the Veeam Backup Enterprise Manager Guide.

Merging Licenses

Merging licenses is an option for customers who have a Perpetual Socket license. If you have a Perpetual Socket license, and want to also protect, for example, your cloud or physical workloads, or work with Veeam plug-ins, you can obtain an instance license and merge it with the socket license.

> **TIP**
>
> Merging licenses is a built-in mechanism that allows you to merge two different license types (socket and instance). If you want to merge several licenses of the same type (socket and socket, instance and instance) to obtain a single license key file, you can use the merge tool in the Customer Portal. For more information on merging licenses in the Customer Portal, see this Veeam KB article. For more information on merge rules and exceptions for the Customer Portal merging process, see the License Key Merge section of the Veeam Licensing Policy.

Under the merged license, the following workloads are processed per socket:

- VMs
- Veeam Plug-ins if the application servers are added to Veeam Backup & Replication
- Veeam Agents if the machines are added to Veeam Backup & Replication.

Other workloads are processed per instance.

License Types Available for Merging

You can merge licenses of the following types:

- Perpetual Socket license and Subscription Instance license
- Perpetual Socket license and Perpetual Instance license.

License Packages Available for Merging

You can merge licenses of the following packages:

Type of license	Essentials Socket	Backup Socket	Suite Socket	ONE Socket
Essentials Instance	×	×	×	×
Backup Instance	✓	✓	✓	×
Suite Instance	✓	✓	✓	×

Type of license	Essentials Socket	Backup Socket	Suite Socket	ONE Socket
ONE Instance	×	×	×	×

Support Terms in Merged Licenses

If the licenses that you want to merge have different support expiration date, the merged license will take the support that expires first.

Merging Licenses

Before you merge licenses, check the following prerequisites:

- The license type and package allow merging.
- The company names are identical in both licenses. Company name check is case sensitive.

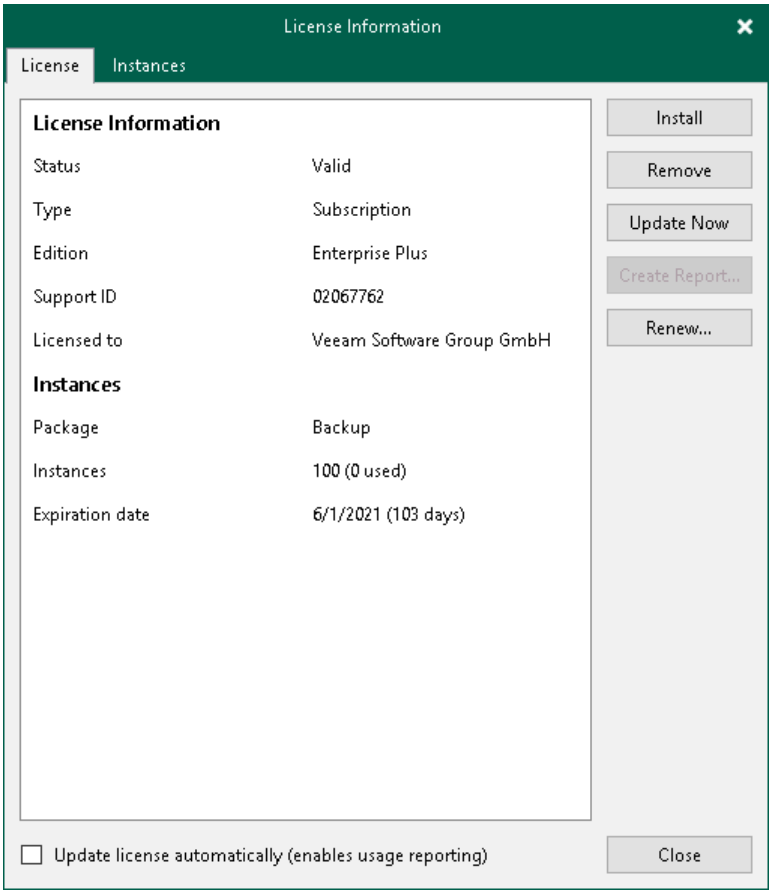
To merge licenses, install a new license over the already installed license. For more information, see Installing License.

> **IMPORTANT**
>
> If the license types do not allow merging, the newly installed license will replace the previous license.

Viewing License Information

You can view details of the installed license in the **License Information** window.

To open the **License Information** window, from the main menu select **License**.



The following details are available for the current license:

- **Status** – license status (*Valid*, *Invalid*, *Expired*, *Not Installed*, *Warning*, *Error*).
- **Type** – license type (*Perpetual*, *Subscription*, *Rental*, *Evaluation*, *NFR*, *Free*).
- **Edition** – license edition (*Community*, *Standard*, *Enterprise*, *Enterprise Plus*).
- **Support ID** – support ID required for contacting Veeam Support.
- **Licensed to** – name of a person or organization to which the license was issued.
- **Cloud Connect Provider** – shows if you can use Veeam Backup & Replication to offer cloud repository as a service and disaster recovery as a service to your customers *(Enterprise, Yes, No)*. For more information on Veeam Cloud Connect, see Veeam Cloud Connect Guide.
- **Package**[1] – Veeam license pack: *Essentials*, *Backup*, *Suite*, *ONE*.
- **Sockets** – number of sockets that you can use to protect workloads.
- **Instances** – number of instances that you can use to protect workloads.
- **Promo instances** – number of additional instances granted by the Promo license.
- **Expiration date** – date when the license expires.

- **Support expiration date** – date when the support expires. Valid for Perpetual Socket and Perpetual Instance licenses.
- **Promo expiration date** – date when the Promo license expires.
- **Total instances including promo** – number of all available instances, regular and promo added up.
- **Capacity** – protected front end capacity (in TiB) for NAS backup.

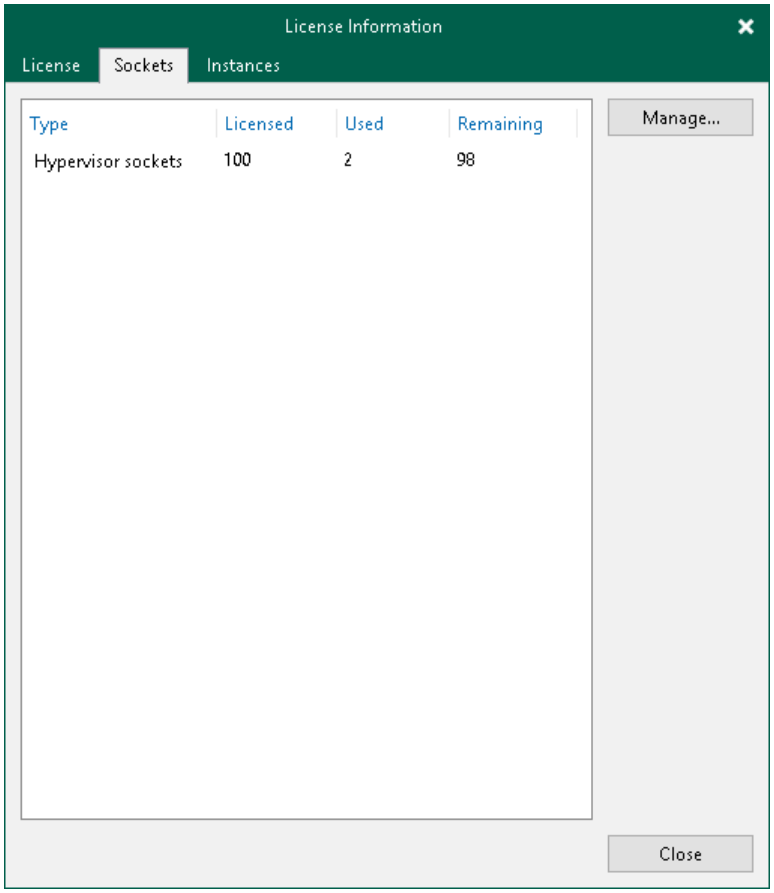
[1] Starting from version 11, the Starter pack license is replaced with Essentials pack. If you were using Starter pack license, you will need either to agree to update the license during the VBR upgrade to version 11 or to download the Essential pack license file from Veeam Customer portal and install it manually.

To enable automatic license update, select the **Update license automatically (enabled usage reporting)** check box. For more information, see Updating License Automatically.

Viewing Information on Sockets

With socket licenses, Veeam Backup & Replication applies a license to the virtualization host on which the processed VMs reside.

To view to which objects the license is currently applied, open the **Sockets** tab.

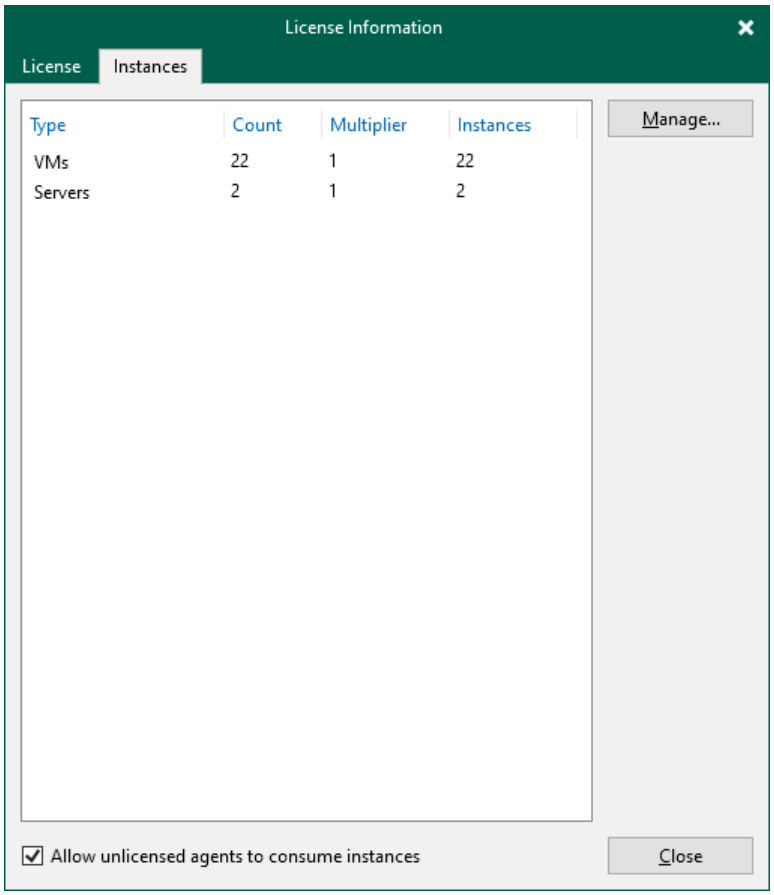


Viewing Information on Instances

With instance licenses, Veeam Backup & Replication applies a license to a protected workload. The number of license instances that a protected workload consumes depends on the workload type and product edition. For details, see Veeam Licensing Policy.

To view to which objects the license is currently applied, open the **Instances** tab.

By default, Veeam Backup & Replication allows Veeam Agents to connect to the Veeam backup server and consume instances in the license. If you do not want Veeam Agents to consume instances, clear the **Allow unlicensed agents to consume instances** check box. For more information on Veeam Agents licensing, see Veeam Agent Management Guide.



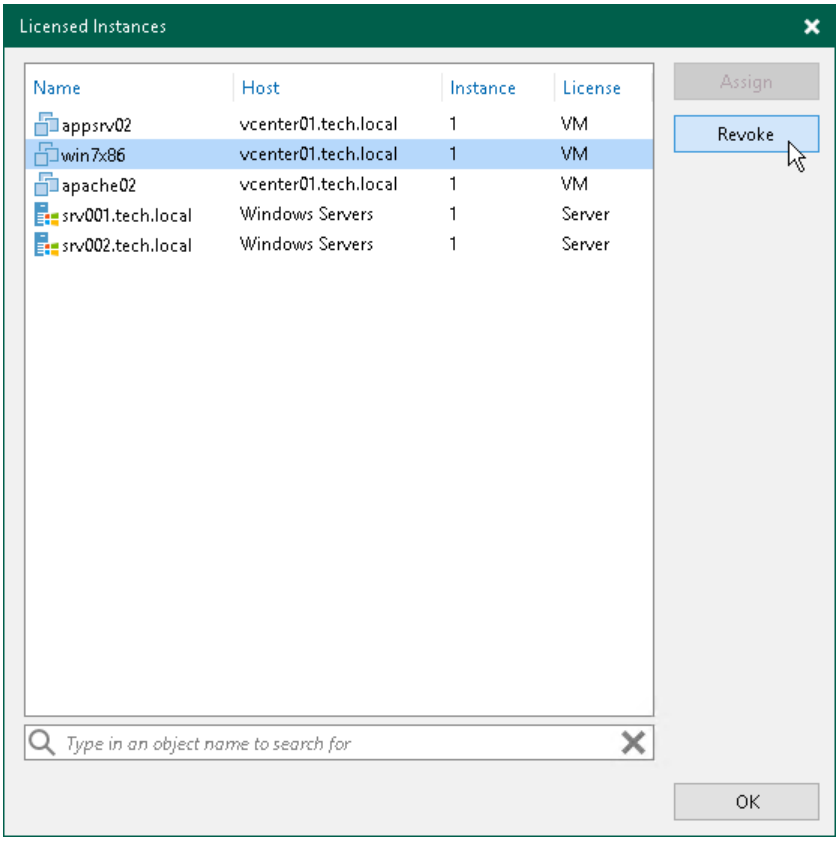
Revoking License

You can revoke licenses from protected workloads or licensed hosts, and re-apply them to other objects that you plan to protect. License revoking can be helpful, for example, if a licensed host goes out of service or you do not want to protect some workloads anymore.

> **NOTE**
>
> If you manually revoke license instances allocated for a NAS share source, the next run of the NAS backup job, which protects this file share, will trigger the recalculation of the NAS share protected size and reallocation of license instances that Veeam Backup & Replication will consume.

To revoke a license, do the following:

1. From the main menu, select **License**.
2. In the **License Information** window:
   - For protected workloads, open the **Instances** tab and click **Manage**.
   - For licensed hosts, open the **Sockets** tab and click **Manage**.
3. In the displayed window, select a protected workload or a licensed host and click **Revoke**.
   Veeam Backup & Replication will revoke the license from the selected object, and the license will be freed for other objects in the backup infrastructure.



Removing License

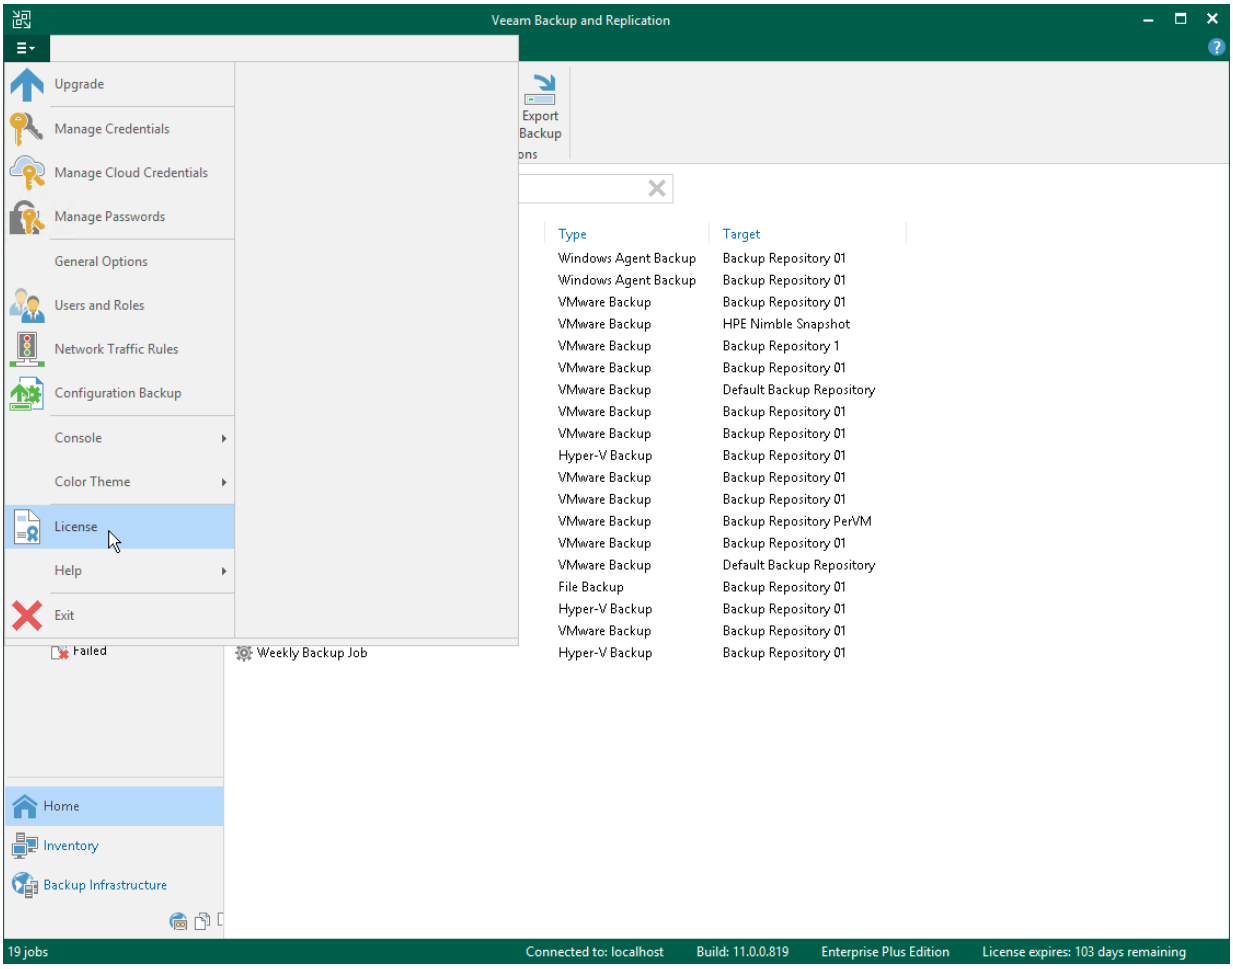
You can remove the installed license. When you remove a license, Veeam Backup & Replication will switch to the Veeam Backup & Replication Community Edition. For more information, see Veeam Backup & Replication Community Edition.

You can also remove a part of merged license. If you do so, Veeam Backup & Replication will operate under the other part of the merged license. For more information, see Merging Licenses.

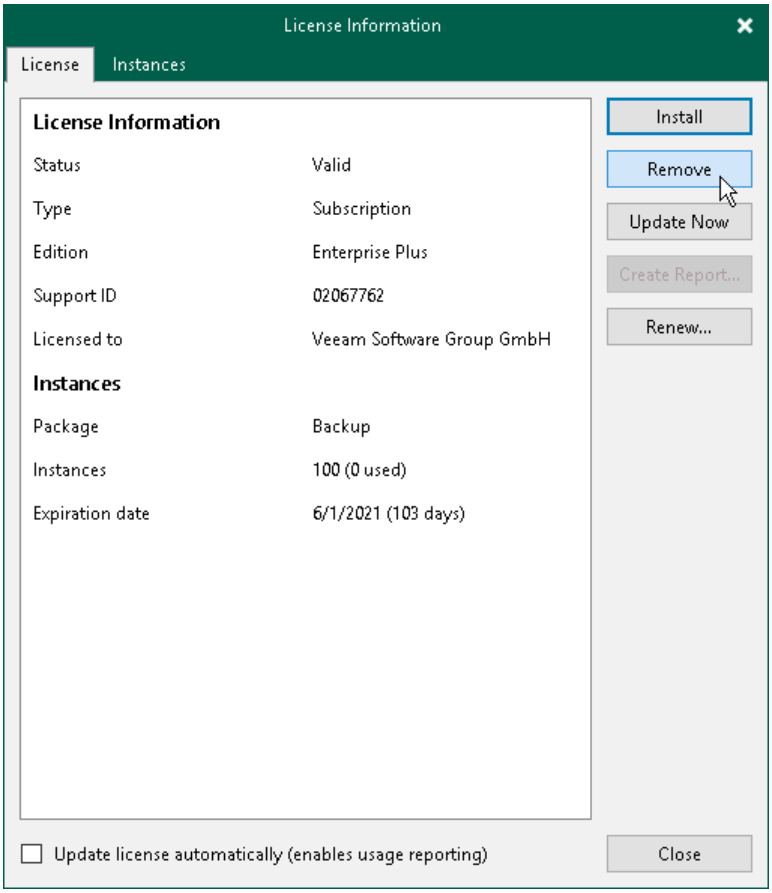
To remove a Promo license, remove the license on top of which it was installed.

To remove a license, do the following:

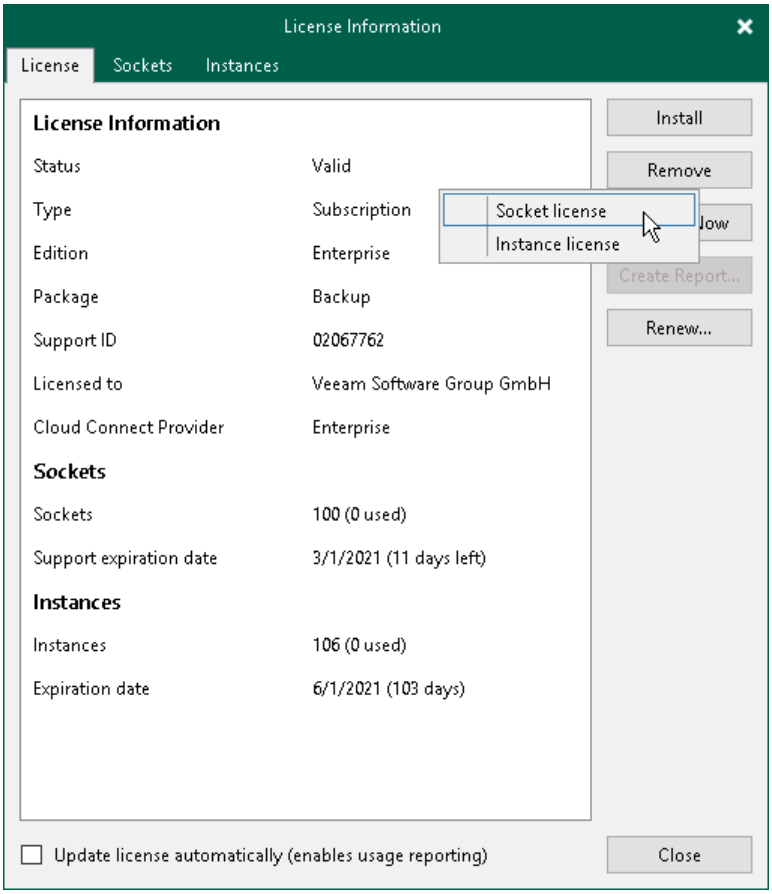
1. From the main menu, select **License**.



2. In the **License Information** window, click **Remove**.



If you have both socket and instance licenses, you will be asked which license you want to remove.



3. Select the license you want to remove and confirm the removal.

Exceeding License Limit

For Veeam Universal Licenses, Veeam Backup & Replication allows you to protect more workloads than covered by the number of instances specified in the license. An increase in the number of protected workloads is allowed throughout the duration of the contract (license key).

The license limit can be exceeded by a number of instances, or a percentage of the total number of instances specified in the license (depends on which number is greater). The exceeding limit varies according to the license type.

License Type	Exceeding Limit	Description
Socket license	Not allowed	Workloads that are exceeding the license limit are not processed.
Veeam Universal License (VUL)	Less than 5 instances (or 5% of the total instance count)	All protected workloads are processed normally, Veeam Backup & Replication does not display a warning message.
	5-10 instances (or 5%-10% of the total instance count)	All protected workloads are processed normally. Once a week when you open the Veeam Backup & Replication console, a warning message is displayed notifying that you are out of compliance with the Veeam Licensing Policy.
	More than 10 instances (or 10% of the total instance count)	Workloads that are exceeding the license limit beyond 10 instances (or 10% of the total instance count) are not processed. Every time you open the Veeam Backup & Replication console, a warning message is displayed notifying that you are out of compliance with the Veeam Licensing Policy.
Rental license	See the Rental License section in the Veeam Cloud Connect Guide.	

For example, you have a Subscription license with 500 instances to protect your workloads. According to the table above, you are allowed to use up to 10 instances or 10% of the total instance count (whichever number is greater) over the license limit. As the number of instances in your license is 500, you are allowed to use additional 50 instances (50 makes *10%* of 500, and 50 is greater than *10*). Consider the following:

- Until the license limit is not exceeded by more than 5% of the total instance count (up to 25 instances), Veeam Backup & Replication processes all protected workloads with no restrictions.
- When the license limit is exceeded by 5%-10% (25 to 50 instances), Veeam Backup & Replication processes protected workloads, and displays a warning message once a week when you open the Veeam Backup & Replication console. In the message, Veeam Backup & Replication provides information on the number of exceeded instances and the number of instances by which the license can be further exceeded.

- If the license limit is exceeded by more than 10% (50 instances and more), Veeam Backup & Replication does not process the workloads exceeding the limit, and displays a warning message every time you open the Veeam Backup & Replication console. In the message, Veeam Backup & Replication provides information on the number of instances by which the license is exceeded.

When the license limit is exceeded, the logs will include the number of instances necessary to finish the job successfully.

> **NOTE**
>
> Mind that capacity-based (per-TB) license limit for NAS backup cannot be exceeded.

License Expiration

When the license expires, Veeam Backup & Replication behaves in the following way depending on the license type:

- Evaluation and NFR licenses: Veeam Backup & Replication will stop processing workloads.
- Paid licenses: Veeam Backup & Replication will switch to the grace period.
- Promo license: Veeam Backup & Replication will remove the granted instances and stop processing workloads for them. Promo license does not have a grace period. Upon expiration of the primary license, the promo license will also expire, regardless of its own expiration date.

Perpetual Socket and Perpetual Instance licenses do not expire. However, such licenses have support expiration date. Veeam Backup & Replication will inform you about the support expiration date.

Grace Period

To ensure a smooth license update and provide sufficient time to install a new license file, Veeam Backup & Replication offers a grace period. Grace period is available for paid licenses.

During the grace period, you can perform all types of data protection and disaster recovery operations. However, Veeam Backup & Replication will inform you about the license expiration when you open the Veeam Backup & Replication console. The license status in the **License Information** window will appear as *Expired*.

You must update your license before the end of the grace period. If you do not update the license, Veeam Backup & Replication stops processing workloads. All existing jobs fail with the *Error* status. However, you will be able to restore machine data from existing backups.

Grace Period Duration

Before the license expires, Veeam Backup & Replication notifies you about soon license expiration.

The number of days for notification and grace period depends on the type of license:

License Type	License Expiration Notification	Grace Period
Subscription	30 days	30 days
Perpetual Instance	14 days before Support expiration date	n/a
Perpetual Socket	14 days before Support expiration date	n/a
Rental	7 days	60 days
Evaluation	30 days	0 days

License Type	License Expiration Notification	Grace Period
NFR	30 days	0 days
Promo	7 days	n/a

Switching to Veeam Backup & Replication *Community Edition*

If you do not want to renew your license, you can switch to the free product version named Veeam Backup & Replication Community (free) Edition. For more information, see Veeam Backup & Replication Community Edition.

To do so, remove your license. For details, see Removing License.

Expiration of Merged Licenses

If you merged licenses with different expiration dates, the merged license will expire on the date that is closer. For example, if you merged a Perpetual license and a Subscription license, the expiration date will be inherited from the Subscription license.

In such case, you can update your Subscription license or continue using the Perpetual license. To continue using the Perpetual license, remove the Subscription license. For more information, see Removing License.

Updating License

To be able to use all data protection and disaster recovery features, you must update your license upon expiry. There are two methods to update the license in Veeam Backup & Replication:

- Update the license manually
- Update the license automatically

> **NOTE**
>
> When updating the license, Veeam Backup & Replication requires internet access to connect to the Veeam License Update Server. If your network is not connected to the internet, instead of update you can download a new license file from my.veeam.com and then install it. For more information on license installation, see Installing License.

Updating License Manually

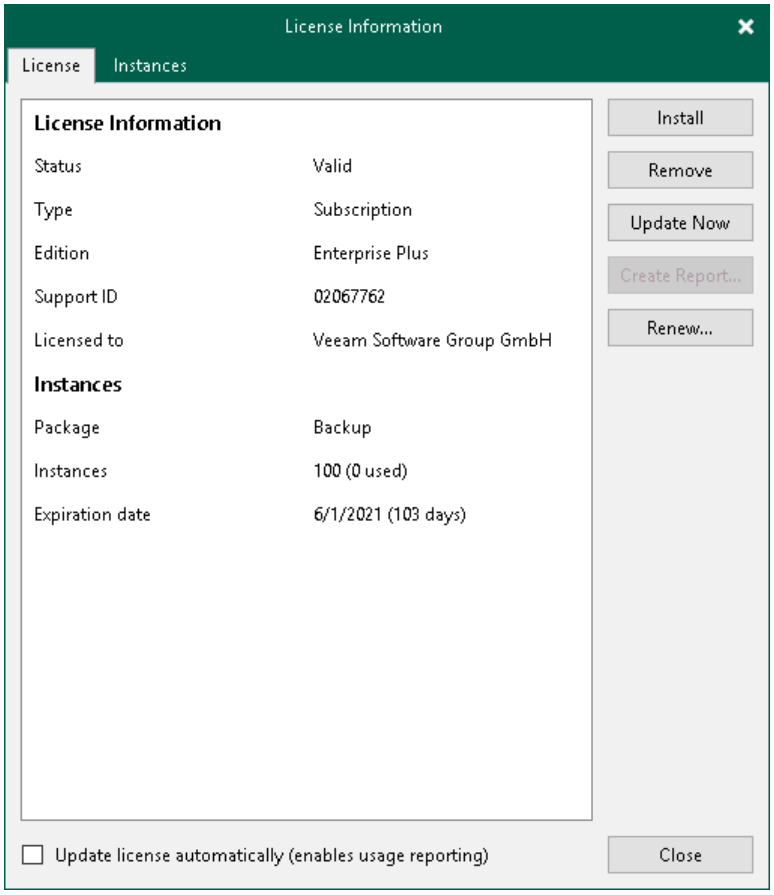
You can update the license manually on demand. When you update the license manually, Veeam Backup & Replication connects to the Veeam License Update Server, downloads a new license from it (if the license is available) and installs it on the backup server.

The new license key differs from the previously installed license key in the license expiration date or support expiration date. If you have obtained a license for a greater number of instances, counters in the new license also display the new number of license instances.

To update the license:

1. From the main menu, select **License**.
2. In the **License Information** window, click **Update Now**.

Statistics on the manual license update process is available under the **System** node in the **History** view. You can double-click the **License key auto-update** job to examine session details for the license update operation.



Manual License Update Results

Manual license update can complete with the following results:

- **Operation is successful**. A new license key is successfully generated, downloaded and installed on the backup server or Veeam Backup Enterprise Manager server.
- **A new license is not required**. The currently installed license key does not need to be updated.
- **The Veeam License Update Server has failed to generate a new license**. Such situation can occur due to some error on the Veeam License Update Server side.

- **Veeam Backup & Replication has received an invalid answer**. Such situation can occur due to connectivity issues between the Veeam License Update Server and Veeam Backup & Replication.
- **Licensing by the contract has been terminated**. In such situation, Veeam Backup & Replication automatically disables automatic license update on the backup server or Veeam Backup Enterprise Manager server.

Updating License Automatically

You can instruct Veeam Backup & Replication to automatically update the license installed on the backup server or Veeam Backup Enterprise Manager server. With automatic license update, you do not need to download and install the license manually each time when you purchase the license extension. If the automatic update option is enabled, Veeam Backup & Replication proactively communicates with the Veeam License Update Server to obtain and install a new license before the current license expires.

Requirements and Limitations for Automatic License Update

- Automatic license update is not available in the Veeam Backup & Replication Community Edition.
- Only licenses that contain a real contract number in the Support ID can be updated with the **Update license key automatically** option.
- If you are managing backup servers with Veeam Backup Enterprise Manager, all license management tasks must be performed in the Veeam Backup Enterprise Manager console. Automatic update settings configured in Veeam Backup Enterprise Manager override automatic update settings configured in Veeam Backup & Replication.

  For example, if the automatic update option is enabled in Veeam Backup Enterprise Manager but disabled in Veeam Backup & Replication, automatic update will be performed anyway. For more information, see the Veeam Backup Enterprise Manager User Guide.

> **NOTE**
>
> Veeam Backup & Replication does not automatically update an existing per-VM or socket license that was obtained for an earlier version of the product to a new instance license.
>
> To overcome this issue, after you upgrade to Veeam Backup & Replication 11, you must obtain in the Veeam Customer Support Portal a new instance license and install it on the backup server manually.

How Automated License Update Works

To update installed licenses automatically, Veeam Backup & Replication performs the following actions:

1. After you enable automatic license update, Veeam Backup & Replication starts sending requests to the Veeam License Update Server on the web (autolk.veeam.com) and checks if a new license key is available. Veeam Backup & Replication sends requests once a week. Communication with the Veeam License Update Server is performed over the HTTPS protocol.
2. Seven days before the expiration date of the current license, Veeam Backup & Replication starts sending requests once a day.
3. When a new license key becomes available, Veeam Backup & Replication automatically downloads it and installs on the backup server or Veeam Backup Enterprise Manager server.

The new license key differs from the previously installed license key in the license expiration date and support expiration date. If you have obtained a license for a greater number of instances, counters in the new license also display the new number of license instances.

Automatic License Update Results

Automatic license update can complete with the following results:

- **Operation is successful**. A new license key is successfully generated, downloaded and installed on the backup server or Veeam Backup Enterprise Manager server.
- **A new license is not required**. The currently installed license key does not need to be updated.
- **The Veeam License Update Server has failed to generate a new license**. Such situation can occur due to some error on the Veeam License Update Server side.
- **Veeam Backup & Replication has received an invalid answer**. Such situation can occur due to connectivity issues between the Veeam License Update Server and Veeam Backup & Replication.
- **Licensing by the contract has been terminated**. In such situation, Veeam Backup & Replication automatically disables automatic license update on the backup server or Veeam Backup Enterprise Manager server.

Automatic Update Retries

If Veeam Backup & Replication fails to update the license, it displays a notification in the session report and sends an email notification to users specified in the global email settings (if global email settings are configured on the backup server). You can resolve the issue, while Veeam Backup & Replication will keep retrying to update the license.

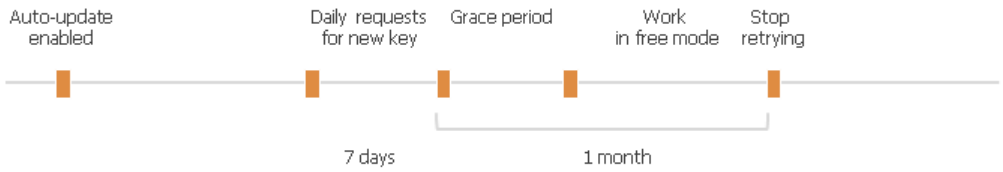
Veeam Backup & Replication retries to update the license key in the following way:

- If Veeam Backup & Replication fails to establish a connection to the Veeam License Update Server, retry takes place every 60 min.
- If Veeam Backup & Replication establishes a connection but you are receiving the "*General license key generation error has occurred*" message, the retry takes place every 24 hours.

The retry period ends one month after the license expiration date or the support expiration date (whichever is earlier). The retry period is equal to the number of days in the month of license expiration. For example, if the license expires in January, the retry period will be 31 day; if the license expires in April, the retry period will be 30 days.

If the retry period is over but the new license has not been installed, the automatic update feature is automatically disabled.

For more information about error cases, see Appendix A. License Update Session Data.

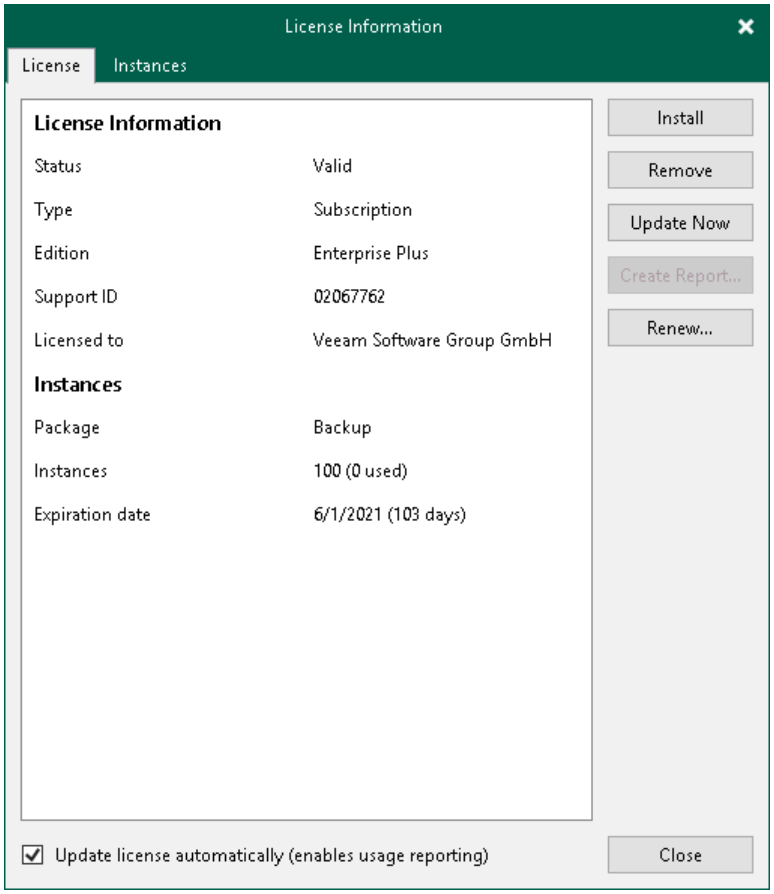


Enabling Automatic License Update

By default, automatic license update is disabled. To enable automatic license update:

1. From the main menu, select **License**.

2. In the **License Information** window, select the **Update license automatically (enabled usage reporting)** check box.



During the installation of a Subscription or Rental license, you will see a pop-up dialog window with a suggestion to enable automatic license update.

Statistics on the automatic license update process is available under the **System** node in the **History** view. You can double-click the **License key auto-update** job to examine session details for the scheduled or ad-hoc automatic license update.

> **NOTE**
>
> [For Rental, Subscription, Perpetual licenses] Enabling license auto update activates Automatic License Usage Reporting.

Automatic License Usage Reporting

When license auto update is enabled for Rental, Subscription, Perpetual licenses, Veeam Backup & Replication performs automatic license usage reporting.

As part of reporting, Veeam Backup & Replication collects statistics on the current license usage and sends it periodically to the Veeam License Update Server. The report provides information about the contract ID, product installation ID, and the maximum number of licensed objects that were managed by Veeam Backup & Replication over the past week (high watermark). The reporting process runs in the background mode, once a week at a random time and day.

The type of reported objects is defined by the product and the installed license. The report can include information about VMs, workstations or servers protected with Veeam backup agents, and so on.

The collected data does not include information on the usage of Veeam Backup & Replication by any individual person identifiable for Veeam, or any data protected by Veeam Backup & Replication.

The collected data allows our back-end system to automatically approve your monthly usage reports as long as they do not deviate from the high watermark value significantly. This helps to keep our report processing costs low, thus allowing us to maintain low rental prices for our solution. Veeam may also use collected data for any other internal business purposes it deems appropriate, including (but not limited to) evaluation, improvement and optimization of Veeam licensing models.

By enabling license auto update you agree with collection, transmission and use of the reporting data. You must not enable license auto update in case you do not agree with such collection, transmission and use.

Getting to Know Veeam Backup & Replication

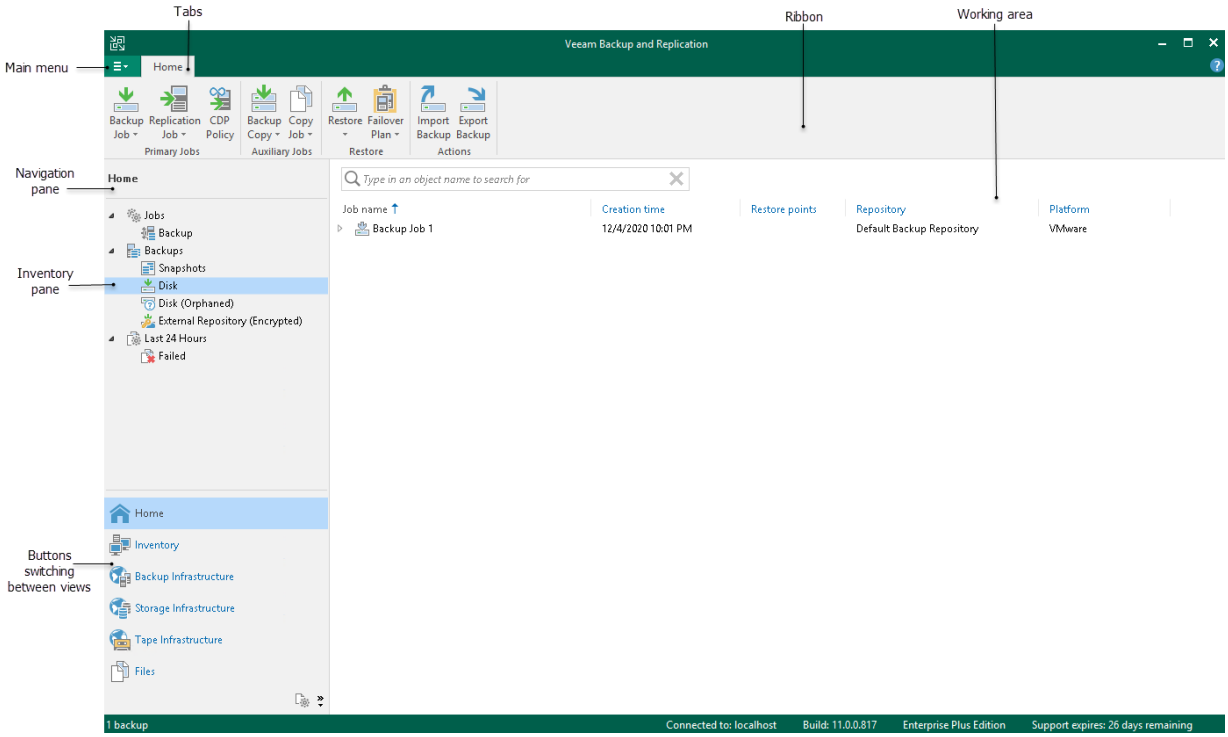
After you install Veeam Backup & Replication, you can get familiar with the product UI, learn about job priorities.

- Veeam Backup & Replication UI
- Job Priorities

Veeam Backup & Replication UI

The user interface of Veeam Backup & Replication is designed to let you quickly find commands that you need and perform data protection and disaster recovery tasks.

- Main Menu
- Navigation Pane
- Ribbon and Tabs
- Views
- Working Area
- Job Filter
- Changing Color Theme



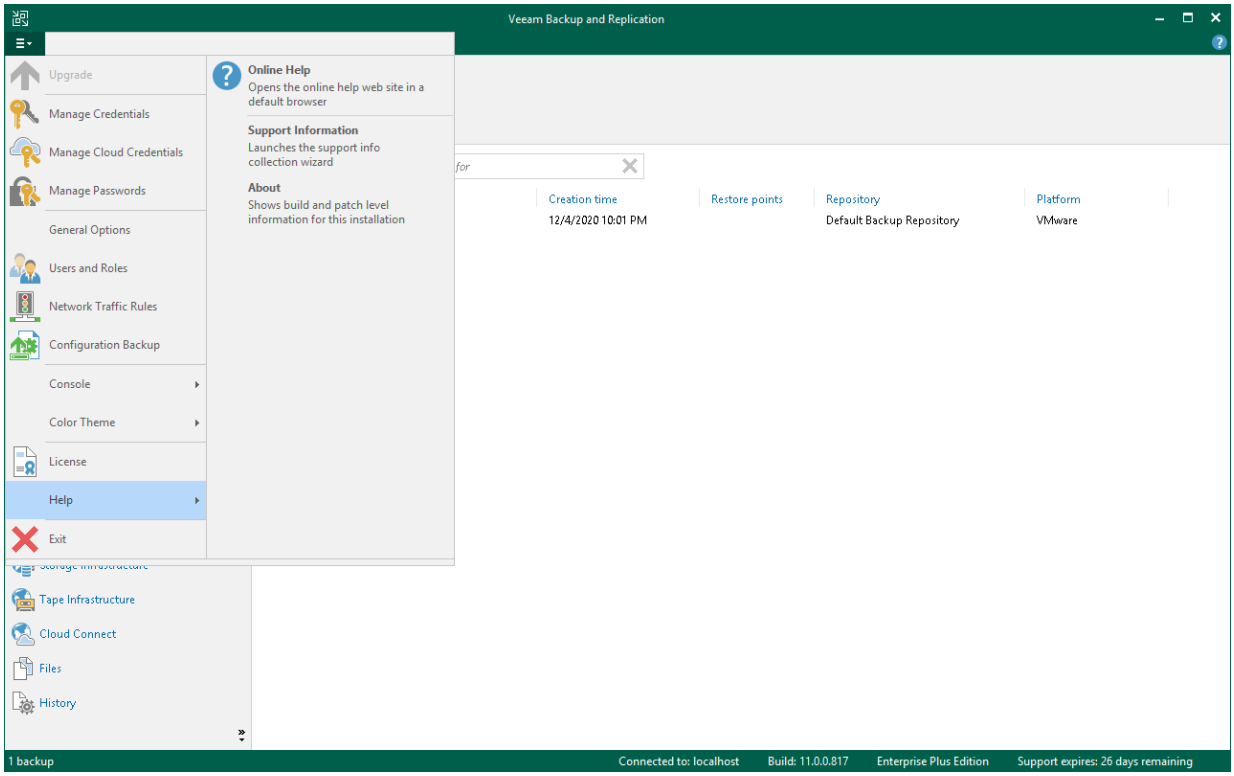
TIP

To open online help, press [F1] in any Veeam Backup & Replication wizard or window. You will be redirected to the corresponding section of the user guide.

Main Menu

The main menu in Veeam Backup & Replication contains commands related to general application settings. You can perform the following operations using the main menu:

- Upgrade backup infrastructure components
- Manage credentials
- Manage cloud credentials
- Manage passwords
- Configure application settings
- Set up user roles
- Configure network traffic rules
- Perform configuration backup and restore
- Start PuTTy and Microsoft PowerShell consoles, and open a remote desktop connection to the backup server
- Change color theme
- Work with licenses
- View Veeam Backup & Replication help and export program logs
- Exit Veeam Backup & Replication



Navigation Pane

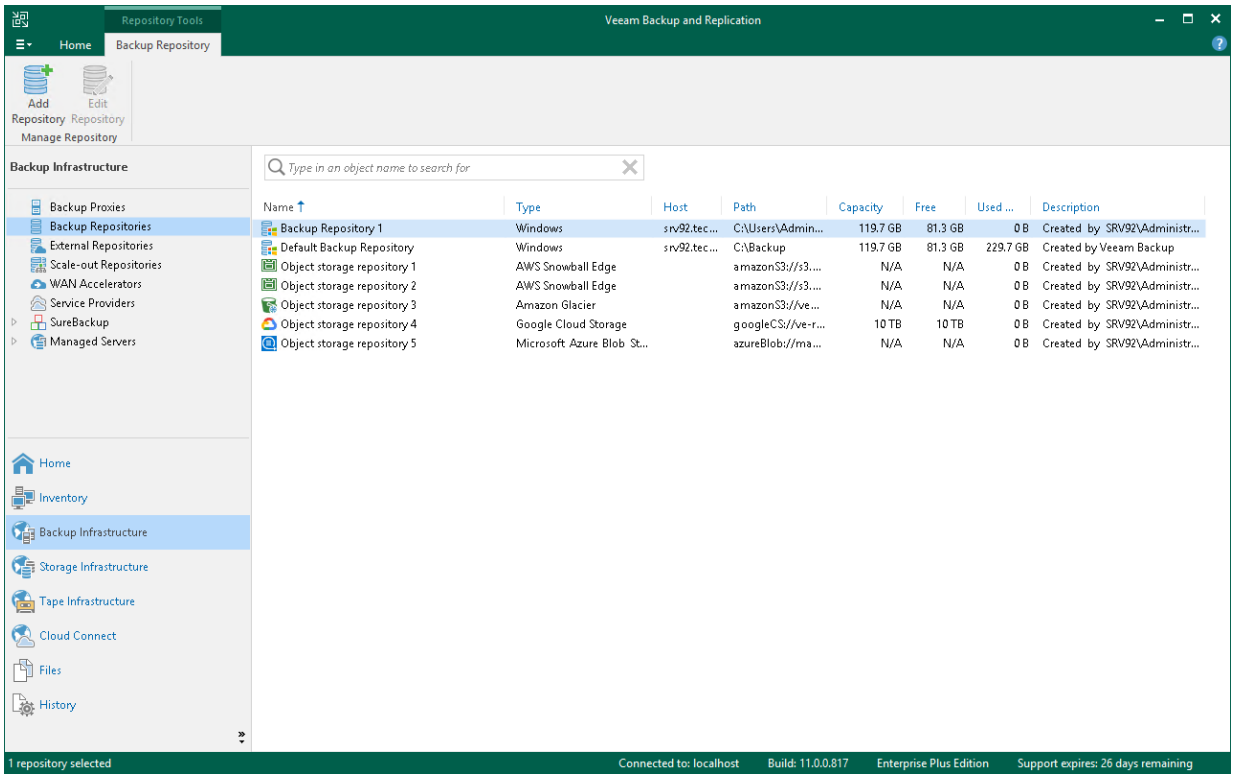
The navigation pane, located on the left of the window, provides centralized navigation and lets you easily access Veeam Backup & Replication items organized in views.

The navigation pane consists of two areas:

- The upper pane, or the inventory pane, displays a hierarchy or list of items relevant for a specific view.

  Items displayed in the inventory pane differ depending on the active view. For example, in the **Backup Infrastructure** view, the inventory pane displays a list of backup infrastructure components — virtualization servers, backup proxies, backup repositories and so on. In the **Inventory** view, the inventory pane displays a list of servers added to the backup infrastructure.

- The lower pane contains a set of buttons that let you switch between views. For more information on views and how to show/hide a view button, see Views.



Ribbon and Tabs

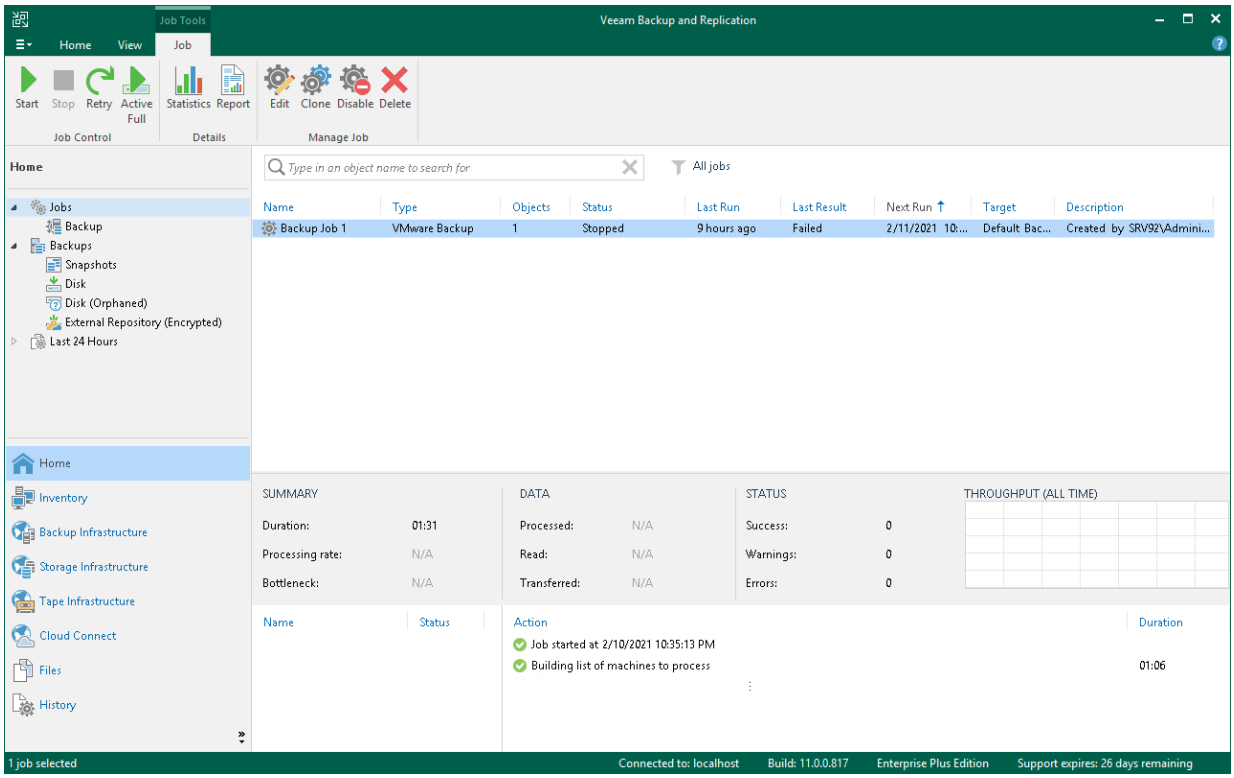
Operation commands in Veeam Backup & Replication are organized in logical groups and displayed under tabs on the ribbon. The ribbon is displayed at the top of the main application window.

On the ribbon, the following tabs are displayed:

- The **Home** tab provides quick access to the most common operations. It lets you configure different types of jobs, perform restore and import operations. This tab is always available, no matter which view is currently active.
- Other tabs contain commands specific for certain items and appear when these items are selected. For example, if you open the **Home** view and select a backup job in the working area, the **Job** tab containing buttons for operations with jobs will appear on the ribbon. If you open the **Files** view and select a file or folder, the **File Tools** tab containing buttons for operations with files will appear on the ribbon.

> **TIP**
>
> Commands for operations with items in Veeam Backup & Replication are also available from the shortcut menu.

You can minimize the ribbon. To do so, right-click anywhere on the ribbon and select **Minimize the Ribbon**. To restore the ribbon, right-click on the minimized ribbon and clear the **Minimize the Ribbon** option.



Views

Veeam Backup & Replication displays its items in views. When you click the button of a specific view in the navigation pane, the view content is displayed in the working area of Veeam Backup & Replication.

Veeam Backup & Replication offers the following views:

- The **Home** view is intended for work with jobs. It also displays a list of created backups and replicas that can be used for various restore operations, and provides statistics for recently performed jobs. For more information about job statistics, see Reporting.
- The **Inventory** view displays the inventory of the virtual infrastructure. The inventory can be presented from different perspectives: **Computer**, **Storage**, **VM Folders, VM Tags** and **vCloud**. You can use this view to work with VMs, and VM containers or groups.
- The **Backup Infrastructure** view displays a list of backup infrastructure components: servers, hosts, backup proxies, backup repositories and so on. You can use this view for backup infrastructure setup — here you can configure backup infrastructure components that will be used for data protection and disaster recovery tasks.
- The **Storage Infrastructure** view displays a list of storage systems, volumes and snapshots. You can use this view to restore VM data from storage snapshots.
- The **Tape Infrastructure** view displays a hierarchy of tape libraries connected to the tape server. You can use this view to archive data to tapes and restore data from tapes.
- The **Cloud Connect Infrastructure** view displays components of the Veeam Cloud Connect infrastructure. This view can be used by SP to manage TLS certificates, configure cloud gateways and create accounts for users who plan to work with cloud resources.
- The **Files** view displays a file tree of servers added to the backup infrastructure. You can use this view for file copying operations.
- The **History** view displays statistics on operations performed with Veeam Backup & Replication.

In some situations, some views may not be displayed. Mind the following:

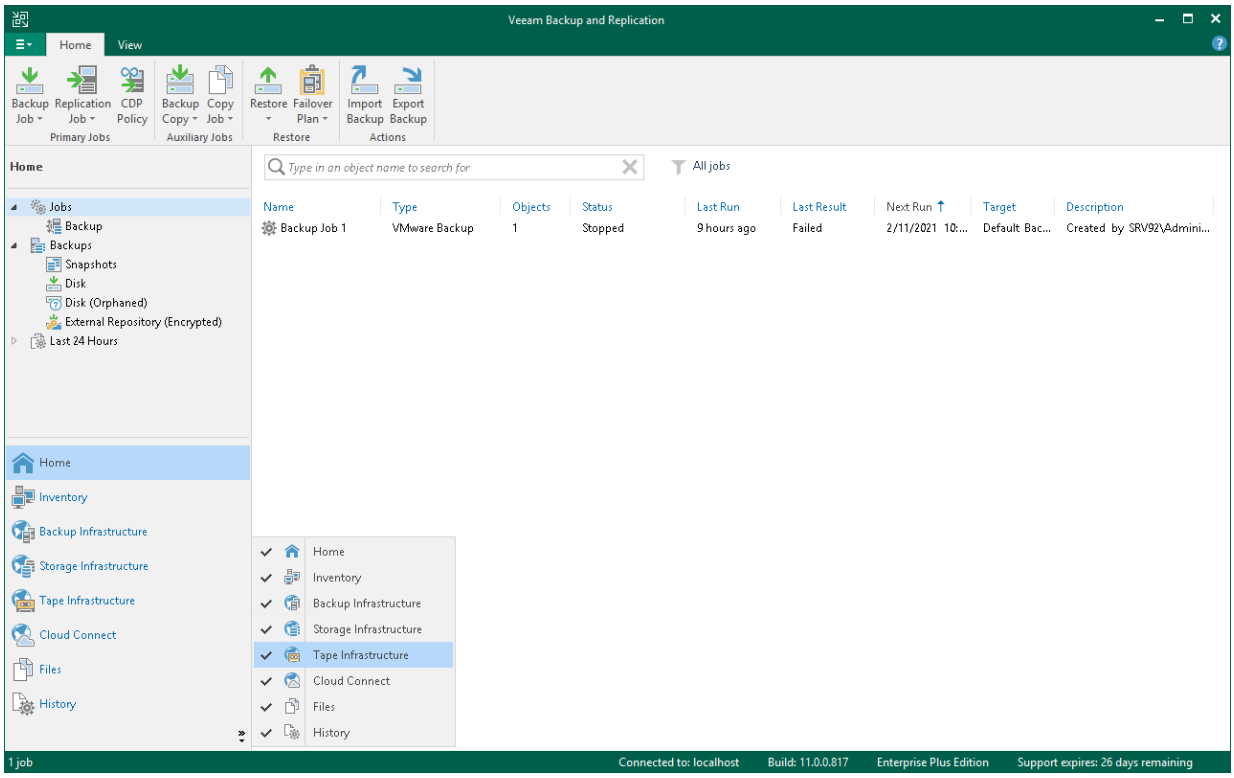
- Right after installation, Veeam Backup & Replication displays only **Backup Infrastructure** and **History** views. To display other views, you must add at least one server or virtualization host to the backup infrastructure.
- Right after installation, Veeam Backup & Replication does not save changes that you make to the navigation pane or views: for example, if you resize panes, display or hide specific views. After you restart the Veeam Backup & Replication console, the main window settings are back to default ones. To save these settings, you must add at least one server or virtualization host to the backup infrastructure.
- To display the **Cloud Connect Infrastructure** view, you must install a valid license that supports the Veeam Cloud Connect functionality.

You can hide views that you do not plan to use. For example, if you do not use tapes for data archiving, you can hide the **Tape Infrastructure** view.

To hide a view:

1. Click the arrow icon (») at the bottom of the navigation pane.

2. Click the view in the list.



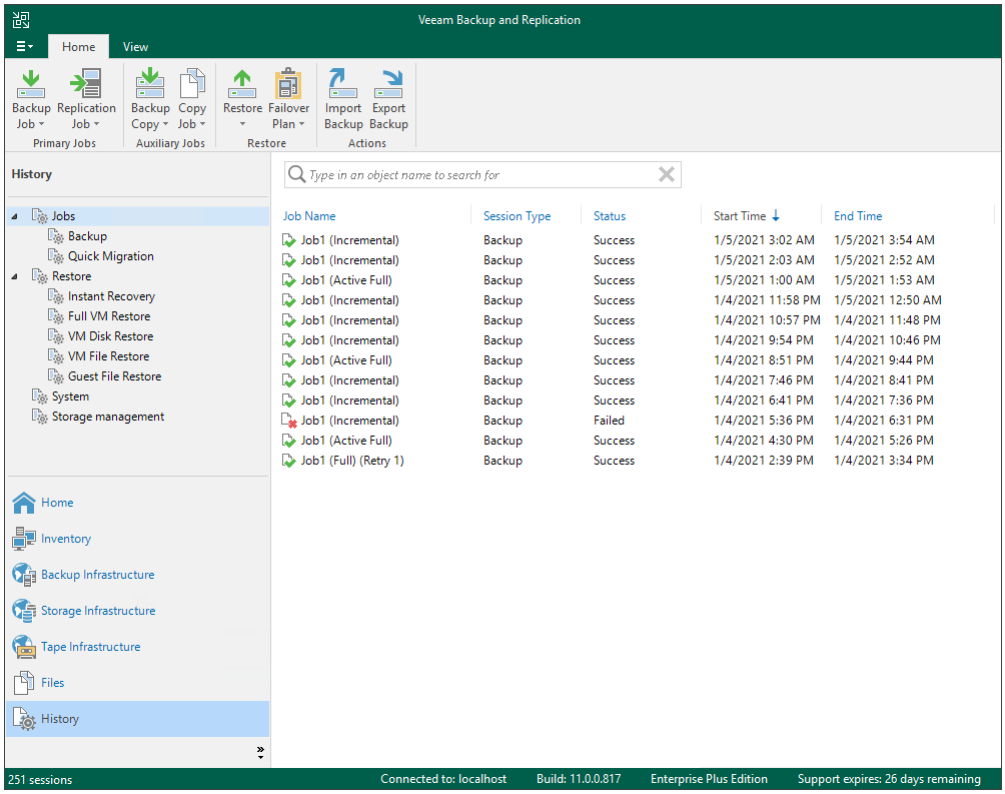
Working Area

The working area of Veeam Backup & Replication displays a list of items relating to a specific view.

The working area looks different depending on the view that is currently active. For example, if you open the **History** view, the working area will display a list of job sessions and restore tasks performed with Veeam Backup & Replication. If you open the **Inventory** view, the working area will display a list of VMs that reside on servers connected to Veeam Backup & Replication.

Every item is described with a set of properties that are presented as column headers. You can click column headers to sort items by a specific property. For example, to sort VMs by the amount of provisioned storage space, click the **Provisioned Size** header.

To hide or display properties, right-click a column header and, in the opened menu, clear or select check boxes near property names.



Job Filter

A job filter allows you to filter jobs by different parameters. For example, you can create a filter that will show only VM backup copy jobs.

Creating Job Filters

To create a filter, do the following:

1. Open the **Home** view.
2. In the inventory pane, select the **Jobs** node. At the top of the working area, click **All jobs** > **Add new**. Alternatively, right-click the **Jobs** node and select **Add view**.
3. In the **Add View** window, specify a filter name. The name must be up to 200 characters long.
4. If you want to show this filter as a subnode of the **Jobs** node in the inventory pane, select the **Add this view to the job management tree** check box.
5. In the **Workload** section, select workloads which jobs must process:
   - **Application** – Veeam Backup & Replication will show Veeam Plug-in jobs and jobs in which SQL transaction log backup or Oracle archive log backup is enabled.
   - **Cloud VM** – Veeam Backup & Replication will show jobs that process VMs stored in clouds.
   - **File** – Veeam Backup & Replication will show jobs that process files.
   - **Server** – Veeam Backup & Replication will show Veeam Agent jobs that process servers. For more information on processed computer types, see Selecting Protected Computer Type.
   - **Workstation** – Veeam Backup & Replication will show Veeam Agent jobs that process workstations. For more information on processed computer types, see Selecting Protected Computer Type.
   - **VM** – Veeam Backup & Replication will show jobs that process VMs.
6. In the **Job type** section, select job types:
   - **Backup** – backup jobs. For example, VM backup jobs or file share backup jobs.
   - **Backup copy** – backup copy jobs. For example, VM backup copy jobs or file share backup copy jobs.
   - **Copy** – copy jobs. For example, file copy jobs or VM copy jobs.
   - **Replication** – replication jobs.
   - **SureBackup** – SureBackup jobs.
   - **Tape** – tape jobs. For example, backup to tape jobs or file to tape jobs. For more information, see the Creating Backup to Tape Jobs and Creating File to Tape Jobs sections in the Veeam Backup & Replication User Guide.
7. In the **Last result** section, select statuses with which jobs must finish: *Success*, *Warning* or *Failed*.

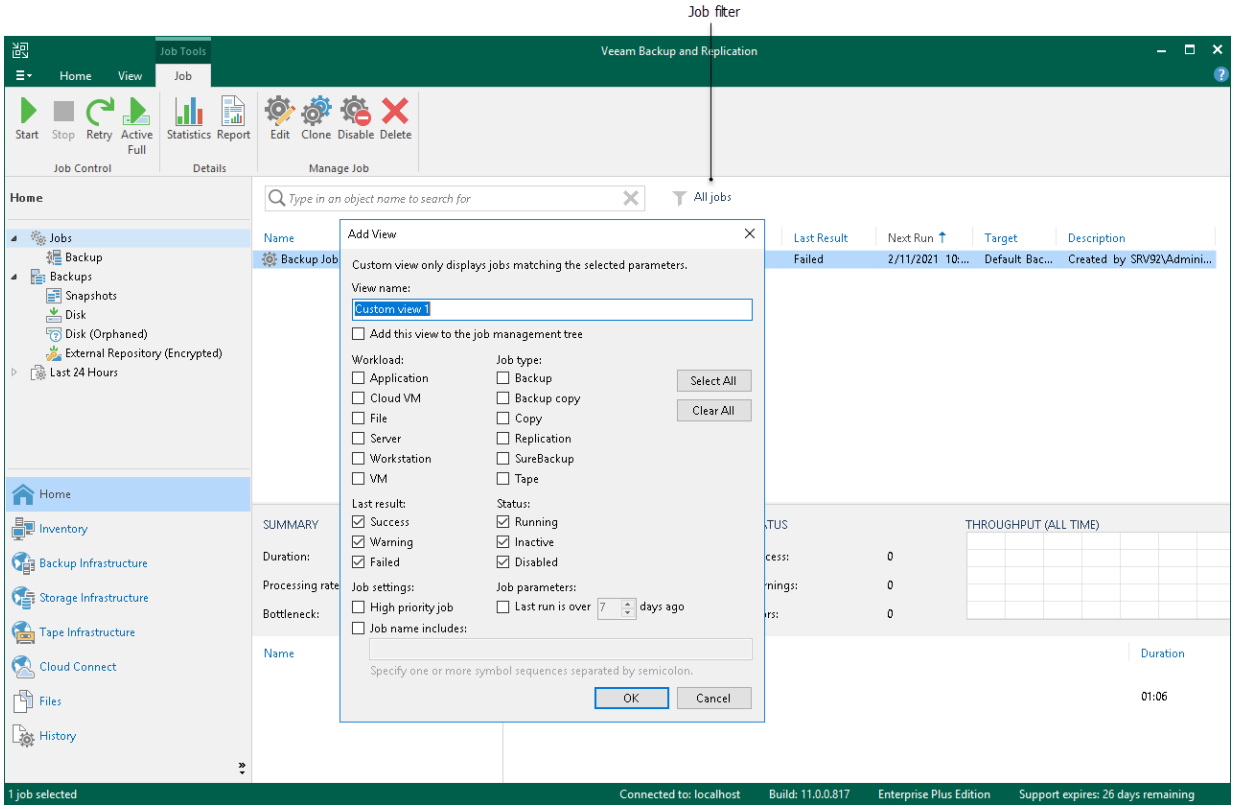
   Jobs that that have never started are considered as *Failed*.
8. In the **Status** section, select states of jobs: *Running*, *Inactive* or *Disabled*.

   Backup copy jobs in the *Idle* state are considered Inactive.

9. If you want to show jobs with manually set **High priority** flag, select the **High priority job** check box. For more information on job priorities, see Job Priorities.
10. If you want to show jobs that were inactive for some period of time, select the **Last run is over N days ago** check box and specify the period in days.
11. If you want to show jobs whose names include specific keywords, select the **Job name includes** check box and enter keywords.

    To show jobs that include any of the specified keywords, separate these keywords by a semicolon without a space. For example, if you enter *"Backup Job;Daily"*, Veeam Backup & Replication will show all jobs that include *"Backup Job"* or *"Daily"* keywords in their names.

> **NOTE**
>
> Only the user who creates filters can access them — that is, other users cannot use these filters.

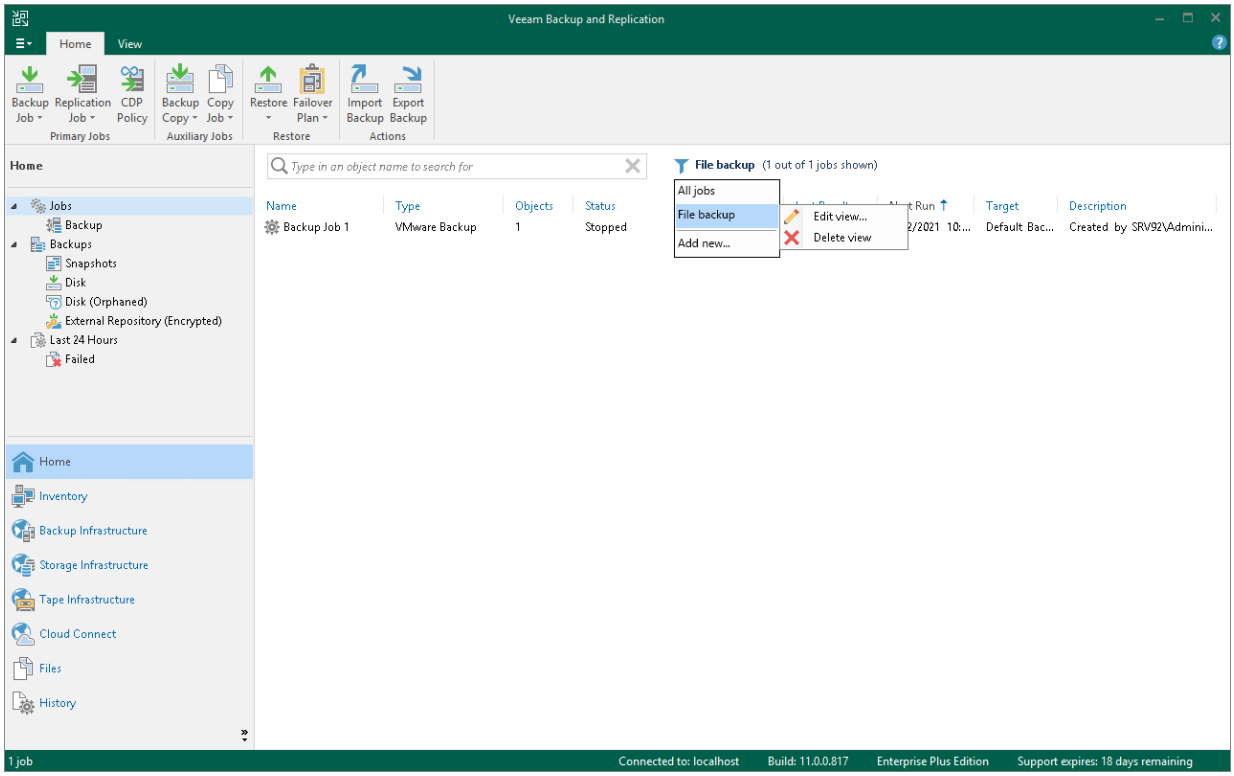


Editing and Deleting Job Filters

To edit or delete a job filter, do the following:

1. Open the **Home** view.
2. In the inventory pane, select the **Jobs** node.
3. At the top of the working area, click the filter icon.
4. Hover the mouse over a job that you want to edit or delete.
5. Click the ellipsis button.
6. Select **Edit view** or **Delete view**.

If you have added the filter as a subnode of the **Jobs** node in the inventory pane, you can right-click the subnode and then click **Edit view** or **Delete view**.



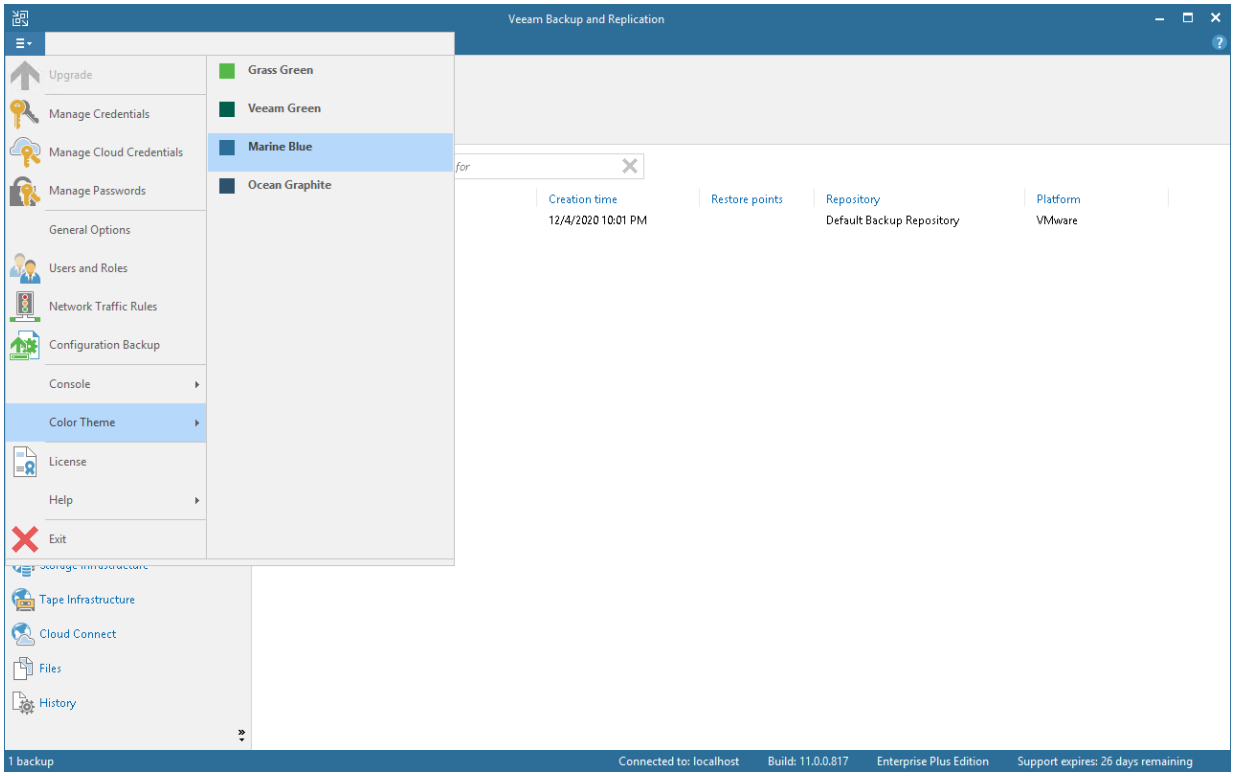
Changing Color Theme

By default, Veeam Backup & Replication uses a '*Veeam Green*' color theme for the UI. If necessary, you can change the color theme. Changing the color theme can be helpful, for example, if you connect to different backup servers from one remote machine on which the Veeam Backup & Replication console is installed. In this case, you will be able to easily differentiate with which backup server you are currently working.

To change the color theme for Veeam Backup & Replication:

1. From the main menu, select **Color Theme**.
2. Choose one of color themes: *Grass Green*, *Veeam Green*, *Marine Blue*, *Ocean Graphite*.

> **NOTE**
>
> Color theme settings are applicable for a specific combination of a backup server and user account. For example, the color theme is initially set to the default one. You log on to the Veeam Backup & Replication console under some user account and change the color theme to **Marine Blue**. If you log on to the same backup server under the same account next time, the color theme will be set to **Marine Blue**. If you log on to the same backup server under another account, Veeam Backup & Replication will use the color theme that was previously set for this account — that is, the default color theme.



Restore Points Icons

Restore points icons display the information about a current state of a restore point. There are 2 types of icons, for backup repository restore points and for scale-out backup repository restore points.

Backup Repository Restore Points

These user interface icons indicate the states of the restore points in all backups.

Icon	State
	Full restore point
	Incremental restore point
	Reverse incremental restore point
	Missing full restore point
	Missing incremental restore point
	Missing reverse incremental restore point

Scale-Out Backup Repository Restore Points

These user interface icons indicate the states of the restore points in the scale-out backup repository.

Icon	State
	Full restore point; on performance tier only
	Full restore point; on performance tier and offloaded to capacity tier
	Full restore point; on capacity tier only
	Full restore point; on capacity tier and partially downloaded to performance tier *(This can happen when download to the performance tier fails)*
	Incremental restore point; on performance tier only
	Incremental restore point; on performance tier and offloaded to capacity tier

Icon	State
	Incremental restore point; on capacity tier only
	Incremental restore point; on capacity tier and partially downloaded to performance tier *(This can happen when download to the performance tier fails)*
	Rollback restore point; on performance tier only
	Rollback restore point; on performance tier and offloaded to capacity tier
	Rollback restore point; on capacity tier only
	Rollback restore point; on capacity tier and partially downloaded to performance tier *(This can happen when download to the performance tier fails)*
	Rollback restore point; on archive tier only

VMware vSphere Icons

These user interface icons represent the hypervisor objects of VMware vSphere workloads and servers, added to the backup infrastructure.

Icon	Description
	Cluster
	Cluster with no license
	Datacenter
	Datastore
	Datastore cluster
	ESX host
	ESX host with no license
	Standby ESX host
	Provider VDC
	VBR Proxy
	Proxy servers folder
	Resource pool
	VBR Tape Server
	vCloud vApp
	Running vCloud vApp
	Paused vCloud vApp
	vCenter

Icon	Description
	vCenter with no license
	VM
	Deleted VM
	Running VM
	Paused VM

Job Priorities

Resources in the backup infrastructure are limited. To make sure that the most crucial jobs are the first to get free resources to provide the reliable data protection, Veeam Backup & Replication uses the system of priorities to allocate resources to different jobs.

The resource scheduler within Veeam Backup & Replication uses several stages to prioritize jobs and provide free resources to them:

1. **Type** – at the first stage, the resource scheduler identifies the priority of the jobs awaiting free resources based on their type:
   a. **Backup restore jobs** – these jobs have the highest priority (**800**) and are the first to get free system resources.
   b. **Continuous data protection jobs** – these jobs have priority (**700**).
   c. **SnapshotDeleter jobs** – these jobs have priority (**600**).
   d. **Quick backup jobs** – these jobs have priority (**500**).
   e. **High priority jobs** – jobs with the enabled **High priority** option have priority (**400**). You can enable the **High priority** option for the following jobs: backup jobs, replication jobs, agent jobs managed by backup server, file share backup jobs.
   f. **Regular backup and replication jobs** – these jobs have priority (**300**).
   g. **Backup copy jobs** – these jobs have priority (**200**).
   h. **Archive jobs** – these jobs have the lowest priority (**100**) and are the last to get free system resources.
2. **Priority** – at the second stage, the resource scheduler identifies the priority of the jobs within each type group from the first stage based on their startup type:
   a. **Scheduled VSS proxy jobs** – the jobs with the configured job schedule and using a VSS proxy have the highest priority (**40**) within the group and are the first to get free system resources.
   b. **Scheduled jobs** – the jobs with the configured job schedule have priority (**30**) within the group.
   c. **Manually started VSS proxy jobs** – the manually started jobs using a VSS proxy have priority (**20**) within the group.
   d. **Manually started jobs** – the manually started jobs have the lowest priority (**10**) within the group and are the last to get free system resources.
3. **Start time** – if the jobs have the same type and priority, resources are first allocated to jobs that were started earlier.

> **NOTE**
>
> Mind that the resource scheduler does not take into account the start time for tape jobs.

You can check the job type and priority of a certain job in service logs. For more information on logs, see Logging.

Deployment

To start working with Veeam Backup & Replication, you must configure a backup server — install Veeam Backup & Replication on a machine that meets the system requirements. To do this, you can use the setup wizard or install the product in the unattended mode.

When you install Veeam Backup & Replication, the Veeam Backup & Replication console is automatically installed on the backup server. If you want to access Veeam Backup & Replication remotely, you can install the Veeam Backup & Replication console on a dedicated machine.

Installing Veeam Backup & Replication

Before you install Veeam Backup & Replication, check prerequisites. Then use the **Veeam Backup & Replication** setup wizard to install the product.

Before You Begin

Before you install Veeam Backup & Replication, check the following prerequisites:

- A machine on which you plan to install Veeam Backup & Replication must meet the system requirements. For more information, see System Requirements.
- A user account that you plan to use for installation must have sufficient permissions. For more information, see Required Permissions.
- Backup infrastructure components communicate with each other over specific ports. These ports must be open. For more information, see Used Ports.
- You must remove Veeam Backup & Replication components of versions that are not supported by the upgrade procedure from the target machine. You may also need to remove earlier versions of other Veeam products and components.
- Before deploying Veeam Backup & Replication, define where the Veeam Backup & Replication server will be located. Depending on what kind of protection are you planning to use, the Veeam Backup & Replication server should be located on the source site or the Disaster Recovery site.
  - **When replication or CDP is used**: If you plan to use replication or Continuous Data Protection (CDP), the Veeam Backup & Replication server should be deployed in the disaster recovery site. In this case, if the production host crashes, Veeam Backup & Replication will automatically fail over to the replica without any manual operations. The source backup infrastructure still can be managed with the same Veeam Backup & Replication server with the help of backup proxies deployed in the source site.
  - **When only backup features are used**: If you plan to use Veeam Backup & Replication for backup jobs only, the backup server should be placed in the production site.

Configuration Database

Before deploying Veeam Backup & Replication, decide which version of Microsoft SQL Server you need to use:

- If you do not prepare an SQL Server in advance, Veeam Backup & Replication will install an Express Edition of SQL Server, which is limited to store 10 GB of configuration data. The Express Edition is enough for evaluation purposes and not very large environments (<500 VMs).
- If your infrastructure is large (more than 500 VMs), you may consider to install a Microsoft SQL Server in advance.
- Also, you must decide do you need the SQL Server installed on the same server as Veeam Backup & Replication or on a remote server.

  Veeam Backup & Replication requires Microsoft SQL Server deployed either locally on the backup server or remotely. If Microsoft SQL Server is not installed, the Veeam Backup & Replication setup will install Microsoft SQL Server Express Edition locally on the backup server. For more information, see System Requirements.

  If Microsoft SQL Server is installed by the previous product version, Veeam Backup & Replication will connect to the existing configuration database, upgrade it (if necessary) and use it for work.

  It is recommended to run the Microsoft SQL Server instance locally to eliminate latency issues. However, in some scenarios a remote Microsoft SQL Server instance can be the better choice:

  - High Availability. SQL Clustering and AlwaysOn Availability Group on external SQL Servers can be used for high availability of the configuration database. To learn about the configuration details, see this Veeam KB article.

    - Licensing. Some enterprises have dedicated virtual clusters for SQL Servers due to licensing constraints. In such cases, you can place the Veeam configuration database on an existing instance to lower the total cost of ownership.

- The usage of Microsoft SQL Server Express Edition is limited by the database size up to 10 GB. If you plan to have larger databases, use other editions of Microsoft SQL Server.

Related Topics

Infrastructure Planning

Step 1. Start Setup Wizard

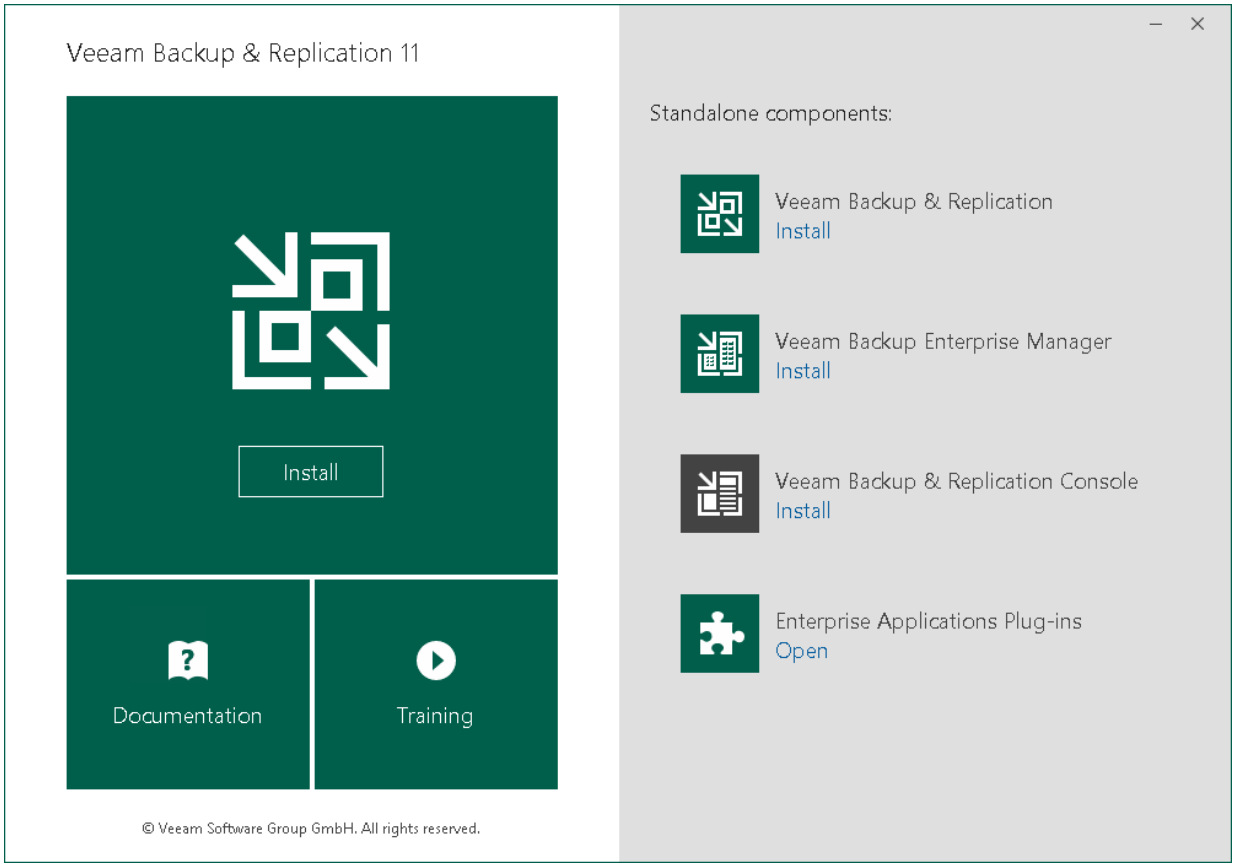
To start the setup wizard:

1. Download the latest version of the Veeam Backup & Replication installation image from the Download Veeam Products page.
2. Mount the installation image to the machine on which you plan to install Veeam Backup & Replication or burn the image file to a flash drive or other removable storage device. If you plan to install Veeam Backup & Replication on a VM, use built-in tools of the virtualization management software to mount the installation image to the VM.

   To extract the content of the ISO, you can also use the latest versions of utilities that can properly extract data from ISOs of large size and can properly work with long file paths.
3. After you mount the image or insert the disk, Autorun will open a splash screen with installation options. If Autorun is not available or disabled, run the `Setup.exe` file from the image or disk.
4. In the **Veeam Backup & Replication** section of the splash screen, click **Install**.

IMPORTANT

It is strongly recommended that you install Veeam Backup & Replication using Autorun or the Setup.exe file. If you run other installation files from the ISO folders, you may miss some components that need to be installed, and Veeam Backup & Replication may not work as expected.



Step 2. Read and Accept License Agreement

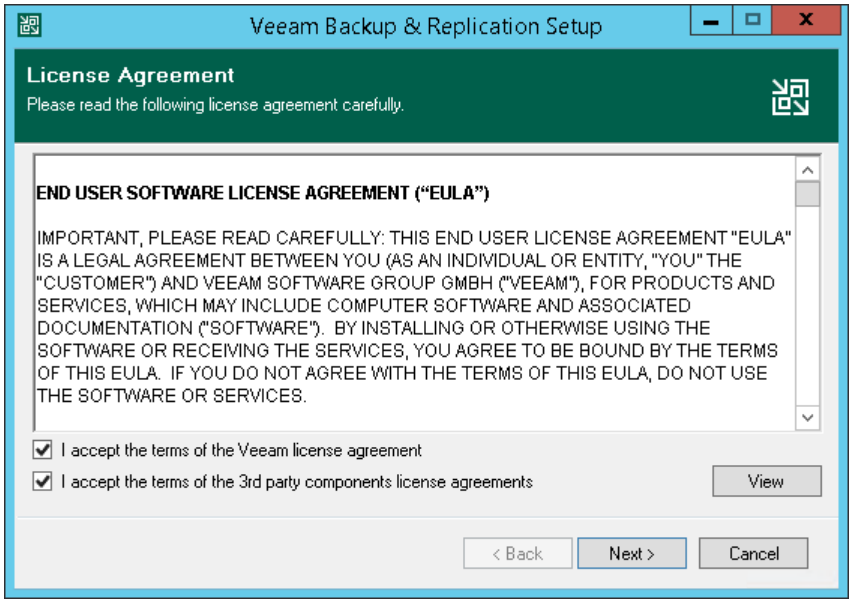
At the **License Agreement** step of the wizard, you must accept the license agreement for Veeam and 3rd party components that Veeam incorporates. If you do not accept the license agreement, you will not be able to pass to the next step of the setup wizard.

1. Read the license agreement.

   To view the license agreement for 3rd party components, click **View**.

2. Select the **I accept the terms of the Veeam license agreement** check box.

3. Select the **I accept the terms of the 3rd party components license agreement** check box.



Step 3. Provide License File

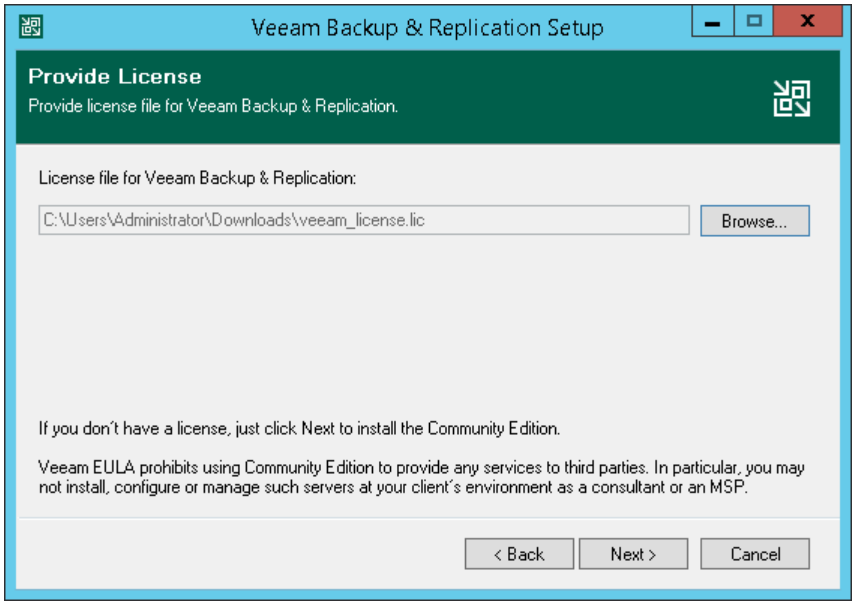
At the **Provide License** step of the wizard, you must specify what license for Veeam Backup & Replication you want to install. For more information, see Licensing.

If a valid license is already installed on the machine, the setup wizard will inform you about it. In this case, you can skip the **Provide License** step and move to the next step of the wizard.

If you do not install a license, the product will operate in the Veeam Backup & Replication Community (free) Edition. For more information, see Veeam Backup & Replication Community Edition.

To install a license:

1. Next to the **License file for Veeam Backup & Replication** field, click **Browse**.
2. Select a valid license file for Veeam Backup & Replication.



Step 4. Review Components and Select Installation Folder

At the **Program Features** step of the wizard, you can check what components the setup wizard will install on the machine and choose the installation folder.

The setup wizard installs the following components:

- Veeam Backup & Replication
- Veeam Backup Catalog (component responsible for storing VM guest OS indexing data)
- Veeam Backup & Replication Console

The setup wizard also installs the following components in the background:

- AWS Plug-In for Veeam Backup & Replication
- Microsoft Azure Plug-In for Veeam Backup & Replication
- Veeam Agent for Linux Redistributable
- Veeam Agent for Mac Redistributable
- Veeam Agent for Microsoft Windows Redistributable
- Veeam Agent for Unix Redistributable
- Veeam Backup PowerShell Module
- Veeam Distribution Server

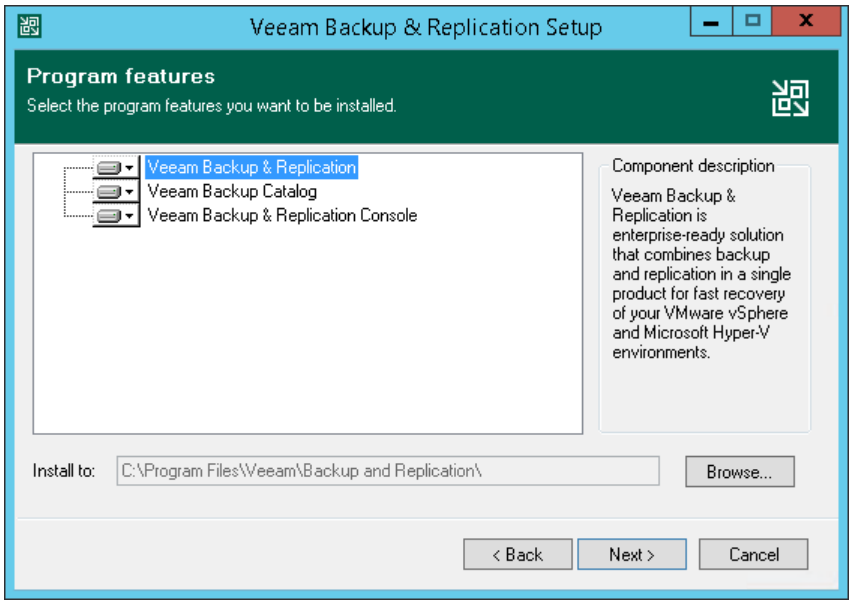
  The distribution server is an architecture component in the Veeam Agent management infrastructure used for deployment of Veeam Agent setup files to protected computers. For more information, see the Veeam Agent Management Infrastructure section in the Veeam Agent Management Guide.
- Veeam Explorer for Microsoft Active Directory
- Veeam Explorer for Microsoft Exchange
- Veeam Explorer for Microsoft OneDrive for Business
- Veeam Explorer for Microsoft SharePoint
- Veeam Explorer for Microsoft SQL Server
- Veeam Explorer for Microsoft SQL Teams
- Veeam Explorer for Oracle
- Veeam Platform Service for Google

These components do not require additional licenses. They are integrated with Veeam Backup & Replication.

To choose the installation folder:

1. On the right of the **Install to** field, click **Browse**.

2. In the **Browse for Folder** window, select the installation folder for the product. The default installation folder is `C:\Program Files\Veeam\Backup and Replication\`.



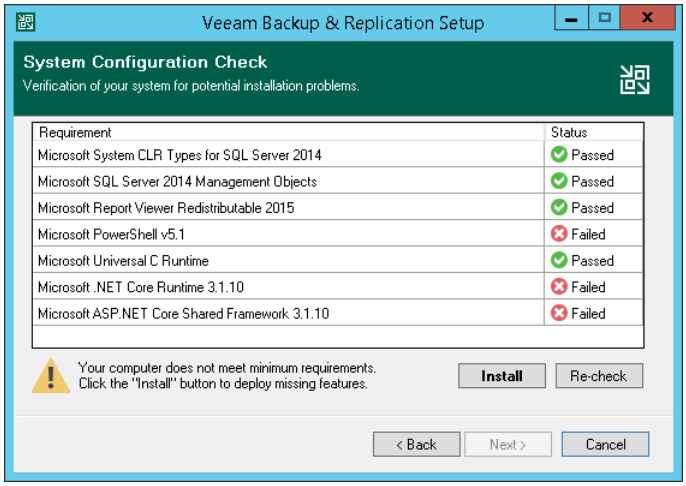
Step 5. Install Missing Software

At the **System Configuration Check** step of the wizard, the setup wizard checks if all prerequisite software is installed on the machine. If required software components are missing, the setup wizard will offer you to install them.

You can install missing components automatically or manually.

- To install missing components automatically, click **Install**. The setup wizard will not interrupt the installation process and install the missing components during the current work session.
- To install missing components manually:
  a. Click **Cancel** and exit the setup wizard.
  b. Install and enable the necessary components manually on the machine.
  c. Start the setup wizard again, pass to the **System Configuration Check** step of the wizard and click **Re-run** to repeat the verification.

> **NOTE**
>
> If all required components are already installed on the machine, the **System Configuration Check** step will be skipped.



Step 6. Specify Installation Settings

At the **Default Configuration** step of the wizard, you can select to install Veeam Backup & Replication with default installation settings or specify custom installation settings.

By default, the setup wizard installs Veeam Backup & Replication with the following settings:

- **Installation folder**: `C:\Program Files\Veeam\Backup and Replication`.
- **Instant recovery cache folder**: the *IRCache* folder on a volume with the maximum amount of free space, for example, `C:\ProgramData\Veeam\Backup\IRCache`.

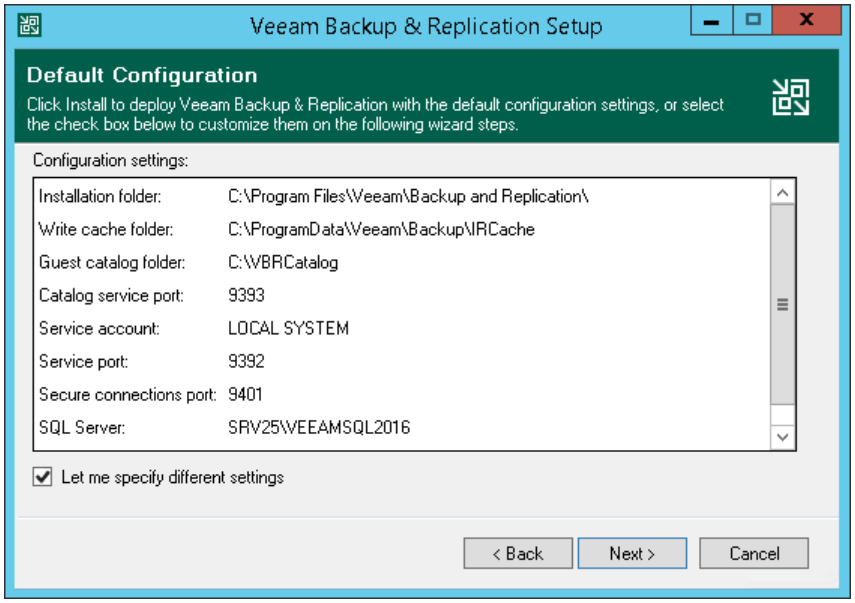
  The IR cache folder stores the write cache for machines that are started from backups during recovery verification or restore operations. Make sure that you have at least 10 GB of free disk space to store the write cache.
- **Guest catalog folder**: the *VBRCatalog* folder on a volume with the maximum amount of free space, for example, `C:\VBRCatalog`.

  The guest catalog folder stores indexing data for VM guest OS files. Indexing data is required for browsing and searching for VM guest OS files inside backups and performing 1-click restore.
- **Catalog service port**: *9393*. The catalog service port is used by the Veeam Guest Catalog Service to replicate catalog data from backup servers to Veeam Backup Enterprise Manager.
- **Service account**: *LOCAL SYSTEM*. The service account is the account under which the Veeam Backup Service runs.
- **Service port**: *9392*. The service port is used by Veeam Backup Enterprise Manager to collect data from backup servers. In addition to it, the Veeam Backup & Replication console uses this service port to connect to the backup server.
- **Secure connections port**: *9401*. The secure connections port is used by the mount server to communicate with the backup server.
- **SQL Server**: *LOCALHOST\VEEAMSQL2012* or *LOCALHOST\VEEAMSQL2016*. During installation, the Veeam Backup & Replication setup installs a new instance of Microsoft SQL Server locally on the backup server:
  - For machines running Microsoft Windows 7 or Microsoft Windows Server 2008 R2, the setup installs Microsoft SQL Server 2012 SP4 Express Edition.
  - For machines running Microsoft Windows Server 2012 or later, the setup installs Microsoft SQL Server 2016 SP2 Express Edition.
- **Database name**: *VeeamBackup*. Veeam Backup & Replication deploys the Veeam Backup & Replication configuration database on the locally installed instance of Microsoft SQL Server.

To use default installation settings:

1. Leave the **Let me specify different settings** check box not selected.
2. Click **Next** to move to the next step.

To use custom installation settings, select the **Let me specify different settings** check box. The setup wizard will include additional steps that will let you configure installation settings.



Step 7. Specify Service Account Settings

At the **Service Account** step of the wizard is available if you have selected to configure installation settings manually.

You can select an account under which you want to run the Veeam Backup Service:

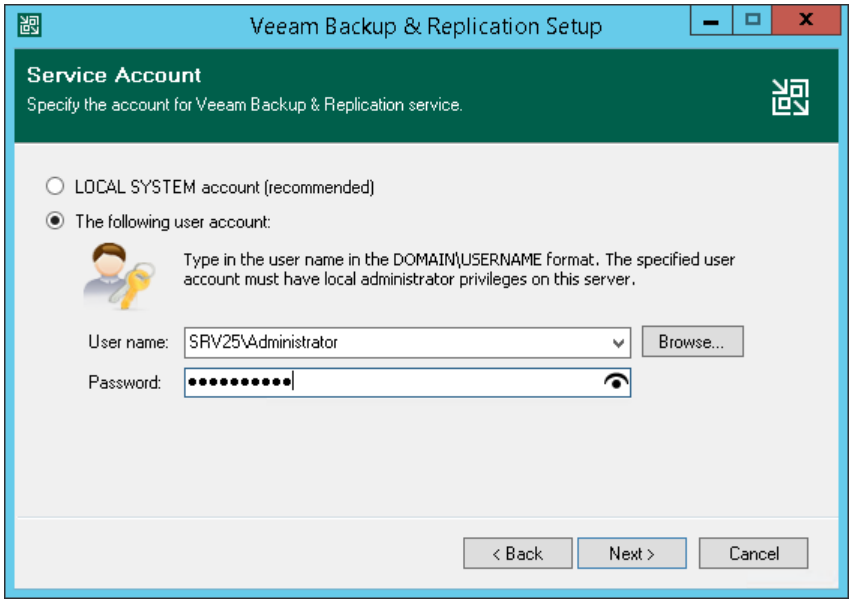
- LOCAL SYSTEM account (recommended, used by default)
- Another user account

The user name of the account must be specified in the *DOMAIN|USERNAME* format.

The user account must have the following rights and permissions:

- The account must be a member of the *Administrators* group on the machine where Veeam Backup & Replication is installed.
- The account must have *db_owner* rights for the configuration database.

Veeam Backup & Replication automatically grants the *Log on as service* right to the specified user account.



Step 8. Select Microsoft SQL Server Instance

The **SQL Server Instance** step of the wizard is available if you have selected to configure installation settings manually.

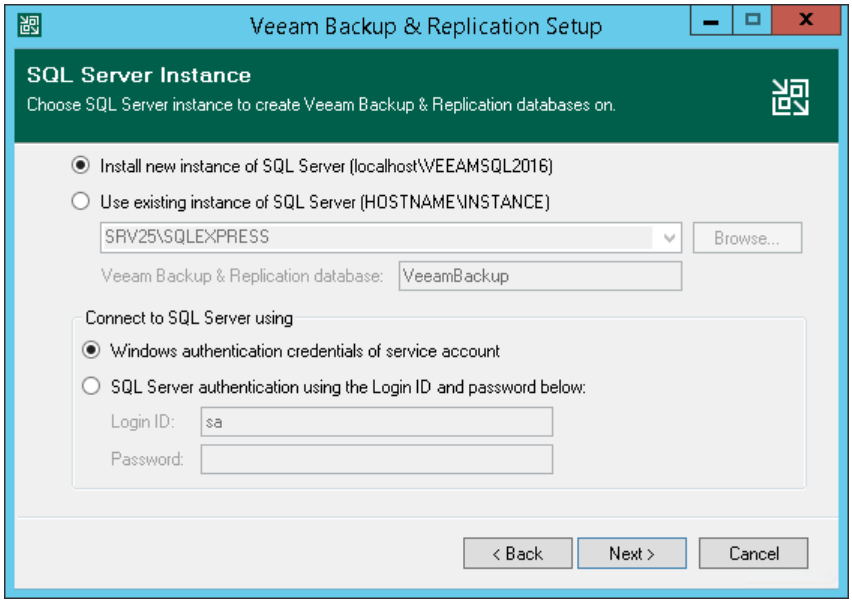
Veeam Backup & Replication stores data about the backup infrastructure, jobs, sessions and other configuration data on a Microsoft SQL Server database. The database instance can be located on a Microsoft SQL Server installed either locally (on the same machine where the backup server is running) or remotely.

> **NOTE**
>
> Mind limitations and considerations in Before You Begin.

You can select a Microsoft SQL Server instance on which you want to deploy the configuration database, and choose the authentication mode.

1. Select a Microsoft SQL Server instance:
   - If a Microsoft SQL Server is not installed locally or remotely, select the **Install new instance of SQL Server** option. The setup will install Microsoft SQL Server locally on the backup server:
     - For machines running Microsoft Windows 7 or Microsoft Windows Server 2008 R2, the setup will install Microsoft SQL Server 2012 SP4 Express Edition.
     - For machines running Microsoft Windows Server 2012 or later, the setup will install Microsoft SQL Server 2016 SP2 Express Edition.
   - If a Microsoft SQL Server is already installed locally or remotely, select the **Use existing instance of SQL Server** option. Enter the instance name in the *HOSTNAME\INSTANCE* format. In the **Database** field, specify a name for the Veeam Backup & Replication configuration database.
2. Select an authentication mode to connect to the Microsoft SQL Server instance: Microsoft Windows authentication or SQL Server authentication. If you select the SQL Server authentication, enter credentials for the Microsoft SQL Server account.

If the configuration database already exists on the Microsoft SQL Server instance (for example, it was created by a previous installation of Veeam Backup & Replication), the setup wizard will notify about it. To connect to the detected database, click **Yes**. If necessary, Veeam Backup & Replication will automatically upgrade the database to the latest version.

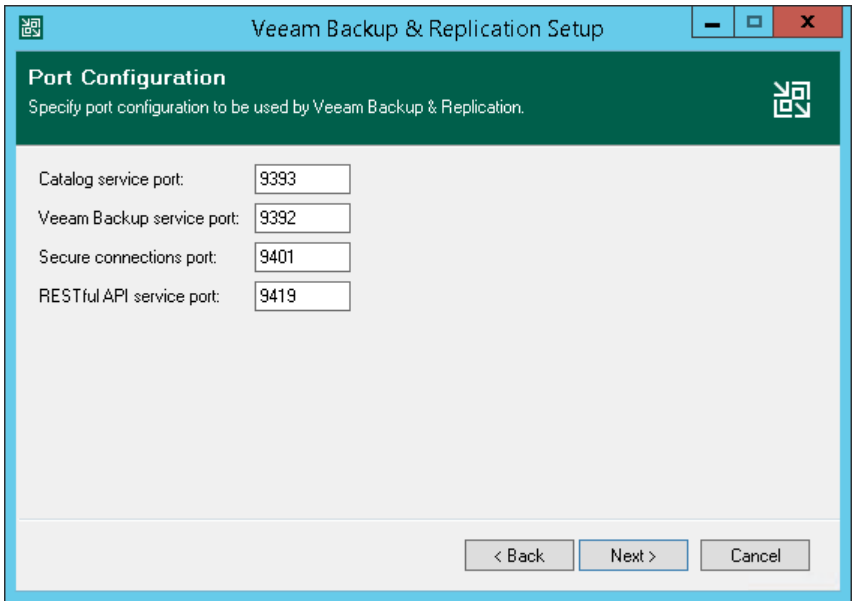


Step 9. Specify Service Ports

The **Port Configuration** step of the wizard is available if you have selected to configure installation settings manually.

You can customize port number values that will be used for communication between backup infrastructure components:

- Catalog service port. The catalog service port is used by the Veeam Guest Catalog Service to replicate catalog data from backup servers to Veeam Backup Enterprise Manager. By default, port 9393 is used.
- Veeam Backup Service port. The service port is used by Veeam Backup Enterprise Manager to collect data from backup servers. In addition to it, the Veeam Backup & Replication console uses this service port to connect to the backup server. By default, port 9392 is used.
- Secure connections port. The secure connections port is used by the mount server to communicate with the backup server. By default, port 9401 is used.
- RESTful API service port. The RESTful API service port is used to communicate with the Veeam Backup & Replication RESTful API. By default, port 9419 is used.



Step 10. Specify Data Locations

The **Data Locations** step of the wizard is available if you have selected to configure installation settings manually.

You can specify where the write cache and indexing data must be stored.

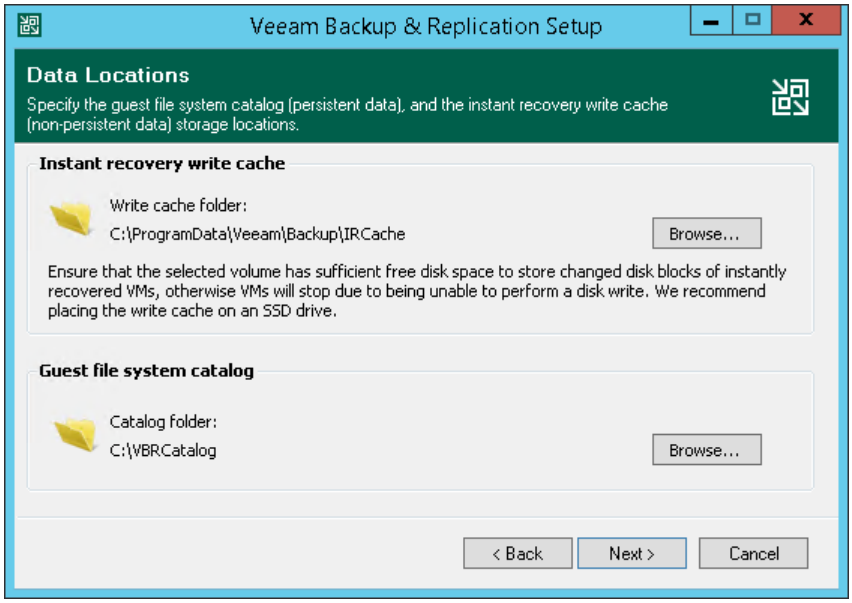
1. [For VMware environments] In the **Instant recovery write cache** section, specify a path to the IR cache folder. The IR cache folder stores the write cache for machines that are started from backups during recovery verification or restore operations. Make sure that you have at least 10 GB of free disk space to store the write cache.

   By default, the setup wizard creates the IR cache folder on a volume with the maximum amount of free space, for example, `C:\ProgramData\Veeam\Backup\IRCache`.

   You do not need to configure this data location for Microsoft Hyper-V environments.

2. In the **Guest file system catalog** section, specify a path to the folder where index files must be stored.

   By default, the setup wizard creates the *VBRCatalog* folder on a volume with the maximum amount of free space, for example: `C:\VBRCatalog`.



Step 11. Verify Installation Settings

At the **Ready to Install** step of the wizard, you can verify Veeam Backup & Replication installation settings and, if required, go back to configuration steps and adjust settings.

By default, the setup wizard installs Veeam Backup & Replication with the following settings:

- **Installation folder**: `C:\Program Files\Veeam\Backup and Replication`.
- **Write cache folder**: the *IRCache* folder on a volume with the maximum amount of free space, for example, `C:\ProgramData\Veeam\Backup\IRCache`.

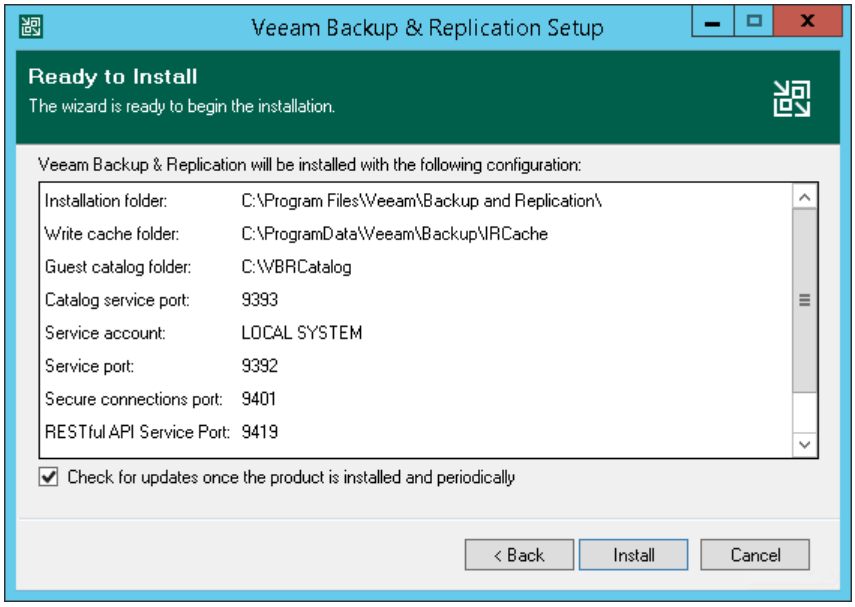
  The IR cache folder stores the write cache for machines that are started from backups during recovery verification or restore operations. Make sure that you have at least 10 GB of free disk space to store the write cache.
- **Guest catalog folder**: the *VBRCatalog* folder on a volume with the maximum amount of free space, for example, `C:\VBRCatalog`.

  The guest catalog folder stores indexing data for VM guest OS files. Indexing data is required for browsing and searching for VM guest OS files inside backups and performing 1-click restore.
- **Catalog service port**: *9393*. The catalog service port is used by the Veeam Guest Catalog Service to replicate catalog data from backup servers to Veeam Backup Enterprise Manager.
- **Service account**: *LOCAL SYSTEM*. The service account is the account under which the Veeam Backup Service runs.
- **Service port**: *9392*. The service port is used by Veeam Backup Enterprise Manager to collect data from backup servers. In addition to it, the Veeam Backup & Replication console uses this service port to connect to the backup server.
- **Secure connections port**: *9401*. The secure connections port is used by the mount server to communicate with the backup server.
- **RESTful API service port**: *9419*. The RESTful API service port is used to communicate with the Veeam Backup & Replication RESTful API.
- **SQL Server**: *LOCALHOST\VEEAMSQL2012* or *LOCALHOST\VEEAMSQL2016*. During installation, the Veeam Backup & Replication setup installs a new instance of Microsoft SQL Server locally on the backup server:
  - For machines running Microsoft Windows 7 or Microsoft Windows Server 2008 R2, the setup installs Microsoft SQL Server 2012 SP4 Express Edition.
  - For machines running Microsoft Windows Server 2012 or later, the setup installs Microsoft SQL Server 2016 SP2 Express Edition.
- **Database name**: *VeeamBackup*. Veeam Backup & Replication deploys the Veeam Backup & Replication configuration database on the locally installed instance of Microsoft SQL Server.

To use displayed installation settings, click **Install**. The installation process will begin.

To adjust installation settings, click **Back** to return to the required installation configuration step and change settings.

To automatically check for Veeam Backup & Replication updates, select the **Check for updates once the product is installed and periodically** check box.

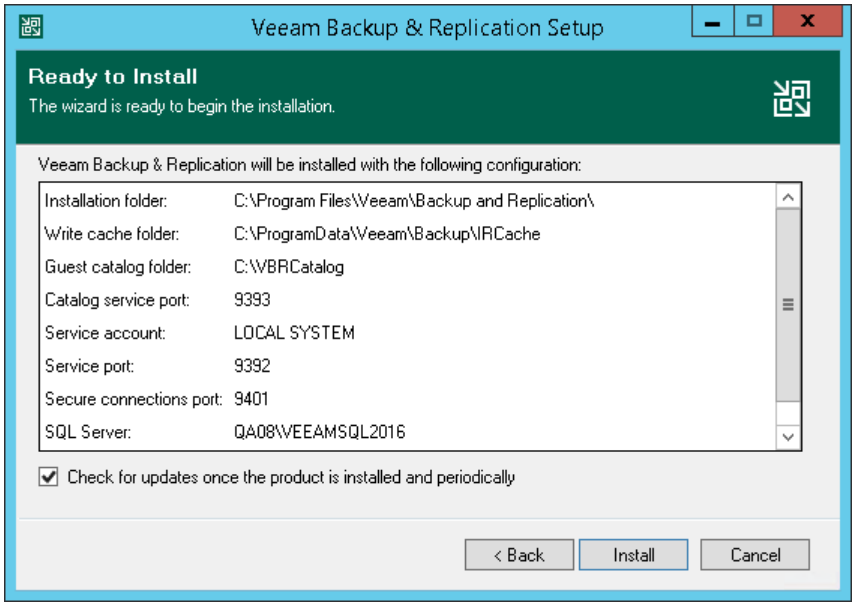


Step 12. Begin Installation

The **Ready to Install** step of the wizard is available if you have selected to configure installation settings manually.

You can review installation settings and start the installation process.

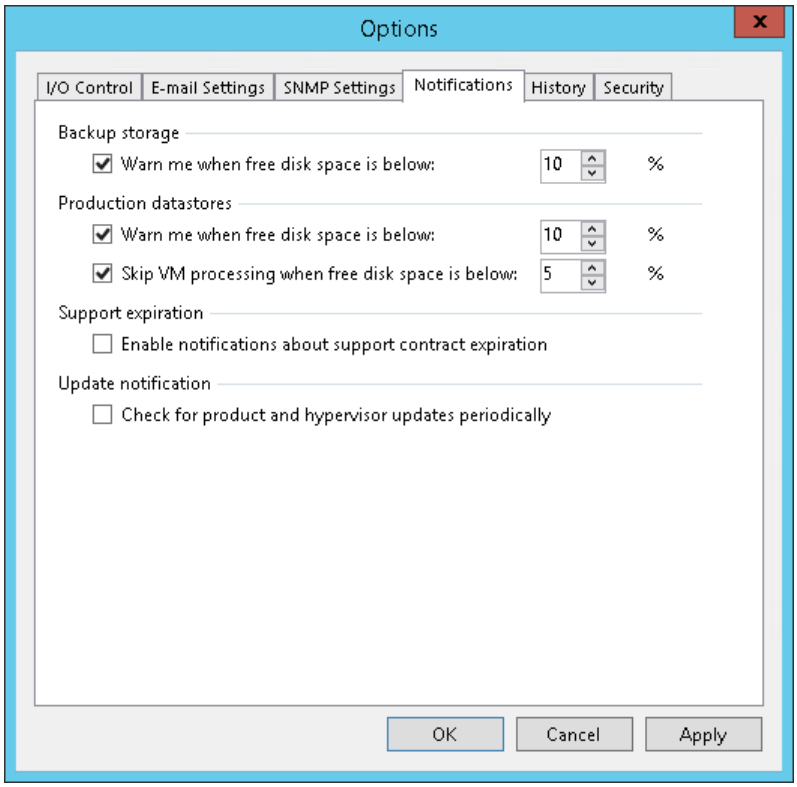
1. If you want Veeam Backup & Replication to periodically check and notify you about product updates, select the **Check for updates once the product is installed and periodically** check box.
2. Click **Install** to begin the installation.
3. Wait for the installation process to complete and click **Finish** to exit the setup wizard.



Step 13. Install Available Updates

It is recommended that you periodically check for Veeam Backup & Replication updates and install them when they are available. Installation of updates lets you make sure that you use the latest version of the product and use its functionality to the full.

You can check for product updates manually or configure Veeam Backup & Replication to automatically notify you about available updates. For more information, see Specifying Other Notification Settings.



TIP

When you install updates for Veeam Backup & Replication, in the update wizard, select the **Update remote components automatically** check box. Veeam Backup & Replication will automatically update its components on all servers added to the backup infrastructure. For more information, see Server Components Upgrade.

Upgrading to Veeam Backup & Replication 11 or 11a

To perform upgrade of Veeam Backup & Replication to version 11 or 11a, you must be running version 9.5 Update 4b (build 9.5.4.2866) or later on the supported operating system (refer to the System Requirements section of this document). To upgrade from previous versions, contact Veeam Customer Support.

Upgrade Checklist

Check the following prerequisites before you upgrade Veeam Backup & Replication:

1. Veeam Backup & Replication 11 or 11a uses the same license file format introduced with version 10, so you can use your existing version 10 license file to install version 11 or any future versions. Your support contract must be active as of the date when the product version you are installing was built.
2. Are you using Veeam Backup Starter? This edition has been discontinued, so Veeam Backup & Replication 11 or 11a will not accept such a license file. Please download a replacement license file from the Customer Portal before upgrading.
3. [For Veeam Backup & Replication 11a (build 11.0.1.1261)] Are you using Server 2019 based ReFS backup repositories? If yes, avoid upgrading them to Server 2022 and/or mounting ReFS volumes from Server 2019 to new Server 2022 installs until you read this thread on Veeam R&D forums. Microsoft has addressed the known regression in the ReFS format upgrade code, the fix is now publicly available.
4. [For Veeam Backup & Replication 11 before 11a (build 11.0.1.1261)] Are you using Cloud Connect? If yes, check with your Cloud Connect service provider if they have already upgraded their system to at least the version you are upgrading to.
5. Are you using Veeam Availability Orchestrator? If yes, mind that Veeam Availability Orchestrator 4.0 is compatible with Veeam Backup & Replication 11, Veeam Availability Orchestrator 5.0 — with Veeam Backup & Replication 11 and 11a. If necessary, upgrade Veeam Availability Orchestrator before upgrading to Veeam Backup & Replication 11 or 11a.
6. Are you using Veeam Backup Enterprise Manager? If yes, start the upgrade procedure with this component. Note that Enterprise Manager 11 supports backup servers version 9.5 Update 4 or later, so you can potentially run both old and new product versions side by side if required.
7. Are you using Veeam ONE to monitor your backup infrastructure? If yes, upgrade it first. Veeam ONE supports monitoring of backup servers version 9.5 Update 4 or later.
8. Are you using Veeam Backup Enterprise Manager server added to Veeam ONE? If yes, first upgrade Veeam ONE, second upgrade Veeam Backup Enterprise Manager, third upgrade Veeam Backup & Replication.
9. Are you using Veeam Backup & Replication within the infrastructure of Nutanix Mine version 2.0.1 or earlier? If yes, upgrade Nutanix Mine to version 3.0, then upgrade Veeam Backup & Replication to version 11a build 11.01.1261 or later.
10. Are you using Veeam Plug-in for Oracle RMAN or Veeam Plug-in for SAP HANA? If yes, you must upgrade Veeam Backup & Replication first, then you can upgrade Veeam Plug-ins. Veeam Backup & Replication 11 or 11a supports Veeam Plug-ins version 11 and 10a (10.0.1.4854).
11. Check if the backup server to be upgraded is installed on the supported operating system version according to System Requirements. If it is not, create a configuration backup, install Veeam Backup & Replication 11 or 11a on the supported OS first, then restore the configuration backup created earlier. For information on how to perform the migration, see this Veeam KB article.

12. Make sure that other servers that you plan to use as backup infrastructure components meet the system requirements listed in the System Requirements section of this document.

13. Make sure that the virtual environment you are going to protect with Veeam Backup & Replication meets the requirements listed in the Platform Support section of this document.

14. Ensure there are no active processes, such as any running jobs and restore sessions. We recommend that you do not stop running jobs and let them complete successfully. Disable any periodic and backup copy jobs, so that they do not start during the upgrade.

15. Make sure that all necessary permissions are granted. For information on permissions, see Required Permissions for VMware vSphere.

16. Do you have backup copy jobs with synthetic GFS full backups? Before the upgrade, make sure that all GFS candidates (incremental restore points created on days when GFS was scheduled and that are expected to be transformed into full GFS restore points) are already transformed into GFS restore points. To force the backup copy job to transform all GFS candidates, you can temporarily decrease the short-term retention to a value less than the number of restore points between the latest restore point and the most recent GFS candidate and then wait till all the candidates are transformed.

    Before Veeam Backup & Replication version 11, Veeam Backup & Replication created GFS candidates on days when GFS was scheduled and only then transformed them into full GFS restore points according to the short-term retention. For more information on how restore points were transformed, see Synthetic Weekly Full Backups. Starting from Veeam Backup & Replication version 11, Veeam Backup & Replication creates GFS restore points according to a new schedule and creates them right on the scheduled days. For more information, see Changes in GFS Retention. After the upgrade, Veeam Backup & Replication no longer transforms previous GFS candidates into full GFS restore points. This means, that all GFS candidates lose their GFS status, they become regular incremental restore points and are deleted according to the short-term retention policy.

17. Are you using backup copy jobs targeted at a backup repository with rotated drives? If yes, mind that after the upgrade, the retention will start taking into account restore points on missing rotated drives. To store the necessary number of restore points on all rotated drives, increase the retention value before the upgrade.

18. [For Veeam Backup & Replication 11a (build 11.0.1.1261)] Are you using integration with Veeam Backup for Microsoft Azure? If yes, after upgrade to 11a, select the existing Microsoft Azure compute account in the Manage Cloud Credentials, click **Edit**, and go through the Microsoft Azure Compute Account to update account permissions. Otherwise, you can face problems when adding an external repository with backups created by Veeam Backup for Microsoft Azure 3.0.

19. Are you using Veeam Agents managed via Veeam Backup & Replication? If you use Veeam Agents v3, they will stop working after upgrading to Veeam Backup & Replication 11 or 11a. In this case, we strongly recommend to immediately upgrade Veeam Agents to v5. If you use Veeam Agents v4, they will continue working after upgrading to Veeam Backup & Replication 11 or 11a, but new features implemented in v11 or v11a will not be supported. In this case, you can upgrade Veeam Agents to v5 later if the support of new features is not critical for you.

Additional considerations for upgrading from version 9.5 Update 4b:

1. Are you using an instance license to protect some of your vSphere or Hyper-V VMs with agent-based backup jobs in presence of a socket license? As an enforcement of the Veeam Licensing Policy, starting from version 10, hosts, where such VMs are running, will consume socket licenses. This may result in your agent-based backup jobs failing after the upgrade due to insufficient socket licenses. Before upgrading to version 11 or 11a, review the **Licensing** section of the What's New in Veeam Backup & Replication 10 document for the detailed description of all changes that happened with version 10 release.

2. The Veeam Backup & Replication 11 or 11a setup wizard will offer to download your version 11 or 11a license automatically. This requires uploading your currently installed license to Veeam servers. If your backup server has no Internet connection, or if you prefer not to have your license uploaded, or in case of license upgrade issues – download your license from the Customer Portal manually. Note that you must have an active maintenance agreement at the time when you access the portal.

Performing Upgrade

To upgrade Veeam Backup & Replication to version 11 or 11a, perform the following steps:

1. Make sure you have checked all the prerequisites listed in the Upgrade Checklist.
2. Download the latest version of the Veeam Backup & Replication ISO from the Veeam Backup & Replication Download page.
3. Make sure the latest run for all existing jobs has completed successfully. Rerun the failed jobs.
4. Ensure there are no running jobs, restore sessions, Instant Recovery sessions and SureBackup jobs. We recommend that you do not stop running jobs and let them complete successfully. Disable any periodic and backup copy jobs temporarily to prevent them from starting during the upgrade.
5. If necessary, perform the configuration backup, as described in Running Configuration Backups Manually.

   Mind that, if you have at least one password created in the Password Manager on the backup server, Veeam Backup & Replication requires that you encrypt the configuration backup. For more information, see Creating Encrypted Configuration Backups.
6. Mount the product ISO and use autorun, or run the `Setup.exe` file.
7. Click the **Upgrade** tile to launch the upgrade wizard.
8. Follow the same steps as described in the Installing Veeam Backup & Replication section. Be sure to select the same SQL database and instance that was used by the previous product version.

   At the **Ready to Install** step of the upgrade wizard, select the **Update remote components automatically** check box to automatically upgrade Veeam Backup & Replication components on all servers added to the backup infrastructure. For more information, see Server Components Upgrade.
9. Wait for the setup program to perform the upgrade.
10. If you use remote backup consoles, upgrade them manually using the product ISO file. Unfortunately, automatic upgrade is not supported this time due to the major version number change.
11. Download and install the latest available update from the Veeam Updates page.
12. Open the Veeam Backup & Replication console. If necessary, the automated upgrade wizard will automatically appear, prompting you to upgrade the product components running on remote servers. Follow the wizard to complete the upgrade process.
13. If some remote servers are unavailable at the time of upgrade, you can run the upgrade wizard at any time later from the main product menu, or by closing and re-opening the Veeam Backup & Replication console. Note that the out-of-date product components cannot be used by jobs until they are updated to the backup server version.
14. If you are using the Virtual Labs functionality, please open settings of each Virtual Lab, and click through the wizard to redeploy each virtual lab with the new proxy appliance version.

15. [For Veeam Backup & Replication 11GA (build 11.0.0.837) with installed cumulative patch P20210319 or later] If you are using Linux servers for your backup infrastructure components, the process of upgrade to 11a will automatically deploy the new persistent data mover only to Linux servers with the VMware Backup Proxy role. To deploy it on other Linux servers, click through the Linux server properties, or use Set-VBRLinux PowerShell cmdlet to mass-deploy. Until you do this, those Linux servers will continue using the legacy run-time data mover to avoid issues with backup repository not meeting the persistent data mover requirements.
16. Enable any scheduled jobs that you have disabled before the upgrade.

Note that immediately after the upgrade, the backup server performance may decrease. This happens due to the maintenance job that optimizes the configuration database. The process may take up to an hour depending on the database size.

> **IMPORTANT**
>
> You must upgrade Veeam components on all remote servers with which the backup server communicates during data protection and disaster recovery tasks. If you do not upgrade components on remote servers, Veeam Backup & Replication jobs will fail. For more information, see Server Components Upgrade.

Unattended Upgrade

To correctly upgrade Veeam Backup & Replication in the unattended mode, perform the following steps:

1. Uninstall the previous version of the product.
2. Install a newer version of the product in the unattended mode. You must connect to the configuration database that was used by the previous product version.

   For information on how to install Veeam Backup & Replication in the unattended mode, see Installing Veeam Backup & Replication in Unattended Mode.

Updating Veeam Backup & Replication 11 or 11a

Apart from major version releases of Veeam Backup & Replication (e.g. *11, 11a*), Veeam Software provides cumulative patches (e.g. *Cumulative Patch P20210525 for v11, Cumulative Patch 20220302 for v11a*). Cumulative patches contain bug fixes, performance enhancements and introduce new features.

Prerequisites

Before you install a cumulative patch for Veeam Backup & Replication 11 or 11a, check the following prerequisites:

- Make sure you have Veeam Backup & Replication 11 (build 11.0.0.837) or 11a (build 11.0.1.1261) of any earlier patch level installed.

  For information on how to upgrade from product version 9.5 Update 4b or later, see Upgrading to Veeam Backup & Replication 11 or 11a.

- Disable all Veeam Backup & Replication jobs and finish all restore processes.
- Close Veeam Backup & Replication processes in the Task Manager.

Performing Update

To install the latest update for Veeam Backup & Replication 11 or 11a, perform the following steps:

1. In case of v11, go to this Veeam KB article. In case of v11a, go to this Veeam KB article.
2. In the **Solution** section of the Veeam KB article, click **DOWNLOAD PATCH**.
3. Run the downloaded file to launch the update wizard.
4. In the update wizard, click **Next**.
5. Select the **Update remote components automatically** and click **Install**.

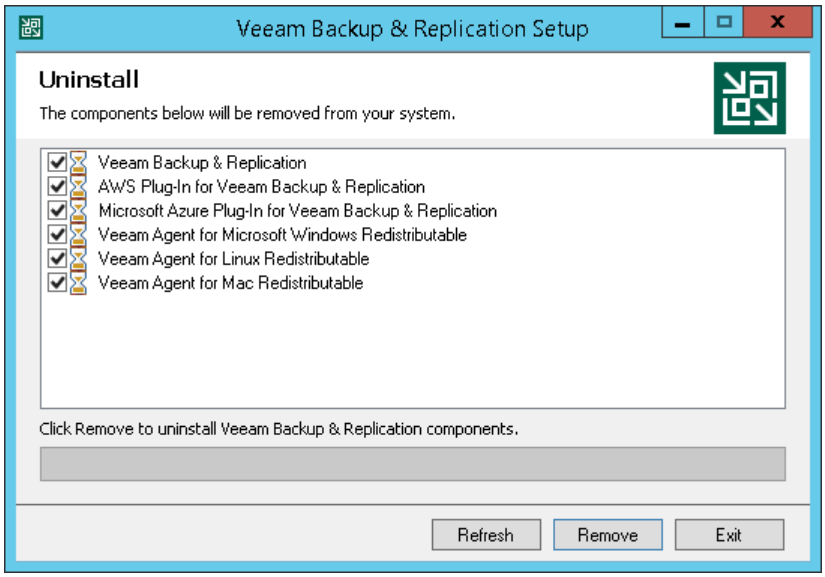
For information on how to update Veeam Backup & Replication in the unattended mode, see Installing Updates in Unattended Mode.

Uninstalling Veeam Backup & Replication

To uninstall Veeam Backup & Replication:

1. From the **Start** menu, select **Control Panel > Programs and Features**.
2. In the programs list, right-click **Veeam Backup & Replication** and select **Uninstall**. If you have Veeam Backup Enterprise Manager installed on this machine, Veeam Backup & Replication will uninstall both components. Wait for the process to complete.
3. If the program list contains additional Veeam Backup & Replication components, right-click the remaining components and select **Uninstall**.

The Veeam Backup & Replication configuration database is not removed during the uninstall process. All configuration data stored in the database remains as well.



Installing Veeam Backup & Replication Console

By default, the Veeam Backup & Replication console is installed on the backup server automatically when you install Veeam Backup & Replication. You do not need to install the console manually.

However, in addition to the default console, you can install the Veeam Backup & Replication console on a dedicated machine to access the backup server remotely. You can install as many remote consoles as you need. For more information, see Backup & Replication Console.

Before you install the Veeam Backup & Replication console, check prerequisites. Then use the **Veeam Backup & Replication Console Setup** wizard to install the console.

Before You Begin

Before you install the Veeam Backup & Replication console, check the following prerequisites:

- The Veeam Backup & Replication console must be of the same version as Veeam Backup & Replication installed on the backup server.
- A machine on which you plan to install the Veeam Backup & Replication console must meet the system requirements. For more information, see System Requirements.
- A user account that you plan to use for installation must have sufficient permissions. For more information, see Required Permissions.
- Backup infrastructure components communicate with each other over specific ports. These ports must be open. For more information, see Used Ports.

Step 1. Start Setup Wizard

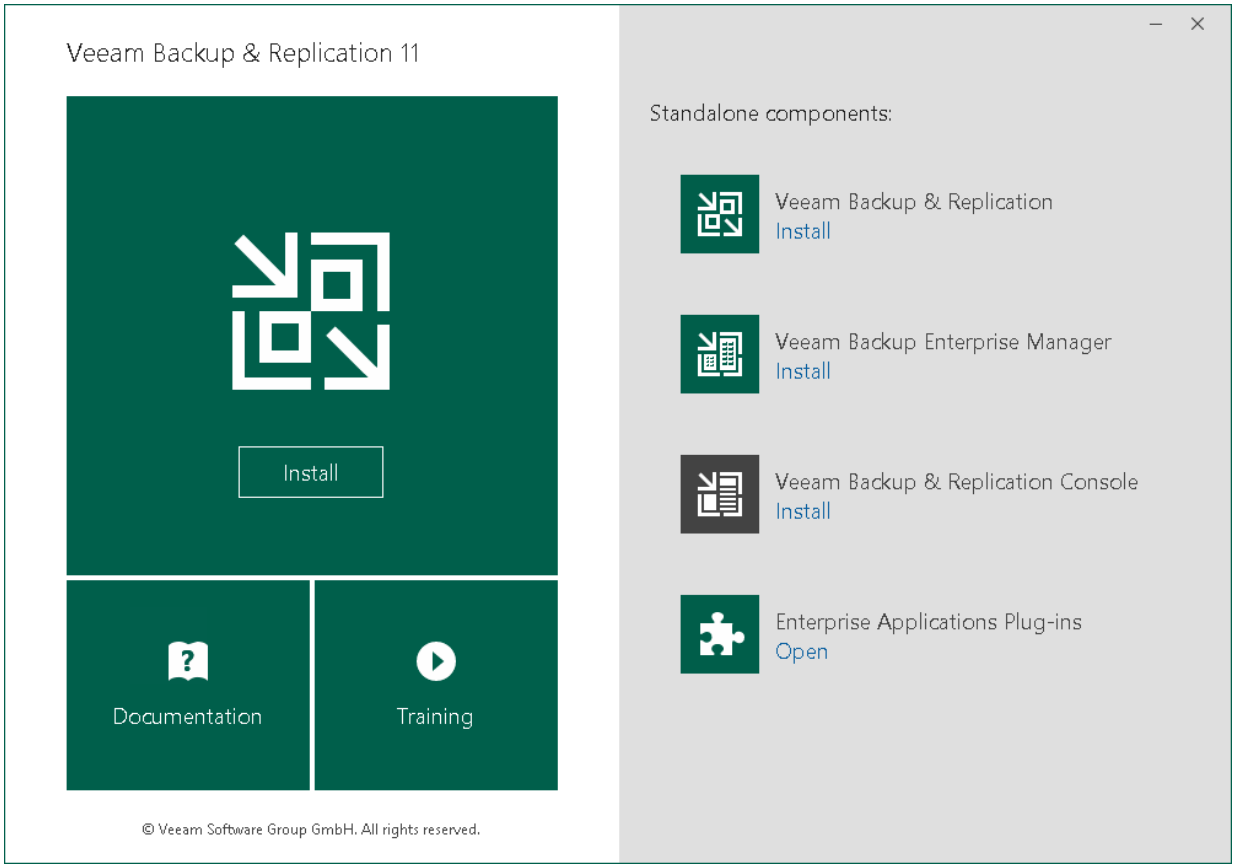
To start the setup wizard:

1. Download the latest version of the Veeam Backup & Replication installation image from www.veeam.com/downloads.html.
2. Use disk image emulation software to mount the installation image to the machine where you plan to install Veeam Backup & Replication or burn the image file to a flash drive or other removable storage device. If you plan to install Veeam Backup & Replication on a VM, use built-in tools of the virtualization management software to mount the installation image to the VM.

   To extract the content of the ISO, you can also use the latest versions of utilities that can properly extract data from ISOs of large size and can properly work with long file paths.
3. After you mount the image or insert the disk, Autorun will open a splash screen with installation options. If Autorun is not available or disabled, run the Setup.exe file from the image or disk.
4. On the splash screen, click **Veeam Backup & Replication Console**.

IMPORTANT

It is strongly recommended that you install the Veeam Backup & Replication console using Autorun or the Setup.exe file. If you run other installation files from the ISO folders, you may miss some components that need to be installed, and Veeam Backup & Replication may not work as expected.



Step 2. Read and Accept License Agreement

At the **License Agreement** step of the wizard, you must accept the license agreement for Veeam and 3rd party components that Veeam incorporates. If you do not accept the license agreement, you will not be able to pass to the next step of the setup wizard.

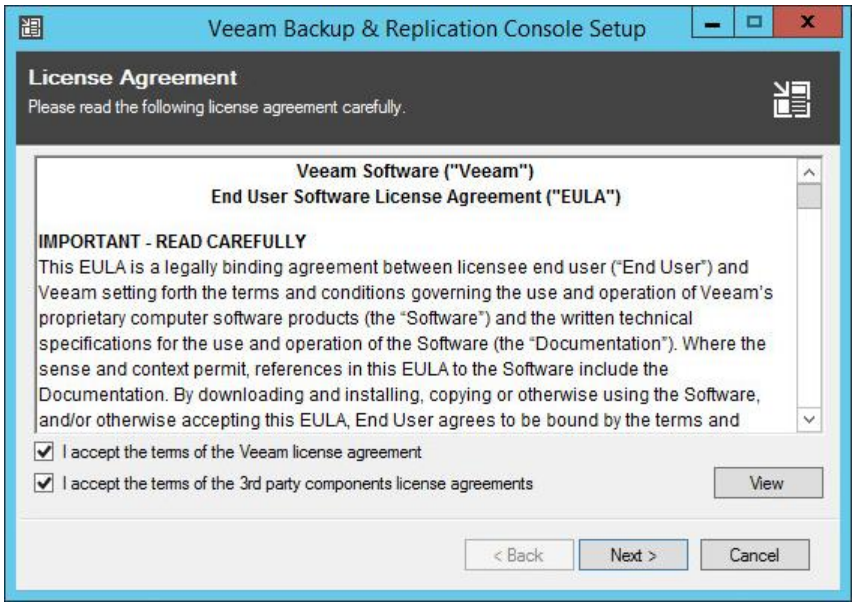
1. Read the license agreement.

   To view the license agreement for 3rd party components, click **View**.

2. Select the **I accept the terms of the Veeam license agreement** check box.

3. Select the **I accept the terms of the 3rd party components license agreement** check box.

4. Click **Next**.

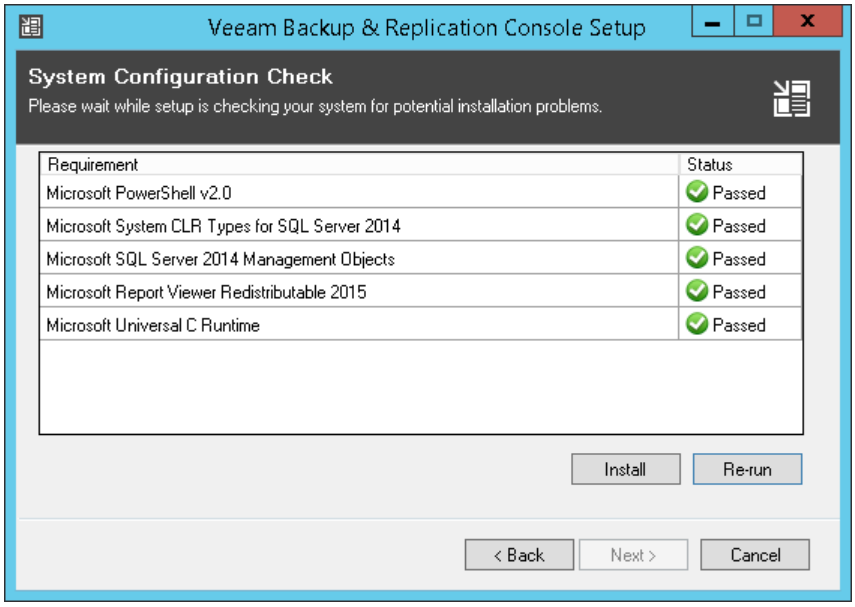


Step 3. Install Missing Software

At the **System Configuration Check** step of the wizard, the setup wizard checks if all prerequisite software is installed on the machine. If required software components are missing, the setup wizard will offer you to install them.

You can install missing components automatically or manually.

- To install missing components automatically, click **Install**. The setup wizard will not interrupt the installation process and install the missing components in the work current session.
- To install missing components manually:
  a. Click **Cancel** and exit the setup wizard.
  b. Install and enable the necessary components manually on the machine.
  c. Start the setup wizard again, pass to the **System Configuration Check** step of the wizard and click **Re-run** to repeat the verification.



Step 4. Specify Installation Settings

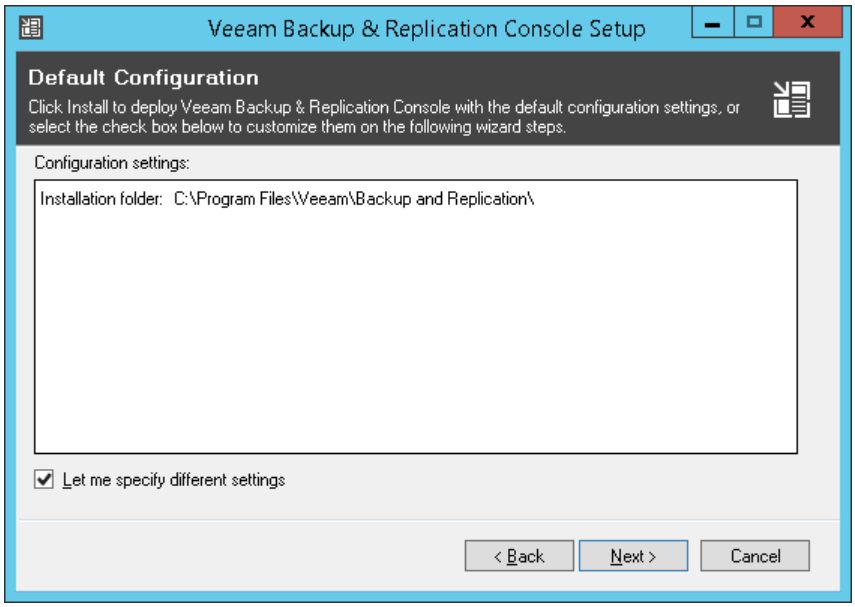
At the **Default Configuration** step of the wizard, you can select to install the Veeam Backup & Replication console with default installation settings or specify custom installation settings.

By default, the setup wizard installs the Veeam Backup & Replication console with the following settings: installation folder — `C:\Program Files\Veeam\Backup and Replication`.

To use default installation settings:

1. Leave the **Let me specify different settings** check box not selected.
2. Click **Install**. The installation process will begin.

To specify custom installation settings, select the **Let me specify different settings** check box. The setup wizard will include additional steps that will let you configure installation settings.

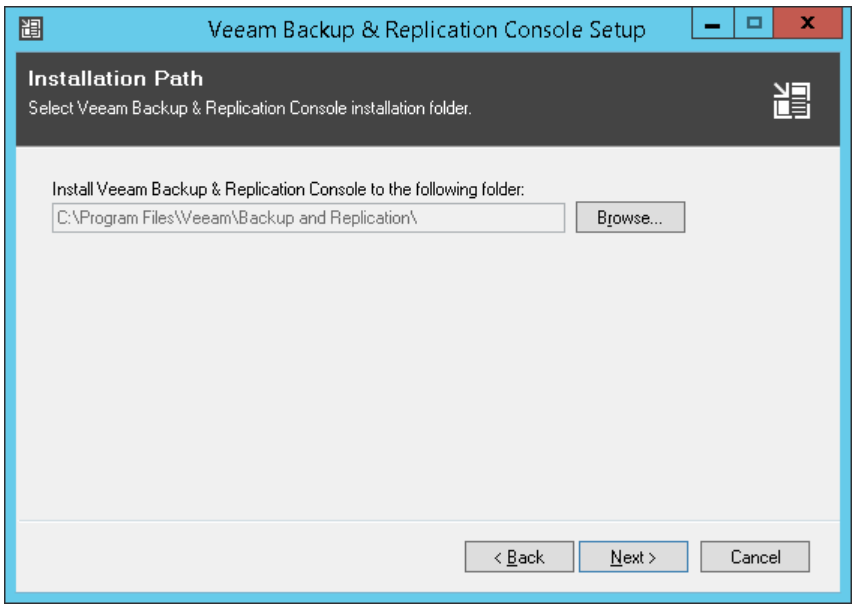


Step 5. Specify Installation Path

The **Installation Path** step of the wizard is available if you have selected to configure installation settings manually.

At the **Installation Path** step of the wizard, you can choose the installation folder for the Veeam Backup & Replication console.

1. On the right of the **Install Veeam Backup & Replication Console to the following folder** field, click **Browse**.
2. In the **Browse for Folder** window, select the installation folder for the product. The default folder is `C:\Program Files\Veeam\Backup and Replication\`.

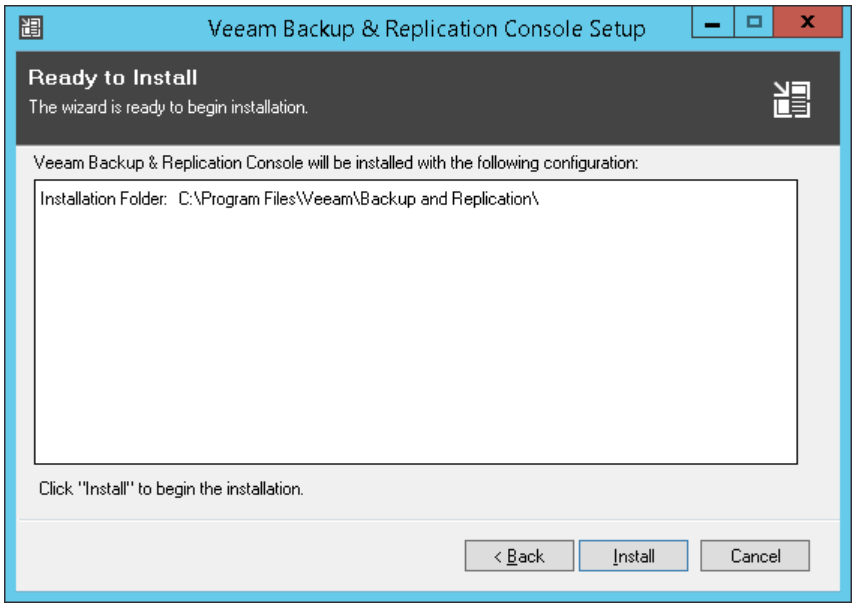


Step 6. Begin Installation

The **Ready to Install** step of the wizard is available if you have selected to configure installation settings manually.

At the **Ready to Install** step of the wizard, you can review the installation settings and start the installation process.

1. Click **Install** to begin the installation.
2. Wait for the installation process to complete and click **Finish** to exit the setup wizard.



Installing Veeam Backup & Replication in Unattended Mode

You can install Veeam Backup & Replication in the unattended mode using the command line interface. The unattended installation mode does not require user interaction. You can use it to automate the installation process in large deployments.

Installation Order

Veeam Backup & Replication components must be installed in the order specified below. The order depends on the type of server that you plan to deploy: backup server or Veeam Backup Enterprise Manager server.

Backup Server

If you want to deploy the backup server (server running Veeam Backup & Replication), you must install components in the following order:

1. Veeam Backup Catalog
2. Veeam Backup & Replication Server
3. Veeam Explorers:
   - Veeam Explorer for Active Directory
   - Veeam Explorer for Exchange
   - Veeam Explorer for Oracle
   - Veeam Explorer for SharePoint and Veeam Explorer for Microsoft OneDrive for Business
   - Veeam Explorer for Microsoft SQL
   - [Optional] Veeam Explorer for Microsoft Teams (required for self-service Veeam Backup for Microsoft Office 365 restores that take place through Cloud Connect Service Provider)
4. [Optional] If you are planning to use the Agent Management feature, install the following components (depending on the management OS):
   - Redistributable package for Veeam Agent for Microsoft Windows
   - Redistributable package for Veeam Agent for Linux
   - Redistributable package for Veeam Agent for Mac

   For more information about Veeam Agents, see the Veeam Agent Management Guide.
5. [Optional] If you are planning to use AWS Plug-in for Veeam Backup & Replication, Microsoft Azure Plug-in for Veeam Backup & Replication or Google Cloud Platform Plug-in for Veeam Backup & Replication, you must install Veeam cloud plug-ins components.
6. [Optional] If you are planning to use Veeam Backup for Nutanix AHV, you must install Veeam Backup for Nutanix AHV components.

Veeam Backup & Replication Console

If you want to deploy the Veeam Backup & Replication console, you must install the Veeam Backup & Replication console component.

Veeam Backup Enterprise Manager Server

If you want to deploy the Veeam Backup Enterprise Manager server (server running Veeam Backup Enterprise Manager), you must install components in the following order:

1. Veeam Backup Catalog
2. Veeam Backup Enterprise Manager

Veeam Cloud Connect Portal

If you want to deploy Veeam Cloud Connect Portal, you must install components in the following order:

1. Veeam Backup Enterprise Manager
2. Veeam Cloud Connect Portal

Before You Begin

Before you start unattended installation, perform the following steps:

1. [For backup server] Manually install Microsoft SQL Server and software components required for the backup server operation. For more information, see **Backup Server** in System Requirements.
2. [For Veeam Explorers] Make sure that the version of a Veeam Explorer that you plan to install matches the version of the Veeam Backup & Replication console on the target machine.
3. [For Veeam Backup Enterprise Manager server] Manually install Microsoft SQL Server and software components required for the operation of the Veeam Backup Enterprise Manager server. For more information, see the System Requirements section of the Enterprise Manager User Guide.
4. Download the Veeam Backup & Replication installation image from the Veeam website. You can burn the downloaded image to a flash drive or mount the image to the target machine using disk image emulation software.
5. Log on to the target machine under the account that has the local Administrator permissions on the machine. For more information, see Required Permissions.
6. Obtain a license file. The license file is required for Veeam Backup Enterprise Manager installation and is optional for Veeam Backup & Replication installation. If you do not specify a path to the license file during Veeam Backup & Replication installation, Veeam Backup & Replication will operate in the Community Edition mode.

Installation Command-Line Syntax

You can install the following Veeam Backup & Replication components in the unattended mode:

- Veeam Backup Catalog
- Veeam Backup & Replication Server
- Veeam Backup & Replication Console
- Veeam Explorer for Microsoft Active Directory
- Veeam Explorer for Microsoft Exchange
- Veeam Explorer for Microsoft SharePoint and Veeam Explorer for Microsoft OneDrive for Business
- Veeam Explorer for Microsoft SQL Server
- Veeam Explorer for Microsoft Teams
- Veeam Explorer for Oracle
- Redistributable Package for Veeam Agent for Linux
- Redistributable Package for Veeam Agent for Mac
- Redistributable Package for Veeam Agent for Microsoft Windows
- Veeam Backup Enterprise Manager
- Veeam Cloud Connect Portal
- Veeam Cloud Plug-Ins
- Veeam Backup for Nutanix AHV

Veeam Backup Catalog

To install Veeam Backup Catalog, use a command with the following syntax:

```
msiexec.exe [/L*v "<path_to_log>"] /qn /i "<path_to_msi>" ACCEPTEULA="1" ACCEPT
_THIRDPARTY_LICENSES="1" [INSTALLDIR="<path_to_installdir >"][VM_CATALOGPATH="<
path_to_catalog_shared_folder>"][VBRC_SERVICE_USER="<Veeam_Guest_Catalog_Servic
e_account>"][VBRC_SERVICE_PASSWORD="<Veeam_Guest_Catalog_Service_account_passwo
rd>"] [VBRC_SERVICE_PORT="<Veeam_Guest_Catalog_Service_port>"]
```

The command has the following parameters:

Option	Parameter	Required	Description
/L	*v logfile	No	Creates an installation log file with the verbose output. Specify a full path to the log file as the parameter value. A setup log file created during the previous installation is cleared. Example: `/L*v "C:\ProgramData\Veeam\Setup\Temp\Logs\Catalog.txt"`
/q	n	Yes	Sets the user interface level to "no", which means no user interaction is needed during installation.
/i	setup file	Yes	Installs Veeam Backup Catalog. Specify a full path to the setup file as the parameter value. Example: `/i "C:\Veeam\VeeamBackupCatalog64.msi"`
ACCEPTEULA	0/1	Yes	Specifies if you want to accept the Veeam license agreement. Specify `1` to accept the license agreement and proceed with installation. Example: `ACCEPTEULA="1"`
ACCEPT_THIRDPARTY_LICENSES	0/1	Yes	Specifies if you want to accept the license agreement for 3rd party components that Veeam incorporates. Specify `1` to accept the license agreement and proceed with installation. Example: `ACCEPT_THIRDPARTY_LICENSES="1"`
INSTALLDIR	path	No	Installs the component to the specified location. By default, Veeam Backup & Replication uses the `Backup Catalog` subfolder in the `C:\Program Files\Veeam\Backup and Replication\` folder. Example: `INSTALLDIR="C:\Catalog\"` The component will be installed to the `C:\Catalog\Backup Catalog` folder.

Option	Parameter	Required	Description
VM_CATALOGPATH	path	No	Specifies a path to the catalog folder where index files must be stored. By default, Veeam Backup & Replication creates the `VBRCatalog` folder on a volume with the maximum amount of free space, for example `C:\VBRCatalog`. Example: `VM_CATALOGPATH="C:\Backup\"` Index files will be stored to the `C:\Backup\VBRCatalog` folder.
VBRC_SERVICE_USER	user	No	Specifies a user account under which the Veeam Guest Catalog Service will run. The account must have full control NTFS permissions on the `VBRCatalog` folder where index files are stored. If you do not specify this parameter, the Veeam Guest Catalog Service will run under the LocalSystem account. Together with the `VBRC_SERVICE_USER` parameter, you must specify the `VBRC_SERVICE_PASSWORD` parameter. Example: `VBRC_SERVICE_USER="BACKUPSERVER\Administrator"`
VBRC_SERVICE_PASSWORD	password	No	This parameter must be used if you have specified the `VBRC_SERVICE_USER` parameter. Specifies a password for the account under which the Veeam Guest Catalog Service will run. Example: `VBRC_SERVICE_PASSWORD="1234"`
VBRC_SERVICE_PORT	port	No	Specifies a TCP port that will be used by the Veeam Guest Catalog Service. By default, port number 9393 is used. Example: `VBRC_SERVICE_PORT="9393"`

Example

Suppose you want to install Veeam Backup Catalog with the following configuration:

- No user interaction
- Path to the MSI file: `E:\Veeam\VeeamBackupCatalog64.msi`
- Installation folder: default
- Catalog folder: default

- Service user account: `VEEAM\Administrator`
- Service user account password: `1243`
- TCP communication port: `9391`

The command to install Veeam Backup Catalog with such configuration will have the following parameters:

```
msiexec.exe /qn /i "E:\Veeam\VeeamBackupCatalog64.msi" ACCEPTEULA="1" ACCEPT_TH
IRDPARTY_LICENSES="1" VBRC_SERVICE_USER="VEEAM\Administrator" VBRC_SERVICE_PASS
WORD="1234" VBRC_SERVICE_PORT="9391"
```

Veeam Backup & Replication Server

To install the Veeam Backup & Replication server, use a command with the following syntax:

```
msiexec.exe [/L*v "<path_to_log>"] /qn /i "<path_to_msi>" ACCEPT_EULA="1" ACCEP
T_THIRDPARTY_LICENSES="1" [INSTALLDIR="<path_to_installdir >"] [VBR_LICENSE_FIL
E="<path_to_license_file>"] [VBR_SERVICE_USER="<Veeam_B&R_Service_account>"] [V
BR_SERVICE_PASSWORD="<Veeam_B&R_Service_account_password>"] [VBR_SERVICE_PORT="
<Veeam_B&R_Service_port>"] [VBR_SECURE_CONNECTIONS_PORT="<SSL_port>"] [VBR_SQLS
ERVER_SERVER="<SQL_server>"] [VBR_SQLSERVER_DATABASE="<database_name>"][VBR_SQL
SERVER_AUTHENTICATION="0"] [VBR_SQLSERVER_USERNAME="<SQL_auth_username>"] [VBR_
SQLSERVER_PASSWORD="<SQL_auth_password>"] [VBR_IRCACHE="<path_to_instant_recove
ry_cache_folder>"] [VBR_CHECK_UPDATES="1"][VBR_AUTO_UPGRADE="1"]
```

NOTE

This command does not install the Veeam Backup & Replication console and the Veeam PowerShell module. To install the Veeam Backup & Replication console, run the command specified in Veeam Backup & Replication Console. The Veeam PowerShell module will be installed with the Veeam Backup & Replication console.

The command has the following parameters:

Option	Parameter	Required	Description
/L	*v logfile	No	Creates an installation log file with the verbose output. Specify a full path to the log file as the parameter value. A setup log file created during the previous installation is cleared. Example: `/L*v "C:\ProgramData\Veeam\Setup\Temp\Logs\Backup.txt"`
/q	n	Yes	Sets the user interface level to "no", which means no user interaction is needed during installation.

Option	Parameter	Required	Description
/i	setup file	Yes	Installs the Veeam Backup & Replication server. Specify a full path to the setup file as the parameter value. Example: `/i "E:\Backup\Server.x64.msi"`
ACCEPT_EULA	0/1	Yes	Specifies if you want to accept the Veeam license agreement. Specify `1` to accept the license agreement and proceed with installation. Example: `ACCEPT_EULA="1"`
ACCEPT_THIRDPARTY_LICENSES	0/1	Yes	Specifies if you want to accept the license agreement for 3rd party components that Veeam incorporates. Specify `1` to accept the license agreement and proceed with installation. Example: `ACCEPT_THIRDPARTY_LICENSES="1"`
INSTALLDIR	path	No	Installs the component to the specified location. By default, Veeam Backup & Replication uses the `Backup` subfolder of the `C:\Program Files\Veeam\Backup and Replication\` folder. Example: `INSTALLDIR="C:\Backup\"` The component will be installed to the `C:\Backup\Backup` folder.
VBR_LICENSE_FILE	license path	No	Specifies a full path to the license file. If you do not specify this parameter, Veeam Backup & Replication will operate in the Community Edition mode. Example: `VBR_LICENSE_FILE="C:\Users\Administrator\Desktop\enterprise - veeam_backup_trial_0_30.lic"`

Option	Parameter	Required	Description
VBR_SERVICE_USER	user	No	Specifies the account under which the Veeam Backup Service will run. The account must have full control NTFS permissions on the `VBRCatalog` folder where index files are stored and the *Database owner* rights for the configuration database on the Microsoft SQL Server where the configuration database is deployed. If you do not specify this parameter, the Veeam Backup Service will run under the LocalSystem account. Together with the `VBR_SERVICE_USER` parameter, you must specify the `VBR_SERVICE_PASSWORD` parameter. Example: `VBR_SERVICE_USER="BACKUPSERVER\Administrator"`
VBR_SERVICE_PASSWORD	password	No	This parameter must be used if you have specified the `VBR_SERVICE_USER` parameter. Specifies a password for the account under which the Veeam Backup Service will run. Example: `VBR_SERVICE_PASSWORD="1234"`
VBR_SERVICE_PORT	port	No	Specifies a TCP port that will be used by the Veeam Backup Service. By default, the port number 9392 is used. Example: `VBR_SERVICE_PORT="9395"`
VBR_SECURE_CONNECTIONS_PORT	port	No	Specifies a port used for communication between the mount server and the backup server. By default, port 9401 is used. Example: `VBR_SECURE_CONNECTIONS_PORT="9402"`
VBR_SQLSERVER_SERVER	SQL server\instance	No	Specifies a Microsoft SQL server and instance on which the configuration database will be deployed. By default, Veeam Backup & Replication uses the `(local)\VEEAMSQL2012` server for machines running Microsoft Windows 7 or Microsoft Windows Server 2008 R2, and `(local)\VEEAMSQL2016` for machines running Microsoft Windows Server 2012 or later. Example: `VBR_SQLSERVER_SERVER="BACKUPSERVER\VEEAMSQL2016_MY"`

Option	Parameter	Required	Description
VBR_SQLSERVER_DATABASE	database	No	Specifies a name for the configuration database. By default, the configuration database is deployed with the `VeeamBackup` name. Example: `VBR_SQLSERVER_DATABASE="VeeamBackup"`
VBR_SQLSERVER_AUTHENTICATION	0/1	No	Specifies if you want to use the SQL Server authentication mode to connect to the Microsoft SQL Server where the Veeam Backup & Replication configuration database is deployed. Specify `1` if you want to use the SQL Server authentication mode. If you do not specify this parameter, Veeam Backup & Replication will connect to the Microsoft SQL Server in the Microsoft Windows authentication mode (default value is `0`). Together with this parameter, you must specify the following parameters: `VBR_SQLSERVER_USERNAME` and `VBR_SQLSERVER_PASSWORD`. Example: `VBR_SQLSERVER_AUTHENTICATION="1"`
VBR_SQLSERVER_USERNAME	user	No	This parameter must be used if you have specified the `VBR_SQLSERVER_AUTHENTICATION` parameter. Specifies a LoginID to connect to the Microsoft SQL Server in the SQL Server authentication mode. Example: `VBR_SQLSERVER_USERNAME="sa"`
VBR_SQLSERVER_PASSWORD	password	No	This parameter must be used if you have specified the `VBR_SQLSERVER_AUTHENTICATION` parameter. Specifies a password to connect to the Microsoft SQL Server in the SQL Server authentication mode. Example: `VBR_SQLSERVER_PASSWORD="1234"`
VBR_IRCACHE	path	No	Specifies the folder to which the instant recovery cache will be stored. By default, Veeam Backup & Replication uses the folder on a volume with the maximum amount of free space, for example, `C:\ProgramData\Veeam\Backup\IRCache\`. Example: `VBR_IRCACHE="C:\ProgramData\Veeam\Backup\IRCache2\"`

Option	Parameter	Required	Description
VBR_CHECK_UPDATES	0/1	No	Specifies if you want Veeam Backup & Replication to automatically check for new product versions and updates. Specify `0` if you do not want to check for updates. If you do not specify this parameter, Veeam Backup & Replication will automatically check for updates (default value is `1`). Example: `VBR_CHECK_UPDATES="0"`
VBR_AUTO_UPGRADE	No/Yes	No	Specifies if you want Veeam Backup & Replication to automatically upgrade existing components in the backup infrastructure. Veeam Backup & Replication performs automatic upgrade after the Veeam Backup Service is started on the backup server. Specify `1` to enable automatic upgrade. Example: `VBR_AUTO_UPGRADE="1"`

Example

Suppose you want to install Veeam Backup & Replication with the following configuration:

- Installation log location: `C:\logs\log1.txt`
- No user interaction
- Path to the MSI file: `E:\Backup\Server.x64.msi`
- Installation folder: `D:\Program Files\Veeam`
- License file location: `C:\License\veeam_license.lic`
- Service user account: `VEEAM\Administrator`
- Service user account password: `1243`
- Service port: default
- TLS port: default
- Configuration database and database name: default
- Path to the instant recovery cache folder: `D:\IRCache`

The command to install Veeam Backup & Replication with such configuration will have the following parameters:

```
msiexec.exe /L*v "C:\logs\log1.txt" /qn /i "E:\Backup\Server.x64.msi" ACCEPT_EU
LA="1" ACCEPT_THIRDPARTY_LICENSES="1" INSTALLDIR="D:\Program Files\Veeam" VBR_L
ICENSE_FILE="C:\License\veeam_license.lic" VBR_SERVICE_USER="VEEAM\Administrato
r" VBR_SERVICE_PASSWORD="1234" VBR_IRCACHE="D:\IRCache"
```

Veeam Backup & Replication Console

To install the Veeam Backup & Replication console, use a command with the following syntax:

```
msiexec.exe [/L*v "<path_to_log>"] /qn /i "<path_to_msi>" ACCEPT_EULA="1" ACCEP
T_THIRDPARTY_LICENSES="1" [INSTALLDIR="<path_to_installdir >"]
```

NOTE

This command also installs the Veeam PowerShell module.

The command has the following parameters:

Option	Parameter	Required	Description
/L	*v logfile	No	Creates an installation log file with the verbose output. Specify a full path to the log file as the parameter value. A setup log file created during the previous installation is cleared. Example: `/L*v "C:\ProgramData\Veeam\Setup\Temp\Logs\Console.txt"`
/q	n	Yes	Sets the user interface level to "no", which means no user interaction is needed during installation.
/i	setup file	Yes	Installs the Veeam Backup & Replication console. Specify a full path to the setup file as the parameter value. Example: `/i "E:\Backup\Shell.x64.msi"`
ACCEPT_EULA	0/1	Yes	Specifies if you want to accept the Veeam license agreement. Specify `1` to accept the license agreement and proceed with installation. Example: `ACCEPT_EULA="1"`

Option	Parameter	Required	Description
ACCEPT_THIRDPARTY_LICENSES	0/1	Yes	Specifies if you want to accept the license agreement for 3rd party components that Veeam incorporates. Specify `1` to accept the license agreement and proceed with installation. Example: `ACCEPT_THIRDPARTY_LICENSES="1"`
INSTALLDIR	path	No	Installs the component to the specified location. By default, Veeam Backup & Replication uses the `Console` subfolder of the `C:\Program Files\Veeam\Backup and Replication\` folder. Example: `INSTALLDIR="C:\Backup\"` The component will be installed to the `C:\Backup\Console` folder.

Example

Suppose you want to install the Veeam Backup & Replication console with the following configuration:

- No user interaction
- Path to the MSI file: `E:\Backup\Shell.x64.msi`
- Installation folder: `C:\Backup`

The command to install the Veeam Backup & Replication console with such configuration will have the following parameters:

```
msiexec.exe /L*v "C:\logs\log1.txt" /qn /i "E:\Backup\Shell.x64.msi" ACCEPT_EUL
A="1" ACCEPT_THIRDPARTY_LICENSES="1" INSTALLDIR="C:\Backup\"
```

Veeam Explorer for Microsoft Active Directory

To install Veeam Explorer for Microsoft Active Directory, use a command with the following syntax:

```
msiexec.exe [/L*v "<path_to_log>"] /qn /i "<path_to_msi>" ACCEPT_EULA="1" ACCEP
T_THIRDPARTY_LICENSES="1"
```

The command has the following parameters:

Option	Parameter	Required	Description
/L	*v logfile	No	Creates an installation log file with the verbose output. Specify a full path to the log file as the parameter value. A setup log file created during the previous installation is cleared. Example: `/L*v "C:\ProgramData\Veeam\Setup\Temp\Logs\VEAD.txt"`
/q	n	Yes	Sets the user interface level to "no", which means no user interaction is needed during installation.
/i	setup file	Yes	Installs Veeam Explorer for Microsoft Active Directory. Specify a full path to the setup file as the parameter value. Example: `/i "C:\Explorers\VeeamExplorerforActiveDirectory.msi"`
ACCEPT_EULA	0/1	Yes	Specifies if you want to accept the Veeam license agreement. Specify `1` to accept the license agreement and proceed with installation. Example: `ACCEPT_EULA="1"`
ACCEPT_THIRDPARTY_LICENSES	0/1	Yes	Specifies if you want to accept the license agreement for 3rd party components that Veeam incorporates. Specify `1` to accept the license agreement and proceed with installation. Example: `ACCEPT_THIRDPARTY_LICENSES="1"`

Veeam Explorer for Microsoft Exchange

To install Veeam Explorer for Microsoft Exchange, use a command with the following syntax:

```
msiexec.exe [/L*v "<path_to_log>"] /qn /i "<path_to_msi>" ACCEPT_EULA="1" ACCEPT_THIRDPARTY_LICENSES="1"
```

The command has the following parameters:

Option	Parameter	Required	Description
/L	*v logfile	No	Creates an installation log file with the verbose output. Specify a full path to the log file as the parameter value. A setup log file created during the previous installation is cleared. Example: `/L*v "C:\ProgramData\Veeam\Setup\Temp\Logs\VEX.txt"`
/q	n	Yes	Sets the user interface level to "no", which means no user interaction is needed during installation.
/i	setup file	Yes	Installs Veeam Explorer for Microsoft Exchange. Specify a full path to the setup file as the parameter value. Example: `/i "C:\Explorers\VeeamExplorerforExchange.msi"`
ACCEPT_EULA	0/1	Yes	Specifies if you want to accept the Veeam license agreement. Specify `1` to accept the license agreement and proceed with installation. Example: `ACCEPT_EULA="1"`
ACCEPT_THIRDPARTY_LICENSES	0/1	Yes	Specifies if you want to accept the license agreement for 3rd party components that Veeam incorporates. Specify `1` to accept the license agreement and proceed with installation. Example: `ACCEPT_THIRDPARTY_LICENSES="1"`

Veeam Explorer for Microsoft SharePoint and Veeam Explorer for Microsoft OneDrive for Business

Veeam Explorer for Microsoft SharePoint is installed together with Veeam Explorer for Microsoft OneDrive for Business from the same setup file.

To install Veeam Explorer for Microsoft SharePoint and Veeam Explorer for Microsoft OneDrive for Business, use a command with the following syntax:

```
msiexec.exe [/L*v "<path_to_log>"] /qn /i "<path_to_msi>" ACCEPT_EULA="1" ACCEP
T_THIRDPARTY_LICENSES="1"
```

The command has the following parameters:

Option	Parameter	Required	Description
/L	*v logfile	No	Creates an installation log file with the verbose output. Specify a full path to the log file as the parameter value. A setup log file created during the previous installation is cleared. Example: `/L*v "C:\ProgramData\Veeam\Setup\Temp\Logs\VESP.txt"`
/q	n	Yes	Sets the user interface level to "no", which means no user interaction is needed during installation.
/i	setup file	Yes	Installs Veeam Explorer for Microsoft SharePoint and Veeam Explorer for Microsoft OneDrive for Business. Specify a full path to the setup file as the parameter value. Veeam Backup & Replication installs both Veeam Explorer for Microsoft SharePoint and Veeam Explorer for Microsoft OneDrive for Business from the same setup file. Example: `/i "C:\Explorers\VeeamExplorerforSharePoint.msi"`
ACCEPT_EULA	0/1	Yes	Specifies if you want to accept the Veeam license agreement. Specify `1` to accept the license agreement and proceed with installation. Example: `ACCEPT_EULA="1"`
ACCEPT_THIRDPARTY_LICENSES	0/1	Yes	Specifies if you want to accept the license agreement for 3rd party components that Veeam incorporates. Specify `1` to accept the license agreement and proceed with installation. Example: `ACCEPT_THIRDPARTY_LICENSES="1"`

Veeam Explorer for Microsoft SQL Server

To install Veeam Explorer for Microsoft SQL Server, use a command with the following syntax:

```
msiexec.exe [/L*v "<path_to_log>"] /qn /i "<path_to_msi>" ACCEPT_EULA="1" ACCEPT_THIRDPARTY_LICENSES="1"
```

The command has the following parameters:

Option	Parameter	Required	Description
/L	*v logfile	No	Creates an installation log file with the verbose output. Specify a full path to the log file as the parameter value. A setup log file created during the previous installation is cleared. Example: `/L*v "C:\ProgramData\Veeam\Setup\Temp\Logs\VESQL.txt"`
/q	n	Yes	Sets the user interface level to "no", which means no user interaction is needed during installation.
/i	setup file	Yes	Installs Veeam Explorer for Microsoft SQL Server. Specify a full path to the setup file as the parameter value. Example: `/i "C:\Explorers\VeeamExplorerforSQL.msi"`
ACCEPT_EULA	0/1	Yes	Specifies if you want to accept the Veeam license agreement. Specify `1` to accept the license agreement and proceed with installation. Example: `ACCEPT_EULA="1"`
ACCEPT_THIRDPARTY_LICENSES	0/1	Yes	Specifies if you want to accept the license agreement for 3rd party components that Veeam incorporates. Specify `1` to accept the license agreement and proceed with installation. Example: `ACCEPT_THIRDPARTY_LICENSES="1"`

Veeam Explorer for Microsoft Teams

To install Veeam Explorer for Microsoft Teams, use a command with the following syntax:

```
msiexec.exe /L*v "C:\ProgramData\Veeam\Setup\Temp\Logs\VET.txt" /qn /i "F:\Expl
orers\VeeamExplorerForTeams.msi" ACCEPT_EULA="1" ACCEPT_THIRDPARTY_LICENSES="1"
```

The command has the following parameters:

Option	Parameter	Required	Description
/L	*v logfile	No	Creates an installation log file with the verbose output. Specify a full path to the log file as the parameter value. A setup log file created during the previous installation is cleared. Example: `/L*v "C:\ProgramData\Veeam\Setup\Temp\Logs\VET.txt"`
/q	n	Yes	Sets the user interface level to "no", which means no user interaction is needed during installation.
/i	setup file	Yes	Installs Veeam Explorer for Microsoft Teams. Specify a full path to the setup file as the parameter value. Example: `/i "C:\Explorers\VeeamExplorerForTeams.msi"`
ACCEPT_EULA	0/1	Yes	Specifies if you want to accept the Veeam license agreement. Specify `1` to accept the license agreement and proceed with installation. Example: `ACCEPT_EULA="1"`
ACCEPT_THIRDPARTY_LICENSES	0/1	Yes	Specifies if you want to accept the license agreement for 3rd party components that Veeam incorporates. Specify `1` to accept the license agreement and proceed with installation. Example: `ACCEPT_THIRDPARTY_LICENSES="1"`

Veeam Explorer for Oracle

To install Veeam Explorer for Oracle, use a command with the following syntax:

```
msiexec.exe [/L*v "<path_to_log>"] /qn /i "<path_to_msi>" ACCEPT_EULA="1" ACCEP
T_THIRDPARTY_LICENSES="1"
```

The command has the following parameters:

Option	Parameter	Required	Description
/L	*v logfile	No	Creates an installation log file with the verbose output. Specify a full path to the log file as the parameter value. A setup log file created during the previous installation is cleared. Example: `/L*v "C:\ProgramData\Veeam\Setup\Temp\Logs\VEO.txt"`
/q	n	Yes	Sets the user interface level to "no", which means no user interaction is needed during installation.
/i	setup file	Yes	Installs Veeam Explorer for Oracle. Specify a full path to the setup file as the parameter value. Example: `/i "C:\Explorers\VeeamExplorerforOracle.msi"`
ACCEPT_EULA	0/1	Yes	Specifies if you want to accept the Veeam license agreement. Specify `1` to accept the license agreement and proceed with installation. Example: `ACCEPT_EULA="1"`
ACCEPT_THIRDPARTY_LICENSES	0/1	Yes	Specifies if you want to accept the license agreement for 3rd party components that Veeam incorporates. Specify `1` to accept the license agreement and proceed with installation. Example: `ACCEPT_THIRDPARTY_LICENSES="1"`

Redistributable Package for Veeam Agent for Linux

To install the redistributable package for Veeam Agent for Linux, use a command with the following syntax:

```
msiexec.exe [/L*v "<path_to_log>"] /qn /i "<path_to_msi>" ACCEPT_EULA="1" ACCEP
T_THIRDPARTY_LICENSES="1"
```

The command has the following parameters:

Option	Parameter	Required	Description
/L	*v logfile	No	Creates an installation log file with the verbose output. Specify a full path to the log file as the parameter value. A setup log file created during the previous installation is cleared. Example: `/L*v "C:\ProgramData\Veeam\Setup\Temp\Logs\VAL.txt"`
/q	n	Yes	Sets the user interface level to "no", which means no user interaction is needed during installation.
/i	setup file	Yes	Installs the redistributable package for Veeam Agent for Linux. Specify a full path to the setup file as the parameter value. Example: `/i "F:\Packages\VALRedist.msi"`
ACCEPT_EULA	0/1	Yes	Specifies if you want to accept the Veeam license agreement. Specify `1` to accept the license agreement and proceed with installation. Example: `ACCEPT_EULA="1"`
ACCEPT_THIRDPARTY_LICENSES	0/1	Yes	Specifies if you want to accept the license agreement for 3rd party components that Veeam incorporates. Specify `1` to accept the license agreement and proceed with installation. Example: `ACCEPT_THIRDPARTY_LICENSES="1"`

Redistributable Package for Veeam Agent for Mac

To install the redistributable package for Veeam Agent for Mac, use a command with the following syntax:

```
msiexec.exe [/L*v "<path_to_log>"] /qn /i "<path_to_msi>" ACCEPT_EULA="1" ACCEPT_THIRDPARTY_LICENSES="1"
```

The command has the following parameters:

Option	Parameter	Required	Description
/L	*v logfile	No	Creates an installation log file with the verbose output. Specify a full path to the log file as the parameter value. A setup log file created during the previous installation is cleared. Example: `/L*v "C:\ProgramData\Veeam\Setup\Temp\Logs\VAM.txt"`
/q	n	Yes	Sets the user interface level to "no", which means no user interaction is needed during installation.
/i	setup file	Yes	Installs the redistributable package for Veeam Agent for Mac. Specify a full path to the setup file as the parameter value. Example: `/i "F:\Packages\VAMRedist.msi"`
ACCEPT_EULA	0/1	Yes	Specifies if you want to accept the Veeam license agreement. Specify `1` to accept the license agreement and proceed with installation. Example: `ACCEPT_EULA="1"`
ACCEPT_THIRDPARTY_LICENSES	0/1	Yes	Specifies if you want to accept the license agreement for 3rd party components that Veeam incorporates. Specify `1` to accept the license agreement and proceed with installation. Example: `ACCEPT_THIRDPARTY_LICENSES="1"`

Redistributable Package for Veeam Agent for Microsoft Windows

To install the redistributable package for Veeam Agent for Microsoft Windows, use a command with the following syntax:

```
msiexec.exe [/L*v "<path_to_log>"] /qn /i "<path_to_msi>" ACCEPT_EULA="1" ACCEPT_THIRDPARTY_LICENSES="1"
```

The command has the following parameters:

Option	Parameter	Required	Description
/L	*v logfile	No	Creates an installation log file with the verbose output. Specify a full path to the log file as the parameter value. A setup log file created during the previous installation is cleared. Example: `/L*v "C:\ProgramData\Veeam\Setup\Temp\Logs\VAW.txt"`
/q	n	Yes	Sets the user interface level to "no", which means no user interaction is needed during installation.
/i	setup file	Yes	Installs the redistributable package for Veeam Agent for Microsoft Windows. Specify a full path to the setup file as the parameter value. Example: `/i "F:\Packages\VAWRedist.msi"`
ACCEPT_EULA	0/1	Yes	Specifies if you want to accept the Veeam license agreement. Specify `1` to accept the license agreement and proceed with installation. Example: `ACCEPT_EULA="1"`
ACCEPT_THIRDPARTY_LICENSES	0/1	Yes	Specifies if you want to accept the license agreement for 3rd party components that Veeam incorporates. Specify `1` to accept the license agreement and proceed with installation. Example: `ACCEPT_THIRDPARTY_LICENSES="1"`

Veeam Backup Enterprise Manager

To install Veeam Backup Enterprise Manager, use a command with the following syntax:

```
msiexec.exe [/L*v "<path_to_log>"] /qn /i "<path_to_msi>" ACCEPT_EULA="1" ACCEP
T_THIRDPARTY_LICENSES="1" [INSTALLDIR="<path_to_installdir >"] VBREM_LICENSE_FI
LE="<path_to_license_file>" [VBREM_SERVICE_USER="<Veeam_EM_Service_account>"][V
BREM_SERVICE_PASSWORD="<Veeam_EM_Service_account_password>"] [VBREM_SERVICE_POR
T="<Veeam_EM_Service_port>"]
[VBREM_SQLSERVER_SERVER="<SQL_server>"] [VBREM_SQLSERVER_DATABASE="<database_na
me>"][VBREM_SQLSERVER_AUTHENTICATION="0"] [VBREM_SQLSERVER_USERNAME="<SQL_auth_
username>"] [VBREM_SQLSERVER_PASSWORD="<SQL_auth_password>"] [VBREM_TCPPORT="<T
CP_port_for_web_site>"] [VBREM_SSLPORT="<SSL_port_for_web_site>”]>"] [VBREM_THU
MBPRINT="<certificate_hash>"] [VBREM_RESTAPISVC_PORT="<TCP_port_for_RestApi_ser
vice>"] [VBREM_RESTAPISVC_SSLPORT="<SSL_port_for_RestApi_service>"][VBREM_CHECK
_UPDATES="1"]
```

The command has the following parameters:

Option	Parameter	Required	Description
/L	*v logfile	No	Creates an installation log file with the verbose output. Specify a full path to the log file as the parameter value. A setup log file created during the previous installation is cleared. Example: `/L*v "C:\ProgramData\Veeam\Setup\Temp\Logs\EM.txt"`
/q	n	Yes	Sets the user interface level to “no”, which means no user interaction is needed during installation.
/i	setup file	Yes	Installs Veeam Backup Enterprise Manager. Specify a full path to the setup file as the parameter value. Example: `/i "C:\Veeam\EnterpriseManager\BackupWeb_x64.msi"`
ACCEPT_EULA	0/1	Yes	Specifies if you want to accept the Veeam license agreement. Specify `1` to accept the license agreement and proceed with installation. Example: `ACCEPT_EULA="1"`

Option	Parameter	Required	Description
ACCEPT_THIRDPARTY_LICENSES	0/1	Yes	Specifies if you want to accept the license agreement for 3rd party components that Veeam incorporates. Specify `1` to accept the license agreement and proceed with installation. Example: `ACCEPT_THIRDPARTY_LICENSES="1"`
INSTALLDIR	path	No	Installs the component to the specified location. By default, Veeam Backup & Replication uses the `Enterprise Manager` subfolder of the `C:\Program Files\Veeam\` folder. Example: `INSTALLDIR="C:\Backup\"` The component will be installed to the `C:\Backup\Enterprise Manager` folder.
VBREM_LICENSE_FILE	license path	Yes	Specifies a full path to the license file. Example: `VBREM_LICENSE_FILE="C:\Users\Administrator\Desktop\enterprise - veeam_backup_trial_0_30.lic"`
VBREM_SERVICE_USER	user	No	Specifies the account under which the Veeam Backup Enterprise Manager Service will run. The account must have full control NTFS permissions on the `VBRCatalog` folder where index files are stored and the *Database owner* rights for the Veeam Backup Enterprise Manager configuration database on the Microsoft SQL Server that you plan to use. If you do not specify this parameter, the Veeam Backup Enterprise Manager Service will run under the LocalSystem account. Together with the `VBREM_SERVICE_USER` parameter, you must specify the `VBREM_SERVICE_PASSWORD` parameter. Example: `VBRC_SERVICE_USER="BACKUPSERVER\Administrator"`
VBREM_SERVICE_PASSWORD	password	No	Specifies a password for the account under which the Veeam Backup Enterprise Manager Service will run. Example: `VBREM_SERVICE_PASSWORD="1234"`

Option	Parameter	Required	Description
VBREM_SERVICE_PORT	Port	No	Specifies a TCP port that will be used by the Veeam Backup Enterprise Manager Service. By default, the port number 9394 is used. Example: `VBREM_SERVICE_PORT ="9394"`
VBREM_SQLSERVER_SERVER	SQL server\instance	No	Specifies a Microsoft SQL Server and instance on which the Veeam Backup Enterprise Manager configuration database will be deployed. By default, Veeam Backup & Replication uses the `(local)\VEEAMSQL2012` server for machines running Microsoft Windows 7 or Microsoft Windows Server 2008 R2, and `(local)\VEEAMSQL2016` for machines running Microsoft Windows Server 2012 or later. Example: `VBREM_SQLSERVER_SERVER="BACKUPSERVER\VEEAMSQL2012_MY"`
VBREM_SQLSERVER_DATABASE	database	No	Specifies a name of the Veeam Backup Enterprise Manager database. By default, the database is deployed with the `VeeamBackupReporting` name. Example: `VBREM_SQLSERVER_DATABASE="VeeamBackupReporting01"`
VBREM_SQLSERVER_AUTHENTICATION	0/1	No	Specifies if you want to use the Microsoft SQL Server authentication mode to connect to the Microsoft SQL Server where the Veeam Backup Enterprise Manager is deployed. Set this parameter to `1` if you want to use the SQL Server authentication mode. If you do not specify this parameter, Veeam Backup Enterprise Manager will connect to the Microsoft SQL Server in the Microsoft Windows authentication mode (default value is `0`). Together with this parameter, you must specify the following parameters: `VBREM_SQLSERVER_USERNAME` and `VBREM_SQLSERVER_PASSWORD`. Example: `VBREM_SQLSERVER_AUTHENTICATION="1"`

Option	Parameter	Required	Description
VBREM_SQLSERVER_USERNAME	user	No	This parameter must be used if you have specified the `VBREM_SQLSERVER_AUTHENTICATION` parameter. Specifies a LoginID to connect to the Microsoft SQL Server in the SQL Server authentication mode. Example: `VBREM_SQLSERVER_USERNAME="sa"`
VBREM_SQLSERVER_PASSWORD	password	No	This parameter must be used if you have specified the `VBR_SQLSERVER_AUTHENTICATION` parameter. Specifies a password to connect to the Microsoft SQL Server in the SQL Server authentication mode. Example: `VBREM_SQLSERVER_USERNAME="1234"`
VBREM_TCPPORT	port	No	Specifies a TCP port that will be used by the Veeam Backup Enterprise Manager website. By default, the port number 9080 is used. Example: `VBREM_TCPPORT="9080"`
VBREM_SSLPORT	port	No	Specifies a port that will be used by the Veeam Backup Enterprise Manager website. By default, the port number 9443 is used. Example: `VBREM_SSLPORT="9443"`
VBREM_THUMBPRINT	hash	No	Specifies the certificate to be used by Veeam Backup Enterprise Manager Service and Veeam RESTful API Service. If this parameter is not specified, a new certificate will be generated by `openssl.exe`. Example: `VBREM_THUMBPRINT="0677d0b8f27caccc966b15d807b41a101587b488"`
VBREM_RESTAPISVC_PORT	port	No	Specifies a TCP port that will be used by the Veeam Backup Enterprise Manager RESTful API Service. By default, the port number 9399 is used. Example: `VBREM_RESTAPISVC_PORT="9399"`

Option	Parameter	Required	Description
VBREM_RESTAPISVC_SSLPORT	port	No	Specifies a port that will be used by the Veeam RESTful API Service. By default, the port number 9398 is used. Example: `VBREM_RESTAPISVC_SSLPORT="9398"`
VBREM_CONFIG_SCHANNEL	0/1	No	Specifies if the TLS 1.2 protocol will be used for secure communication with the Veeam Backup Enterprise Manager website.
VBREM_CHECK_UPDATES	0/1	No	Specifies if you want Veeam Backup Enterprise Manager to automatically check for new product versions and updates. Specify `0` if you do not want to check for updates. If you do not specify this parameter, Veeam Backup Enterprise Manager will automatically check for updates (default value is `1`). Example: `VBREM_CHECK_UPDATES="0"`

Example

Suppose you want to install Veeam Backup Enterprise Manager with the following settings:

- Installation log location: `C:\logs\log1.txt`
- No user interaction
- Path to the MSI file: `E:\Veeam\EnterpriseManager\BackupWeb_x64.msi`
- Installation folder: `D:\Program Files\Veeam`
- License file location: `C:\License\veeam_license.lic`
- Service user account: `VEEAM\Administrator`
- Service user account password: `1243`
- Service port: default
- Microsoft SQL Server database: `BACKUPSERVER\VEEAMSQL2012_MY`
- Database name: `VeeamReporting01`
- TCP and TLS ports: default
- Certificate: default
- TCP port for RESTful API: `9396`
- TLS port for RESTful API: `9397`
- No check for updates

The command to install Veeam Backup Enterprise Manager with such configuration will have the following parameters:

```
msiexec.exe /L*v "C:\logs\log1.txt" /qn /i "E:\Veeam\EnterpriseManager\BackupWeb_x64.msi" ACCEPT_EULA="1" ACCEPT_THIRDPARTY_LICENSES="1" INSTALLDIR="D:\Program Files\Veeam" VBREM_LICENSE_FILE="C:\License\veeam_license.lic" VBREM_SERVICE_USER="VEEAM\Administrator" VBREM_SERVICE_PASSWORD="1234" VBREM_SQLSERVER_SERVER="BACKUPSERVER\VEEAMSQL2012_MY" VBREM_SQLSERVER_DATABASE="VeeamReporting01" VBREM_RESTAPISVC_PORT="9396" VBREM_CHECK_UPDATES="0"
```

Veeam Cloud Connect Portal

Veeam Cloud Connect Portal requires Veeam Backup Enterprise Manager of the same version to be installed on the target machine.

To install Veeam Cloud Connect Portal, use a command with the following syntax:

```
msiexec.exe [/L*v "<path_to_log>"] /qn /i "<path_to_msi>" ACCEPT_EULA="1" ACCEPT_THIRDPARTY_LICENSES="1" [INSTALLDIR="<path_to_installdir >"] VBCP_SSLPORT="<SSL_port">
```

The command has the following parameters:

Option	Parameter	Required	Description
/L	*v logfile	No	Creates an installation log file with the verbose output. Specify a full path to the log file as the parameter value. A setup log file created during the previous installation is cleared. Example: `/L*v "C:\ProgramData\Veeam\Setup\Temp\Logs\CloudPortal.txt"`
/q	n	Yes	Sets the user interface level to "no", which means no user interaction is needed during installation.
/i	setup file	Yes	Installs the Veeam Cloud Connect Portal. Specify a full path to the setup file as the parameter value. Example: `/i "C:\Cloud portal\BackupCloudPortal_x64.msi"`
ACCEPT_EULA	0/1	Yes	Specifies if you want to accept the Veeam license agreement. Specify `1` to accept the license agreement and proceed with installation. Example: `ACCEPT_EULA="1"`

Option	Parameter	Required	Description
ACCEPT_THIRDPARTY_LICENSES	0/1	Yes	Specifies if you want to accept the license agreement for 3rd party components that Veeam incorporates. Specify `1` to accept the license agreement and proceed with installation. Example: `ACCEPT_THIRDPARTY_LICENSES="1"`
INSTALLDIR	path	No	Installs the component to the specified location. By default, Veeam Backup & Replication uses the `CloudPortal` subfolder of the `C:\Program Files\Veeam\Backup and Replication\` folder. Example: `INSTALLDIR="C:\Backup\"`. The component will be installed to the `C:\Backup\CloudPortal` folder
VBCP_SSLPORT	port	No	Specifies a port that will be used by the Veeam Cloud Connect Portal website. By default, the port number 6443 is used. Example: `VBREM_SSLPORT="7443"`

Example

Suppose you want to install Veeam Cloud Connect Portal with the following configuration:

- No user interaction
- Path to the MSI file: `E:\Cloud portal\BackupCloudPortal_x64.msi`
- Installation folder: `C:\Backup`
- TLS port: default

The command to install Veeam Cloud Connect Portal with such configuration will have the following parameters:

```
msiexec.exe /qn /L*v "C:\logs\log1.txt" /qn /i "E:\Cloud portal\BackupCloudPortal_x64.msi" ACCEPT_EULA="1" ACCEPT_THIRDPARTY_LICENSES="1" INSTALLDIR="C:\Backup\"
```

Veeam Cloud Plug-Ins

You can install the following cloud plug-ins in the unattended mode:

- AWS Plug-in for Veeam Backup & Replication

- Microsoft Azure Plug-in for Veeam Backup & Replication
- Google Cloud Platform Plug-in for Veeam Backup & Replication

Note that if you want to manage cloud plug-ins in the Veeam Backup & Replication interface, you must install both the cloud plug-in service and the cloud plug-in UI.

To install cloud plug-ins, use a command with the following syntax:

```
msiexec /qn /i <msi> ACCEPT_THIRDPARTY_LICENSES=1 ACCEPT_EULA=1
```

The command has the following parameters:

Option	Parameter	Required	Description
/L	*v logfile	No	Creates an installation log file with the verbose output. Specify a full path to the log file as the parameter value. A setup log file created during the previous installation is cleared. Example: `/L*v "C:\ProgramData\Veeam\Setup\Temp\Logs\Backup.txt"`
/q	n	Yes	Sets the user interface level to "no", which means no user interaction is needed during installation.
/i	setup file	Yes	Installs the Veeam Backup & Replication server. Specify a full path to the setup file as the parameter value. Example: `/i "C:\Veeam\Server.x64.msi"`

Option	Parameter	Required	Description
msi	• AWSPlugin.msi • AWSPluginUI.msi • MicrosoftAzurePlugin.msi • MicrosoftAzurePluginUI.msi • GCPPlugin.msi • GCPPluginUI.msi	Yes	`AWSPlugin.msi`: Installs the AWS Plug-in for Veeam Backup & Replication services. `AWSPluginUI.msi`: Installs the AWS Plug-in for Veeam Backup & Replication UI. `MicrosoftAzurePlugin.msi`: Installs the Microsoft Azure Plug-in for Veeam Backup & Replication services. `MicrosoftAzurePluginUI.msi`: Installs the Microsoft Azure Plug-in for Veeam Backup & Replication UI. `GCPPlugin`.msi: Installs the Google Cloud Platform Plug-in for Veeam Backup & Replication services. `GCPPluginUI`.msi: Installs the Google Cloud Platform Plug-in for Veeam Backup & Replication UI.
ACCEPT_EULA	0/1	Yes	Specifies if you want to accept the Veeam license agreement. Specify `1` to accept the license agreement and proceed with installation. Example: `ACCEPT_EULA="1"`
ACCEPT_THIRDPARTY_LICENSES	0/1	Yes	Specifies if you want to accept the license agreement for 3rd party components that Veeam incorporates. Specify `1` to accept the license agreement and proceed with installation. Example: `ACCEPT_THIRDPARTY_LICENSES="1"`
INSTALLDIR	path	No	Installs the component to the specified location. By default, Veeam Backup & Replication uses the `Backup` subfolder of the `C:\Program Files\Veeam\Backup and Replication\` folder. Example: `INSTALLDIR="C:\Backup\"` The component will be installed to the `C:\Backup\Backup` folder.

Examples

Installing AWS Plug-in for Veeam Backup & Replication Service and UI

To install the AWS Plug-in for Veeam Backup & Replication service and the AWS Plug-in for Veeam Backup & Replication UI to Veeam Backup & Replication, perform the following steps:

1. Run the following command to install the AWS Plug-in for Veeam Backup & Replication service.

```
msiexec /qn /i AWSPlugin.msi ACCEPT_THIRDPARTY_LICENSES=1 ACCEPT_EULA=1
```

2. Run the following command to install the AWS Plug-in for Veeam Backup & Replication UI.

```
msiexec /qn /i AWSPluginUI.msi ACCEPT_THIRDPARTY_LICENSES=1 ACCEPT_EULA=1
```

Veeam Backup for Nutanix AHV

You can install the Nutanix AHV Plug-in in the unattended mode. Note that if you want to manage AHV VM backup in the Veeam Backup & Replication console, you must install both the Nutanix AHV Plug-in service (`NutanixAHVPlugin.msi`) and the UI components of Nutanix AHV Plug-in (`NutanixAHVPluginUI.msi`).

To install Nutanix AHV Plug-in, use a command with the following syntax:

```
msiexec /qn /i <msi> ACCEPT_THIRDPARTY_LICENSES=1 ACCEPT_EULA=1
```

The command has the following parameters:

Option	Parameter	Required	Description
/L	*v logfile	No	Creates an installation log file with the verbose output. Specify a full path to the log file as the parameter value. A setup log file created during the previous installation is cleared. Example: `/L*v "C:\ProgramData\Veeam\Setup\Temp\Logs\Backup.txt"`
/q	n	Yes	Sets the user interface level to "no", which means no user interaction is needed during installation.

Option	Parameter	Required	Description
/i	setup file	Yes	Installs the Veeam Backup & Replication server. Specify a full path to the setup file as the parameter value. Example: `/i "C:\Veeam\Server.x64.msi"`
msi	• NutanixAHVPlugin.msi • NutanixAHVPluginUI.msi	Yes	`NutanixAHVPlugin.msi`: Installs the Nutanix AHV Plug-in services. `NutanixAHVPluginUI.msi`: Installs UI components of the Nutanix AHV Plug-in.
ACCEPT_EULA	0/1	Yes	Specifies if you want to accept the Veeam license agreement. Specify `1` to accept the license agreement and proceed with installation. Example: `ACCEPT_EULA="1"`
ACCEPT_THIRDPARTY_LICENSES	0/1	Yes	Specifies if you want to accept the license agreement for 3rd party components that Veeam incorporates. Specify `1` to accept the license agreement and proceed with installation. Example: `ACCEPT_THIRDPARTY_LICENSES="1"`
INSTALLDIR	path	No	Installs the component to the specified location. By default, Veeam Backup & Replication uses the `Backup` subfolder of the `C:\Program Files\Veeam\Backup and Replication\` folder. Example: `INSTALLDIR="C:\Backup\"` The component will be installed to the `C:\Backup\Backup` folder.

Examples

Installing Nutanix AHV Plug-in Service and UI

To install the Nutanix AHV Plug-in service and UI components to Veeam Backup & Replication, perform the following steps:

1. Run the following command to install the Nutanix AHV Plug-in service.

```
msiexec /qn /i NutanixAHVPlugin.msi ACCEPT_THIRDPARTY_LICENSES=1 ACCEPT_EU
LA=1
```

2. Run the following command to install UI components of Nutanix AHV Plug-in.

```
msiexec /qn /i NutanixAHVPluginUI.msi ACCEPT_THIRDPARTY_LICENSES=1 ACCEPT_
EULA=1
```

Installing Updates in Unattended Mode

Veeam Backup & Replication updates can be installed in the unattended mode.

To install a Veeam Backup & Replication update, perform the following steps:

1. Download the update installation archive and extract the executable file.
2. Install the update on the backup server.

> **IMPORTANT!**
>
> The script that installs Veeam Backup & Replication updates must be run with elevated privileges (run as Administrator).

Step 1. Download and Extract Executable File

Download and extract the executable file for update installation:

1. Download the installation archive for the Veeam Backup & Replication update from the Release Notes page for a certain update/patch. For example for v11 updates, download the latest patch from this Veeam KB article, for v11a updates — from this Veeam KB article.
2. Extract the executable file from the downloaded archive.
3. Save the extracted file locally on the backup server where you plan to install the update, or place it in a network shared folder.

Alternatively, you can get the update/patch executable file from the **Updates** folder on the downloaded ISO image.

Step 2. Install Update

To install the Veeam Backup & Replication update on the backup server, use the following command syntax:

```
%patch% [/silent][/noreboot][/log <log_path>][VBR_AUTO_UPGRADE="1"]
```

The command has the following parameters:

%patch%	path	Yes	Specifies a path to the update installation file on the backup server or in a network shared folder. Example: `C:\Temp\VeeamBackup&Replication_11.0.1.1261_20220302.exe`
silent	—	Yes	Sets the user interface level to "no", which means no user interaction is needed during installation.

noreboot	—	No	Suppresses reboot if reboot is required during the Veeam Backup & Replication update installation.
log	path	No	Specifies a full path to the log file for the Veeam Backup & Replication update installation. Example: `C:\Logs\veeam.log`
VBR_AUTO_UPGRADE	0/1	No	Specifies if you want Veeam Backup & Replication to automatically upgrade existing components in the backup infrastructure. Veeam Backup & Replication performs automatic upgrade after the Veeam Backup Service is started on the backup server. Specify `1` to enable automatic upgrade. Example: `VBR_AUTO_UPGRADE="1"`

For example:

You want to install the Veeam Backup & Replication update with the following options:

- Path to the update installation file: `C:\Temp\VeeamBackup&Replication_11.0.1.1261_20220302.exe`
- Silent install: enabled
- Noreboot: enabled
- Path to the log file: `C:\Logs\veeam.log`
- Components auto upgrade: enabled

The command to install the Veeam Backup & Replication update will be the following:

```
C:\Temp\VeeamBackup&Replication_11.0.1.1261_20220302.exe /silent /noreboot /log C
:\Logs\veeam.log VBR_AUTO_UPGRADE="1"
```

Installation Results

You can use the last exit code to verify if the installation process has completed successfully.

- In `cmd.exe`, use the `%ERRORLEVEL%` variable to check the last exit code.
- In Microsoft Windows PowerShell, use the `$LastExitCode` variable to check the last exit code.

Veeam Backup & Replication does not provide any confirmation about the results of automatic components upgrade. To check if components have been successfully upgraded, use the Veeam Backup & Replication console.

Backup Infrastructure

Veeam Backup & Replication is a modular solution that lets you build a scalable backup infrastructure for environments of different sizes and configuration. The installation package of Veeam Backup & Replication includes a set of components that you can use to configure the backup infrastructure. Some components are mandatory and provide core functionality; some components are optional and can be installed to provide additional functionality for your business and deployment needs. You can co-install Veeam Backup & Replication components on the same machine, physical or virtual, or you can set them up separately for a more scalable approach.

Backup Infrastructure Components

The Veeam backup infrastructure comprises a set of components. Some components can be deployed with the help of the setup file. Other components can be deployed using the Veeam Backup & Replication console.

Backup Server

The backup server is a Windows-based physical or virtual machine on which Veeam Backup & Replication is installed. It is the core component in the backup infrastructure that fills the role of the "configuration and control center". The backup server performs all types of administrative activities:

- Coordinates backup, replication, recovery verification and restore tasks
- Controls job scheduling and resource allocation
- Is used to set up and manage backup infrastructure components as well as specify global settings for the backup infrastructure

In addition to its primary functions, a newly deployed backup server also performs the roles of the default backup proxy and the backup repository (it manages data handling and data storing tasks).

Backup Server Services and Components

The backup server uses the following services and components:

- **Veeam Backup Service** is a Windows service that coordinates all operations performed by Veeam Backup & Replication such as backup, replication, recovery verification and restore tasks. The Veeam Backup Service runs under the LocalSystem account or account that has the local Administrator permissions on the backup server.
- **Veeam Broker Service** interacts with the virtual infrastructure to collect and cache the virtual infrastructure topology. Jobs and tasks query information about the virtual infrastructure topology from the broker service, which accelerates job and task performance.
- **Veeam Guest Catalog Service** manages guest OS file system indexing for VMs and replicates system index data files to enable search through guest OS files. Index data is stored in the Veeam Backup Catalog — a folder on the backup server. The Veeam Guest Catalog Service running on the backup server works in conjunction with search components installed on Veeam Backup Enterprise Manager and (optionally) a dedicated Microsoft Search Server.
- **Mount Service** mounts backups and replicas for file-level access, browsing the VM guest file system and restoring VM guest OS files and application items to the original location.
- **Backup Proxy Services** In addition to dedicated services, the backup server runs a set of Veeam Data Movers. For details, see Backup Proxy.
- **Veeam CDP Coordinator Service** communicates with vCenter, assigns continuous data protection (CDP) tasks and manages the infrastructure components involved in CDP. For more information on CDP, see Continuous Data Protection (CDP).
- **Veeam Backup & Replication Configuration Database** stores data about the backup infrastructure, jobs, sessions and other configuration data. The database instance can be located on an SQL Server installed either locally (on the same machine where the backup server is running) or remotely. For more information, see the following guidelines: Veeam Backup & Replication Best Practices.

  Veeam Backup & Replication maintains the configuration database. Veeam Backup & Replication runs the *DatabaseMaintenance* system job once a week and when the Veeam Backup Service is restarted. The job updates the database internal statistics, defragments indexes and clears unused data. For details, see the `Job.DatabaseMaintenance` log file in the `%ProgramData%\Veeam\Backup` folder.
- **Veeam Backup & Replication Console** provides the application user interface and allows user access to the application functionality.

- **Veeam Backup PowerShell Module** is an extension for Microsoft Windows PowerShell that adds a set of cmdlets to allow users to perform backup, replication and recovery tasks through the command-line interface of PowerShell or run custom scripts to fully automate operation of Veeam Backup & Replication. For more information on the Microsoft Windows PowerShell module, see the Getting Started section of Veeam PowerShell Reference.

Credentials Manager

You can use the Credentials Manager to create and maintain a list of credentials records that you plan to use to connect to components in the backup infrastructure.

The Credentials Manager lets you create the following types of credentials records:

- Standard account
- Linux account (user name and password)
- Linux private key (Identity/Pubkey)

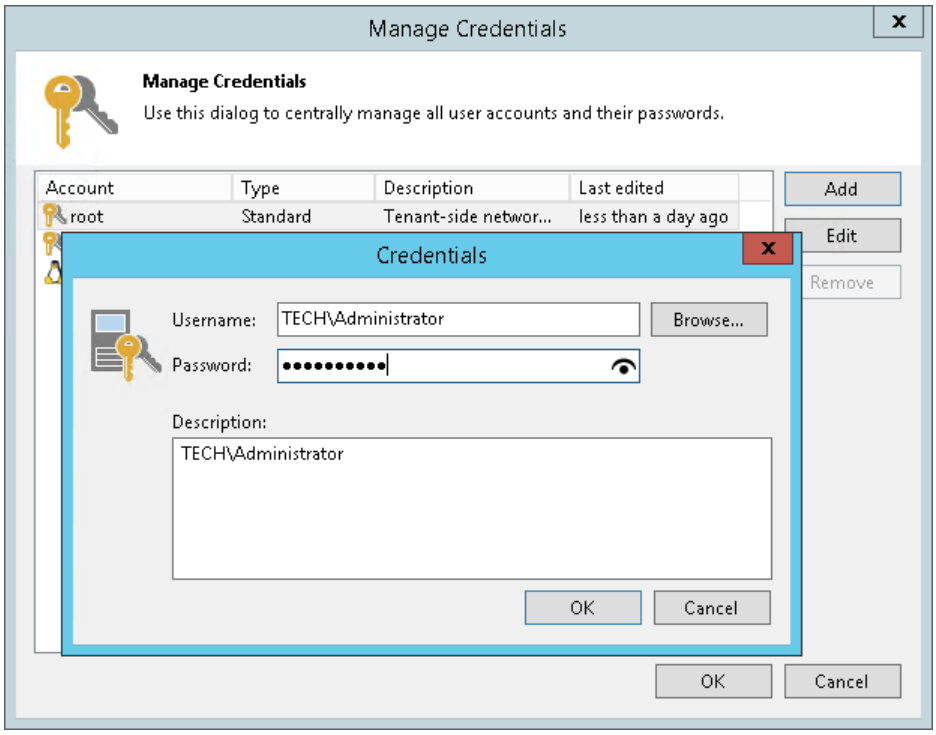
Standard Accounts

You can create a credentials record for an account that you plan to use to connect to a Microsoft Windows server, a VM running Microsoft Windows OS or a storage system.

To create a new standard credentials record:

1. From the main menu, select **Manage Credentials**.
2. Click **Add** > **Standard account**.
3. In the **Username** field, enter a user name for the account that you want to add. You can also click **Browse** to select an existing user account.
4. In the **Password** field, enter a password for the account that you want to add. To view the entered password, click and hold the eye icon on the right of the field.

5. In the **Description** field, enter a description for the created credentials record. As there can be a number of similar account names, for example, *Administrator*, it is recommended that you provide a meaningful unique description for the credentials record so that you can distinguish it in the list. The description is shown in brackets, following the user name.



Linux Accounts (User Name and Password)

You can create a credentials record for the account that you plan to use to connect to a Linux server or VM running Linux OS.

To create a new credentials record with a user name and password for a Linux server:

1. From the main menu, select **Manage Credentials**.
2. Click **Add** > **Linux account**.
3. In the **Username** field, enter a user name for the account that you plan to add.
4. In the **Password** field, enter a password for the account that you want to add. To view the entered password, click and hold the eye icon on the right of the field.
5. In the **SSH port** field, specify the SSH port over which you want to connect to a Linux server. By default, port 22 is used.
6. If you specify data for a non-root account that does not have root permissions on a Linux server, you can use the **Non-root account** section to grant sudo rights to this account.
   a. To provide a non-root user with root account privileges, select the **Elevate account privileges automatically** check box.
   b. To add the user account to `sudoers` file, select the **Add account to the sudoers file** check box. In the **Root password** field, enter the password for the root account.

      If you do not enable this option, you will have to manually add the user account to the `sudoers` file.

c. If you plan to use the account to connect to Linux servers where the `sudo` command is not available or fails, you have an option to use the `su` command instead. To enable the su command, select the **Use "su" if "sudo" fails** check box and in the **Root password** field, enter the password for the root account.

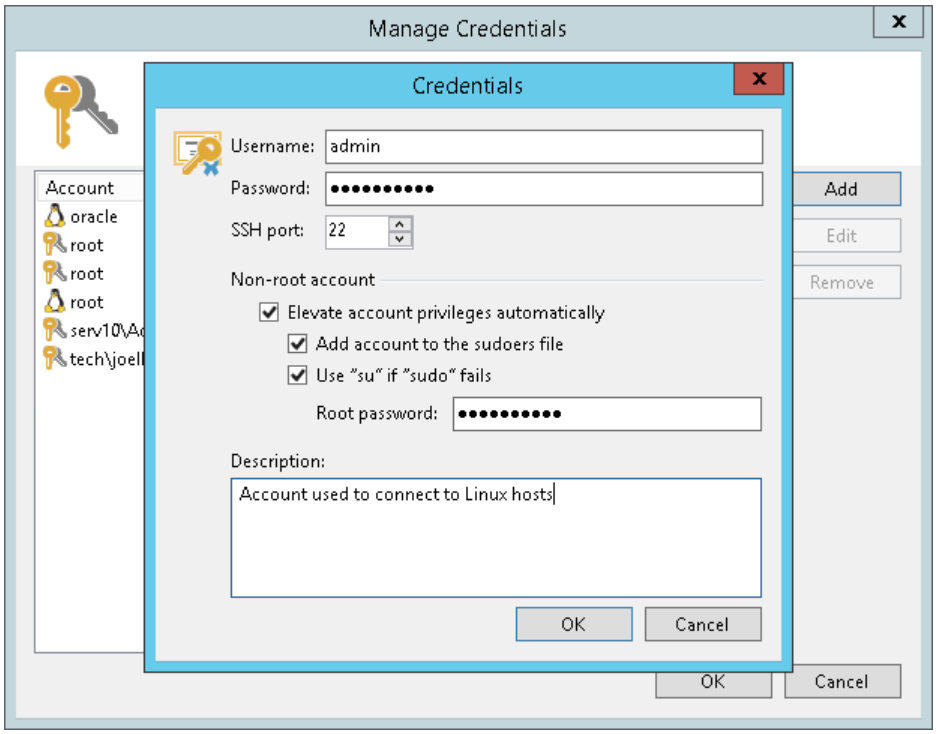
   Veeam Backup & Replication will first try to use the `sudo` command. If the attempt fails, Veeam Backup & Replication will use the `su` command.

7. In the **Description** field, enter a description for the created credentials record. As there can be a number of similar account names, for example, *Root*, it is recommended that you provide a meaningful unique description for the credentials record so that you can distinguish it in the list. The description is shown in brackets, following the user name.

IMPORTANT

Mind the following:

- You can create a separate user account intended for work with Veeam Backup & Replication on a Linux-based VM, grant root privileges to this account and specify settings of this account in the Credentials Manager. It is recommended that you avoid additional commands output for this user (like messages echoed from within ~/.bashrc or command traces before execution) because they may affect Linux VM processing.
- Cases when root password is required to elevate account rights to root using sudo are no longer supported.



Linux Private Keys (Identity/Pubkey)

You can log on to a Linux server or VM running Linux OS using the Identity/Pubkey authentication method. The Identity/Pubkey authentication method helps protect against malicious applications like keyloggers, strengthens the security level and simplifies launch of automated tasks.

To use the Identity/Pubkey authentication method, you must generate a pair of keys — a public key and private key:

- Public key is stored on Linux servers to which you plan to connect from the backup server. The key is kept in a special `authorized_keys` file containing a list of public keys.
- Private key is stored on the client machine — backup server. The private key is protected with a passphrase. Even if the private key is intercepted, the eavesdropper will have to provide the passphrase to unlock the key and use it.

For authentication on a Linux server, the client must prove that it has the private key matching the public key stored on the Linux server. To do this, the client generates a cryptogram using the private key and passes this cryptogram to the Linux server. If the client uses the "correct" private key for the cryptogram, the Linux server can easily decrypt the cryptogram with a matching public key.

Veeam Backup & Replication has the following limitations for the Identity/Pubkey authentication method:

- Veeam Backup & Replication does not support keys that are stored as binary data, for example, in a file of DER format.
- Veeam Backup & Replication supports only keys whose passphrase is encrypted with algorithms supported by PuTTY:
  - AES (Rijndael): 128-bit, 192-bit and 256-bit CBC or CTR (SSH-2 only)
  - Blowfish: 128-bit CBC
  - Triple-DES: 168-bit CBC
- Veeam Backup & Replication supports the DSA and RSA algorithms with the following key formats:
  - PEM
  - Private
  - Private-openSSH
  - SSH-com
  - PKCS8 (RSA only)
- [Starting from Veeam Backup & Replication 11a (build 11.0.1.1261)] the following key formats are also supported for the following algorithms:
  - RSA: PEM, PKCS8, RFC4716 (private-openssh-new), private-openssh, private, sshcom
  - DSA: PEM, RFC4716 (private-openssh-new), private-openssh, private, sshcom
  - ECDSA: PEM, PKCS8, RFC4716 (private-openssh-new), private-openssh, private
  - EDDSA (ED25519): RFC4716 (private-openssh-new), private

> **IMPORTANT**
>
> If you use VMware VIX/vSphere Web Services, Veeam Backup & Replication does not support usage of public keys for guest processing on Linux guest servers.

To add a credentials record using the Identity/Pubkey authentication method:

1. Generate a pair of keys using a key generation utility, for example, ssh-keygen.

   Note that keys generated as ED448 as are not supported.

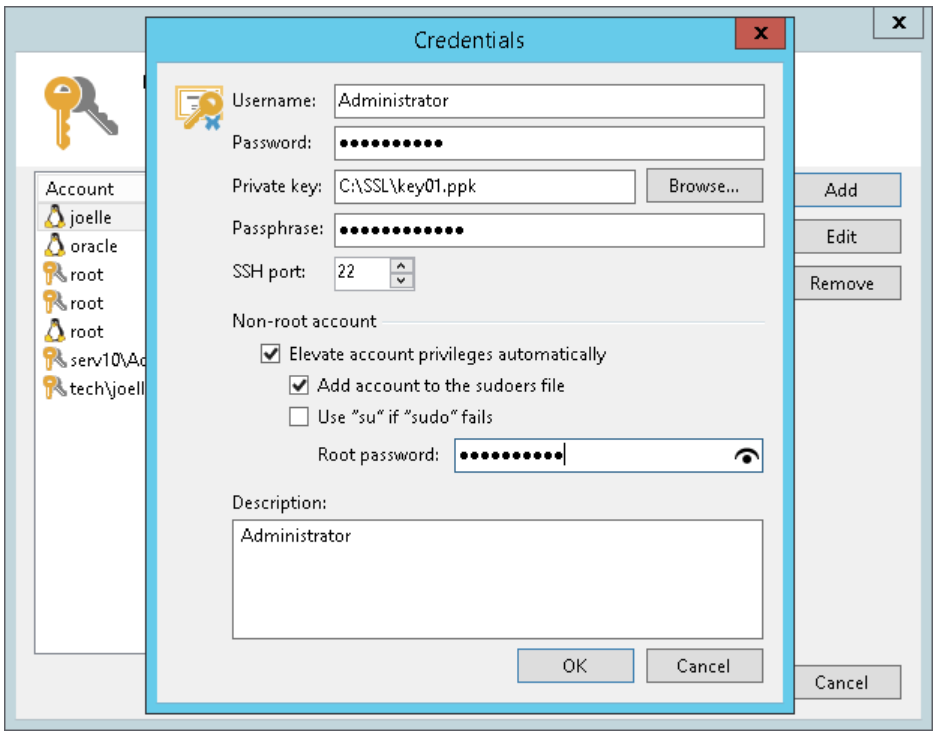
2. Place the public key on a Linux server. To do this, add the public key to the `authorized_keys` file in the `.ssh/` directory in the home directory on the Linux server.
3. Place the private key in some folder on the backup server or in a network shared folder.
4. In Veeam Backup & Replication, from the main menu select **Manage Credentials**.
5. Click **Add** > **Linux private key**.
6. In the **Username** field, specify a user name for the created credentials record.
7. In the **Password** field, specify the password for the user account. The password is required in all cases except when you use root or a user with enabled *NOPASSWD:ALL* setting in */etc/sudoers*.
8. In the **Private key** field, enter a path to the private key or click **Browse** to select a private key.
9. In the **Passphrase** field, specify a passphrase for the private key on the backup server. To view the entered passphrase, click and hold the eye icon on the right of the field.
10. In the **SSH port** field, specify a number of the SSH port that you plan to use to connect to a Linux server. By default, port 22 is used.
11. If you specify data for a non-root account that does not have root permissions on a Linux server, you can use the **Non-root account** section to grant sudo rights to this account.
    a. To provide a non-root user with root account privileges, select the **Elevate specified account to root** check box.
    b. To add the user account to sudoers file, select the **Add account to the sudoers file automatically** check box. In the **Root password** field, enter the password for the root account.

       If you do not enable this option, you will have to manually add the user account to the sudoers file.
    c. When registering a Linux server, you have an option to failover to using the su command for distros where the sudo command is not available.

       To enable the failover, select the **Use "su" if "sudo" fails** check box and in the **Root password** field, enter the password for the root account.
12. In the **Description** field, enter a description for the created credentials record. As there can be a number of similar account names, for example, *Root*, it is recommended that you supply a meaningful unique description for the credentials record so that you can distinguish it in the list. The description is shown in brackets, following the user name.

IMPORTANT

Cases when root password is required to elevate account rights to root using sudo are no longer supported.



Editing and Deleting Credentials Records

You can edit or delete credentials records that you have created.

To edit a credentials record:

1. From the main menu, select **Manage Credentials**.
2. Select the credentials record in the list and click **Edit**.
3. If the credentials record is already used for any component in the backup infrastructure, Veeam Backup & Replication will display a warning. Click **Yes** to confirm your intention.
4. Edit settings of the credentials record as required.

To delete a credentials record:

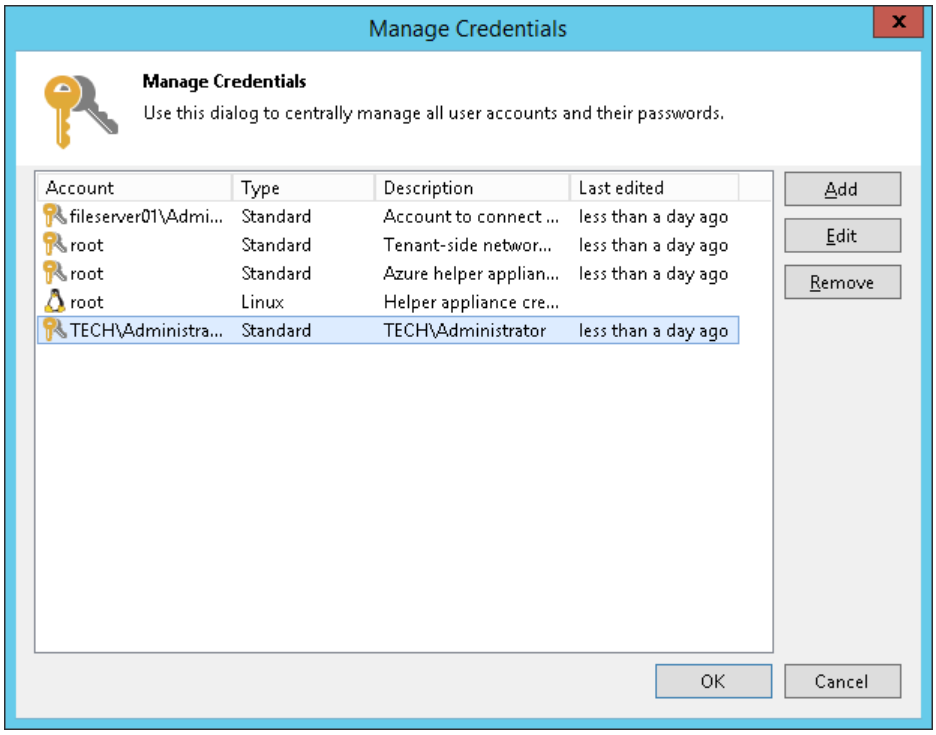
1. From the main menu, select **Manage Credentials**.
2. Select the credentials record in the list and click **Remove**. You cannot delete a record that is already used for any component in the backup infrastructure.

NOTE

The Credentials Manager contains 3 system credentials records:

- A credentials record for the Veeam FLR appliance
- A credentials record for the tenant-side network extension appliance
- A credentials record for Microsoft Azure helper appliance.

You cannot delete these credentials records. However, you can edit them: change a password and record description.



Cloud Credentials Manager

You can use the Cloud Credentials Manager to create and maintain a list of credentials records that you plan to use to connect to cloud services.

The Cloud Credentials Manager lets you create the following types of credentials records:

- Veeam Cloud Connect Accounts
- Access Keys for AWS Users
- Microsoft Azure Storage Accounts
- Microsoft Azure Compute Accounts
- Microsoft Azure Stack Hub Compute Accounts
- Google Cloud Accounts
- Google Cloud Service Accounts

Veeam Cloud Connect Accounts

You can add a credentials record for a tenant account — an account that you plan to use to connect to a service provider (SP).

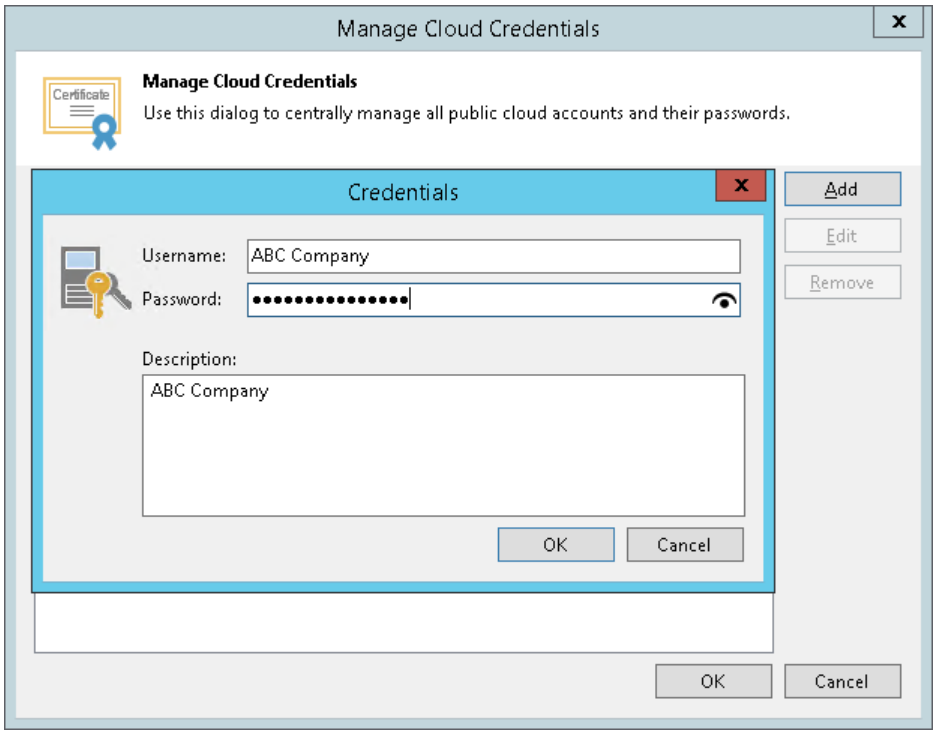
Before you add a credentials record, the SP must register a tenant account on the SP Veeam backup server. Tenants without accounts cannot connect to the SP and use Veeam Cloud Connect resources. For more information, see the Registering Tenant Accounts section in the Veeam Cloud Connect Guide.

To create a credentials record for a tenant account:

1. From the main menu, select **Manage Cloud Credentials**.
2. Click **Add** > **Veeam Cloud Connect service provider account**.
3. In the **Username** field, enter a user name for the account that the SP has provided to you.

   > **NOTE**
   >
   > If the SP used vCloud Director to allocate replication resources to you, you must enter a user name for the vCloud Director tenant account in the following format: *Organization*|*Username*. For example: *TechCompanyOrg*|*Administrator*.

4. In the **Password** field, enter a password for the account that the SP has provided to you. To view the entered password, click and hold the eye icon on the right of the field.
5. In the **Description** field, enter a description for the created credentials record.



Access Keys for AWS Users

You can create a record for credentials that you plan to use to connect to AWS.

To access AWS resources, you can use *Identity and Access Management (IAM) user* credentials or *AWS account root user* credentials. However, AWS recommends that you use the IAM user credentials. For details, see the AWS Account Root User Credentials vs. IAM User Credentials section in the AWS General Reference. The permissions that you must provide to the user are listed in the Required Permissions section.

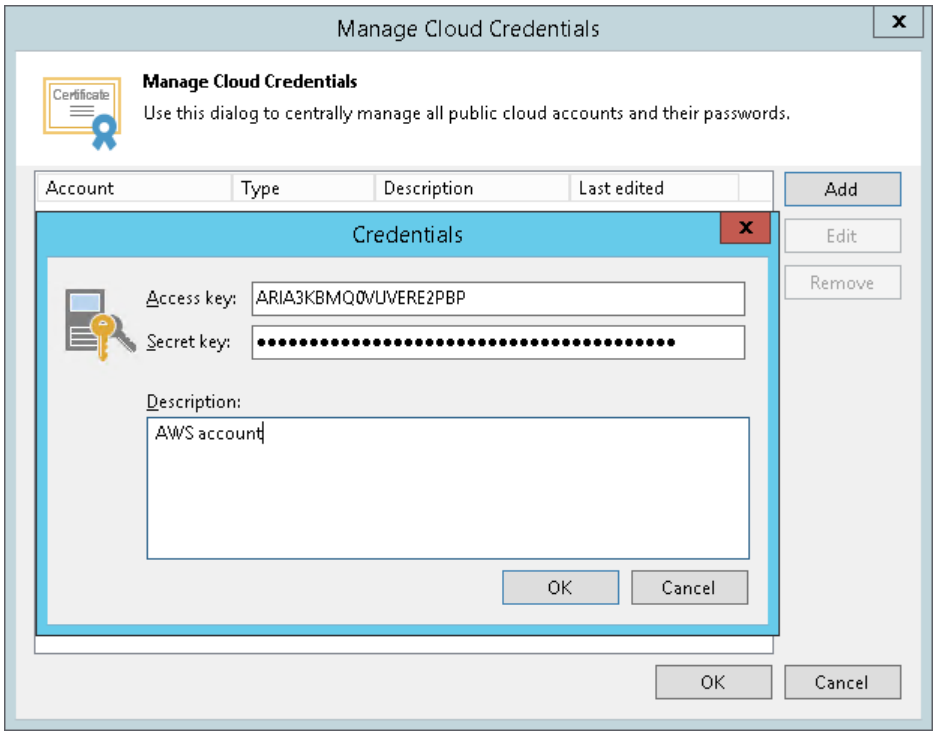
Since Veeam Backup & Replication uses AWS CLI commands to perform operations in AWS, instead of a user name and password you must specify an AWS access key. AWS access keys are long-term user credentials that consists of two parts: an access key ID and a secret access key. For details, see the Managing Access Keys for IAM Users section in the AWS IAM User Guide.

To create a credentials record:

1. From the main menu, select **Manage Cloud Credentials**.
2. Click **Add** > **AWS access key**.
3. In the **Access key** field, enter an access key ID.
4. In the **Secret key** field, enter a secret access key. To view the entered secret key, click and hold the eye icon on the right of the field.
5. In the **Description** field, enter a description for the created credentials record.

IMPORTANT

It is recommended that the user whose credentials you plan to use to connect to AWS S3 has administrative permissions — access to all AWS S3 actions and resources.



Microsoft Azure Storage Accounts

You can create a credentials record for an account that you plan to use to connect to Microsoft Azure Blob storage.

The following types of storage accounts are supported.

Storage account type	Supported services	Supported performance tiers	Supported access tiers
General-purpose V2 BlobStorage	Blob (block blobs only)	Standard	Cool: for infrequently accessed data. Hot: for frequently accessed data. Archive: for rarely accessed data. Can be set only on blob-level. Supported in Archive Tier object storage systems. Veeam Backup & Replication will use the access tier you select as the default one.
General-purpose V1	Blob (block blobs only)	Standard	N/A
BlockBlobStorage	Blob (block blobs only)	Premium	N/A

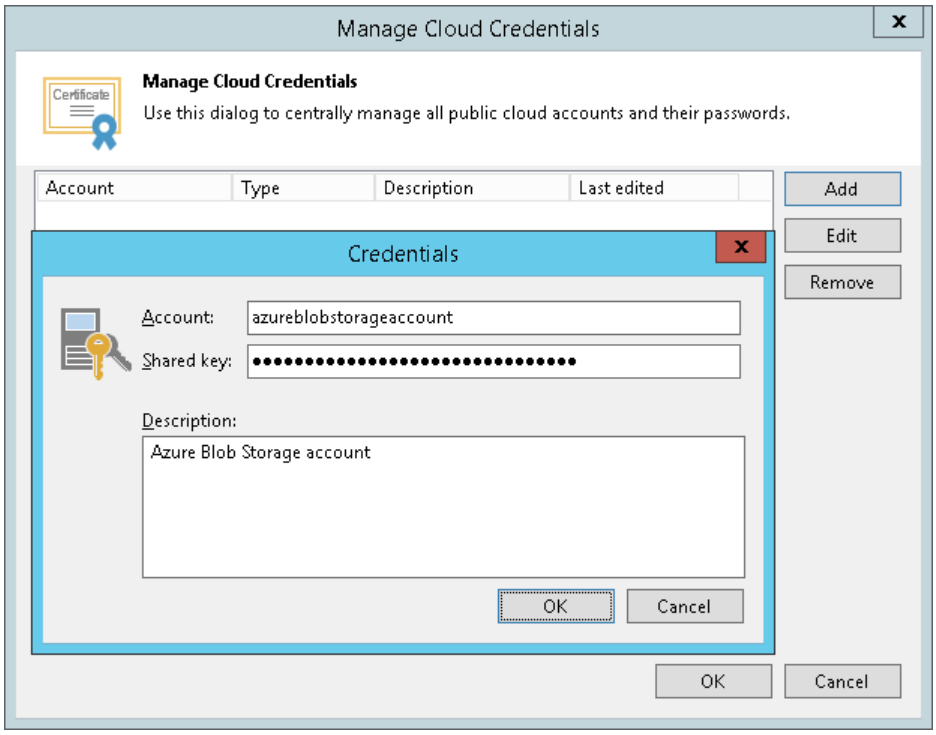
For more information about the types of storage accounts in Azure, see Microsoft Docs.

To create a record for a Microsoft Azure storage account:

1. From the main menu, select **Manage Cloud Credentials**.
2. Click **Add** > **Microsoft Azure storage account**.
3. In the **Account** field, enter the storage account name.
4. In the **Shared key** field, enter the storage account shared key. To view the entered key, click and hold the eye icon on the right of the field.
5. In the **Description** field, enter an optional description for the credentials record.

TIP

If you do not have a Microsoft Azure storage account, you can create it in the Azure portal, as described in the Azure Storage Documentation.



Microsoft Azure Compute Accounts

To restore workloads to Microsoft Azure or add Azure archive storage, you must add a Microsoft Azure compute account to Veeam Backup & Replication. When you add a Microsoft Azure compute account, Veeam Backup & Replication imports information about subscriptions and resources associated with this account. During the restore process, Veeam Backup & Replication accesses these resources and uses them to register new VMs in Microsoft Azure.

If necessary, you can add multiple Microsoft Azure compute accounts to Veeam Backup & Replication. In this case, Veeam Backup & Replication will import information about all subscriptions and resources associated with the provided accounts, and you will be able to use these resources for restore.

Information about subscriptions and resources is saved to the Veeam Backup & Replication configuration database. You can re-import this information at any time.

To add a Microsoft Azure compute account, use the **Microsoft Azure Compute Account** wizard.

Before You Begin

Before you add a Microsoft Azure compute account to Veeam Backup & Replication, check the following prerequisites:

- At the Subscription step of the **Microsoft Azure Compute Account** wizard, you specify how Veeam Backup & Replication will access Microsoft Azure. If you plan to use an Azure Active Directory (AD) application, mind the following:
  - Make sure that you already have a user account in Microsoft Azure AD. This user account must have privileges to register applications. This can be the *Global Administrator* privileges assigned to the user or the **Users can register applications** option enabled for the user in Azure portal. For details, see Microsoft Azure Docs.
  - You must install Microsoft Azure PowerShell version 6.13.2 on the machine running the Veeam Backup & Replication console version 11a (build 11.0.1.1261) or later. For the machine running the Veeam Backup & Replication console version prior 11a (build 11.0.1.1261), install Microsoft Azure PowerShell version 5.1.1. If the installed Microsoft Azure PowerShell version differs, you will not be able to add a Microsoft Azure compute account.

    If you do not have Microsoft Azure PowerShell on the machine, Veeam Backup & Replication will prompt you to install it at the **Subscription** step of the **Microsoft Azure Compute Account** wizard.
- At the Subscription step of the **Microsoft Azure Compute Account** wizard, you specify how Veeam Backup & Replication will access Microsoft Azure. If you plan to use an existing Azure AD user account, mind the following:
  - The Azure account must have the *Contributor* role privileges for the selected subscription.

    [Restore to Microsoft Azure] If the *Contributor* role cannot be used, you can create a custom role with minimal permissions. To learn how to create a custom role, see Creating Custom Role for Azure Account.
  - If you have more than one Azure Active Directory tenant associated with your account, you will need to specify which tenant to use. For information on how to do it, see this Veeam KB article.
  - Only subscriptions that belong to the selected tenant will be added.
  - In this scenario, you cannot add an account with enabled Multi-Factor Authentication (MFA). You must disable MFA for the required account. App passwords are not supported.
- [For Microsoft Server OS] The Protected Mode must be switched off in the Internet Explorer settings. Otherwise, you will not be able to log on to Microsoft Azure when passing through the **Microsoft Azure Compute Account** wizard.

  If you do not want to switch off the Protected Mode for security reasons, you can add the following sites to the list of trusted hosts in **Internet Options** > **Secure** settings in Internet Explorer or in **Control Panel** > **Network and Internet**:
  - https://login.live.com
  - https://login.microsoftonline.com
  - https://secure.aadcdn.microsoftonline-p.com
  - https://auth.gfx.ms
  - about:security_veeam.backup.shell.exe

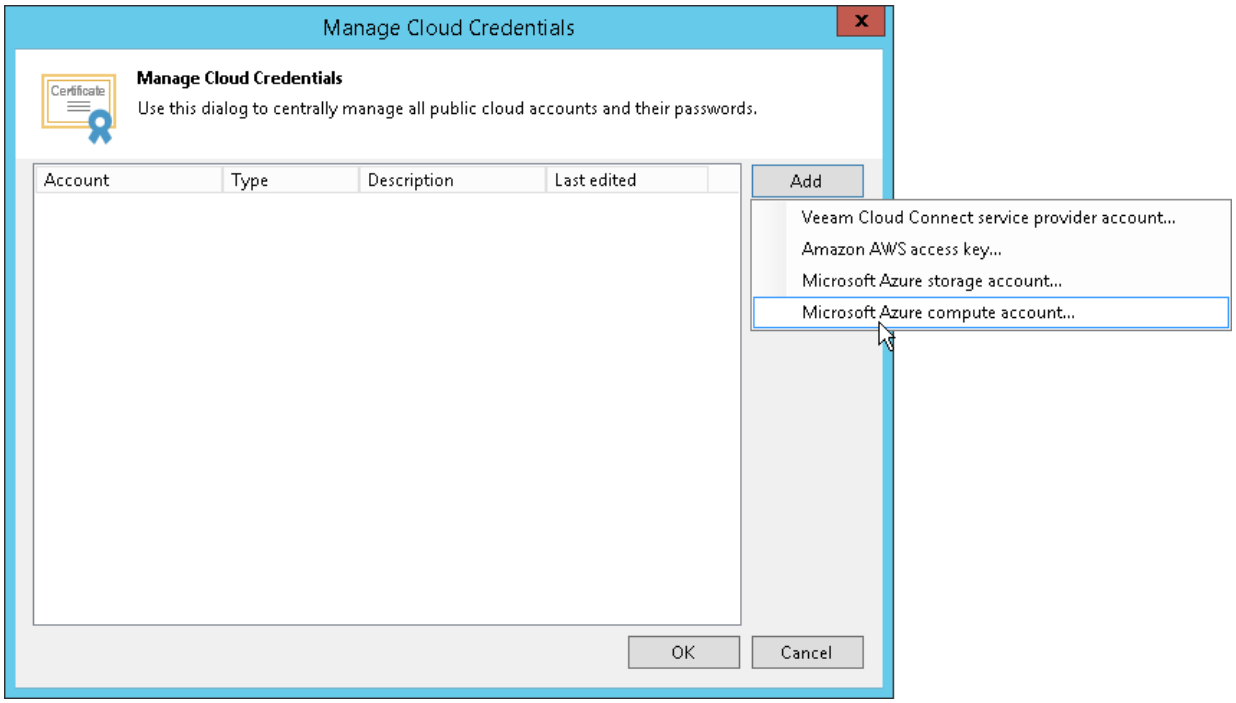
  You may need to additionally disable the Internet Explorer Enhanced Security Configuration in Server Manager.

- On the backup server, you must set the correct time according to the timezone where the backup server is located. Otherwise, you may not be able to add a Microsoft Azure user account to Veeam Backup & Replication.
- When the Internet access is possible only through HTTP/HTTPS proxy, you must configure the proxy settings for the Local System account or account under which the Veeam Backup Service is running. For more information, see this Microsoft article.

Step 1. Launch Microsoft Azure Compute Account Wizard

To launch the **Microsoft Azure Compute Account** wizard, do the following:

1. From the main menu, select **Manage Cloud Credentials**.
2. In the **Manage Cloud Credentials** window, click **Add and select Microsoft Azure compute account**.

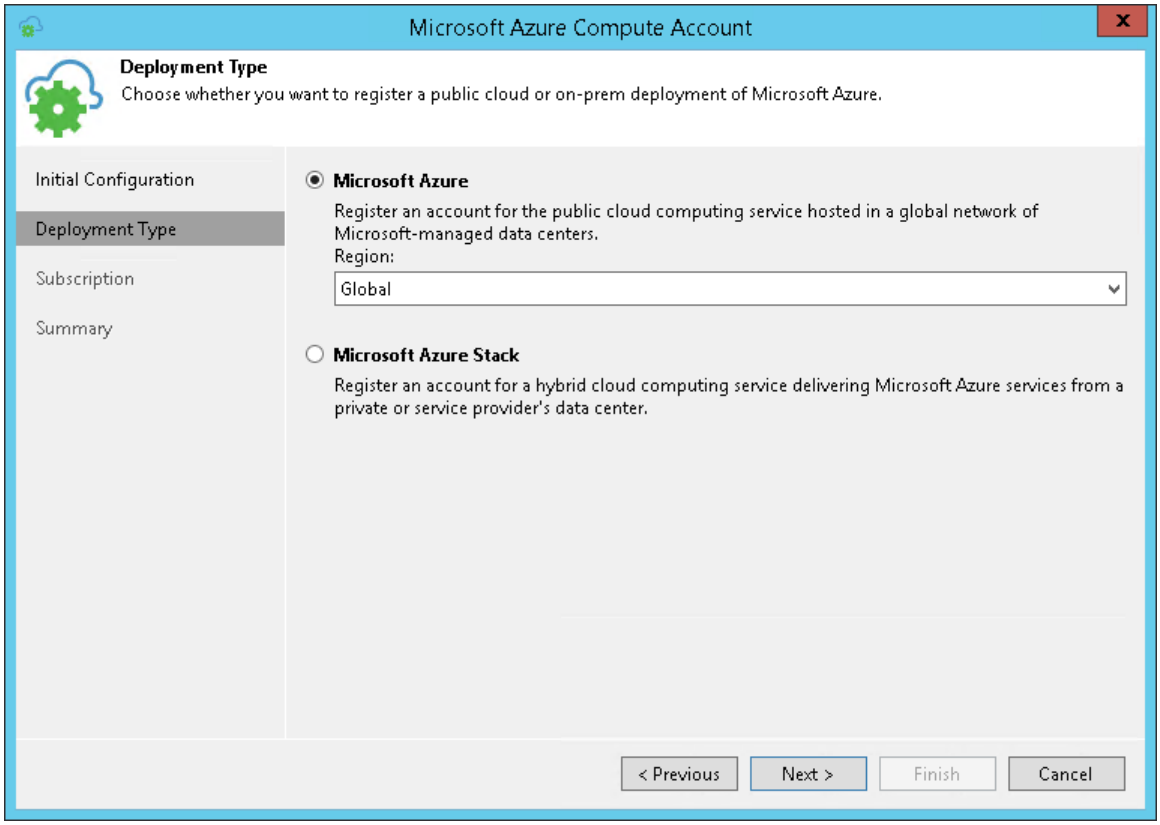


Step 2. Select Deployment Type and Region

To select a deployment type and a region where your Microsoft Azure compute account is located:

1. At the **Initial Configuration** step of the wizard, click **Next**.
2. At the **Deployment Type** step of the wizard, select **Microsoft Azure**.
3. From the **Region** drop-down list, select a Microsoft Azure region.

> **IMPORTANT**
>
> Veeam Backup & Replication no longer supports Azure Germany since Microsoft closed Microsoft Cloud Deutschland.



Step 3. Select Access Type

At the **Subscription** step of the wizard, specify how Veeam Backup & Replication will access Microsoft Azure:

- Using an Azure AD application
- Using an Azure AD user account

Using Azure AD Application

If you select to use an Azure AD application, Veeam Backup & Replication registers a new AD application in Microsoft Azure. Veeam Backup & Replication will use this application to communicate with Azure. For more information on AD applications, see Microsoft Azure Docs.

The created Azure AD application is assigned the *Owner* role privileges for the subscription that will be used for restore to Microsoft Azure. Starting from Veeam Backup & Replication version 11a (build 11.0.1.1261), the Azure AD application is also assigned the *Key Vault Crypto User* role privileges. For more information on roles, see Microsoft Azure Docs.

To use an Azure AD application, do the following:

1. Select **Create a new account**.
2. Click the **Configure account** link.

   Veeam Backup & Replication checks if Microsoft Azure PowerShell is installed on the machine that runs the Veeam Backup & Replication console. If Microsoft Azure PowerShell is not installed, Veeam Backup & Replication will display a warning.

   If Microsoft Azure PowerShell is not installed, do the following:

   a. In the warning window, click **this link**. Veeam Backup & Replication will launch the Microsoft Azure Powershell installation wizard. Follow the steps of the installation wizard to set up the Microsoft Azure PowerShell on the machine.

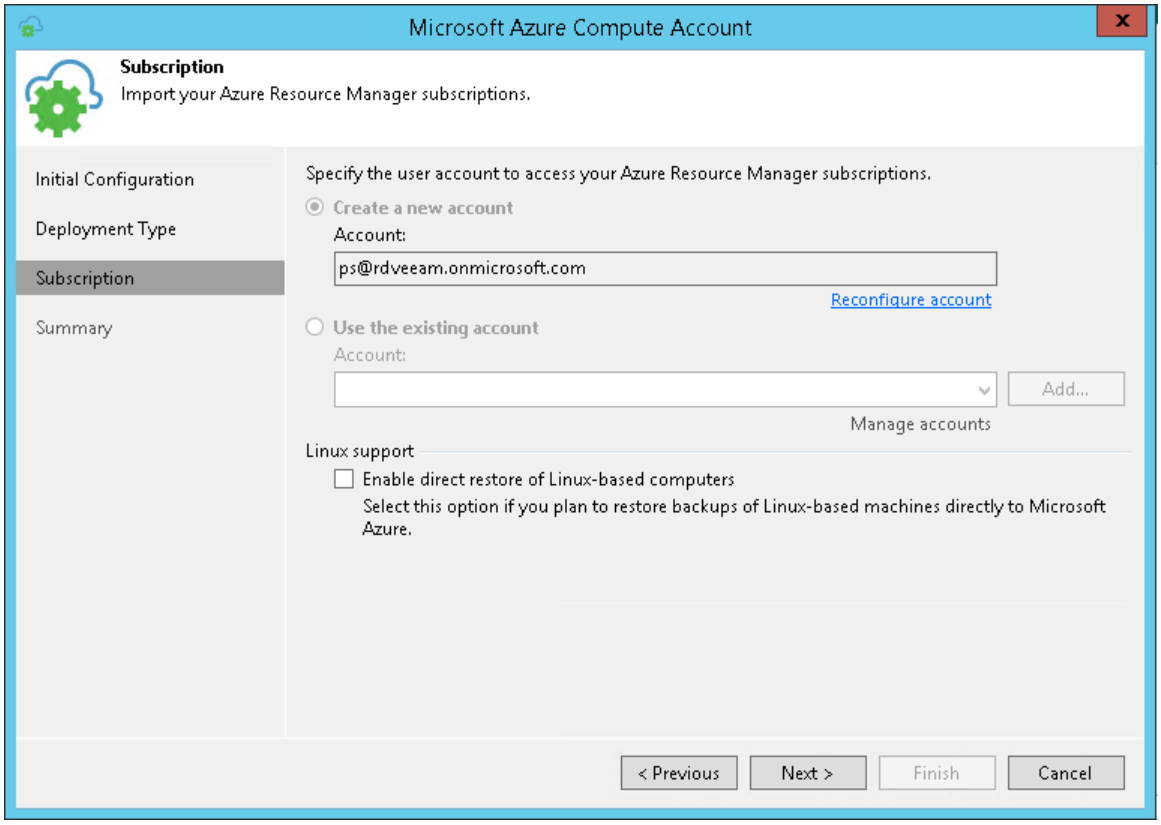
   b. After the installation process is complete, close the Veeam Backup & Replication console. In rare cases,you will need to restart the machine to allow the Veeam Backup & Replication console to detect the newly installed version of Microsoft Azure PowerShell.

   c. Open the Veeam Backup & Replication console and pass through the **Microsoft Azure Compute Account** wizard once again.

   d. Click the **Configure account** link once again.

3. You will be prompted to log in to the Microsoft Azure portal in a separate browser window. Enter credentials of an existing Azure user account in this window. Note that this account must have privileges described in Before You Begin.

   Veeam Backup & Replication will retrieve information about subscriptions and resources associated with the specified account and will create the AD application in this account.



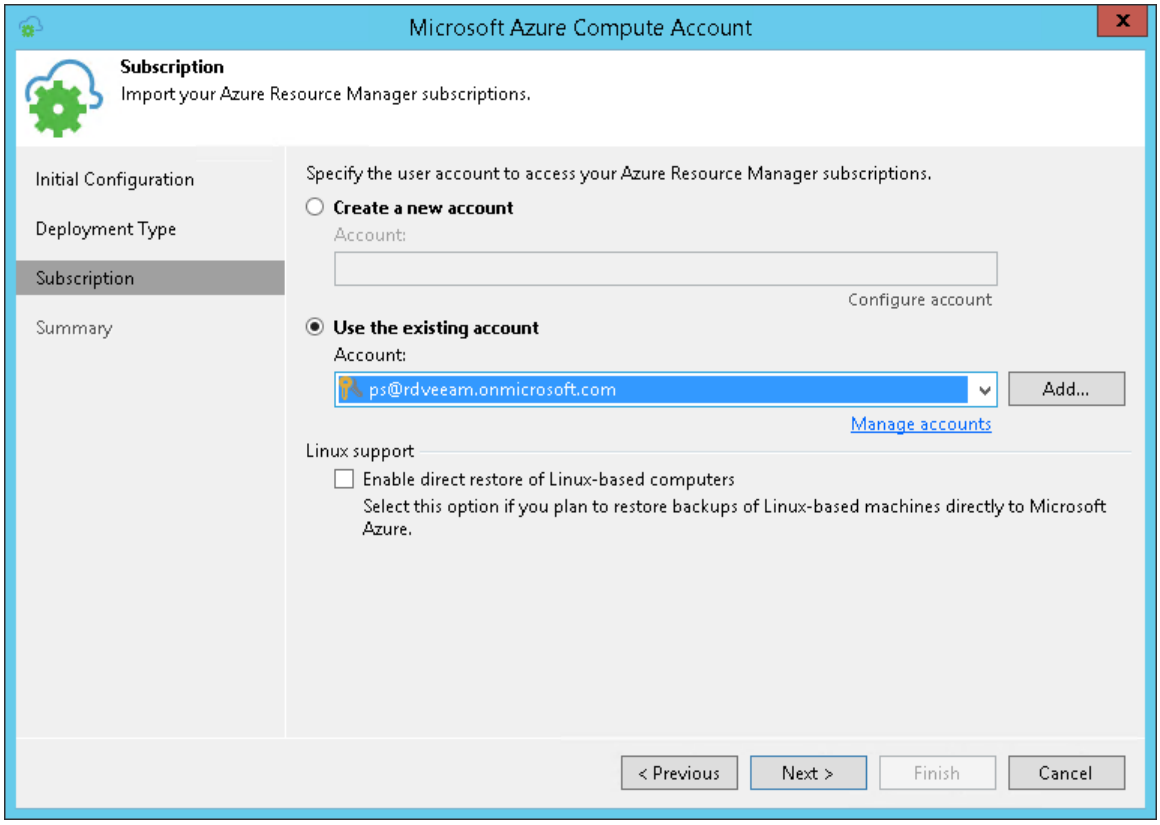
Using Azure AD User Account

To use an existing Azure AD user account:

1. Select **Use the existing account.**

2. From the **Account** drop-down list, select an Azure AD user account you want to use.
Veeam Backup & Replication imports information about subscriptions and resources associated with this account. Note that this account must have privileges described in Before You Begin.

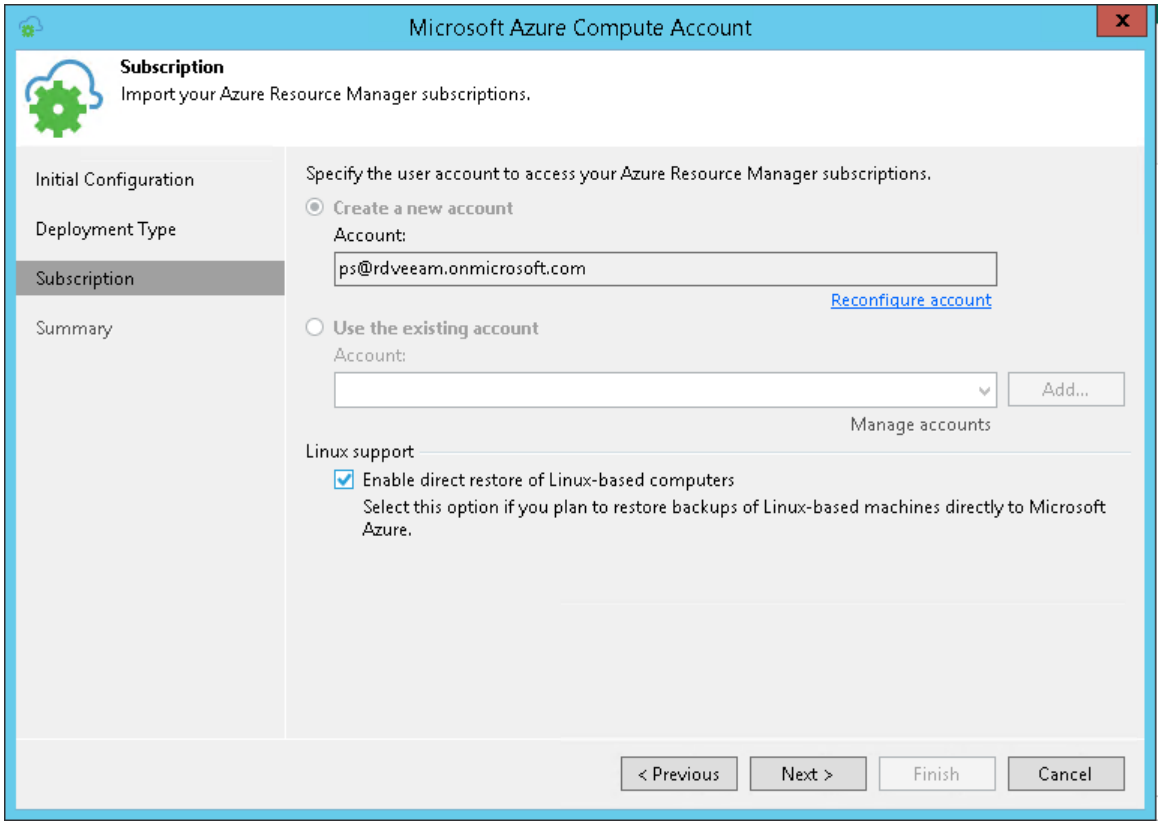
   If you have not added credentials beforehand, click the **Manage accounts** link or **Add** to add the necessary credentials.



Step 4. Enabling Linux Restore

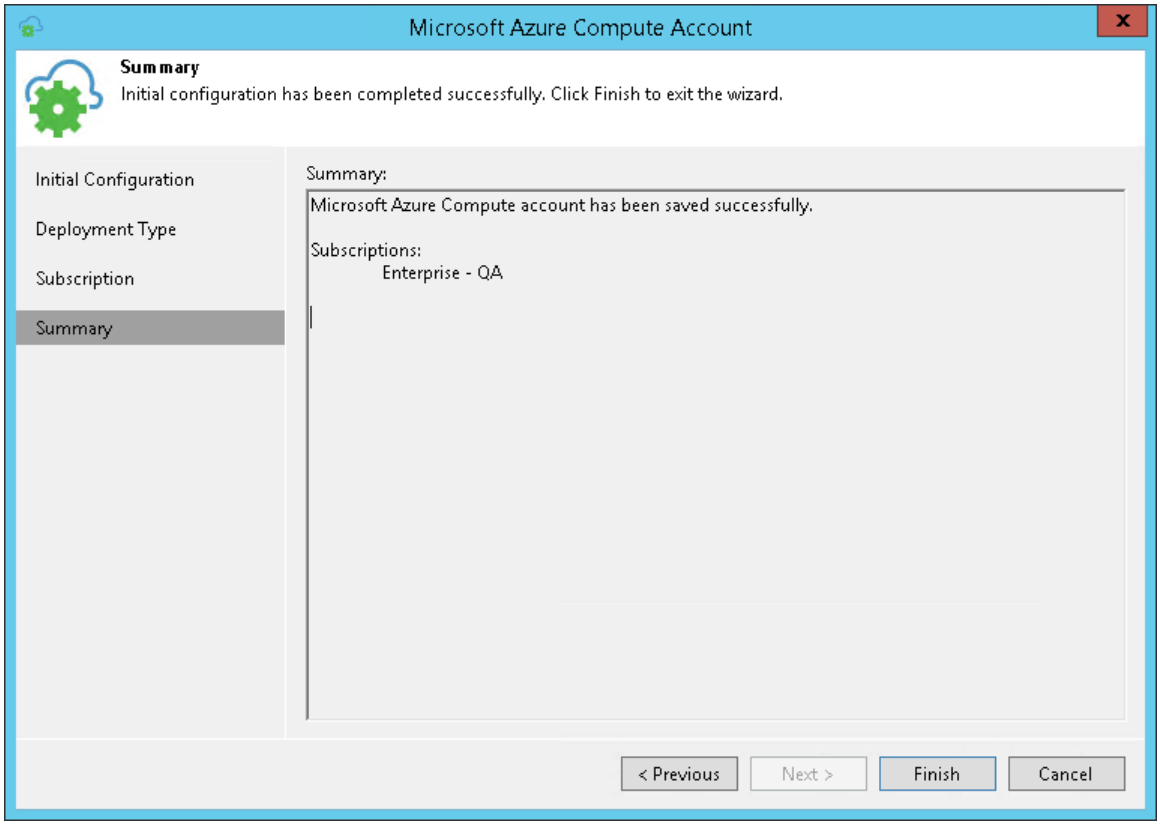
At the **Subscription** step of the wizard, specify whether you plan to restore Linux-based machines. If you are planning to restore them, select the **Enable restore of Linux-based computers** check box.

Veeam Backup & Replication will deploy a helper appliance in Microsoft Azure and use it for restore of Linux machines. For more information about helper appliances and how to setup up them, see Configuring Helper Appliances.



Step 5. Finish Working with Wizard

At the **Summary** step of the wizard, review details of configured settings and click **Finish** to close the wizard.



Microsoft Azure Stack Hub Compute Accounts

To restore workloads to Microsoft Azure Stack Hub, you must add an Azure Stack Hub compute account to Veeam Backup & Replication. When you add an Azure Stack Hub compute account, Veeam Backup & Replication imports information about subscriptions and resources associated with this account. During the restore process, Veeam Backup & Replication accesses these resources and uses them to register new VMs in Azure Stack Hub.

If necessary, you can add different user accounts to Veeam Backup & Replication. In this case, Veeam Backup & Replication will import information about all subscriptions and resources associated with provided accounts, and you will be able to use these resources for restore.

Information about subscriptions and resources is saved to the Veeam Backup & Replication configuration database. You can re-import this information at any time.

Prerequisites

Before restoring workloads to Microsoft Azure Stack Hub, mind the following:

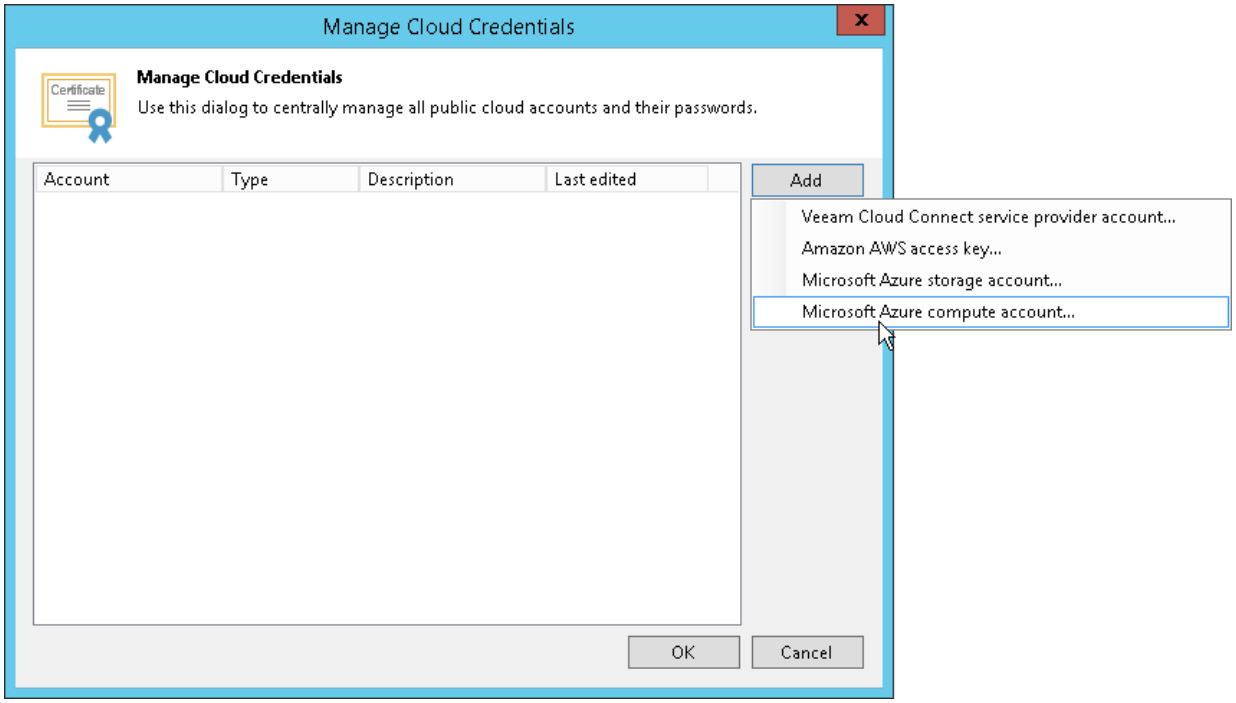
- [For restore to Microsoft Azure Stack Hub version 1808 and later] You must configure the backup server. For more information, see this Veeam KB article.
- Azure Stack Hub tenant user account must have the *Owner* role privileges for the required subscription.

  [For restore to Microsoft Azure] If the *Owner* role cannot be used, you can create a custom role with minimal permissions. To learn how to create a custom role, see Creating Custom Role for Azure Account.
- Accounts with the enabled Multi-Factor Authentication (MFA) are not supported. You must disable MFA for the required account. App passwords are not supported.

Adding Microsoft Azure Stack Hub Account

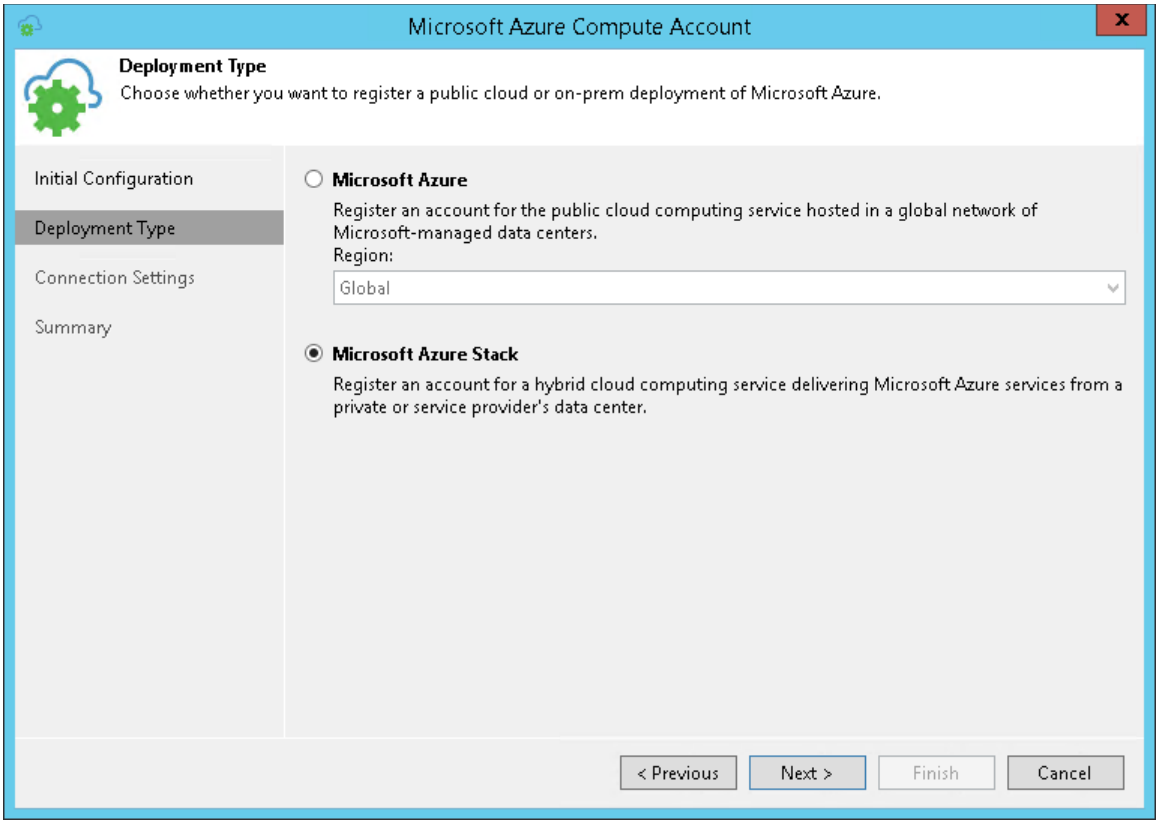
To add a Microsoft Azure Stack Hub compute account, do the following.

1. From the main menu, select **Manage Cloud Credentials**.
2. In the **Manage Cloud Credentials** window, click **Add and select Microsoft Azure compute account**.



3. At the **Initial Configuration** step of the wizard, click **Next**.

4. At the **Deployment Type** step of the wizard, select **Microsoft Azure Stack and click Next.**



5. At the **Connection Settings** step of the wizard, do the following:

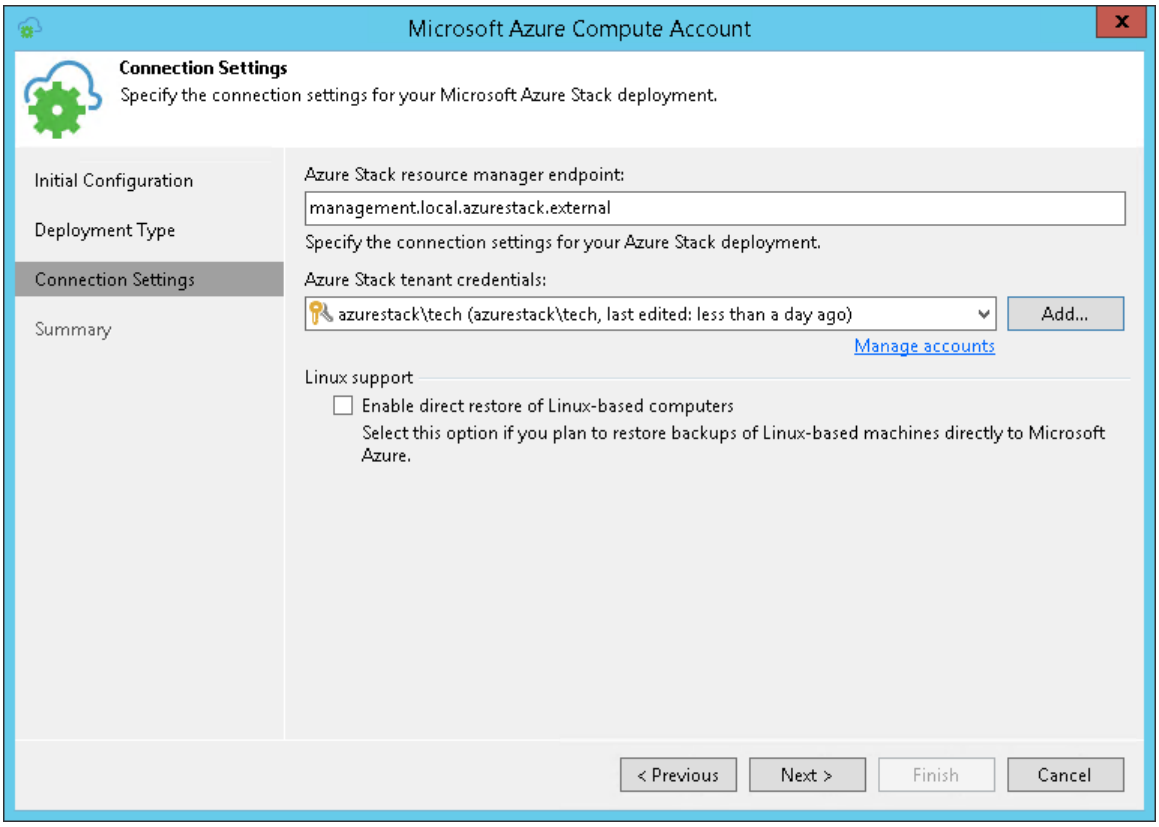
   a. In the **Azure Stack resource manager endpoint** field, specify the virtual IP address of Azure Resource Manager in the following format: `management.<region>.<FQDN>`.

      To learn about Azure Stack Hub virtual IP addresses, see Microsoft Docs.

   b. From the **Azure Stack tenant account** drop-down list, select the necessary account.

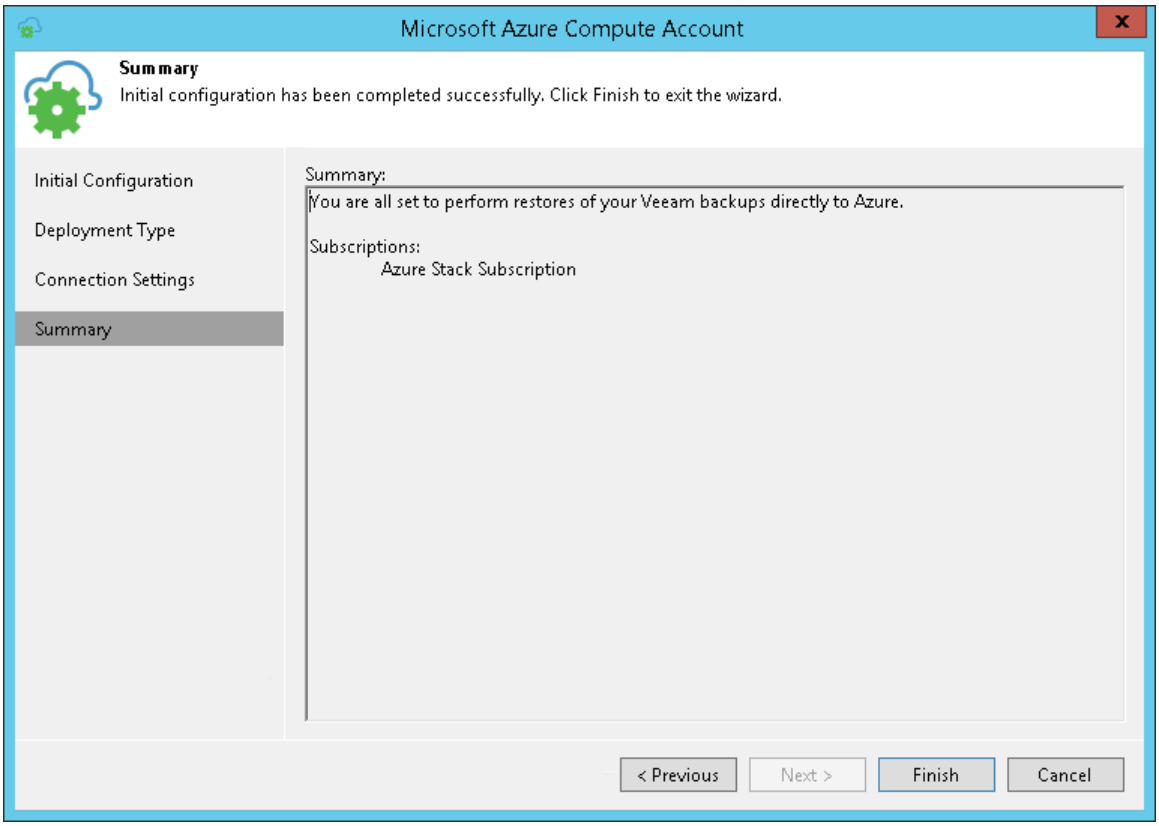
      If you have not added credentials beforehand, click the **Manage accounts** link or click **Add** on the right to add an Azure Stack Hub tenant user credentials.

> **NOTE**
>
> Only subscriptions that belong to the directory of the selected account will be added.



6. If you want to restore Linux-based computers, select the **Enable restore of Linux-based computers** check box. Veeam Backup & Replication will deploy a helper appliance in Microsoft Azure. The helper appliance will be used to restore Linux workloads. For more information about helper appliance setup, see Configuring Helper Appliances.

7. At the **Summary** step of the wizard, review details of configured settings and click **Finish** to close the wizard.



Google Cloud Accounts

You can add to the backup infrastructure a credentials record for the Google Cloud account that will be used to connect to Google Cloud Storage. Veeam Backup & Replication uses these credentials to add a Google Cloud storage repository to the backup infrastructure as an external repository.

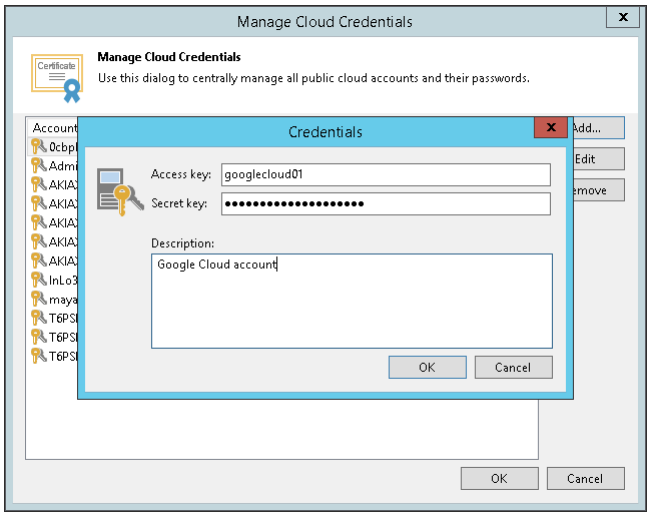
To create a record for a Google Cloud account:

1. From the main menu, select **Manage Cloud Credentials**.
2. Click **Add** > **Google Cloud access key**.
3. In the **Access Key** and **Secret key** fields, specify the Hash-based Message Authentication Code (HMAC) key associated with the Google Cloud account. Veeam Backup & Replication will use the HMAC key to authenticate requests to the Google Cloud Storage. To view the entered key, click and hold the eye icon on the right of the field. For more information on Google Cloud accounts, see the Google Cloud documentation.

   If you have not created the HMAC key beforehand, you can do one of the following:

   - Create the HMAC key in the Google Cloud console, as described in the Google Cloud documentation.
   - Create the HMAC key in the Veeam Backup for Google Cloud Platform console, as described in the Veeam Backup for Google Cloud Platform User Guide.

4. In the **Description** field, enter an optional description for the credentials record.



Google Cloud Platform Service Accounts

> **NOTE**
> This feature is available in Veeam Backup & Replication starting from version 11a (build 11.0.1.1261).

You can create a record for credentials that you plan to use to connect to Google Compute Engine within Google Cloud Platform. This Google Cloud Platform service account is used by Veeam Backup & Replication to perform direct restore to Google Compute Engine and backup and restore operations available with Google Cloud Platform Plug-in for Veeam Backup & Replication. For more information on the latter, see the Integration with Veeam Backup for Google Cloud Platform Guide.

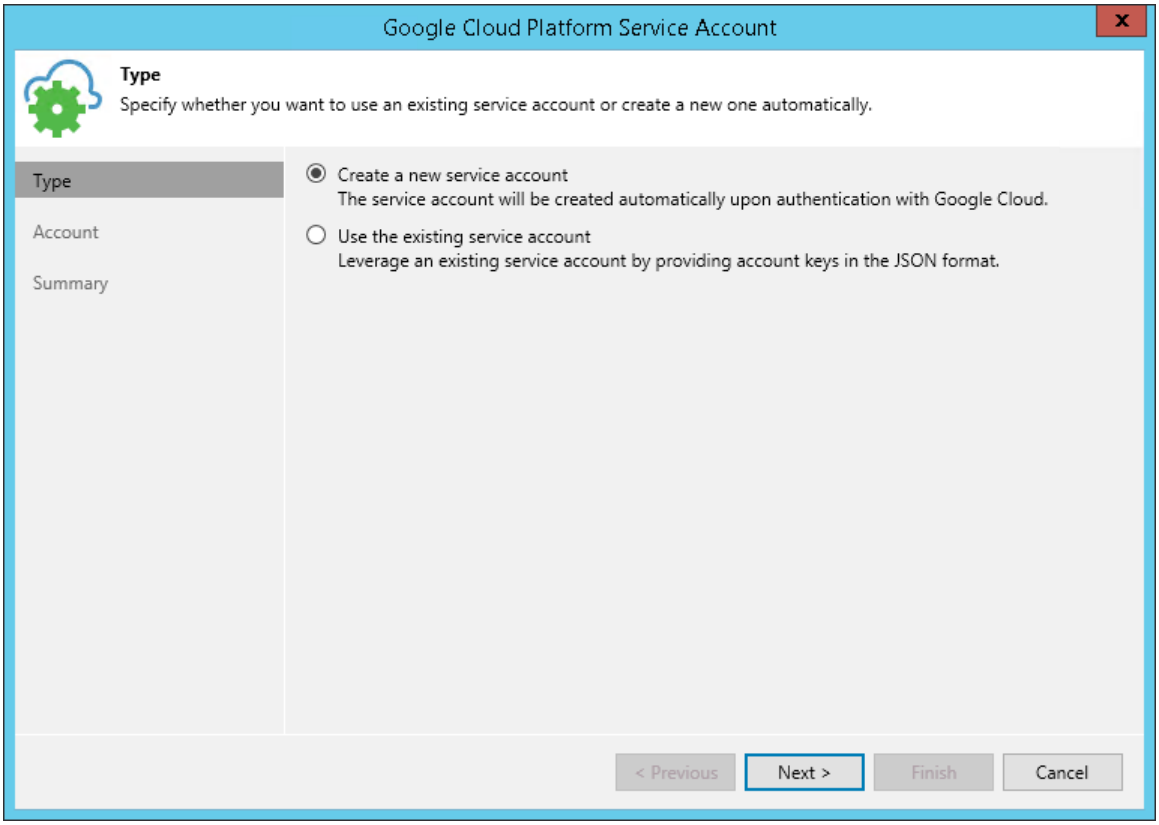
To create a credentials record for a Google Cloud Platform service account:

1. From the main menu, select **Manage Cloud Credentials**.
2. Click **Add** > **Google Cloud Platform service account**.
3. At the **Type** step of the wizard, select if you want to create a new service account automatically or use an existing service account.

NOTE

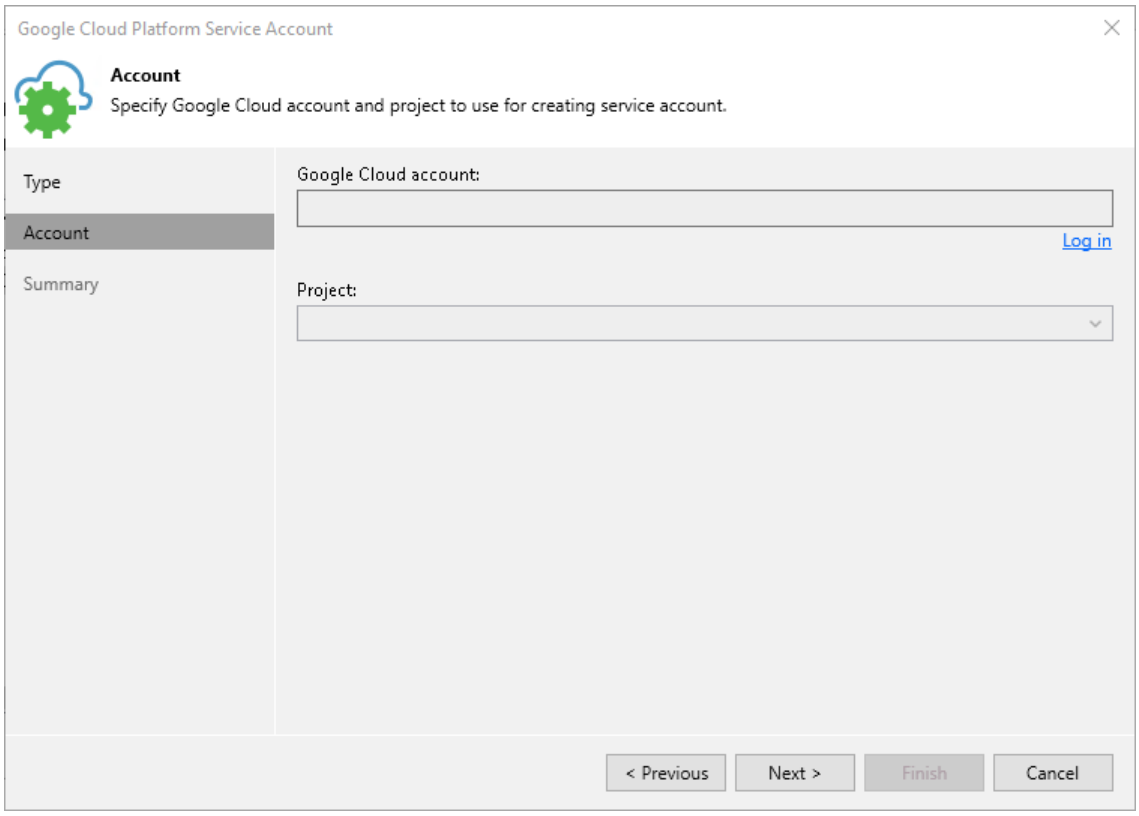
If you select **Create a new service account**, the created service account will be granted the Owner IAM role with a wide scope of permissions and capabilities. If you want to limit the list of permissions granted to the service account, create a user-managed service account, as described in the Google Cloud documentation, with the limited set of permissions:

- For the information on permissions required to restore to Google Compute Engine, see Google Compute Engine IAM User Permissions.
- For the information on permissions required to deploy GCP Plug-in for Veeam Backup & Replication, see the Required Permissions section in the Integration with Veeam Backup for Google Cloud Platform Guide.



4. At the **Account** step of the wizard, specify credentials required for accessing the service account:
   - If you have selected **Create a new service account**, do the following:
     i. Log into your Google Cloud account. Read and accept the Google Terms of Service and the Google Privacy Policy.
     ii. Allow Veeam Backup & Replication to access your Google account. After that, Veeam Backup & Replication can manage your Identity and Access Management (IAM) policies, and see, edit, configure and delete your Google Cloud Platform data.

iii. Return to the wizard and select the project with which you want the created service account to work.

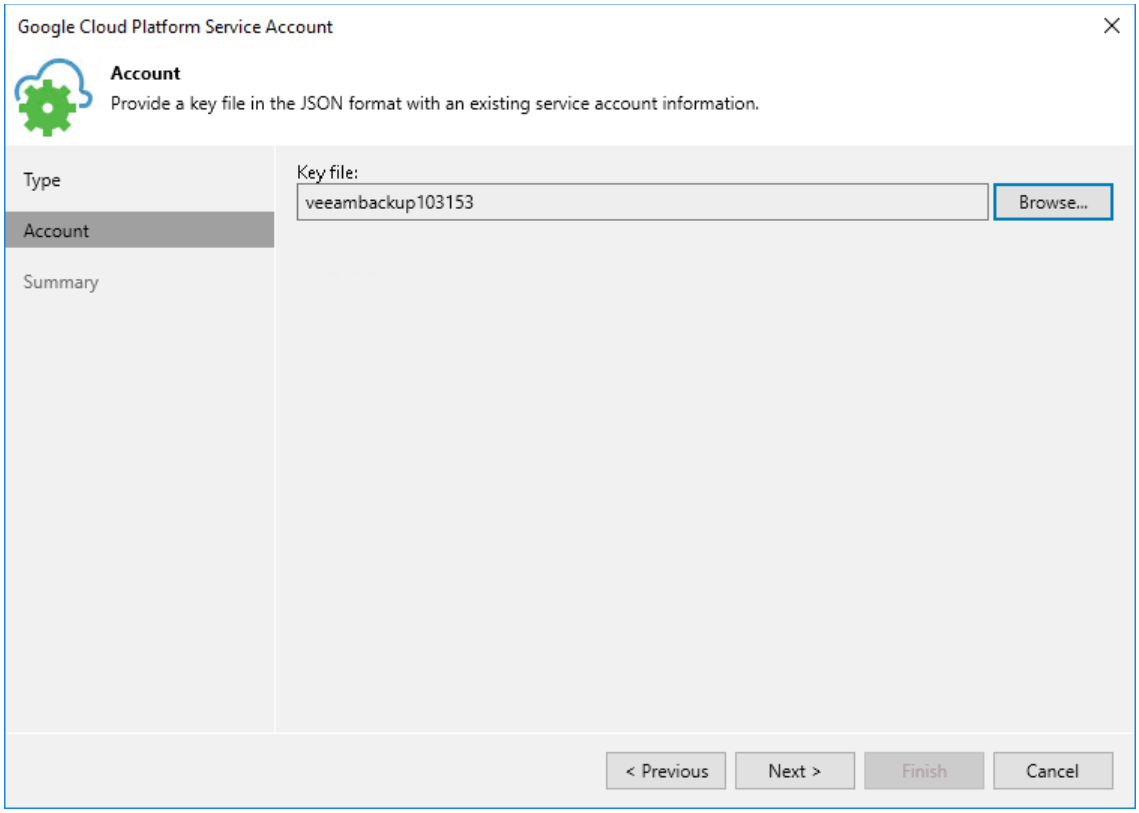


- If you have selected **Use the existing service account**, do the following:

  i. Download the service account key in the JSON format, created as described in Google Cloud documentation.

  > **IMPORTANT**
  >
  > Depending on the scenarios that the service account will be used for, make sure that the service account meets all requirements and limitations.
  >
  > For restoring virtual workloads from backups to Google Cloud, mind the requirements and limitations listed in Restore to Google Compute Engine.

ii. At the **Account** step of the wizard, select the downloaded service account key.



5. At the **Summary** step of the wizard, review details of the configured account and click Finish to close the wizard.

Editing and Deleting Credentials Records

You can edit or delete existing cloud credentials records.

To edit a credentials record:

1. From the main menu, select **Manage Cloud Credentials**.
2. Select the credentials record in the list and click **Edit**.
3. Edit settings of the credentials record as required.

To delete a credentials record:

1. From the main menu, select **Manage Credentials**.
2. Select the credentials record in the list and click **Remove**. You cannot delete a record that is already used for any component in the backup infrastructure.

> **TIP**
>
> You can use the Cloud Credentials Manager to change the password for a tenant account provided by the SP. For more information, see the Changing Password for Tenant Account section in the Veeam Cloud Connect Guide.

Password Manager

You can use the Password Manager to create and maintain a list of passwords that you plan to use for data encryption. Password management can be helpful in the following situations:

- You want to create new passwords. You can use one password per job or share the same password between several jobs on the backup server.
- You want to edit an existing password, for example, change its hint, or delete a password.

> **TIP**
>
> Periodical change of passwords is a security best practice. You can create new passwords as often as you need based on your company security needs and regulatory requirements.

Creating Passwords

You can use the Password Manager to create one or more passwords.

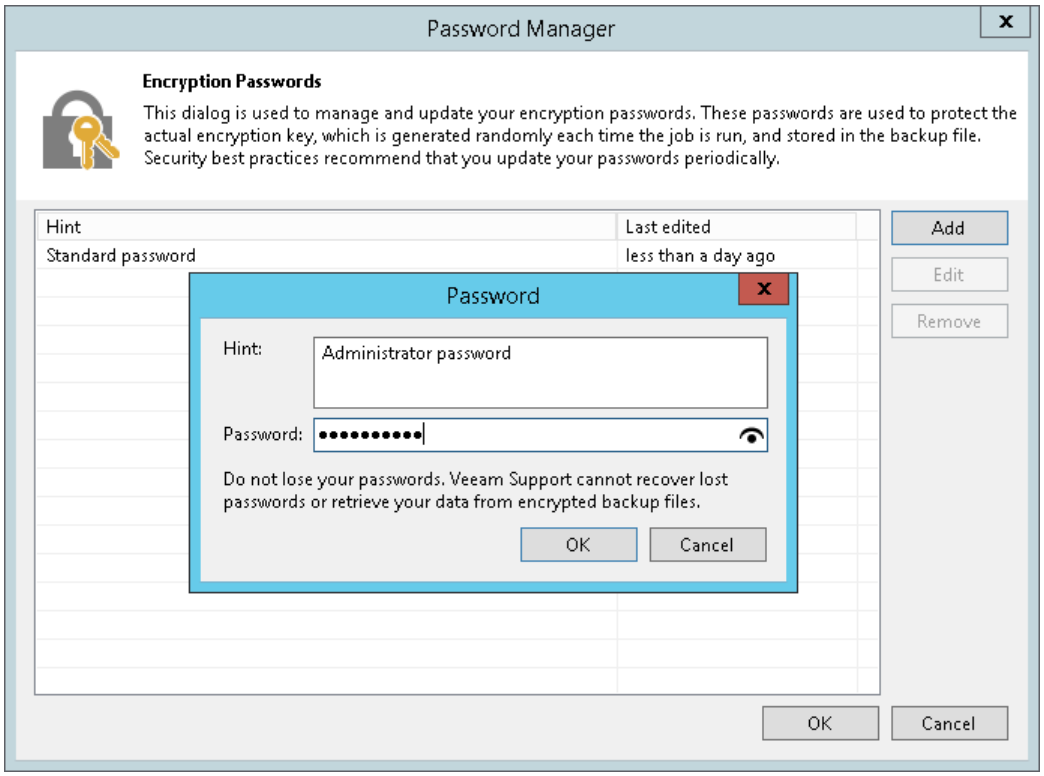
To create a new password:

1. From the main menu, select **Manage Passwords**. Alternatively, you can use job properties to create a new password:
   a. Open the **Home** view.
   b. In the inventory pane, select **Jobs**.
   c. In the working area, right-click the backup or backup copy job and select **Edit**.
   d. At the **Storage** step of the wizard (for backup job) or **Target** step of the wizard (for backup copy job), click **Advanced**.
   e. Click the **Storage** tab.
   f. In the **Encryption** section of the **Advanced Setting** window, select the **Enable backup file encryption** check box and click the **Manage passwords** link.

   Veeam Backup & Replication will open the Password Manager.
2. In the Password Manager, click **Add**.
3. In the **Hint** field, specify a hint for the created password. It is recommended that you provide a meaningful hint that will help you recall the password. The password hint is displayed when you import an encrypted file on the backup server and access this file.
4. In the **Password** field, enter a password. To view the entered password, click and hold the eye icon on the right of the field.

IMPORTANT

Always save a copy of the password you create in a secure place. If you lose the password, you will not be able to restore it.



Editing Passwords

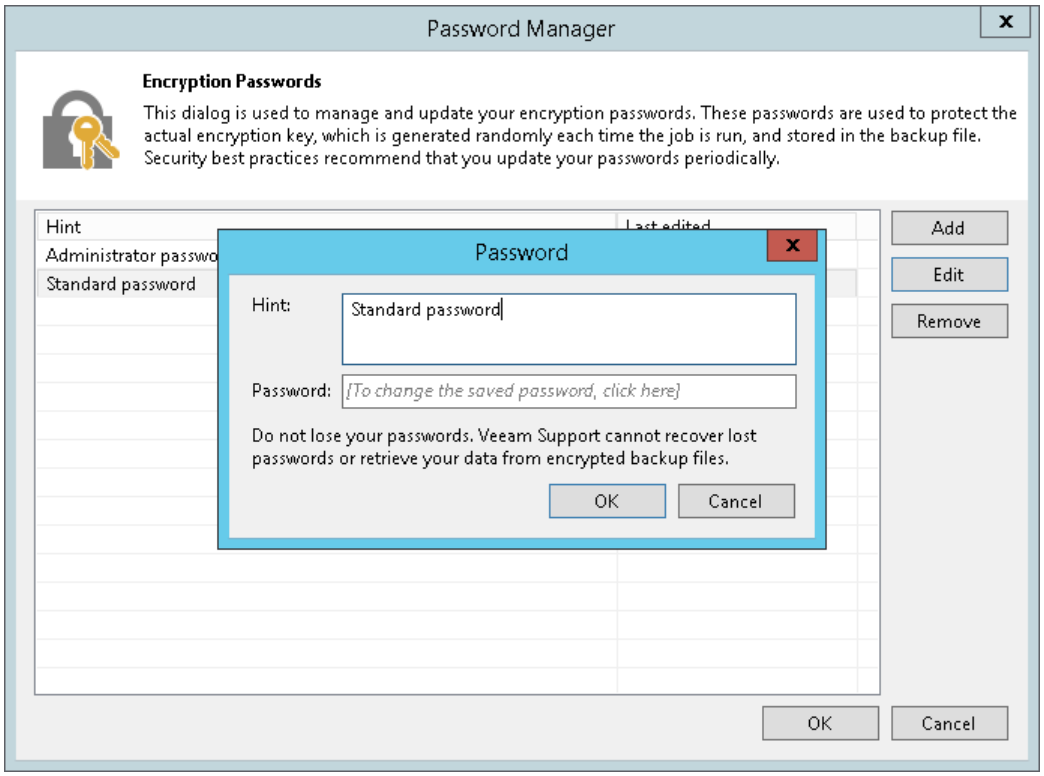
You can edit passwords you have created using the Password Manager.

To edit a password:

1. From the main menu, select **Manage passwords**. Alternatively, you can use job properties to edit the password:
   a. Open the **Home** view.
   b. In the inventory pane, select **Jobs**.
   c. In the working area, right-click the backup or backup copy job and select **Edit**.
   d. At the **Storage** step of the wizard (for backup job) or **Target** step of the wizard (for backup copy job), click **Advanced**.
   e. Click the **Storage** tab.
   f. In the **Encryption** section of the **Advanced Setting** window, select the **Enable backup file encryption** check box and click the **Manage passwords** link.

   Veeam Backup & Replication will open the Password Manager.
2. In the Password Manager, select the password and click **Edit**.
3. Edit the password data: hint and password, as required.

After you edit the password, you do not need to perform any other actions. Veeam Backup & Replication will start using the changed password after a job runs for the next time. Before the job run, the old password is still used. For more information on which password to use when you restore data, see Restoring Data from Encrypted Backups.



Deleting Passwords

You can delete passwords using the Password Manager.

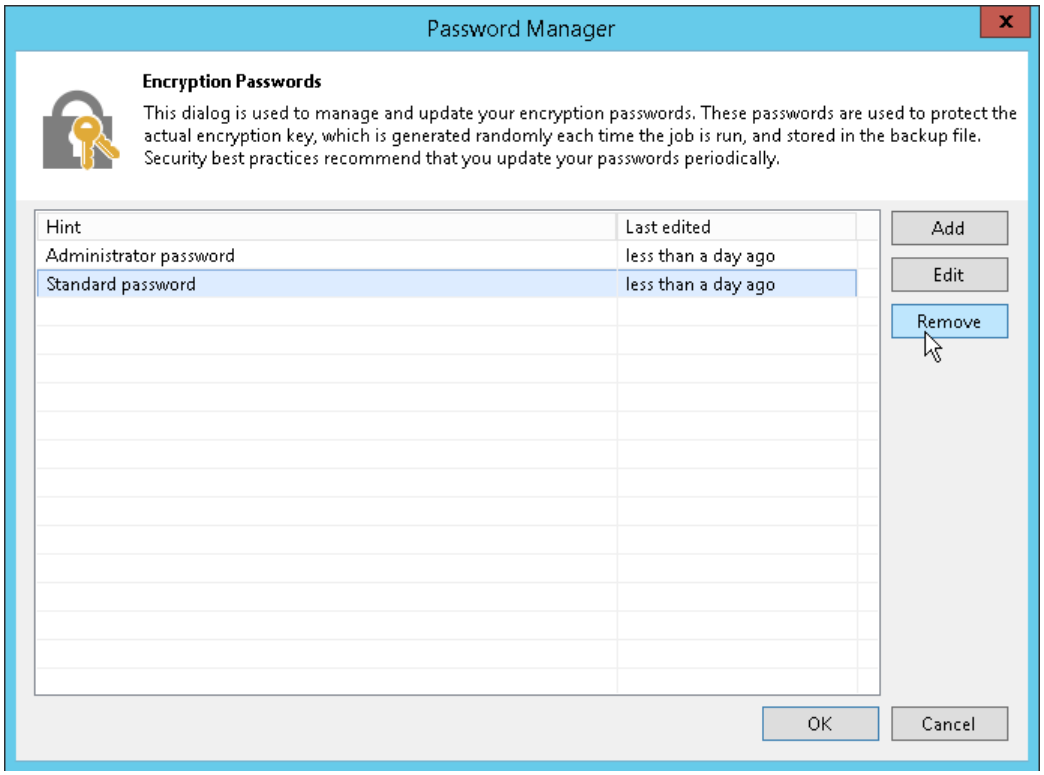
You cannot remove a password that is currently used by any job on the backup server. To remove such password, you first need to delete a reference to this password in the job settings.

To delete a password:

1. From the main menu, select **Manage passwords**. Alternatively, you can use job properties to delete passwords:
   a. Open the **Home** view.
   b. In the inventory pane, select **Jobs**.
   c. In the working area, right-click the backup or backup copy job and select **Edit**.
   d. At the **Storage** step of the wizard (for backup job) or **Target** step of the wizard (for backup copy job), click **Advanced**.
   e. Click the **Storage** tab.
   f. In the **Encryption** section of the **Advanced Setting** window, select the **Enable backup file encryption** check box and click the **Manage passwords** link.

   Veeam Backup & Replication will open the Password Manager.

2. In the Password Manager, select the password and click **Remove**.



Backup & Replication Console

The Veeam Backup & Replication console is a client-side component that provides access to the backup server. The console allows you to log in to Veeam Backup & Replication and perform all kinds of data protection and disaster recovery operations on the backup server.

The console does not have a direct access to the backup infrastructure components and configuration database. Such data as user credentials, passwords, roles and permissions are stored on the backup server side. To access this data, the console needs to connect to the backup server and query this information periodically during the work session.

To make users work as uninterrupted as possible, the remote console maintains the session for 5 minutes if the connection is lost. If the connection is re-established within this period, you can continue working without re-logging to the console.

Backup & Replication Console Deployment

The console is installed locally on the backup server by default. You can also use it in a standalone mode — install the console on any machine and access Veeam Backup & Replication remotely over the network.

You can install as many remote consoles as you need so that multiple users can access Veeam Backup & Replication simultaneously. Veeam Backup & Replication prevents concurrent modifications on the backup server. If several users are working with Veeam Backup & Replication at the same time, the user who saves the changes first has the priority. Other users will be prompted to reload the wizard or window to get the most recent information about the changes in the configuration database.

If you have multiple backup servers in the infrastructure, you can connect to any of them from the same console. For convenience, you can save several shortcuts for these connections.

> **IMPORTANT**
>
> You cannot use the same console to connect to backup servers with different versions of Veeam Backup & Replication. Mind this if you have more than one backup server in your backup environment, and these backup servers run different versions of Veeam Backup & Replication. For example, if one of your backup servers run version 9.5 Update 3, and another backup server runs version 9.5 Update 4, you will need to use 2 separate consoles for connecting to these servers.

The console supports automatic update. Every time you connect to the backup server locally or remotely, the console checks for updates. If the backup server has updates installed, the console will be updated automatically.

Mind the following:

- Upgrade to another Veeam Backup & Replication major product version is not supported. If you upgrade Veeam Backup & Replication to another major version, you must upgrade the console to the same version manually. Automatic upgrade is not supported for Preview, Beta or RTM versions of Veeam Backup & Replication.
- Downgrade of the console is not supported. If the console is of a higher version than the backup server (for example, you have upgraded the console manually), the connection to the server will fail.

If other Veeam Backup & Replication components, such as Veeam Cloud Connect Portal or Veeam Backup Enterprise Manager, are installed on the machine where the console runs, these components will also be upgraded.

Backup & Replication Console Components

When you install a remote console on a machine, Veeam Backup & Replication installs the following components:

- Veeam Backup PowerShell Module
- Veeam Explorer for Microsoft Active Directory
- Veeam Explorer for Microsoft Exchange
- Veeam Explorer for Oracle
- Veeam Explorer for Microsoft SQL
- Veeam Explorer for Microsoft SharePoint
- Veeam Explorer for Microsoft OneDrive for Business
- Mount server

Backup & Replication Console User Access Rights

To log in to Veeam Backup & Replication using the console, the user must be added to the Local Users group on the backup server or a group of domain users who have access to the backup server. The user can perform the scope of operations permitted by his or her role in Veeam Backup & Replication. For more information, see Assigning Roles to Users.

Requirements for Backup & Replication Console

A machine on which you install the Veeam Backup & Replication console must meet the following requirements:

- The machine must meet the system requirements. For more information, see System Requirements.
- The remote console can be installed on a Microsoft Windows machine (physical or virtual).
- If you install the console remotely, you can deploy it behind NAT. However, the backup server must be outside NAT. The opposite type of deployment is not supported: if the backup server is deployed behind NAT and the remote console is deployed outside NAT, you will not be able to connect to the backup server.

Limitations for Backup & Replication Console

The Veeam Backup & Replication console has the following limitations:

- You cannot perform restore from the configuration backup using the remote console.
- The machines on which the remote console is installed are not added to the list of managed servers automatically. For this reason, you cannot perform some operations, for example, import backup files that reside on the remote console machine or assign roles of backup infrastructure components to this machine. To perform these operations, you must add the remote console machine as a managed server to Veeam Backup & Replication. For more information, see Managing Servers.

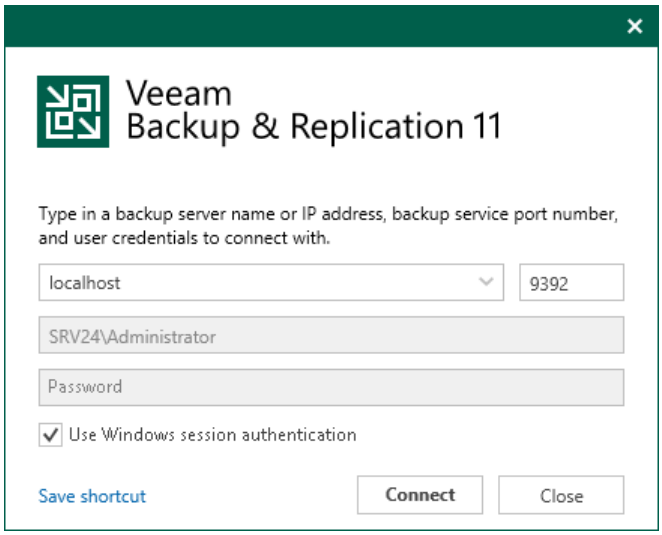
Logging in to Veeam Backup & Replication

To log in to Veeam Backup & Replication, you must open the Veeam Backup & Replication console and specify connection settings to access the backup server.

1. To open the Veeam Backup & Replication console, do one of the following:
   - Double-click the console icon on the desktop.
   - From the Microsoft Windows **Start** menu, select **All Programs** > **Veeam** > **Veeam Backup & Replication Console**.
   - Use the Microsoft Windows search to find the **Veeam Backup & Replication Console** program on the computer.
2. In the **Server** field, type the name or IP address of the backup server or select it from the list of recent connections. By default, the console connects to the backup server installed locally — localhost.
3. In the **Port** field, enter the port over which you want to connect to the backup server. The port number is set at the **Port Configuration** step of the setup wizard for Veeam Backup & Replication. By default, port 9392 is used.
4. In the **Username** and **Password** fields, enter credentials of the user account that you want to use to connect to the backup server. The user account must be added to the Local Users group on the backup server or a group of domain users who have access to the backup server.

   You can also select the **Use Windows session authentication** check box. In this case, you will log in to Veeam Backup & Replication using the account under which you are currently logged in to Microsoft Windows.
5. To create a shortcut for the connection, click **Save shortcut**. You can create as many shortcuts as you need.

> **NOTE**
>
> If you create a shortcut for a connection, the credentials for this connection will be stored in the Windows Credentials Manager. The credentials are saved after the first successful login.



Virtualization Servers and Hosts

You can add the following types of servers and hosts to the backup infrastructure:

- VMware vSphere Server
- VMware vCloud Director
- Microsoft Windows Server
- Linux Server

> **NOTE**
>
> We recommend that only one instance of a server or host is present in the backup infrastructure at a time. Do not add the same server or host multiple times, for example, by a DNS name and IP address, this can cause unexpected behavior.

You can add physical machines and VMs to the backup infrastructure and assign different roles to them. The table below describes which roles can be assigned to the different types of servers.

Server Type	Source Host	Target Host	Backup Proxy	Backup Repository
VMware vSphere Server (standalone ESXi host or vCenter Server)	●	●	○	○
VMware vCloud Director	●	○	○	○
Microsoft Windows server	○	○	●	●
Linux server	○	○	●	●

NOTE

You can also add the following servers:

- Hyper-V servers. For more information on how to add servers, see the Adding Microsoft Hyper-V Servers section in the Veeam Backup & Replication User Guide for Microsoft Hyper-V.
- Veeam Backup for AWS servers. For more information on how to add servers, see the Deployment and Configuration section in the Integration with Veeam Backup for AWS Guide.
- Veeam Backup for Microsoft Azure servers. For more information on how to add servers, see the Deployment and Configuration section in the Integration with Veeam Backup for Microsoft Azure Guide.
- Veeam Backup for Google Cloud Platform servers. For more information on how to add servers, see the Deployment and Configuration section in the Integration with Veeam Backup for Google Cloud Platform Guide.
- Nutanix AHV clusters. For more information on how to add AHV clusters, see the Deployment section in the Veeam Backup for Nutanix AHV User Guide.

Related Topics

- Veeam Data Movers
- Rescanning Servers
- Editing Server Settings
- Removing Servers

Adding VMware vSphere Servers

You must add to the backup infrastructure VMware vSphere servers that you plan to use as source and target for backup, replication and other activities.

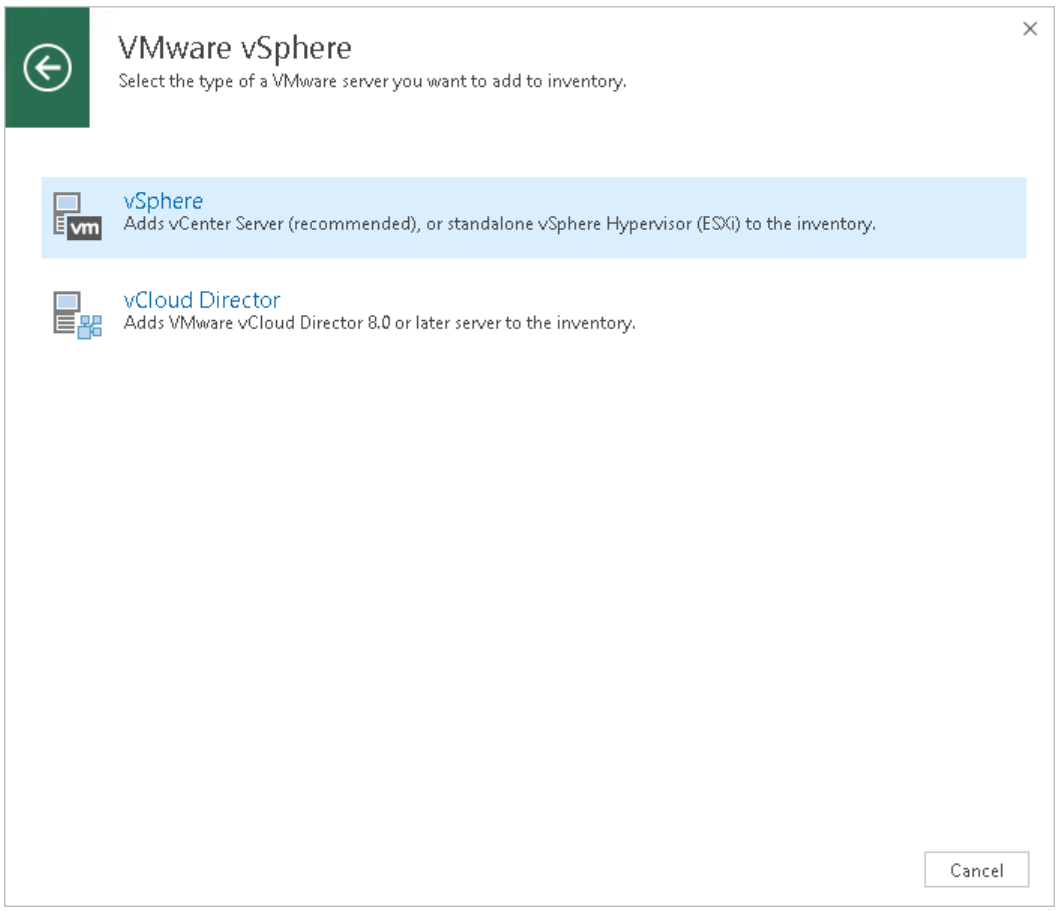
You can add vCenter Servers and ESXi hosts. If an ESXi host is managed by a vCenter Server, it is recommended that you add the vCenter Server, not a standalone ESXi host. If you move VMs between ESXi hosts managed by the vCenter Server, you will not have to re-configure jobs in Veeam Backup & Replication.
Veeam Backup & Replication will automatically locate migrated VMs and continue processing them as usual.

To add a VMware vSphere server, use the **New VMware Server** wizard.

Step 1. Launch New VMware Server Wizard

To launch the **New VMware Server** wizard, do one of the following:

- Open the **Backup Infrastructure** view. In the inventory pane, select the **Managed Servers** node and click **Add Server** on the ribbon or right-click the **Managed Servers** node and select **Add Server.** In the **Add Server** window, click **VMware vSphere** > **vSphere**.
- Open the **Inventory** view, in the inventory pane select the **VMware vSphere** node and click **Add Server** on the ribbon. You can also right-click the **VMware vSphere** node and select **Add Server**.



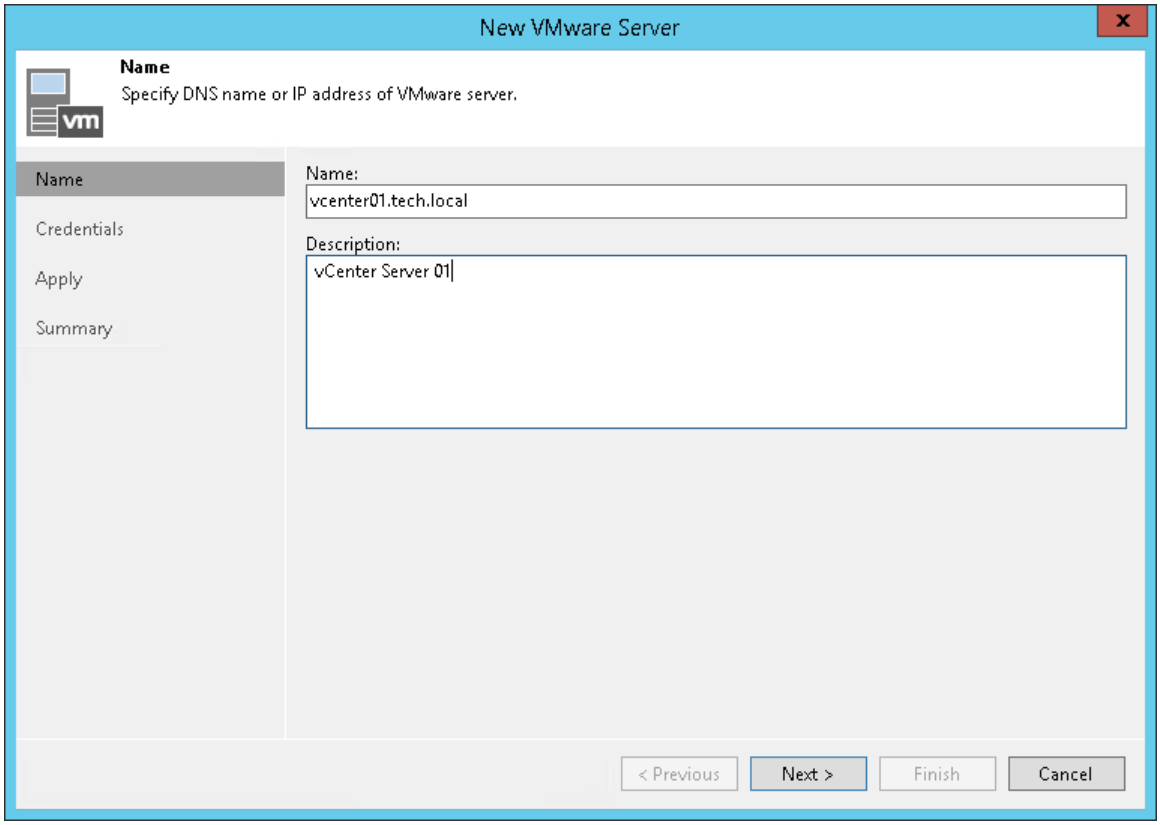
Step 2. Specify Server Name or Address

At the **Name** step of the wizard, specify an address and description for the VMware vSphere server.

1. Enter a full DNS name or IP address of the vCenter Server or standalone ESXi host.

   If you add a VMware Cloud on AWS vCenter Server, use its Fully Qualified Domain Name (FQDN). Make sure the name you specify ends with <vmc.vmware.com>.

2. Provide a description for future reference. The default description contains information about the user who added the server, date and time when the server was added.



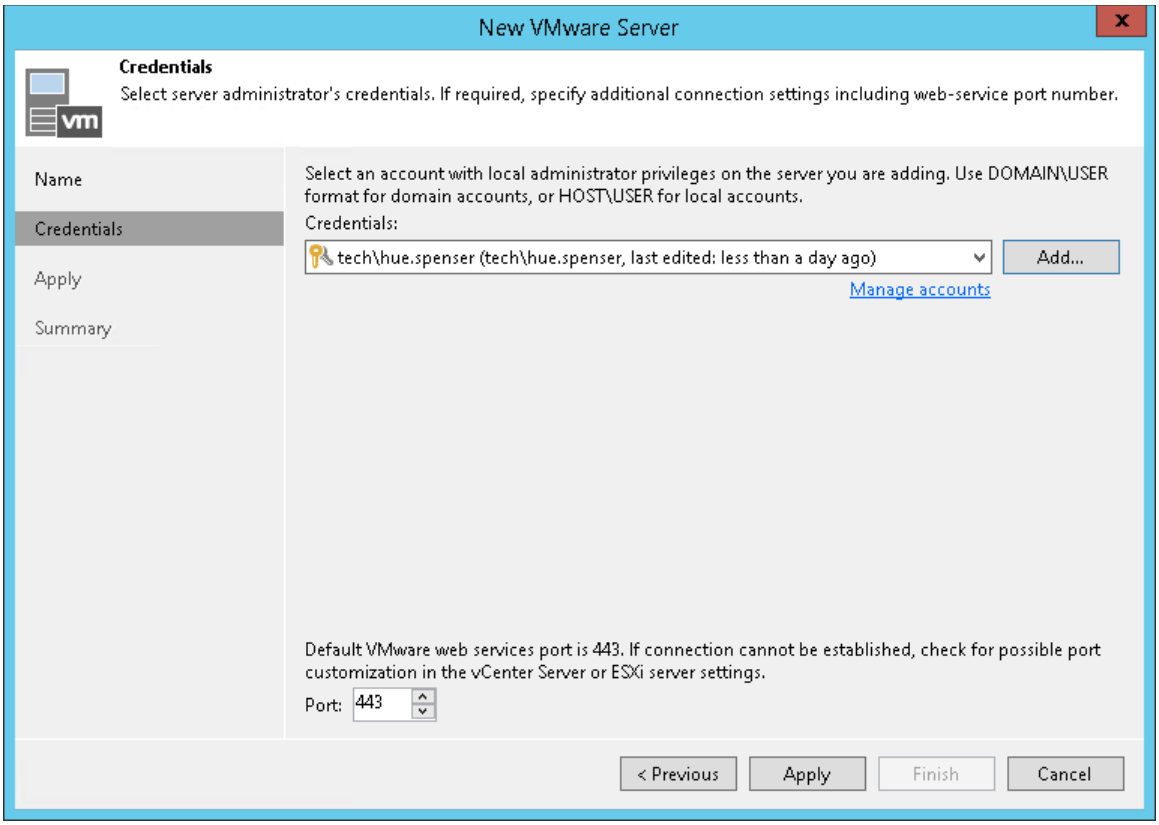
Step 3. Specify Credentials

At the **Credentials** step of the wizard, specify credentials and port settings for the VMware vSphere server.

1. From the **Credentials** list, select credentials with the required permissions. For more information, see Required Permissions.

   If you have not set up credentials beforehand, click the **Manage accounts** link or click **Add** on the right to add the credentials. For more information, see Managing Credentials.

2. By default, Veeam Backup & Replication uses port 443 to communicate with vCenter Servers and ESXi hosts. If a connection with the vCenter Server or ESXi host over this port cannot be established, you can customize the port number in vCenter Server/ESXi host settings and specify the new port number in the **Port** field.



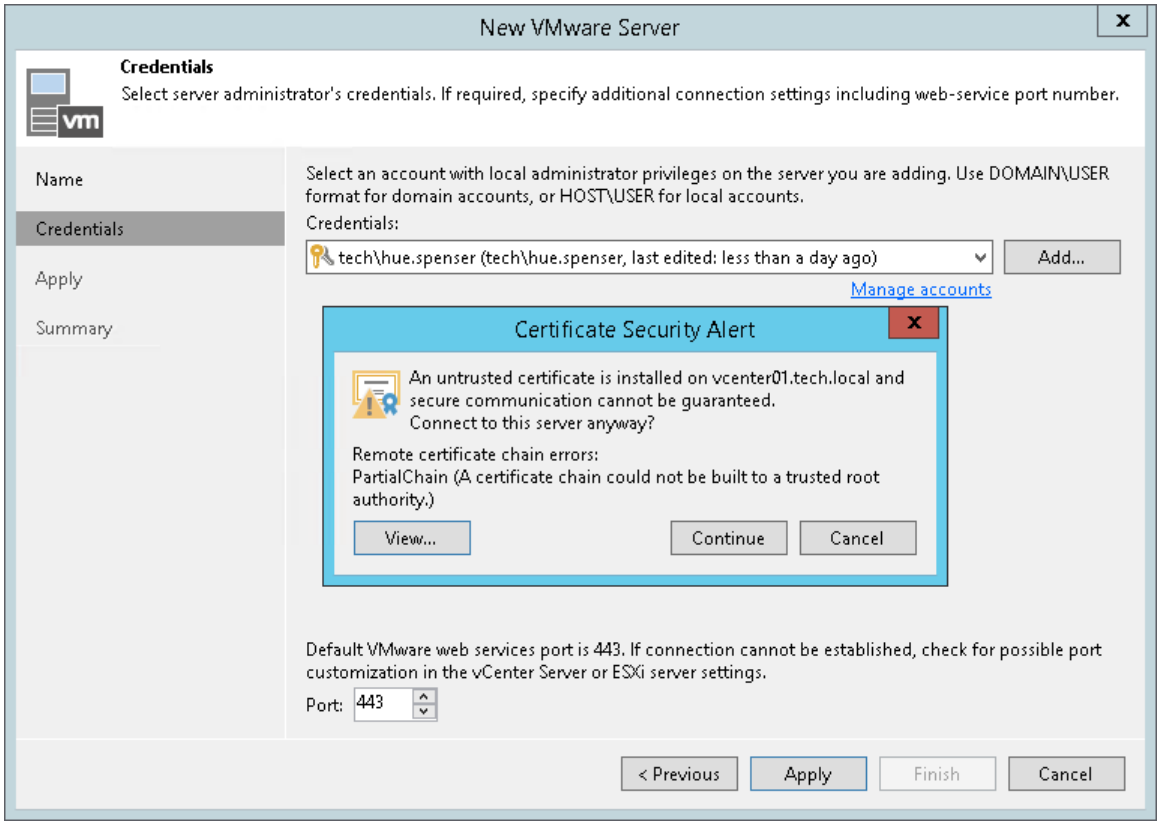
3. When you add a vCenter Server or ESXi host, Veeam Backup & Replication saves to the configuration database a thumbprint of the TLS certificate installed on the vCenter Server or ESXi host. During every subsequent connection to the server, Veeam Backup & Replication uses the saved thumbprint to verify the server identity and avoid the man-in-the-middle attack. For details on managing TLS Certificates, see TLS Certificates.

   If the certificate installed on the server is not trusted, Veeam Backup & Replication displays a warning.

   - To view detailed information about the certificate, click **View**.
   - If you trust the server, click **Continue**.
   - If you do not trust the server, click **Cancel**. Veeam Backup & Replication will display an error message, and you will not be able to connect to the server.

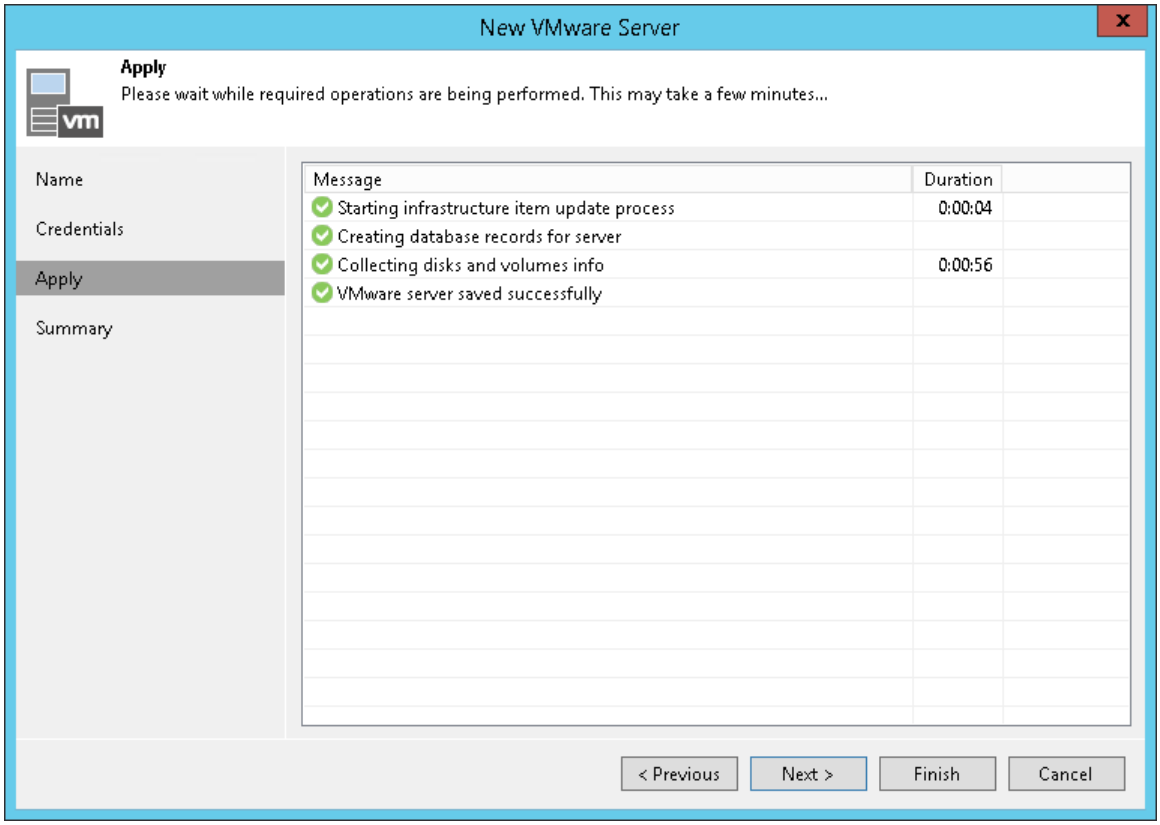
NOTE

If you update the certificate on the server, you must acknowledge the new certificate in the server connection settings. To do this, in the **Backup Infrastructure** view open the server settings, pass through the **Edit Server** wizard and click **Trust** to acknowledge the key.



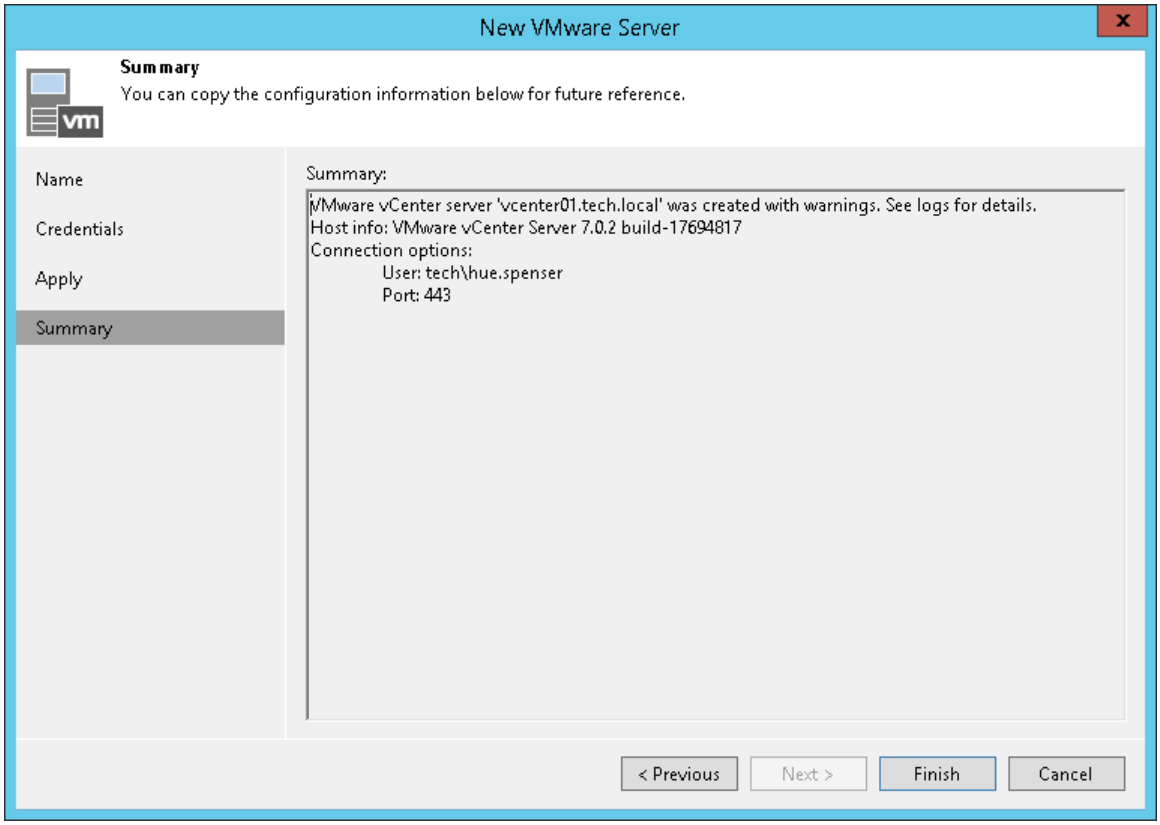
Step 4. Apply Settings

At the **Apply** step of the wizard, wait till Veeam Backup & Replication installs and configures all the required components. Click **Next** to complete the adding of the server.



Step 5. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the VMware vSphere server and click **Finish** to exit the wizard.



Adding VMware vCloud Director

To work with vApps and VMs managed by vCloud Director, you must add vCloud Director to the backup infrastructure.

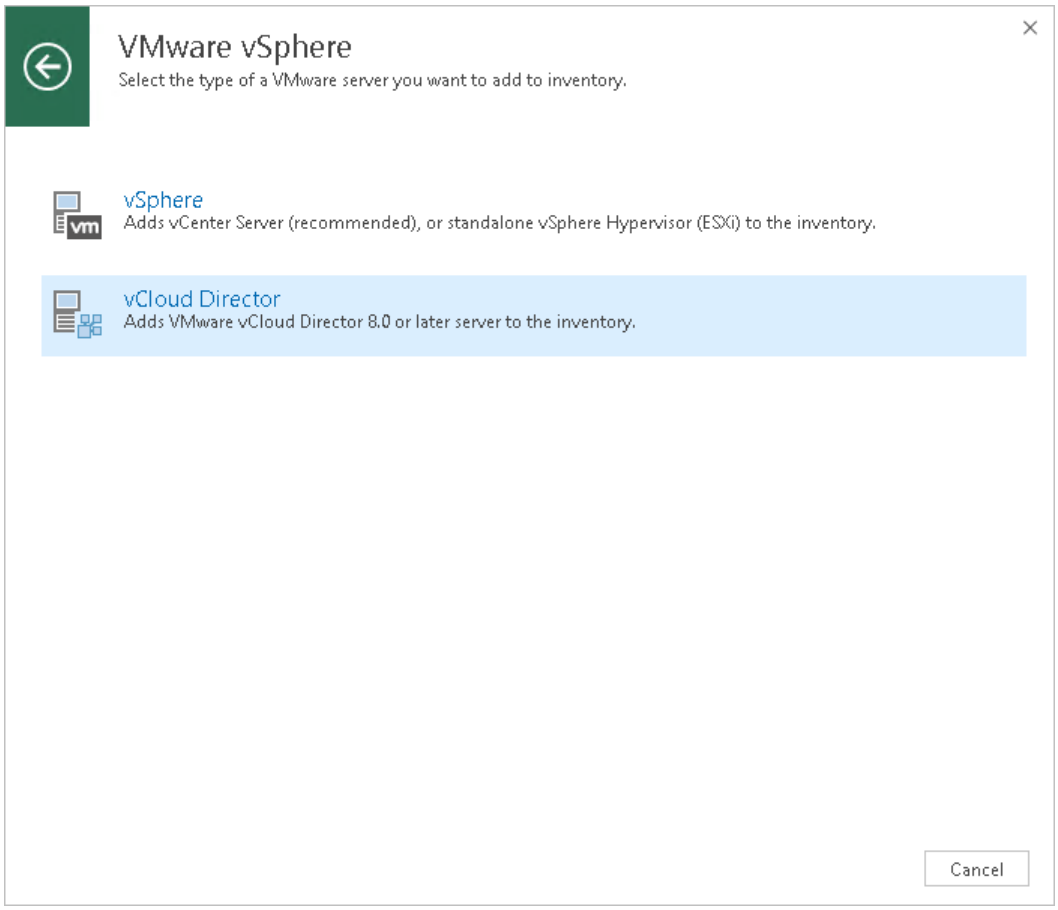
When you add vCloud Director to the backup infrastructure, the vCloud Director hierarchy is displayed under the **vCloud Director** node in the **Inventory** view. You can work with VMs managed by vCloud Director directly in the Veeam Backup & Replication console.

To add the vCloud Director server, use the **New VMware vCloud Director** wizard.

Step 1. Launch New VMware vCloud Director Server Wizard

To launch the **New VMware vCloud Director Server** wizard, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, right-click the **Managed Servers** node and select **Add Server**. Alternatively, you can click **Add Server** on the ribbon.
3. In the **Add Server** window, click **VMware vSphere** > **vCloud Director**.

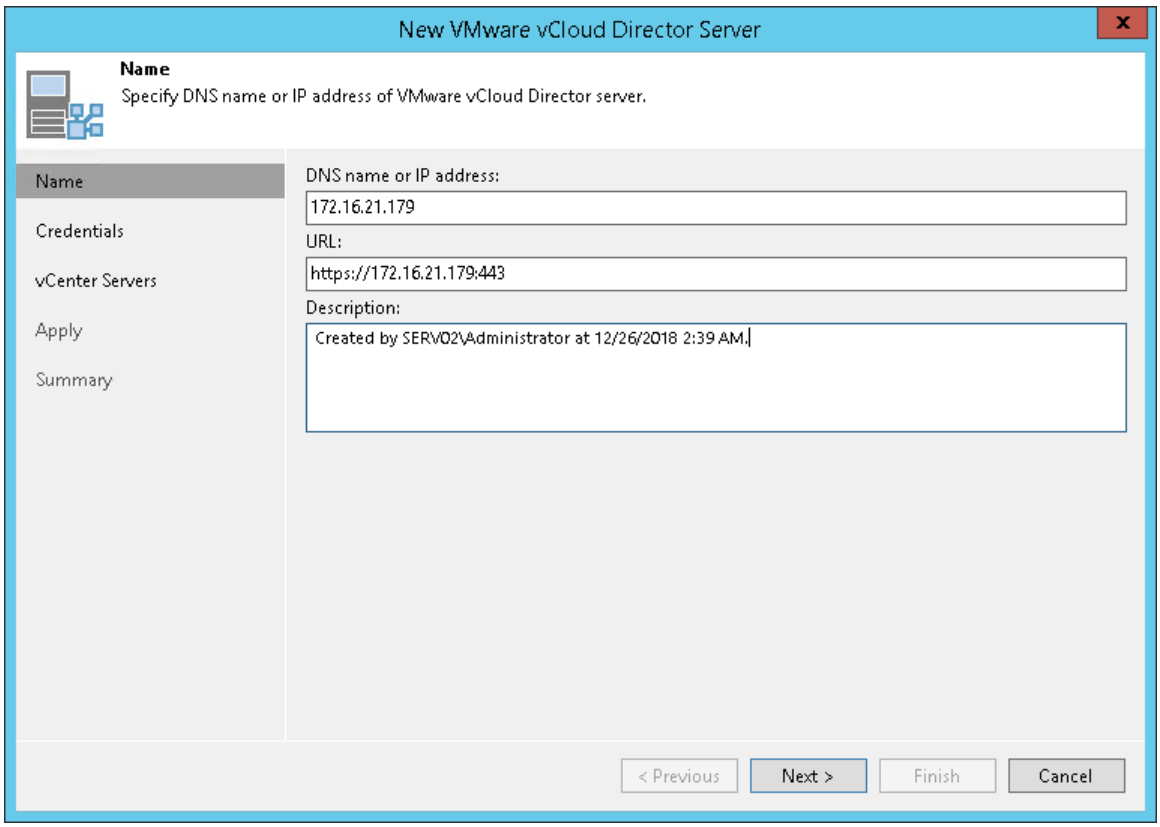


Step 2. Specify Server Name or Address

At the **Name** step of the wizard, specify connection settings for vCloud Director. If the vCloud Director infrastructure comprises several cells, you can specify connection settings for any cell in the vCloud Director hierarchy.

1. In the **DNS name or IP address** field, enter a full DNS name or IP address of the vCloud Director server or any cell in the vCloud Director infrastructure.
2. In the **URL** field, enter a URL of the vCloud Director server. By default, Veeam Backup & Replication uses the following URL: `https://<vcdservername>:443`,

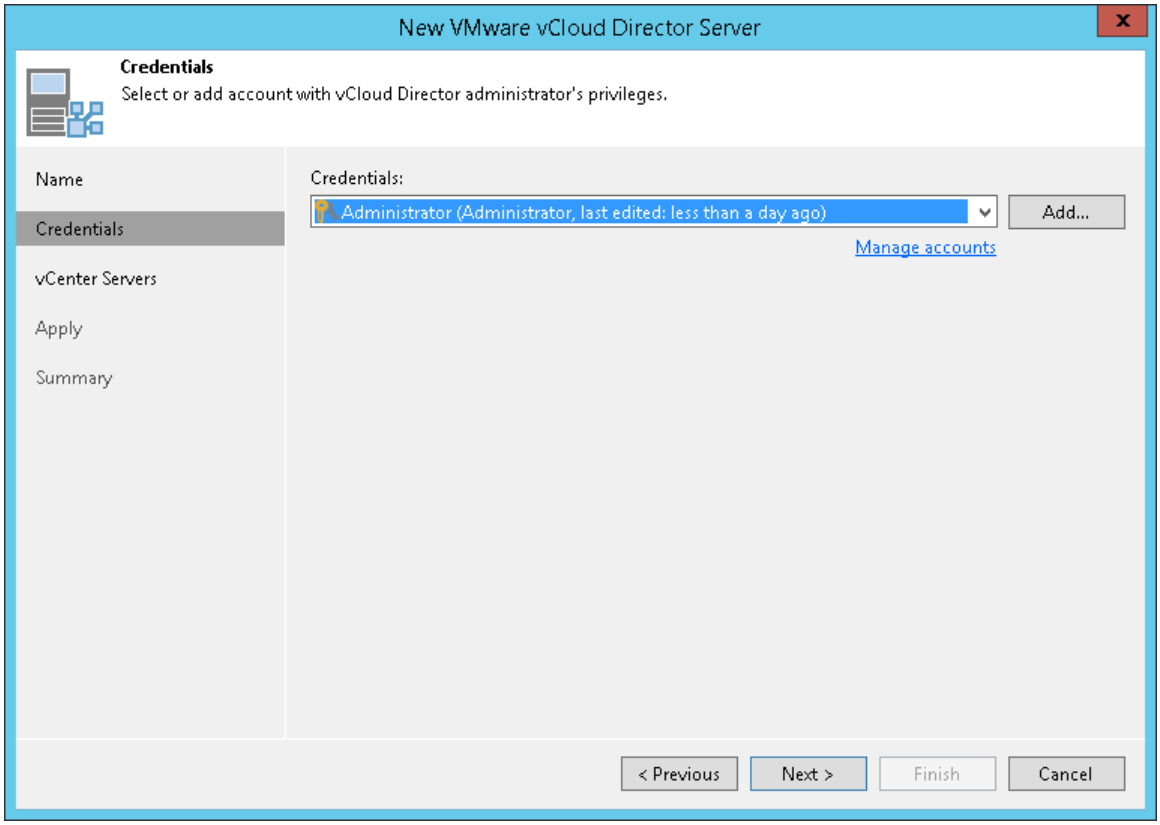
   where `<vcdservername>` is the name or IP address of the vCloud Director server that you have specified in the field above and `443` is the default port for communication with vCloud Director.
3. In the **Description** field, provide a description for future reference. The default description contains information about the user who added the server, date and time when the server was added.



Step 3. Specify VMware vCloud Director Credentials

At the **Credentials** step of the wizard, specify credentials to connect to vCloud Director.

From the **Credentials** list, select credentials for the account that has system administrator privileges on vCloud Director (you cannot use the organization administrator account to add vCloud Director). If you have not set up credentials beforehand, click the **Manage accounts** link at the bottom of the list or click **Add** on the right to add the credentials. For more information, see Managing Credentials.

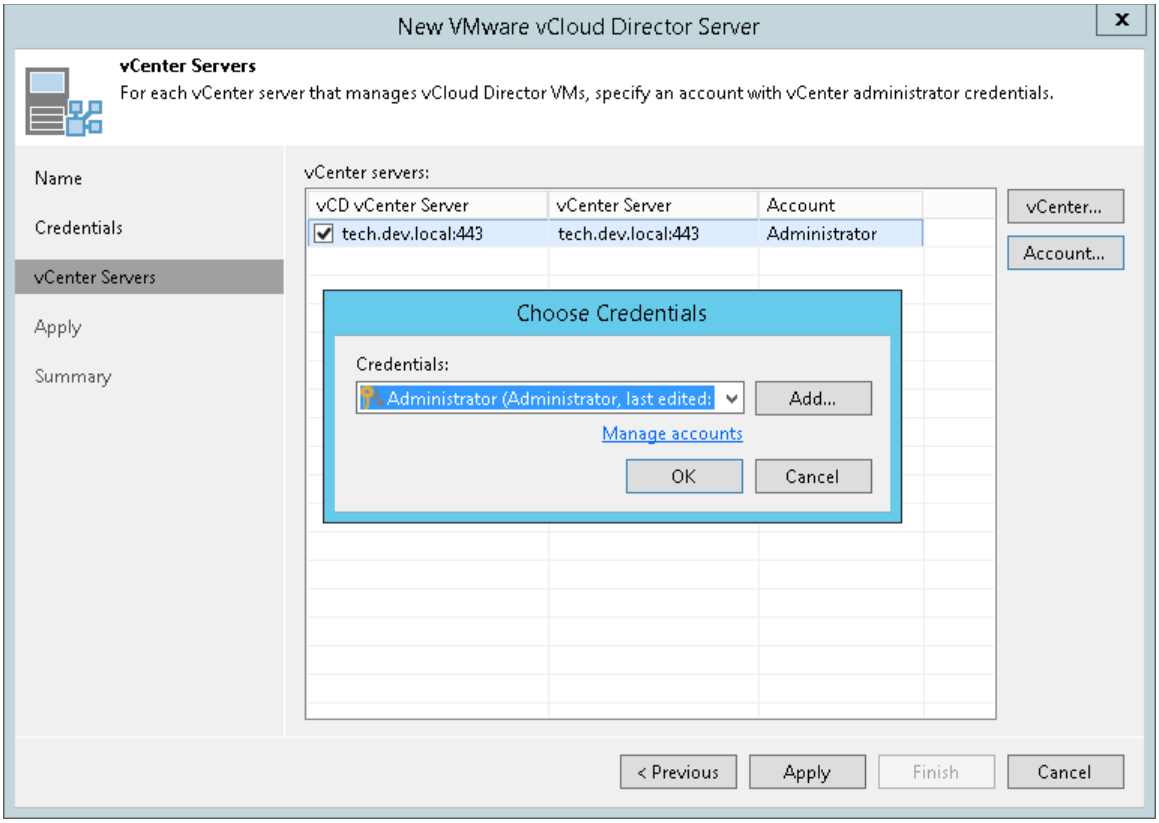


Step 4. Specify Credentials for Underlying vCenter Servers

At the **vCenter Servers** step of the wizard, specify credentials for every vCenter Server added to vCloud Director. If the vCenter Server is already added to the backup infrastructure, you do not need to specify credentials for it once again. Veeam Backup & Replication will automatically detect the credentials you provided when adding this vCenter Server and use them.

1. From the **vCenter servers** list, select a vCenter Server.
2. Click **Account** on the right and select credentials to connect to the vCenter Server. By default, Veeam Backup & Replication uses the same credentials that you have specified for vCloud Director at the previous step of the wizard.

   If you have not set up the credentials beforehand, click the **Manage accounts** link at the bottom of the list or click **Add** on the right to add the credentials. For more information, see Managing Credentials.
3. Veeam Backup & Replication automatically detects a port used to communicate with the vCenter Server. If necessary, you can change the connection port for the vCenter Server. Click **vCenter** on the right and adjust the port number.

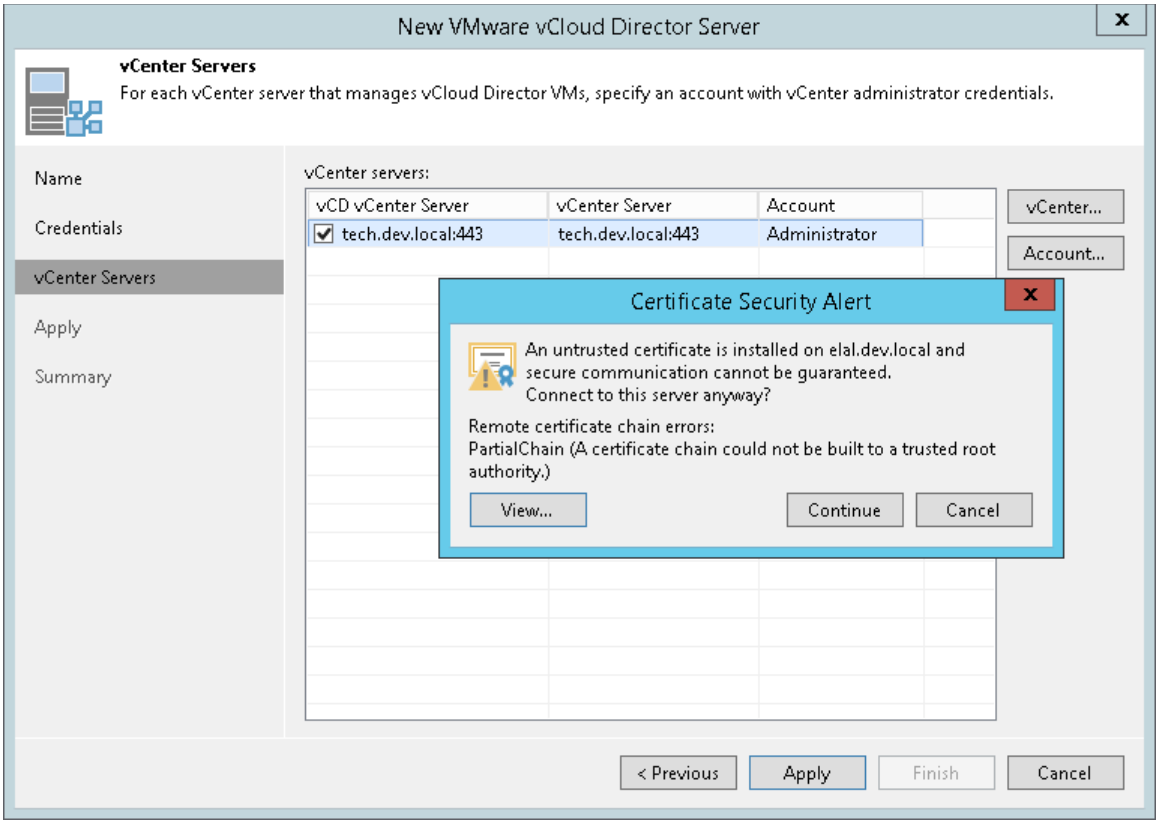


4. When you add a vCenter Server, Veeam Backup & Replication saves a thumbprint of the TLS certificate installed on the vCenter Server to the configuration database. During every subsequent connection to the server, Veeam Backup & Replication uses the saved thumbprint to verify the server identity and avoid the man-in-the-middle attack. For details on managing TLS Certificates, see TLS Certificates.

   If the certificate installed on the server is not trusted, Veeam Backup & Replication displays a warning.

   - To view detailed information about the certificate, click **View**.
   - If you trust the server, click **Continue**.

- If you do not trust the server, click **Cancel**. Veeam Backup & Replication will display an error message, and you will not be able to connect to the server.



5. Repeat steps 1-3 for all vCenter Servers added to vCloud Director.

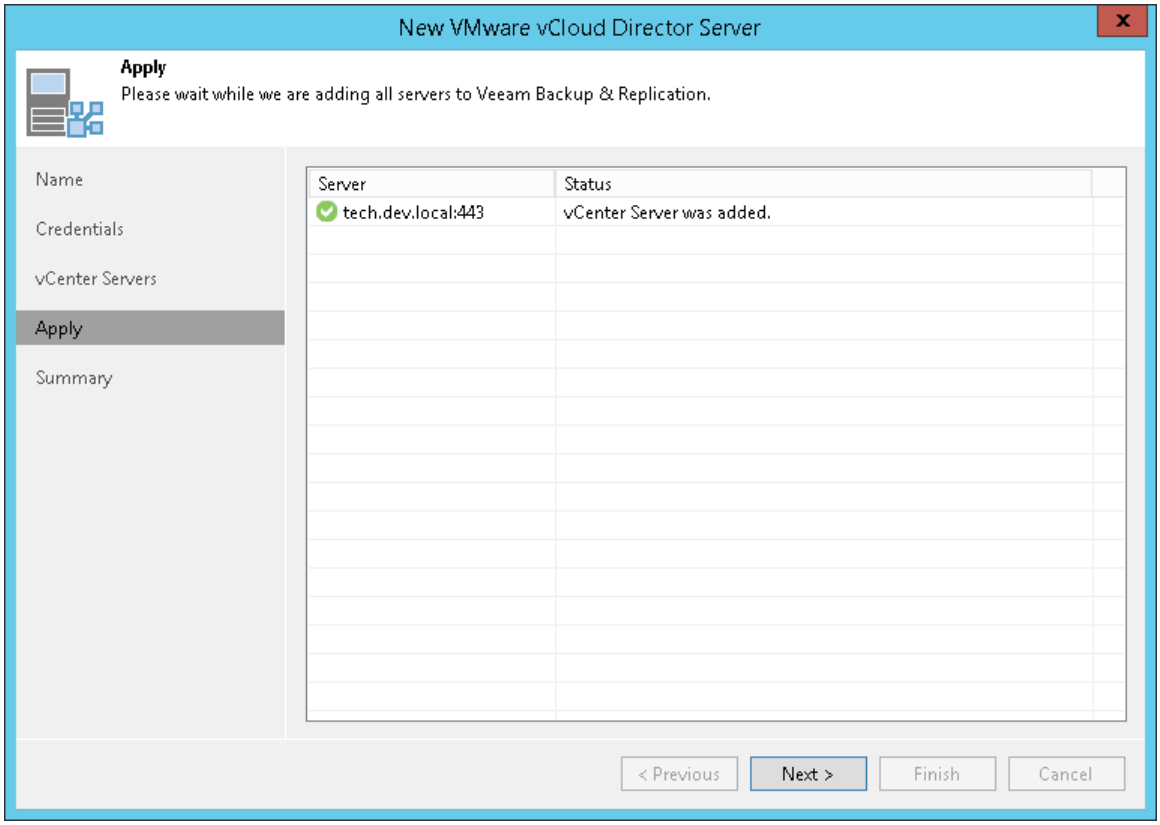
> **NOTE**
>
> After you update the certificate on the server, you must acknowledge the new certificate in the server connection settings. To do this, in the **Backup Infrastructure** view open the server settings, pass through the **Edit Server** wizard and click **Trust** to acknowledge the key.

Step 5. Finish Working with Wizard

At the **Apply** step of the wizard, complete the procedure of vCloud Director adding.

1. Review details of vCloud Director.
2. Click **Next**, and then click **Finish** to exit the wizard.

If vCenter Servers underlying vCloud Director are already added to the backup infrastructure, they will not be added for the second time. Veeam Backup & Replication will create associations with the vCenter Servers and display them in the vCloud Director hierarchy.



Adding Microsoft Windows Servers

You must add to the backup infrastructure Microsoft Windows servers that you plan to use as backup infrastructure components and servers that you plan to use for various types of restore operations.

Before you add a Microsoft Windows server, check prerequisites. Then use the **New Windows Server** wizard to add the server.

Before You Begin

Before you add a Microsoft Windows server to the backup infrastructure, check network connection settings of this server.

- File and printer sharing must be enabled in network connection settings of the added Microsoft Windows server. On every connected Microsoft Windows server, Veeam Backup & Replication deploys two components:
  - Veeam Installer Service

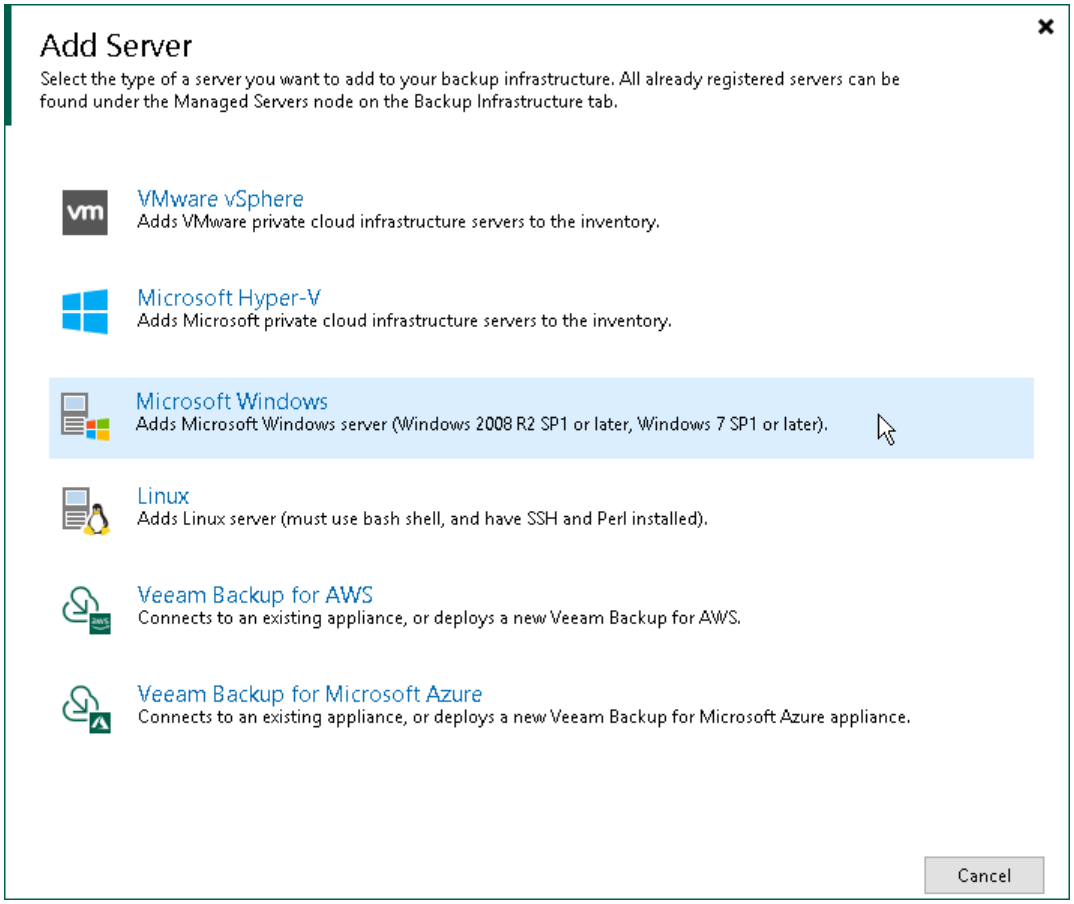
- Veeam Data Mover

If file and printer sharing is not enabled, Veeam Backup & Replication will fail to deploy these components.

Step 1. Launch New Windows Server Wizard

To launch the **New Windows Server** wizard, do one of the following:

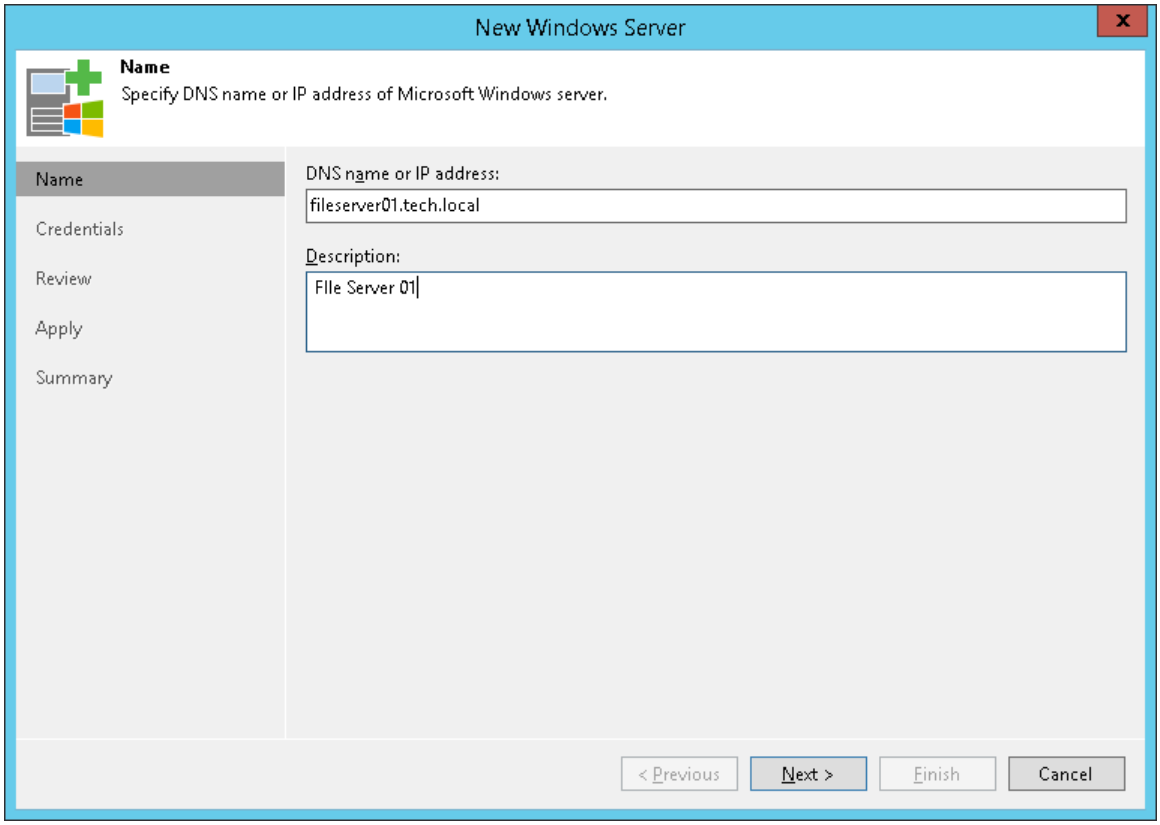
- Open the **Backup Infrastructure** or **Files** view, in the inventory pane select the **Microsoft Windows** node and click **Add Server** on the ribbon.
- Open the **Backup Infrastructure** view. In the inventory pane, select the **Managed Servers** node and click **Add Server** on the ribbon or right-click the **Managed Servers** node and select **Add Server.** In the **Add Server** window, select **Microsoft Windows**.



Step 2. Specify Server Name or Address

At the **Name** step of the wizard, specify an address and description for the Microsoft Windows server.

1. Enter a full DNS name or IP address of the Microsoft Windows server.
2. Provide a description for future reference. The default description contains information about the user who added the server, date and time when the server was added.



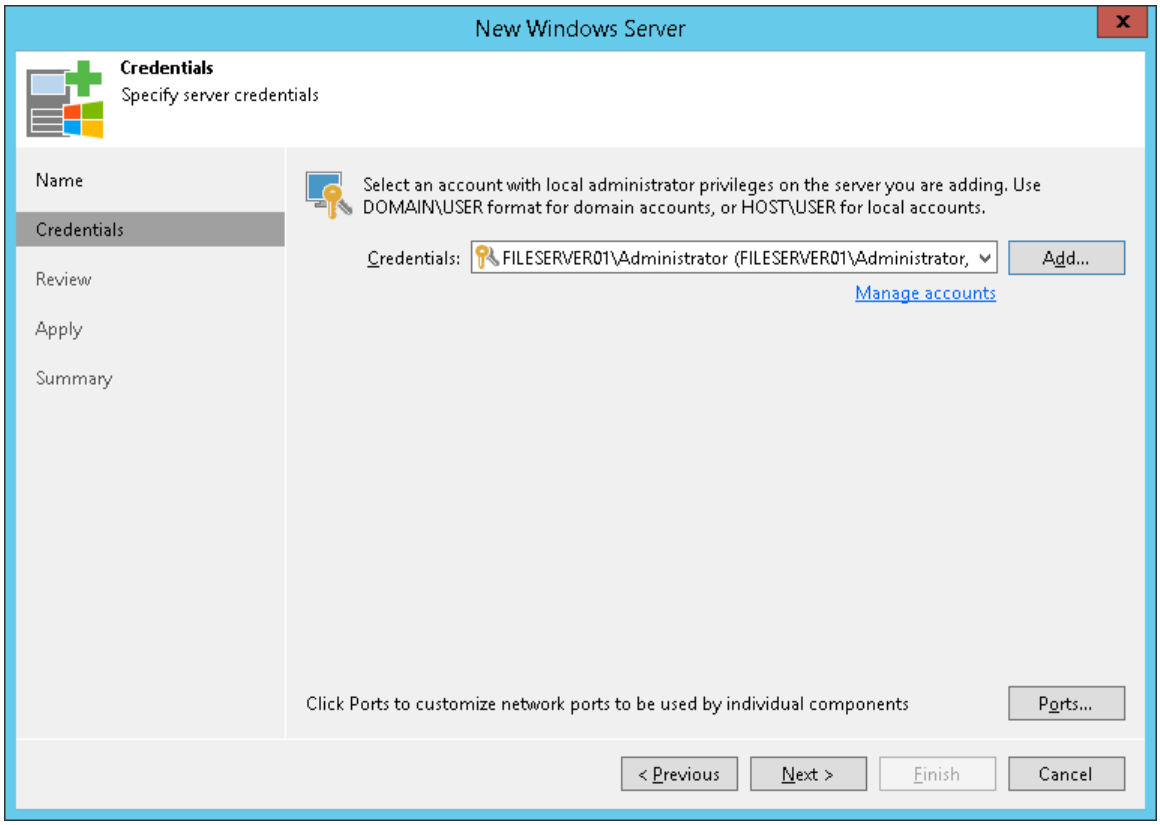
Step 3. Specify Credentials

At the **Credentials** step of the wizard, specify credentials for the Microsoft Windows server.

1. From the **Credentials** list, select credentials for the account that has privileges described in Required Permissions.

   If you have not set up credentials beforehand, click the **Manage accounts** link or click **Add** on the right to add the credentials. For more information, see Managing Credentials.

   Veeam Backup & Replication will use the provided credentials to deploy its components on the added server.



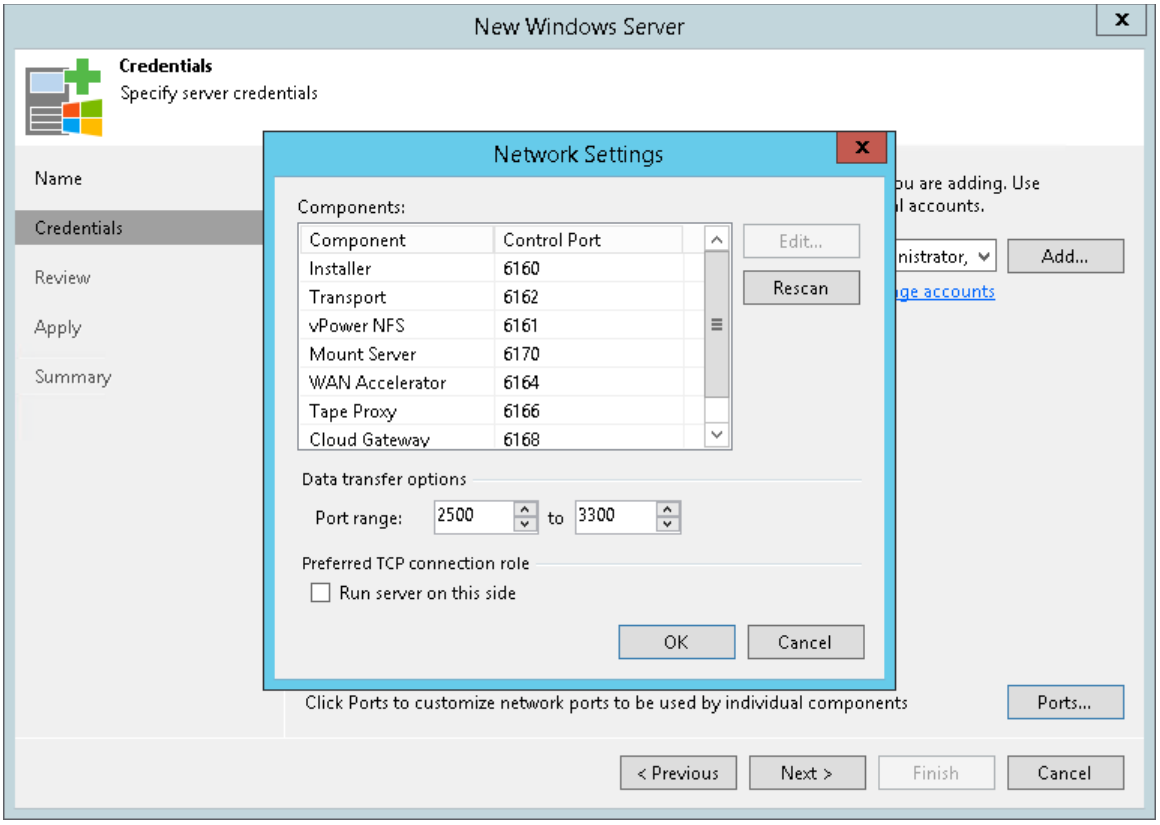
2. To customize network ports used by Veeam Backup & Replication components, click **Ports**. For default ports used by the Veeam Backup & Replication components, see Used Ports.

   If necessary, adjust the port numbers.

3. In the **Data transfer options** section of the **Network Settings** window, specify connection settings for file copy operations. Provide a range of ports that will be used as transmission channels between the source server and target server (one port per task). By default, Veeam Backup & Replication uses port range 2500-3300.

   If the virtual environment is not large and data traffic will not be significant, you can specify a smaller range of ports, for example, 2500-2510 to run 10 concurrent tasks at the same time. Keep in mind that Veeam Backup & Replication processes each VM disk as a separate task.

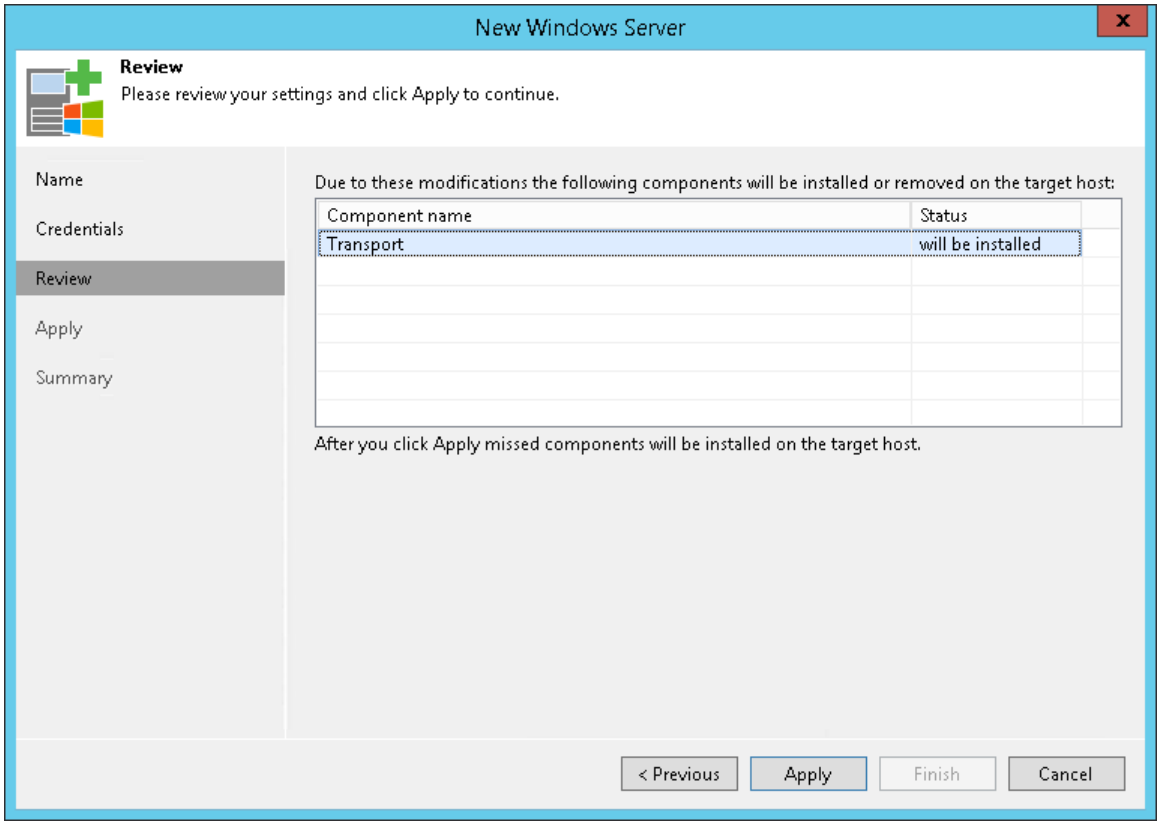
4. If the Microsoft Windows server is deployed outside NAT, in the **Preferred TCP connection role** section select the **Run server on this side** check box. In the NAT scenario, the outside client cannot initiate a connection to the server on the NAT network. As a result, services that require initiation of the connection from outside can be disrupted. With this option selected, you will be able to overcome this limitation and initiate a 'server-client' connection — that is, a connection in the direction of the Microsoft Windows server.



Step 4. Review Components

At the **Review** step of the wizard, review what Veeam Backup & Replication components are already installed on the server and what components will be installed.

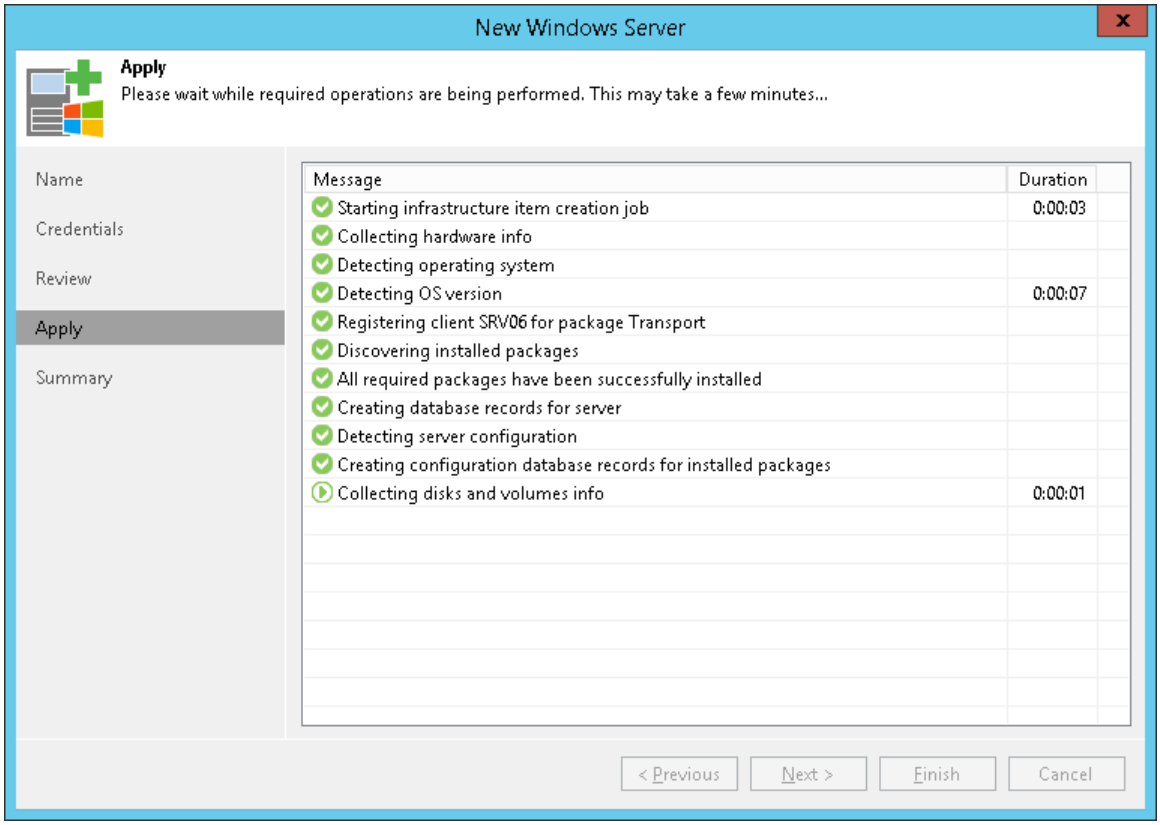
1. Review the components.
2. Click **Apply** to add the Microsoft Windows server to the backup infrastructure.



Step 5. Finish Working with Wizard

At the **Apply** step of the wizard, complete the procedure of Microsoft Windows server adding.

1. Review details of the Microsoft Windows server.
2. Click **Next**, then click **Finish** to exit the wizard.



Adding Linux Servers

You must add to the backup infrastructure Linux servers that you plan to use as backup repositories and servers that you plan to use for various types of restore operations.

Before you add a Linux server, check prerequisites. Then use the **New Linux Server** wizard to add the server.

Before You Begin

Before you add a Linux server to the Veeam Backup & Replication infrastructure, check the following prerequisites.

If you add a server using single-use credentials, it will not be available through SSH connection.

Linux Firewalls

When you add a Linux server to the backup infrastructure, Veeam Backup & Replication automatically opens ports used by the Veeam Data Mover on the Linux server. Generally, Veeam Backup & Replication automatically open ports for most of popular firewalls (iptables, ufw, firewall-cmd). However, if for some reason the ports are not opened, you can open the ports manually. You can also specify these ports at the SSH Connection step of the **New Linux Server** wizard. Note that ports are opened dynamically: if 10 concurrent jobs are running, Veeam Backup & Replication opens ports 2500-2510.

If you use the `firewalld` tool, you can configure firewall rules to open ports only in necessary zones. By default, Veeam Backup & Replication opens ports in all active `firewalld` zones. If your firewall is configured for different zones, and you want to minimize security holes, you can configure Veeam Backup & Replication to open the ports only for certain zones. To do this, perform the following:

1. On the helper host or target Linux host, create the `/etc/VeeamNetConfig` file and define the following parameter:

   ```
   FirewalldZones=zone_name_1, zone_name_2
   ```

   where `zone_name_1, zone_name_2` is a list of zone names where the ports must be open. Veeam Backup & Replication will skip the zones that are not in this list.

2. [Only for helper host] If you select a Linux host that is already added to the Veeam Backup & Replication infrastructure, you should also add required zones to the `/opt/veeam/transport/VeeamTransportConfig` file.

   ```
   FirewalldZones=zone_name_1, zone_name_2
   ```

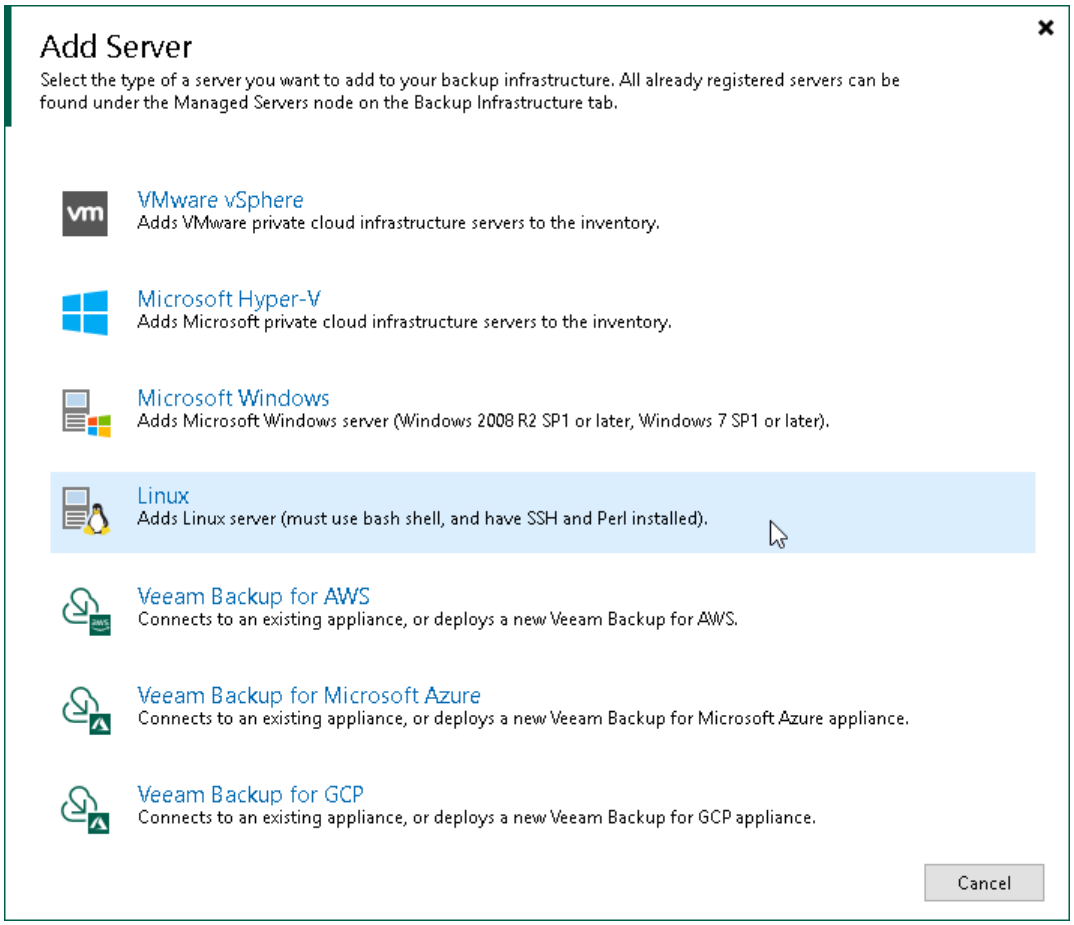
   NOTE

   Veeam Backup & Replication opens the port 2500 in all zones even if you have specified the required zones in configuration files.

Step 1. Launch New Linux Server Wizard

To launch the **New Linux Server** wizard, do the following:

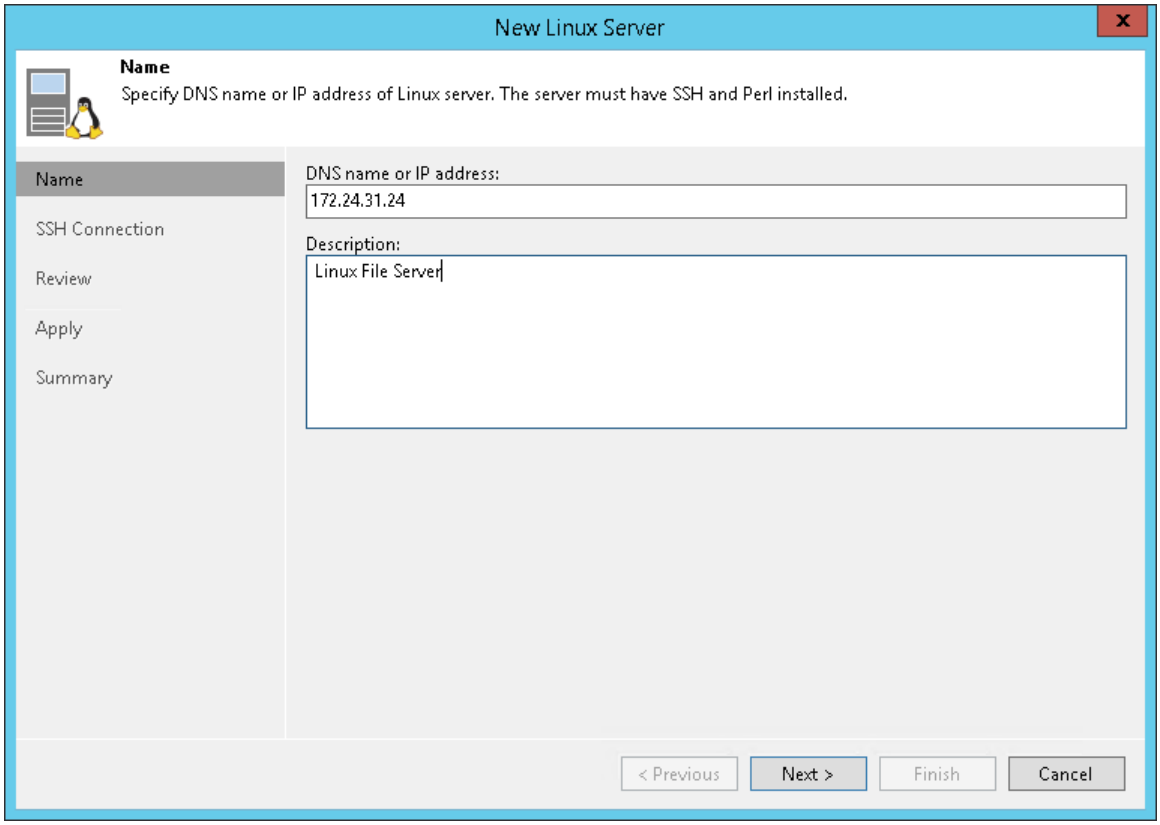
1. Open the **Backup Infrastructure** view.
2. In the inventory pane, right-click the **Managed Servers** node and select **Add Server**. Alternatively, you can click **Add Server** on the ribbon.
3. In the **Add Server** window, select **Linux**.



Step 2. Specify Server Name or Address

At the **Name** step of the wizard, specify an address and description for the Linux server.

1. Enter a full DNS name or IP address of the Linux server.
2. Provide a description for future reference. The default description contains information about the user who added the server, date and time when the server was added.



Step 3. Specify Credentials and SSH Settings

At the **SSH Connection** step of the wizard, specify credentials for the Linux server and additional SSH connection settings.

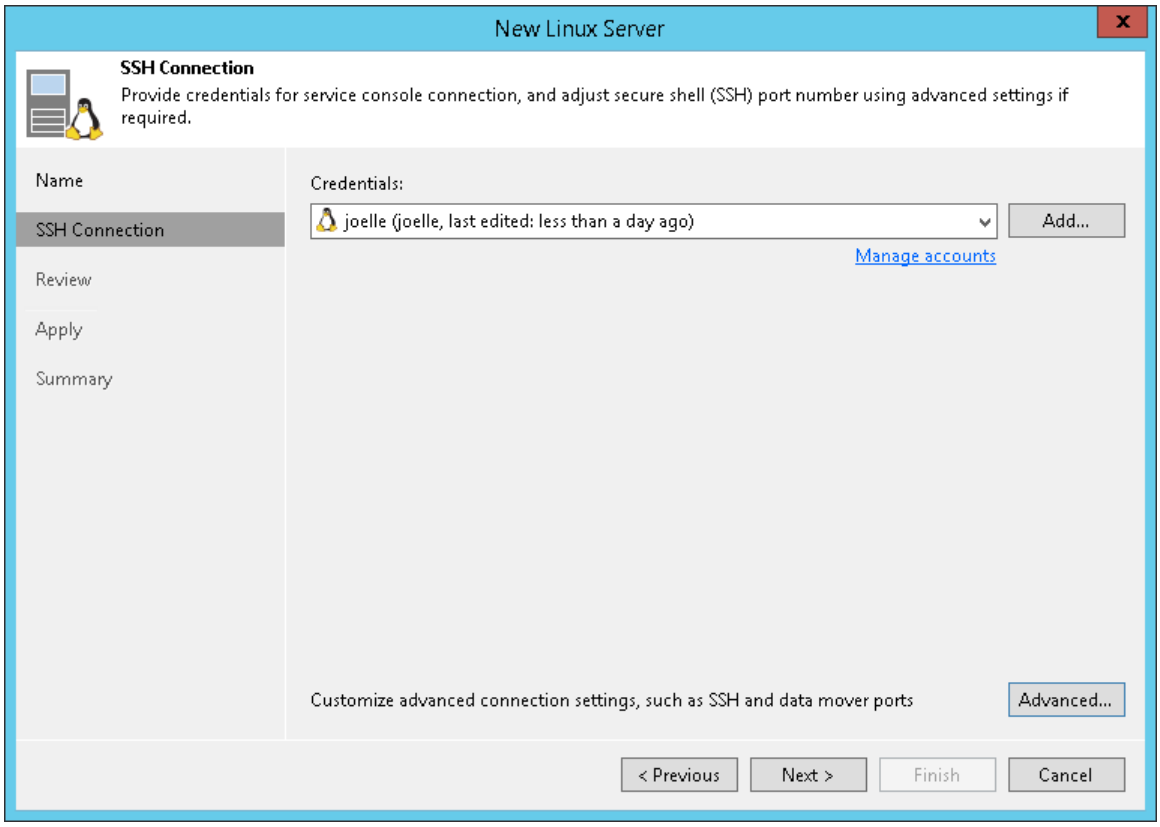
1. From the **Credentials** list, select credentials for the account that has permissions described in Required Permissions. You can select a credentials record that uses the password authentication method or credentials record that uses the Identity/Pubkey authentication method.

   If you have not set up credentials beforehand, click the **Manage accounts** link or click **Add** on the right to add the credentials. For more information, see Managing Credentials.

   To build a hardened repository, click **Add** and select **Single-use credentials for hardened repository**. Veeam Backup & Replication does not store these credentials, they are used only for deploying Veeam Data Mover to the server. These credentials reduce the rights for the Veeam Data Mover. Single-use, or temporary, credentials is a recommended option for a hardened repository, but you can also use persistent credentials. In this case, the rights for the transport service will be reduced at the Configure Backup Repository Settings step of the **Adding Backup Repository** wizard. User with root permissions cannot use single-use credentials.

   In the **Credentials** window, specify username, password, SSH port and select the **Use "su" if "sudo" fails** check box. For more information, see Deploying Hardened Repository.

> **NOTE**
>
> If you add a Linux server with single-use credentials, the folder with the repository must be accessible for accounts with user permissions (and not only root).



2. To configure SSH settings, click **Advanced**. This option becomes available after you have entered your credentials. In the **SSH Settings** window:

   a. In the **Service console connection** section, specify an SSH timeout. By default, the SSH timeout is set to 20000 ms. If a task targeted at the Linux server is inactive after the specified timeout, Veeam Backup & Replication will automatically terminate the task.

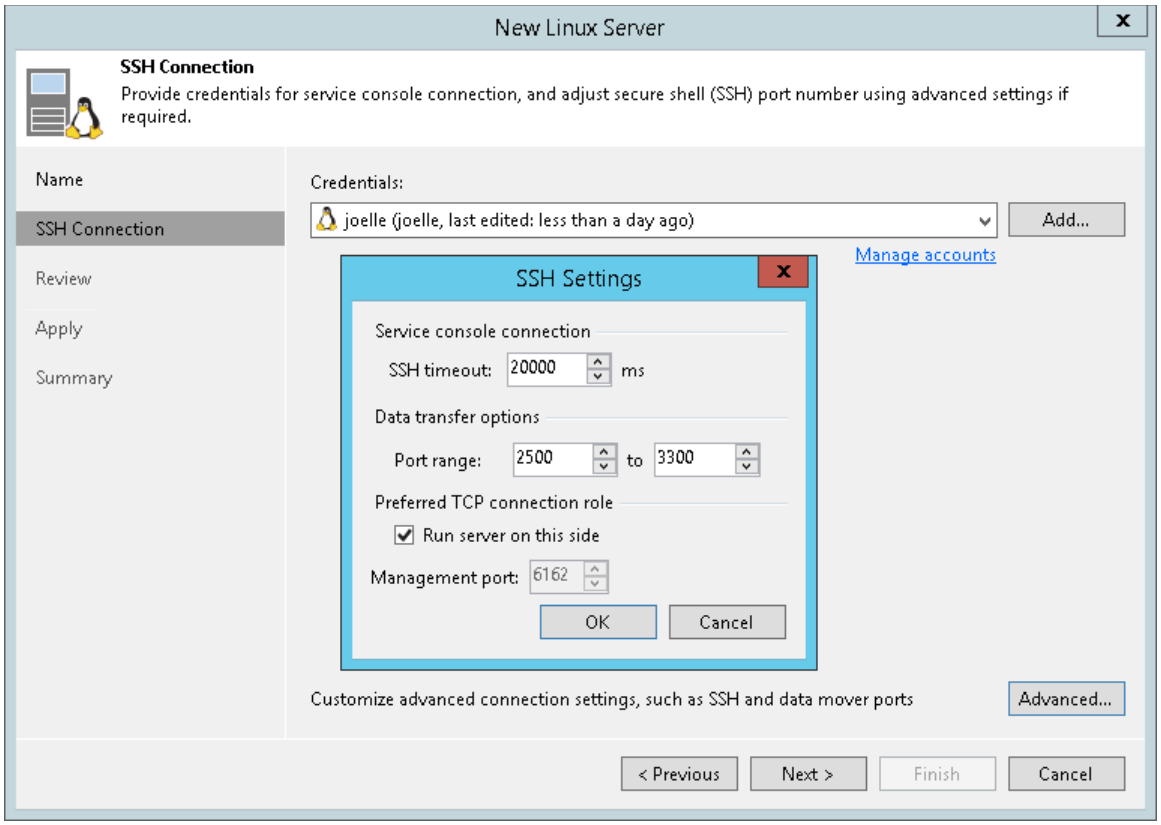
   b. In the **Data transfer options** section, specify connection settings for file copy operations. Provide a range of ports that will be used as transmission channels between the source host and target host (one port per task). By default, Veeam Backup & Replication uses port range 2500-3300. If the virtual environment is not large and data traffic will not be significant, you can specify a smaller range of ports, for example, 2500-2510 to run 10 concurrent tasks at the same time.

   Port 6162 is opened by default. It is a port used by Veeam Data Mover.

   > **NOTE**
   >
   > If you want to open these ports only for certain `firewalld` zones, you can specify the required zones in the configuration files. For instructions, see the Before You Begin section.

   c. [For the Linux server deployed outside NAT] In the **Preferred TCP connection role** section, select the **Run server on this side** check box. In the NAT scenario, the outside client cannot initiate a connection to the server on the NAT network. As a result, services that require initiation of the connection from outside can be disrupted. With this option selected, you will be able to overcome this limitation and initiate a 'server-client' connection — that is, a connection in the direction of the Linux server.

You can also change the SSH port over which you want to connect to the Linux server. For this, click the **Manage accounts** link and edit the account used to connect to the Linux server.

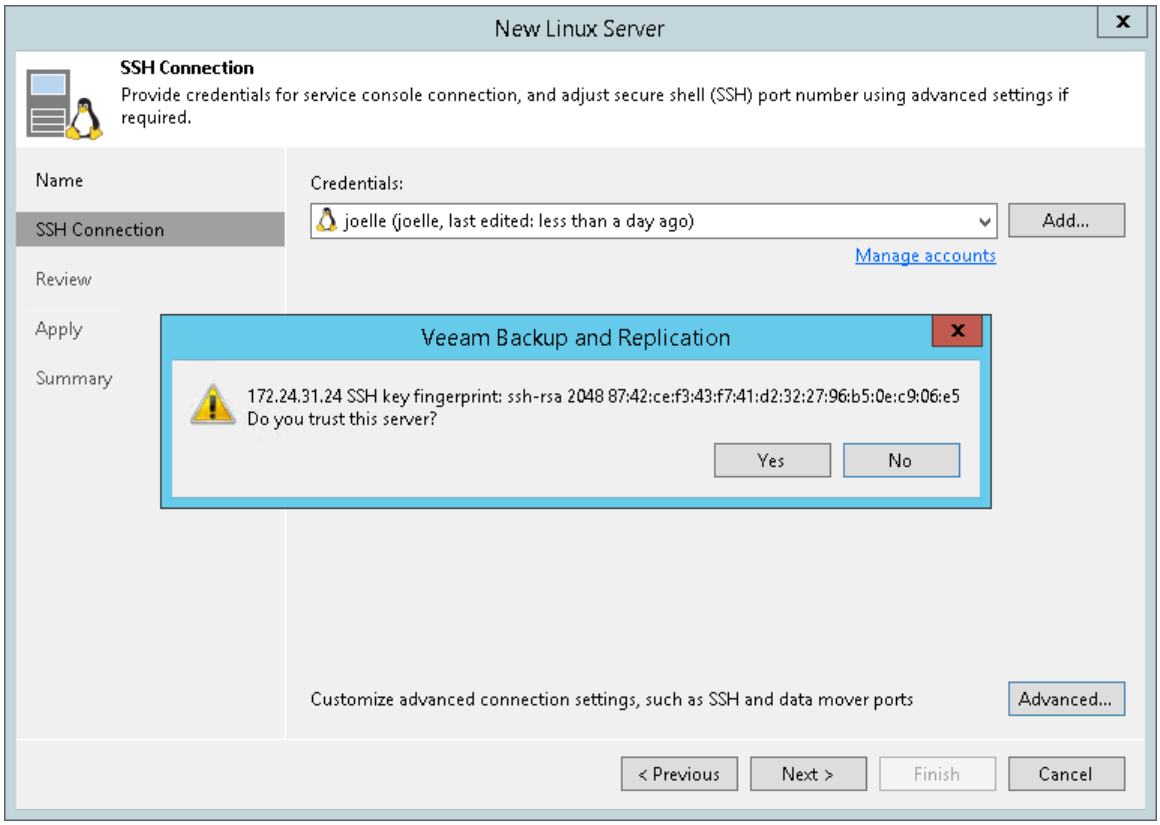


3. When you add a Linux server, Veeam Backup & Replication saves a fingerprint of the Linux host SSH key to the configuration database. During every subsequent connection to the server, Veeam Backup & Replication uses the saved fingerprint to verify the server identity and avoid the man-in-the-middle attack.

   To let you identify the server, Veeam Backup & Replication displays the SSH key fingerprint:

   - If you trust the server and want to connect to it, click **Yes**.
   - If you do not trust the server, click **No**. Veeam Backup & Replication will display an error message, and you will not be able to connect to the server.

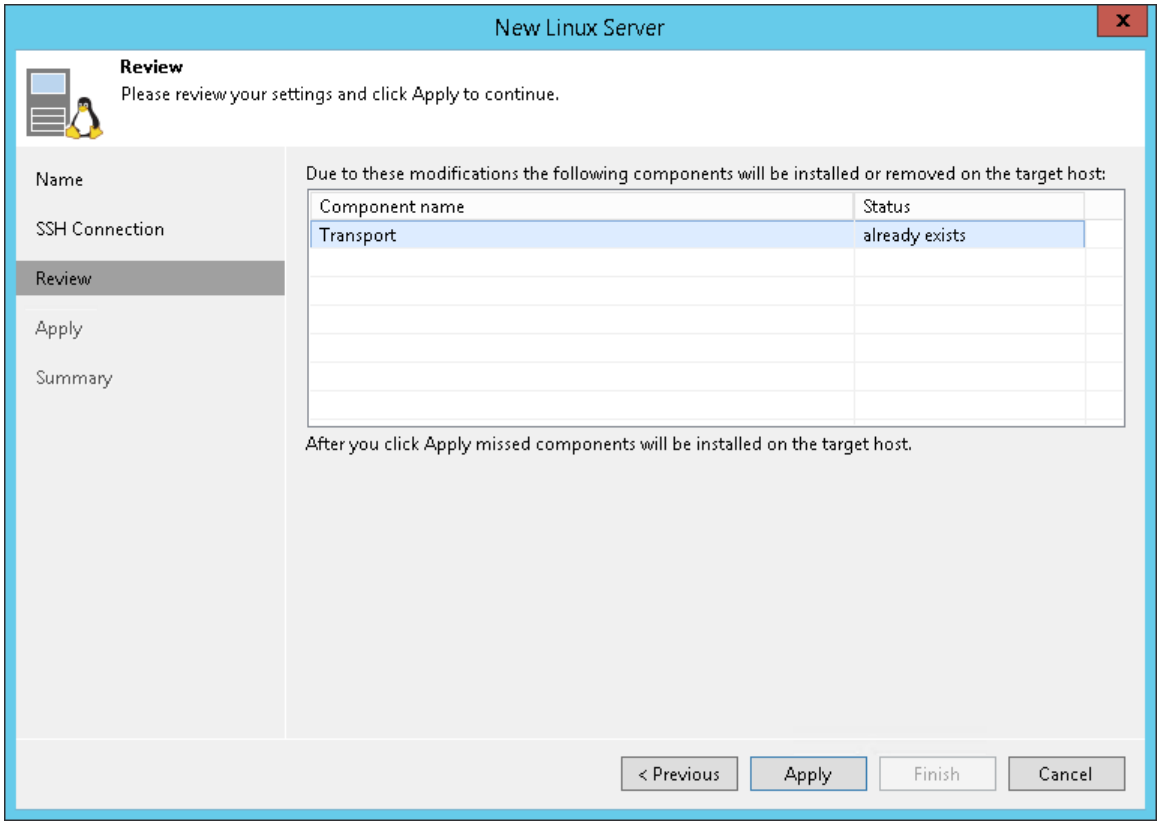
> **NOTE**
>
> If you update the SSH key on the server, you must acknowledge the new key in the server connection settings. To do this, in the **Backup Infrastructure** view open the server settings, pass through the **Edit Server** wizard and click **Trust** to acknowledge the new key.



Step 4. Review Components

At the **Review** step of the wizard, review what Veeam Backup & Replication components are already installed on the server and what components will be installed.

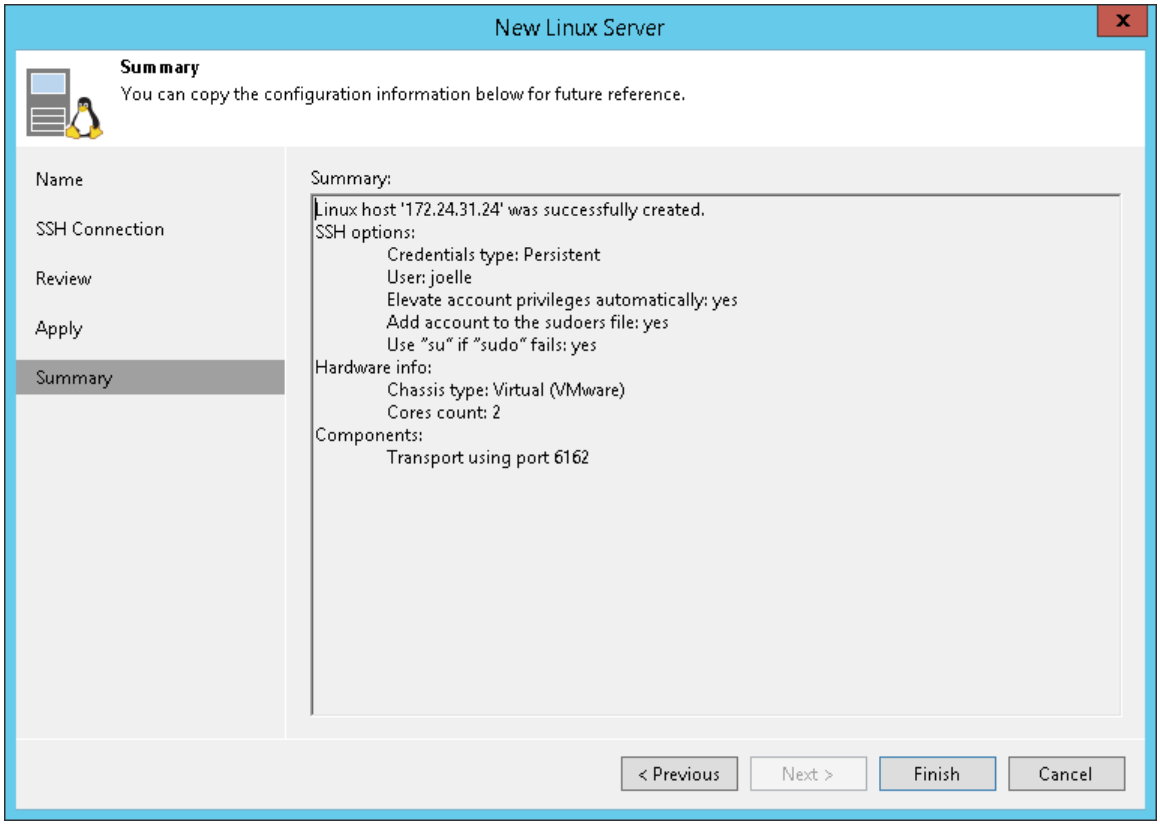
1. Review the components.
2. Click **Apply** to add the Linux server to the backup infrastructure.



Step 5. Finish Working with Wizard

At the **Summary** step of the wizard, complete the procedure of Linux server adding.

1. Review details of the Linux server.
2. Click **Next**, then click **Finish** to exit the wizard.



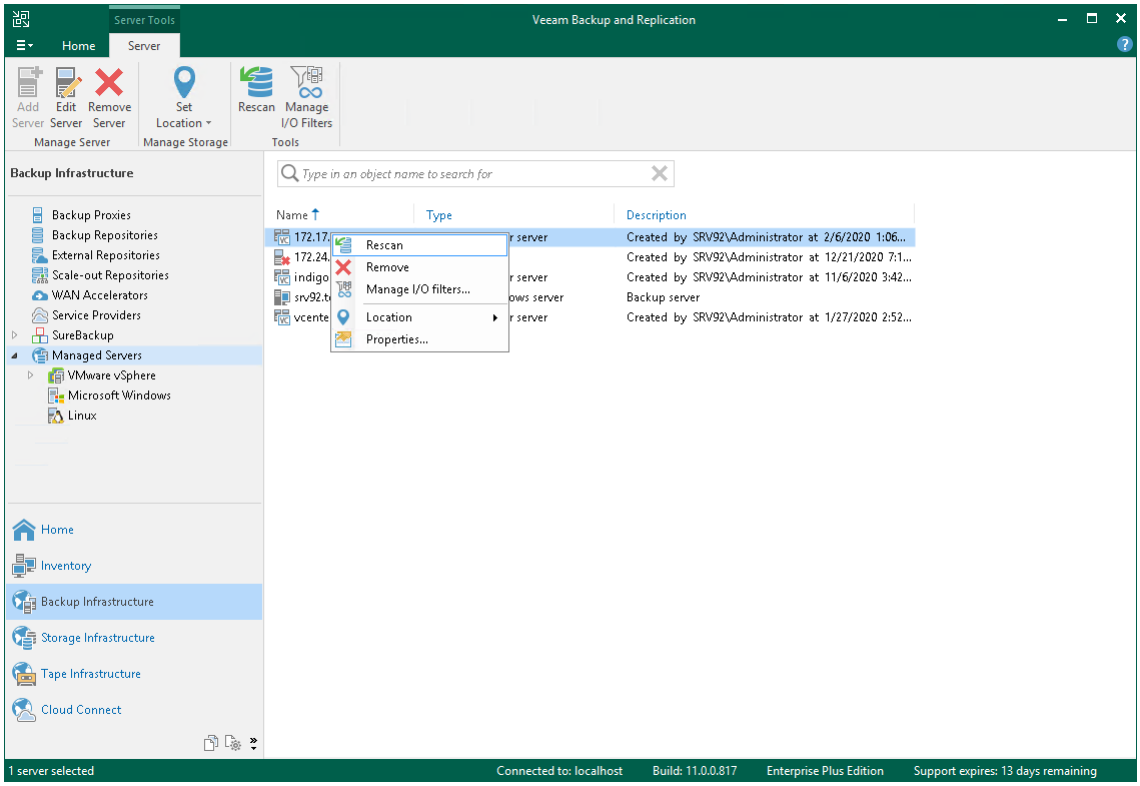
Rescanning Servers

In some cases, you may need to rescan hosts or servers in the backup infrastructure. The rescan operation may be required if you have added or removed new disks and volumes to/from the host or server and want to display actual information in Veeam Backup & Replication. During the rescan operation, Veeam Backup & Replication retrieves information about disks and volumes that are currently connected to a host or server and stores this information to the configuration database.

Veeam Backup & Replication automatically performs a rescan operation every 4 hours. You can also start the rescan operation manually:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Managed servers**.

3. In the working area, select the server or host and click **Rescan** on the ribbon. Alternatively, you can right-click the server or host and select **Rescan**.

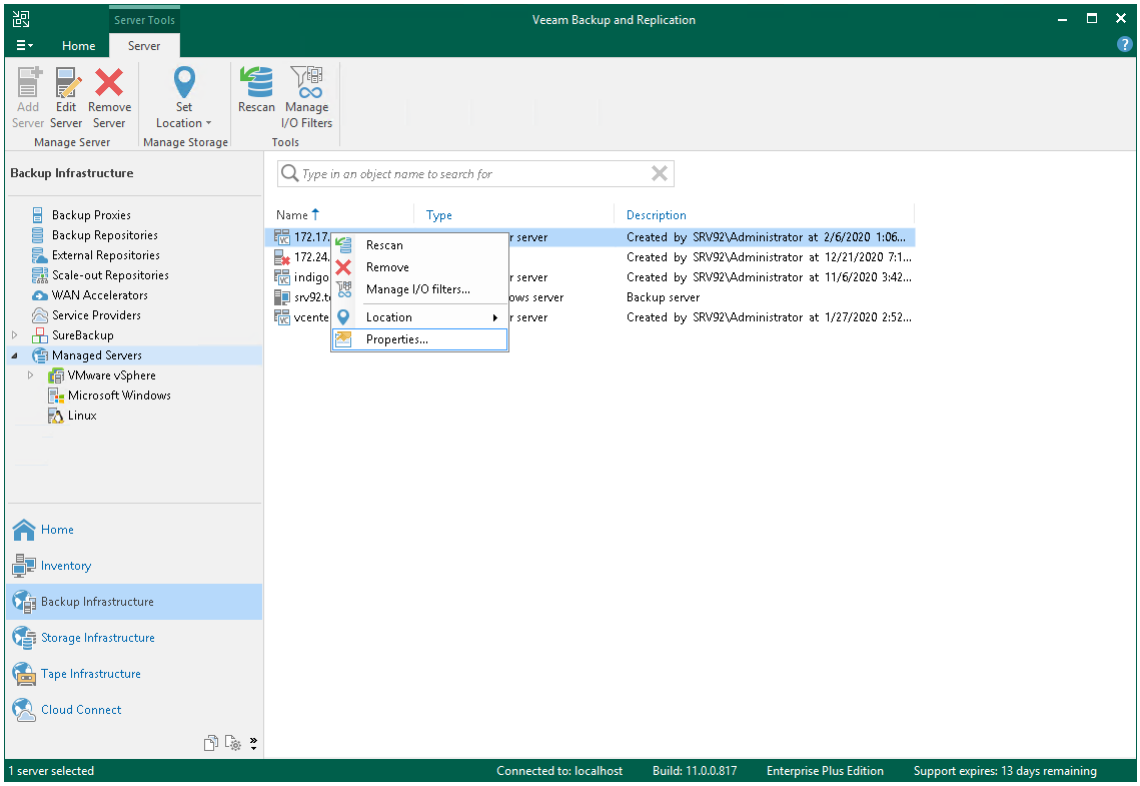


Editing Server Settings

To edit settings of a server in the backup infrastructure:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Managed servers**.
3. In the working area, select the server and click **Edit Server** on the ribbon or right-click the server and select **Properties**.

4. You will follow the same steps as you have followed when adding the server. Edit server settings as required.



Removing Servers

If you do not plan to use some server anymore, you can remove it from the backup infrastructure.

You cannot remove a server that has any dependencies. For example, you cannot remove a server that is referenced by a backup or replication job, performs the role of a backup proxy or backup repository. To remove such server, you will need to delete all referencing jobs and roles first.

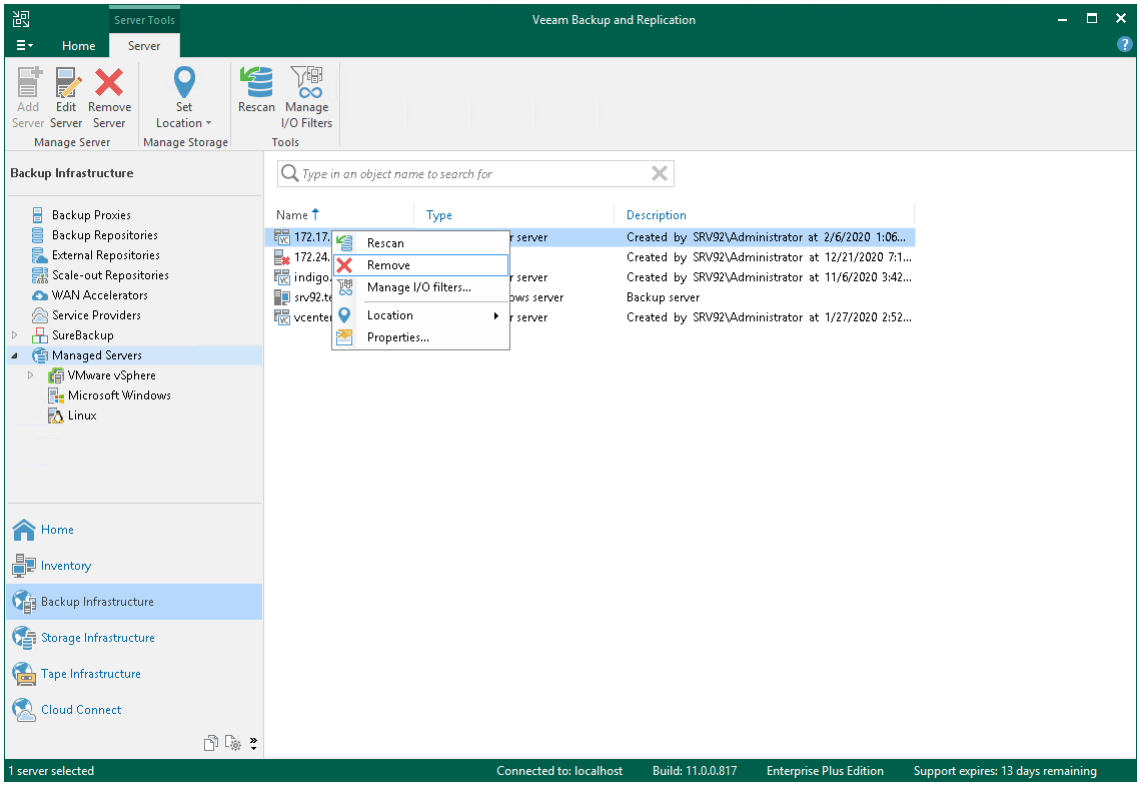
When you remove a server that is used as a target host or backup repository, backup files and replica files are not removed from disk. You can easily import these files later to Veeam Backup & Replication if needed.

> **NOTE**
>
> When you remove vCloud Director from the backup infrastructure, vCenter Servers added to vCloud Director are not removed. To remove the vCenter Server, in the inventory pane expand the **vCenter Servers** node, right-click the vCenter Server and select **Remove**.
>
> You cannot remove vCenter Servers added to vCloud Director until the vCloud Director server is removed from the backup infrastructure.

To remove a server from the backup infrastructure:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Managed servers**.

3. In the working area, select the server and click **Remove Server** on the ribbon or right-click the server and select **Remove**.



Backup Proxy

A backup proxy is an architecture component that sits between the backup server and other components of the backup infrastructure. While the backup server administers tasks, the proxy processes jobs and delivers backup traffic.

Basic backup proxy tasks include the following:

- Retrieving VM data from the production storage
- Compressing
- Deduplicating
- Encrypting
- Sending it to the backup repository (for example, if you run a backup job) or another backup proxy (for example, if you run a replication job)

A backup proxy can operate as a data mover in NAS backup and transfer data from the source file share to the backup repository. For more information about using a backup proxy in NAS backup, see NAS Backup Support.

Backup Proxy Transport Modes

Depending on your backup architecture, a backup proxy can use one of the following data transport modes:

- Direct storage access
- Virtual appliance
- Network

If the VM disks are located on the storage system and the storage system is added to the Veeam Backup & Replication console, the backup proxy can also use the Backup from Storage Snapshots mode.

You can explicitly select the transport mode or let Veeam Backup & Replication automatically choose the mode. For details, see the Transport Modes and Configuring Backup Proxy for Storage Snapshots sections in the Veeam Backup & Replication User Guide.

Backup Proxy Deployment

By default, the role of the proxy is assigned to the backup server itself. However, this is sufficient only for small installations with low traffic load. For large installations, it is recommended to deploy dedicated backup proxies.

To optimize performance of several concurrent jobs, you can use several backup proxies. In this case, Veeam Backup & Replication will distribute the backup workload between available backup proxies. You can deploy backup proxies both in the primary site and in remote sites.

To deploy a proxy, you need to add a Windows-based or Linux-based server to Veeam Backup & Replication and assign the role of the backup proxy to the added server. For requirements and limitations that backup proxies have, see Requirements and Limitations for Backup Proxies.

Backup Proxy Services and Components

Backup proxies run light-weight services that take a few seconds to deploy. Deployment is fully automated. Veeam Backup & Replication installs the following components and services:

- **Veeam Installer Service** is an auxiliary service that is installed and started on any Windows server once it is added to the list of managed servers in the Veeam Backup & Replication console. This service analyzes the system, installs and upgrades necessary components and services depending on the role selected for the server.
- **Veeam Data Mover** is a component that performs data processing tasks on behalf of Veeam Backup & Replication, such as retrieving source VM data, performing data deduplication and compression, and storing backed-up data on the target storage.

In This Section

- Requirements and Limitations for Backup Proxies
- Transport Modes
- Adding VMware Backup Proxies
- Editing Backup Proxy Settings
- Disabling and Removing Backup Proxies

Requirements and Limitations for Backup Proxies

Before you assign the role of a backup proxy, check the following prerequisites and limitations.

Connection to Storage

The following list shows possible connections between the machine and storage that keeps backups of this machine. The first connection is the most efficient, the last one is the least efficient.

- A machine used as the backup proxy should have direct access to the storage on which VMs reside or the storage where VM data is written. This way, the backup proxy will retrieve data directly from the datastore bypassing LAN.
- The backup proxy can be a VM with HotAdd access to VM disks on the datastore. This type of proxy also enables LAN-free data transfer between the host and the backup proxy.
- If neither of the above scenarios is possible, you can assign the role of the backup proxy to a machine whose network is located closer to the source or the target storage to which the proxy will connect. In this case, VM data will be transported over LAN using the NBD protocol.

General Requirements and Limitations

- The role of a backup proxy can be assigned to the following machines:
  - Physical or virtual Microsoft Windows machine
  - Physical or virtual Linux machine
- The machine must meet the system requirements. For more information, see System Requirements.

- The account that you specify when adding a server must have permissions described in Required Permissions.
- You must add the machine to the Veeam Backup & Replication console as a managed server.
- If you back up proxies that use the Virtual appliance (Hot-Add) mode to process VM data, the change block tracking mechanism (CBT) will be disabled. For more information on CBT, see Changed Block Tracking.
- A backup proxy should be as close to the source data as possible with a high bandwidth connection. Consider a good connection between proxy and repository.

Requirements and Limitations for Linux Backup Proxy

In addition to the general requirements and limitations, the following ones apply to Linux backup proxies:

- Linux backup proxies require SSH connection.

  Therefore, you cannot assign the role of a backup proxy to a Linux server added with single-use credentials. For more information, see Requirements and Limitations for Veeam Data Movers and Specify Credentials and SSH Settings step of the **Adding Linux Servers** wizard.
- The `disk.EnableUUID` parameter of the Linux server must be set to the `TRUE` value in the VMware vSphere client.
- Linux backup proxies cannot be used with VMware Cloud on AWS. This is because VDDK settings required by VMware cannot be enabled on Linux backup proxies.
- Linux backup proxies that use virtual appliance (Hot-Add) transport mode do not support the VM copy scenario.
- For integration with storage systems: Linux backup proxies do not support the NFS protocol.
- Linux backup proxies cannot be used for guest interaction.
- For Direct SAN access, note that Linux backup proxies must have the Open-iSCSI initiator enabled.
- For Direct NFS access, mind the following:
  - Linux backup proxies must have NFS client package installed.
  - Debian-based backup proxies must have the `nfs-common` package installed.
  - RHEL-based backup proxies must have the `nfs-utils` package installed.
- See recommendations for backup proxy parameters in the Veeam Backup & Replication Best Practice Guide.

Transport Modes

A transport mode is a method that is used by the Veeam Data Mover to retrieve VM data from the source and write VM data to the target. Job efficiency and time required for job completion greatly depend on the transport mode.

For data retrieval, Veeam Backup & Replication offers the following modes (starting from the most efficient):

- Direct storage access
- Virtual Appliance (HotAdd)
- Network

The Veeam Data Mover responsible for data retrieval runs on a backup proxy. Correspondingly, the transport mode can be defined in the settings of the backup proxy that performs the job.

When you configure backup proxy settings, you can manually select a transport mode, or let Veeam Backup & Replication select the most appropriate mode automatically. If you use automatic mode selection, Veeam Backup & Replication will scan backup proxy configuration and its connection to the VMware vSphere infrastructure to choose the optimal transport mode. If several transport modes are available for the same backup proxy, Veeam Backup & Replication will choose the mode in the following order: Direct storage access > Virtual appliance > Network.

The selected transport mode is used for data retrieval. For writing data to the target, Veeam Backup & Replication picks the transport mode automatically, based on the configuration of the backup proxy and transport mode limitations.

Veeam Backup & Replication leverages VMware vStorage APIs for Data Protection (VADP) for all transport modes except for backup from storage snapshots, the direct NFS transport mode and the virtual appliance transport mode. VADP can be used for VMware vSphere starting from version 4.

Applicability and efficiency of each transport mode primarily depends on the type of datastore used by the source host — local or shared, and on the backup proxy type — physical or virtual. The table below shows recommendations for installing the backup proxy, depending on the storage type and desired transport mode.

<table>
<tr><th>Production Storage Type</th><th>Direct Storage Access</th><th>Virtual Appliance</th><th>Network Mode</th></tr>
<tr><td>Fiber Channel (FC) SAN</td><td>Install a backup proxy on a physical server with a direct FC access to the SAN.</td><td rowspan="4">Install a backup proxy on a VM running on an ESXi host connected to the storage device.</td><td rowspan="8">This mode is not recommended on 1 Gb Ethernet but works well with 10 Gb Ethernet.

Install a backup proxy on any machine on the storage network.</td></tr>
<tr><td>iSCSI SAN</td><td rowspan="2">Install a backup proxy on a physical or virtual machine.</td></tr>
<tr><td>NFS Storage</td></tr>
<tr><td>Shared SAS</td><td>Install a backup proxy on a physical server with a direct SAS access to the SAN.</td></tr>
<tr><td>vSAN</td><td rowspan="3">Not supported.</td><td>Install a backup proxy on a VM running on an ESXi host connected to the VSAN storage device.</td></tr>
<tr><td>VVol</td><td>Install a backup proxy on a VM running on an ESXi host connected to the VVol storage.</td></tr>
<tr><td>Local Storage</td><td>Install a backup proxy on a VM on every ESXi host.</td></tr>
</table>

Production Storage Type	Direct Storage Access	Virtual Appliance	Network Mode

NOTE

If you use VMware Cloud on AWS, the only available transport mode is Virtual appliance. We recommend you install a backup proxy on a VM running on an ESXi host connected to the VSAN storage device.

Direct Storage Access

In the Direct storage access mode, Veeam Backup & Replication reads/writes data directly from/to the storage system where VM data or backups are located. This mode comprises two transport modes:

- Direct SAN access
- Direct NFS access

Direct SAN Access

The Direct SAN access transport mode is recommended for VMs whose disks are located on shared VMFS SAN LUNs that are connected to ESXi hosts over FC, FCoE, iSCSI, and on shared SAS storage.

In the Direct SAN access transport mode, Veeam Backup & Replication leverages VMware VADP to transport VM data directly from and to FC, FCoE and iSCSI storage over the SAN. VM data travels over the SAN, bypassing ESXi hosts and the LAN. The Direct SAN access transport method provides the fastest data transfer speed and produces no load on the production network.

The Direct SAN access transport mode can be used for all operations where the backup proxy is engaged:

- Backup
- Replication
- VM copy
- Quick migration
- Entire VM restore
- VM disk restore
- Replica failback

Requirements for the Direct SAN Access Mode

To use the Direct SAN access transport mode, make sure that the following requirements are met:

- It is strongly recommended that you assign the role of a backup proxy working in the Direct SAN access mode to a physical machine. If you assign this role to a VM, the backup proxy performance may not be optimal.
- A backup proxy using the Direct SAN access transport mode must have a direct access to the production storage using a hardware or software HBA. If a direct SAN connection is not configured or not available when a job or task starts, the job or task will fail.

- SAN storage volumes presented as VMware datastores must be exposed to the OS of the backup proxy that works in the Direct SAN access transport mode.

  The volumes must be visible in Disk Management but must not be initialized by the OS. Otherwise, the VMFS filesystem will be overwritten with NTFS, and volumes will become unrecognizable by ESXi hosts. To prevent volumes initialization, Veeam Backup & Replication automatically sets the SAN Policy within each proxy to Offline Shared.

- [For restore operations] A backup proxy must have write access to LUNs where VM disks are located.

Limitations for the Direct SAN Access Mode

- The Direct SAN access transport mode can be used to restore only thick VM disks.
- You cannot use the Direct SAN access mode in the following cases:
  - For VMs residing on vSAN. You can use Virtual appliance and Network transport modes to process such VMs. For details on vSAN restrictions, see VDDK release notes. For example, release notes for VDDK 7.0.3.
  - If at least one VM disk is located on a VVol.
  - For Veeam Cloud Connect Replication because in this scenario Veeam Backup & Replication always creates VM replicas with thin disks.
  - For incremental restore due to VMware limitations. Either disable CBT for VM virtual disks for the duration of the restore process or select another transport mode for incremental restore.
  - For backing up VM templates.
- Veeam Backup & Replication uses the Direct SAN access transport mode to read and write VM data only during the first session of the replication job. During subsequent replication job sessions, Veeam Backup & Replication will use the Virtual appliance or Network transport mode on the target side. The source side proxy will keep reading VM data from the source datastore in the Direct SAN access transport mode.

  Veeam Backup & Replication writes VM data to the target datastore in the Direct SAN access transport mode only if disks of a VM replica are thick-provisioned. If disks are thin-provisioned, Veeam Backup & Replication will write VM data in the Network or Virtual appliance mode. By default, Veeam Backup & Replication replicates VM disks in the thin format. To write VM data to the target datastore in the Direct SAN access transport mode, select to convert VM disks to the thick format at the **Destination** step of the replication job wizard.
- IDE and SATA disks can be processed in the Direct SAN access transport mode.

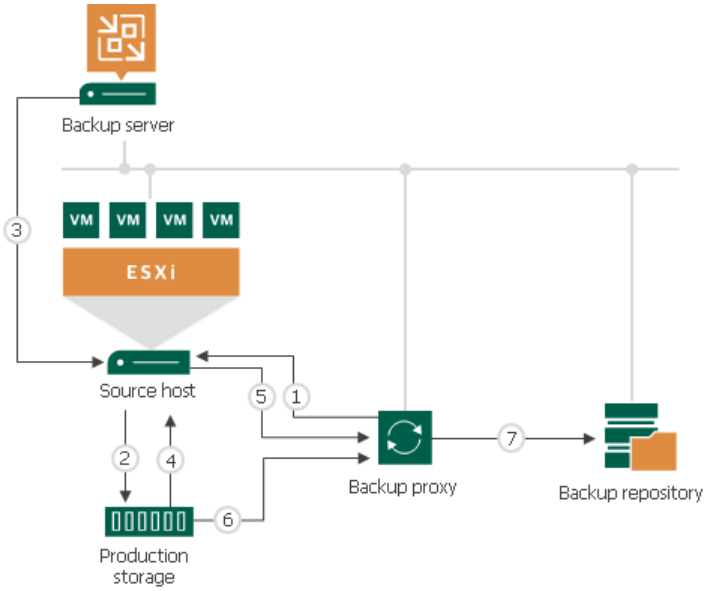
Data Backup in Direct SAN Access Mode

To retrieve VM data blocks from a SAN LUN during backup, the backup proxy uses metadata about the layout of VM disks on the SAN.

Data backup in the Direct SAN access transport mode includes the following steps:

1. The backup proxy sends a request to the ESXi host to locate the necessary VM on the datastore.
2. The ESXi host locates the VM.
3. Veeam Backup & Replication triggers VMware vSphere to create a VM snapshot.
4. The ESXi host retrieves metadata about the layout of VM disks on the storage (physical addresses of data blocks).

5. The ESXi host sends metadata to the backup proxy.
6. The backup proxy uses metadata to copy VM data blocks directly from the source storage over the SAN.
7. The backup proxy processes copied data blocks and sends them to the target.



Data Restore in Direct SAN Access Mode

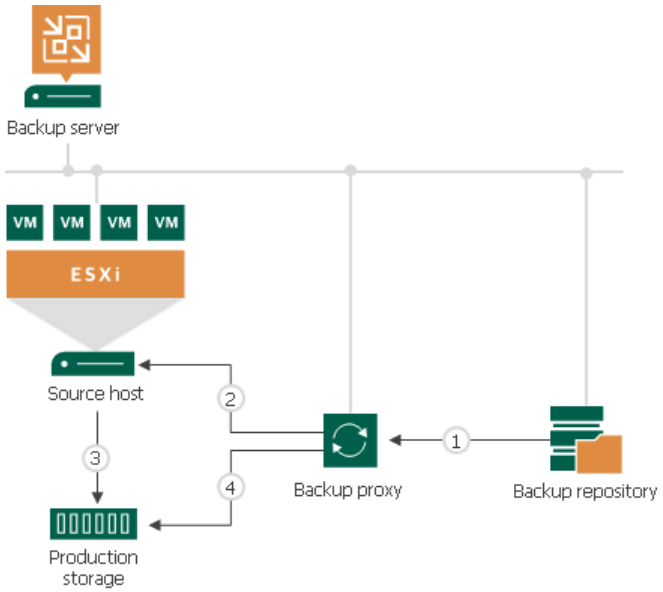
The Direct SAN access transport mode can be used to restore VMs with thick disks. To restore VMs with thin disks, you can use the Virtual appliance or Network mode.

If you plan to process VMs that have both thin and thick disks, you can enable the Direct SAN access mode. However, Veeam Backup & Replication will use the Network transport mode to restore disks of these VMs. If you want to use Direct SAN access mode, restore all VM disks as thick.

Data restore in the Direct SAN access transport mode includes the following steps:

1. The backup proxy retrieves data blocks from the backup repository or a datastore in the target site.
2. The backup proxy sends a request to the ESXi host in the source site to restore data to a necessary datastore.
3. The ESXi host in the source site allocates space on the datastore.

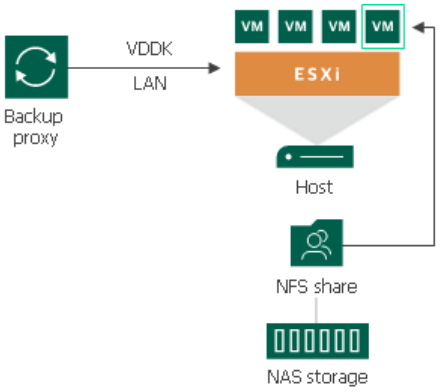
4. Data blocks are written to the datastore.



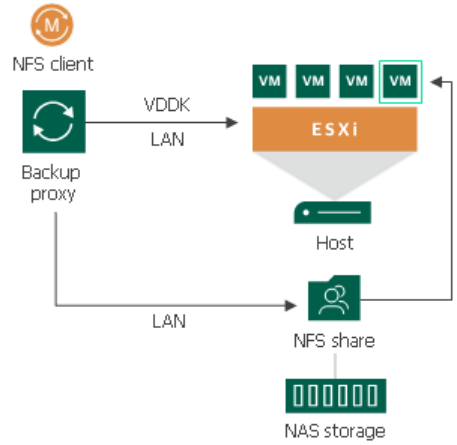
Direct NFS Access

The Direct NFS access is a recommended transport mode for VMs whose disks are located on NFS datastores.

The Direct NFS access mode provides an alternative to the Network mode. When Veeam Backup & Replication processes VM data in the Network mode, it uses VMware VDDK to communicate with the ESXi host. This produces additional load on the ESXi host.



In the Direct NFS access mode, Veeam Backup & Replication bypasses the ESXi host and reads/writes data directly from/to NFS datastores. To do this, Veeam Backup & Replication deploys its native NFS client on the backup proxy and uses it for VM data transport. VM data still travels over LAN but there is no load on the ESXi host.



The Direct NFS access mode can be used for all operations where the backup proxy is engaged:

- Backup
- Replication
- Quick migration
- VM copy
- Entire VM restore
- VM disk restore
- Replica failback

Requirements for the Direct NFS Access Mode

- Direct NFS access mode can be used in VMware vSphere environments running NFS version 3 and 4.1.
- The backup proxy used for VM data processing must have access to the NFS datastores where VM disks are located. For more information, see Backup Proxy for Direct NFS Access Mode.
- If NFS volumes are mounted on the ESXi host under names, not IP addresses, the volume names must be resolved by DNS from the backup proxy.

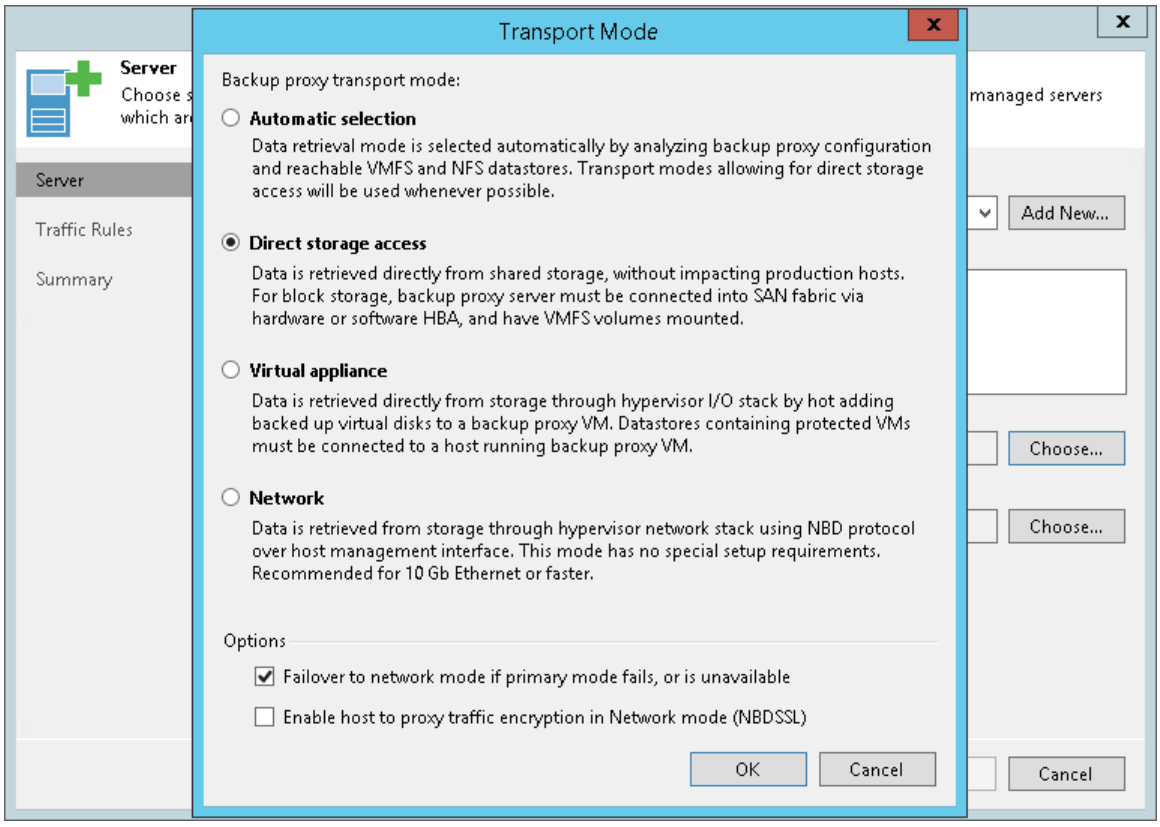
Limitations for Direct NFS Access Mode

- Veeam Backup & Replication cannot parse delta disks in the Direct NFS access mode. For this reason, the Direct NFS access mode has the following limitations:
  - The Direct NFS access mode cannot be used for VMs that have at least one snapshot.
  - Veeam Backup & Replication uses the Direct NFS transport mode to read and write VM data only during the first session of the replication job. During subsequent replication job sessions, the VM replica will already have one or more snapshots. For this reason, Veeam Backup & Replication will use another transport mode to write VM data to the datastore on the target side. The source side proxy will keep reading VM data from the source datastore in the Direct NFS transport mode.

- If you enable the **Enable VMware tools quiescence** option in the job settings, Veeam Backup & Replication will not use the Direct NFS transport mode to process running Microsoft Windows VMs that have VMware Tools installed. The Direct NFS transport mode is not used because during VM quiescence VMware creates a snapshot with two delta disks per virtual disk.
- If a VM has some disks that cannot be processed in the Direct NFS access mode, Veeam Backup & Replication processes these VM disks in the Network transport mode.

Backup Proxy for Direct NFS Access Mode

Veeam Backup & Replication deploys its NFS agent on every backup proxy when you assign the backup proxy role to a Microsoft Windows server (physical or virtual). To instruct the backup proxy to use the Direct NFS access mode, you must choose the **Automatic selection** or **Direct storage access** option in the backup proxy settings.



To read and write data in the Direct NFS transport mode, the backup proxy must meet the following requirements:

1. The backup proxy must have access to the NFS datastore.
2. The backup proxy must have *ReadOnly/Write* permissions and root access to the NFS datastore.

Backup Proxy Selection

Veeam Backup & Replication selects backup proxies working in the Direct NFS access transport mode by the following rules:

- If you instruct Veeam Backup & Replication to select a backup proxy automatically for a job or task, Veeam Backup & Replication picks a backup proxy with the minimum number of hops to the NFS datastore. If there are several backup proxies with the equal number of hops in the backup infrastructure, Veeam Backup & Replication picks the least busy backup proxy in the backup infrastructure.

  If all backup proxies with the minimum number of hops are busy at the moment, Veeam Backup & Replication waits until these backup proxies are free. Veeam Backup & Replication does not pick a backup proxy that has a greater number of hops to the NFS datastore and works in the Direct NFS access or Virtual appliance transport mode.

- If you select one or more backup proxies explicitly for a job or task, Veeam Backup & Replication does not regard the number of hops to the NFS datastore. Veeam Backup & Replication picks the least busy backup proxy working in the Direct NFS access transport mode.

  If all backup proxies working in the Direct NFS access transport mode are busy, Veeam Backup & Replication waits until these backup proxies are free. Veeam Backup & Replication does not pick a backup proxy working in the Virtual appliance transport mode.

To detect the number of hops from a backup proxy to the NFS datastore, Veeam Backup & Replication uses the host discovery process. During host discovery, Veeam Backup & Replication obtains information about the number of hops, checks to which NFS datastores the backup proxy has access and what permissions the backup proxy has on NFS datastores.

The host discovery process rescans all Microsoft Windows machines to which the backup proxy role is assigned. The process starts automatically every 4 hours. Host discovery is also triggered when you change the transport mode settings and choose to use the Direct storage access for the backup proxy.

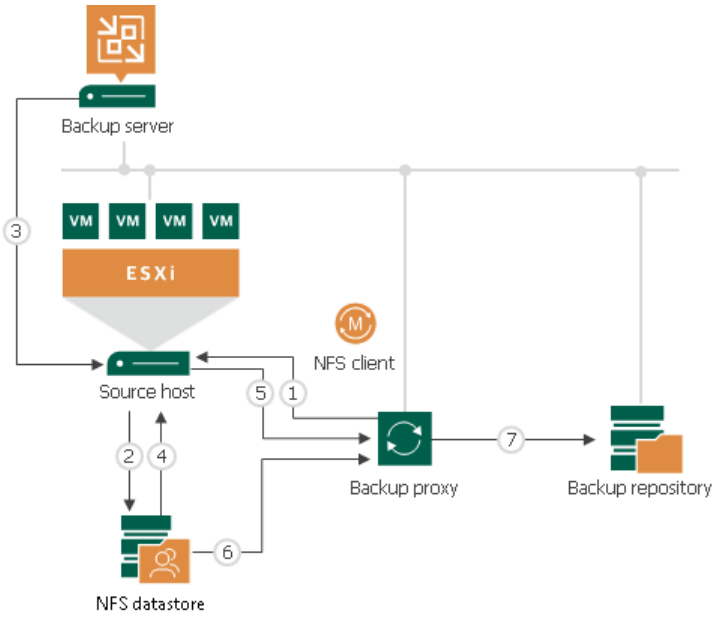
If necessary, you can start the host discovery process manually. To do this, perform the **Rescan** operation for a machine to which the backup proxy role is assigned.

Data Backup in Direct NFS Access Mode

Data backup in the Direct NFS access transport mode is performed in the following way:

1. The backup proxy sends a request to the ESXi host to locate a VM on the NFS datastore.
2. The ESXi host locates the VM.
3. Veeam Backup & Replication triggers VMware vSphere to create a VM snapshot.
4. The ESXi host retrieves metadata about the layout of VM disks on the storage (physical addresses of data blocks).
5. The ESXi host sends metadata to the backup proxy.
6. The backup proxy uses metadata to copy VM data blocks directly from the NFS datastore over LAN.

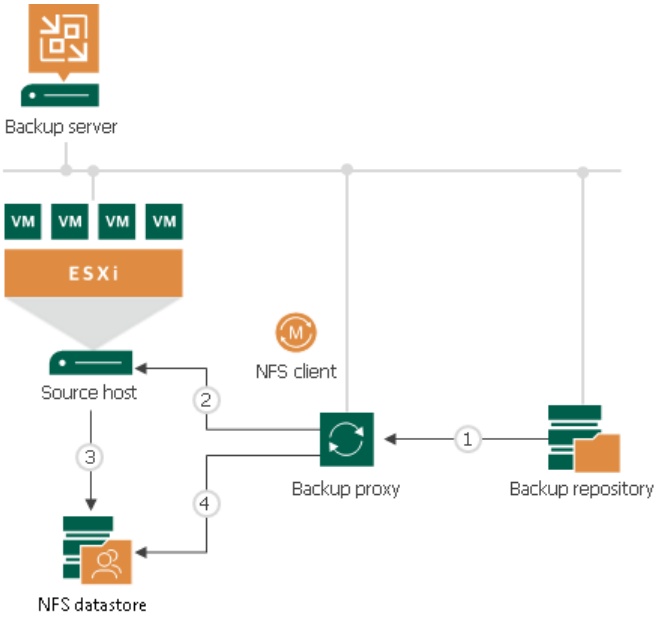
7. The backup proxy processes copied data blocks and sends them to the target over LAN.



Data Restore in Direct NFS Access Mode

Data restore in the Direct NFS access transport mode is performed in the following way:

1. The backup proxy retrieves data blocks from the backup repository or a datastore in the target site.
2. The backup proxy sends a request to the ESXi host to restore data to an NFS datastore.
3. The ESXi host allocates space on the NFS datastore.
4. Data blocks obtained from the backup proxy are written to the NFS datastore over LAN.



Virtual Appliance (HotAdd)

The Virtual appliance mode is not so efficient as the Direct storage access mode but provides better performance than the Network mode. The Virtual appliance mode is recommended if the role of a backup proxy is assigned to a VM.

In the Virtual appliance mode, Veeam Backup & Replication uses the VMware SCSI HotAdd capability that allows attaching devices to a VM while the VM is running. During backup, replication or restore disks of the processed VM are attached to the backup proxy. VM data is retrieved or written directly from/to the datastore, instead of going through the network.

The Virtual appliance transport mode can be used for all operations where the backup proxy is engaged:

- Backup
- Replication
- VM copy
- Quick migration
- Entire VM restore
- VM disk restore
- Replica failback

Requirements for the Virtual Appliance mode

To use the Virtual appliance transport mode, make sure that the following requirements are met:

- The role of a backup proxy must be assigned to a VM.
- The backup proxy and processed VMs must reside in the same datacenter.
- The ESXi host on which the backup proxy is deployed must have access to the datastore hosting disks of VMs that you plan to process.
- [For NFS 3.0] If you plan to process VMs that store disks on the NFS datastore, you must configure Veeam Backup & Replication to use the proxy on the same host as VMs. This is required due to an issue described in this VMware KB article. For more information on how to configure the proxy, see this Veeam KB article.

  As an alternative, you can use ESXi 6.0 or higher and NFS 4.1.
- The backup proxy must have the latest version of VMware Tools installed. Note that the backup server installed on a VM can also perform the role of the backup proxy that uses Virtual appliance transport mode. In this case, make sure the backup server has the latest version of VMware Tools installed.
- SCSI 0:X controller must be present on a backup proxy. In the opposite case, VM data processing in the Virtual appliance transport mode will fail.

Limitations for the Virtual Appliance mode

- If a backup proxy used to process a source VM resides on a VMFS 3 datastore, it must be formatted with proper block size to be able to mount the largest virtual disk of hot-added VMs:
  - 1 MB block size — 256 GB maximum file size
  - 2 MB block size — 512 GB maximum file size

- 4 MB block size — 1024 GB maximum file size
- 8 MB block size — 2048 GB maximum file size

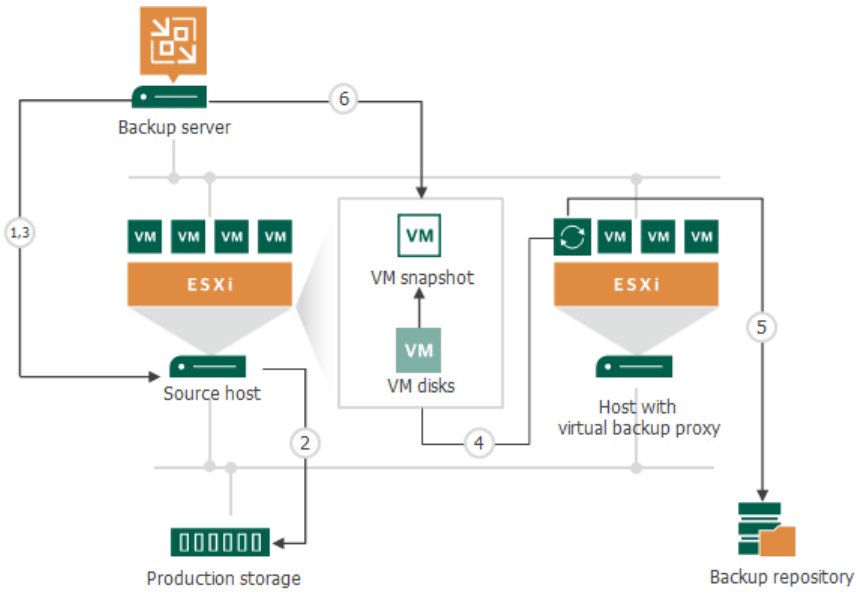
This limitation does not apply to VMFS-5 volumes that always have 1 MB file block size.

- For vSphere 5.5 and later the maximum supported VMDK size is 62 TB.
- [For Microsoft Windows proxy] Before running a data protection task, Veeam Backup & Replication disables the volume automount feature, and it remains disabled after the data protection task is completed.
- Backup and restore of IDE disks in the Virtual appliance mode is not supported.
- Backup and restore of SATA disks in the Virtual appliance mode is supported if you use VMware vSphere 6.0 and later.

Data Backup and Restore in Virtual Appliance Mode

The process of data retrieval in the Virtual appliance transport mode includes the following steps:

1. The backup server sends a request to the ESXi host to locate the necessary VM on the datastore.
2. The ESXi host locates the VM.
3. Veeam Backup & Replication triggers VMware vSphere to create a VM snapshot.
4. VM disks are attached (hot-added) to the backup proxy.
5. Veeam Backup & Replication reads data directly from disks attached to the backup proxy.
6. When the VM processing is complete, VM disks are detached from the backup proxy and the VM snapshot is deleted.



The process of data restore in the Virtual appliance mode works in a similar manner. VM disks from the backup are attached to the backup proxy and Veeam Backup & Replication transports VM data to the target datastore. After the restore process is finished, VM disks are detached from the backup proxy.

ESXi host interacts with VMware Cloud on AWS through VMware vCenter. Veeam Backup & Replication performs backup through the networkless Virtual appliance (HotAdd) mode.

Virtual Appliance Mode for VMs on VSAN

To transport data of VMs residing on VSAN in the Virtual appliance mode, you must assign the backup proxy role to a VM.

The backup proxy VM must meet the following requirements:

- The backup proxy VM must reside on any of ESXi hosts connected to a VSAN cluster.

  Veeam Backup & Replication will retrieve data of processed VMs over the I/O stack of the ESXi host on which the backup proxy is deployed.

- Disks of the backup proxy VM must reside on the VSAN datastore.

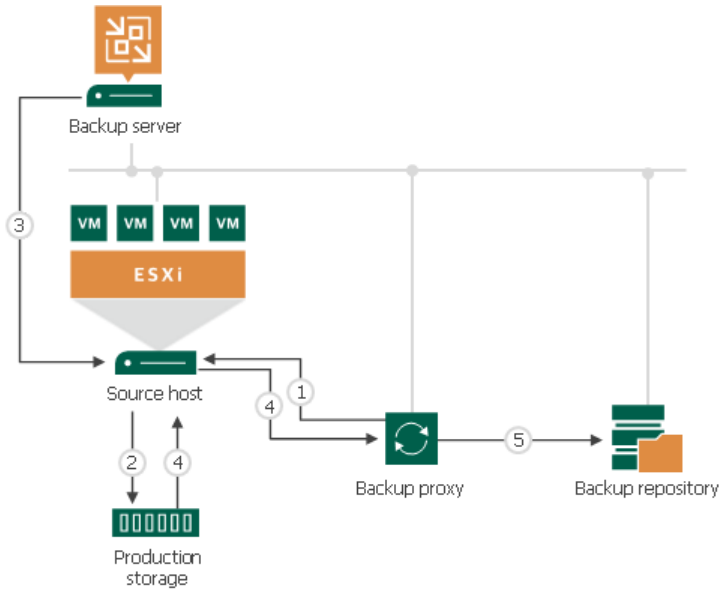
If you have several backup proxies on ESXi hosts in the VSAN cluster, Veeam Backup & Replication chooses the most appropriate backup proxy to reduce the backup traffic on the VSAN cluster network. To choose a backup proxy, Veeam Backup & Replication checks HDDs directly attached to every ESXi host and calculates the amount of VM data on these HDDs. The preference is given to the ESXi host that has a direct access to an HDD with the maximum amount of VM data. This approach helps reduce workload on the ESXi I/O stack during data transport.

> **NOTE**
>
> Even if disks of a VM are located on a host where the backup proxy is deployed, VSAN traffic may still be observed between hosts in the cluster. This behavior depends on the VSAN cluster itself and cannot be modified in Veeam Backup & Replication.

Network Mode

The Network mode can be used with any infrastructure configuration. In this mode, data is retrieved through the ESXi host over LAN using the Network Block Device protocol (NBD).

The Network mode has low data transfer speed over LAN. To take the load off the LAN, Veeam Backup & Replication provides two alternative modes: Direct Storage Access and Virtual Appliance. However, the Network mode is the only applicable mode when the backup proxy role is assigned to a physical machine and the host uses local storage. Also, the Network mode can be the best choice if you have a large virtual environment with hundreds of small VMs, with 10 Gb Ethernet networks and with a small change rate.



The process of data retrieval in Network mode includes the following steps:

1. The backup proxy sends a request to the ESXi host on which the processed VM is registered to locate the VM on the datastore.
2. The ESXi host locates the processed VM on the datastore.
3. Veeam Backup & Replication instructs VMware vSphere to create a VMware vSphere VM snapshot.
4. ESXi host copies VM data blocks from the source storage and sends them to the backup proxy over LAN.

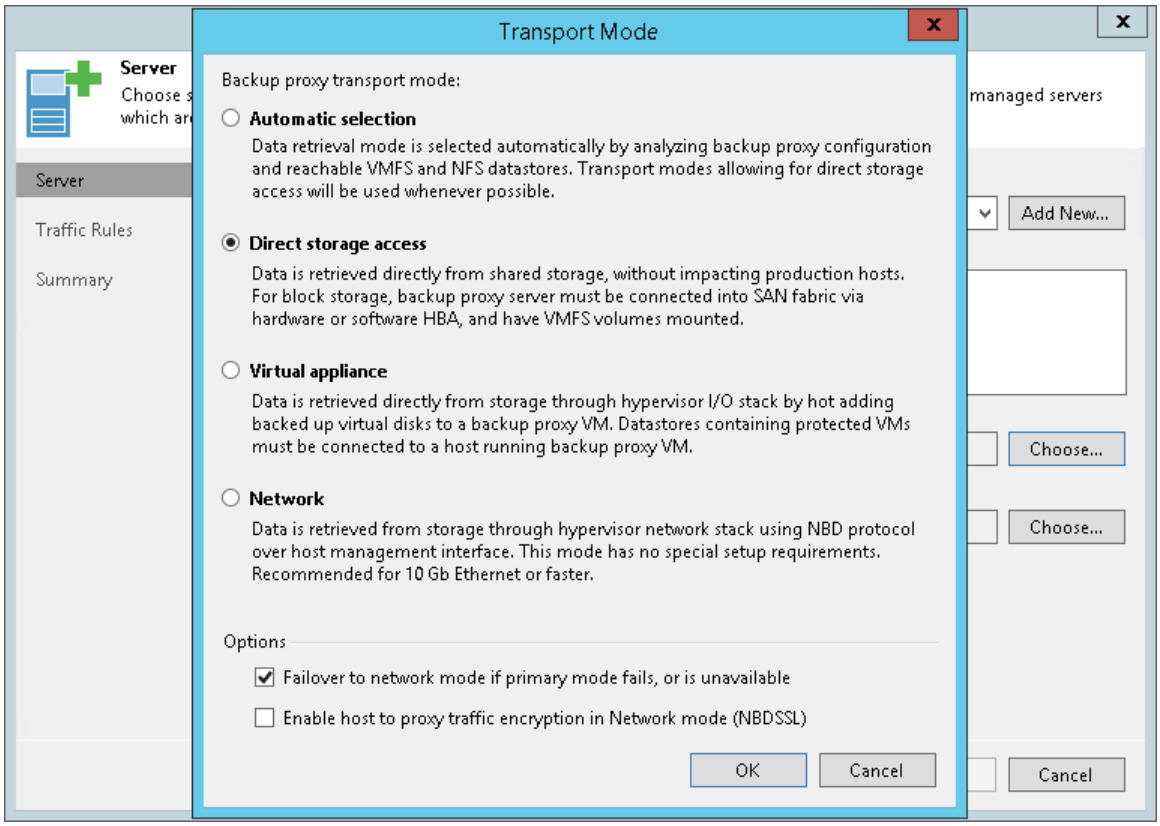
   Note that the real data transfer speed may be significantly less than the available speed. This is because the backup proxy and the ESXi host communicate over the ESXi management network.
5. The backup proxy sends the data to target.

Veeam Backup & Replication processes VM disks in parallel. If VM disks are located on different storage types (for example, on the SAN and local storage), Veeam Backup & Replication uses different transport modes to process VM disks. In such scenario, it is strongly recommended that you select the **Failover to network mode if primary mode fails, or is unavailable** option when configuring the mode settings for the backup proxy.

Failover to Network Mode

You can instruct Veeam Backup & Replication to switch to the Network transport mode and transfer VM data over LAN if the primary transport mode — Direct storage access or Virtual appliance — fails during the job session. This option is enabled by default to ensure that jobs and tasks can be completed successfully in any situation.

Note that data transport over LAN puts additional load on your production network and may potentially affect performance if you accomplish data protection and disaster recovery tasks in business hours.



Adding VMware Backup Proxies

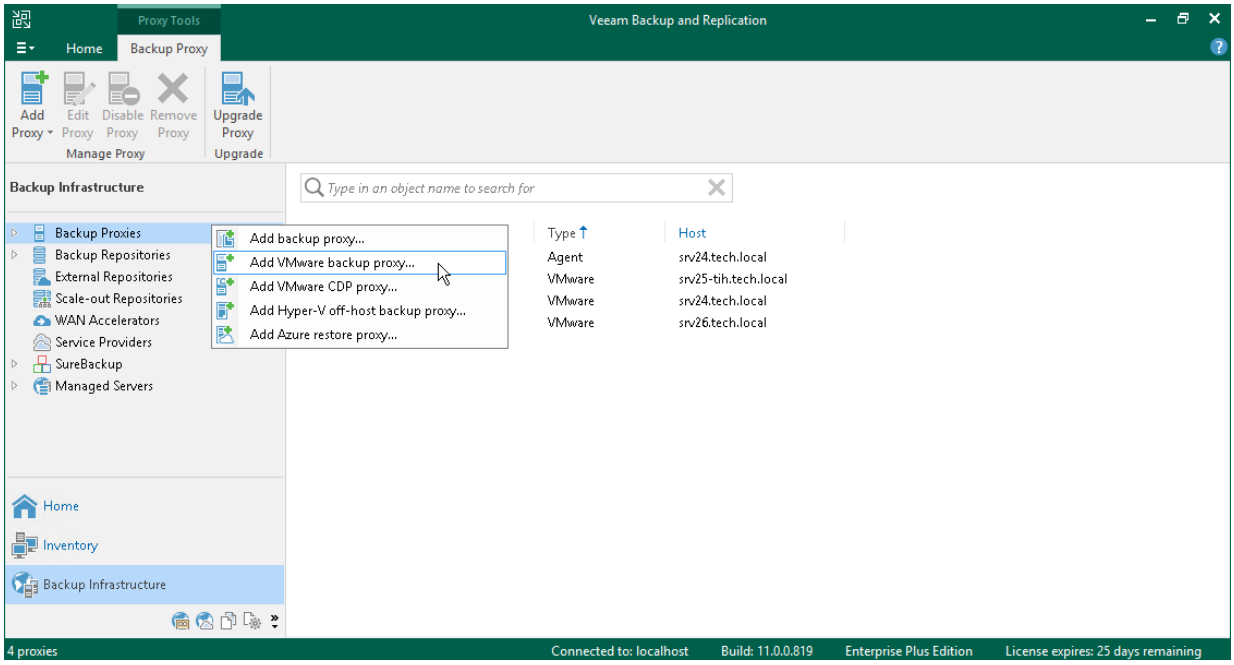
You can configure one or more backup proxies in the backup infrastructure.

Before adding a backup proxy, check prerequisites. Then use the **New VMware Proxy** wizard.

Step 1. Launch New VMware Proxy Wizard

To launch the **New VMware Proxy** wizard, do either of the following:

- Open the **Backup Infrastructure** view, in the inventory pane select the **Backup Proxies** node, click **Add Proxy** on the ribbon and select **VMware**.
- Open the **Backup Infrastructure** view, in the inventory pane right-click the **Backup Proxies** node and select **Add VMware Backup Proxy**.



Step 2. Choose Server

At the **Server** step of the wizard, specify server settings for the backup proxy:

1. From the **Choose server** list, select a Microsoft Windows or Linux server to which you want to assign the backup proxy role.

   If the server is not added to the backup infrastructure, click **Add New** to open the **Add Server** wizard. For more information, see Adding Microsoft Windows Servers or Adding Linux Servers.

   > **NOTE**
   >
   > If a user account specified for the Linux server does not have root or elevated to root permission, Veeam Backup & Replication will direct you to the **Edit Linux Server** wizard. In this wizard, you can change the user account.

2. In the **Proxy description** field, provide a description. The default description contains information about the user who added the backup proxy, date and time when the backup proxy was added.

3. By default, Veeam Backup & Replication analyzes the backup proxy configuration, defines to which datastores it has access and automatically selects the best transport mode depending on the type of connection between the backup proxy and datastores.

   For Microsoft Windows backup proxy, you can manually select the data transport mode. Click **Choose** on the right of the **Transport mode** field. In the opened window, select one of the following modes: *Direct storage access*, *Virtual appliance* or *Network*. For more information, see Transport Modes.

4. In the **Options** section of the **Transport Mode** window, specify additional options for the selected transport mode:

   - [For the Direct storage access and Virtual appliance transport modes] If the primary transport mode fails during the job session, Veeam Backup & Replication will automatically fail over to the Network transport mode. To disable failover, clear the **Failover to network mode if primary mode fails, or is unavailable** check box.
   - [For the Network mode] You can choose to transfer VM data over an encrypted TLS connection. To do this, select the **Enable host to proxy traffic encryption in Network mode (NBDSSL)** check box. Traffic encryption puts more stress on the CPU of an ESXi host but ensures secure data transfer.

   > **NOTE**
   >
   > In some cases, the backup proxy may not be able to use some transport modes due to known limitations. For more information, see Transport Modes.

5. In the **Connected datastores** field, specify datastores to which the backup proxy has a direct SAN or NFS connection. By default, Veeam Backup & Replication automatically detects all datastores that the backup proxy can access.

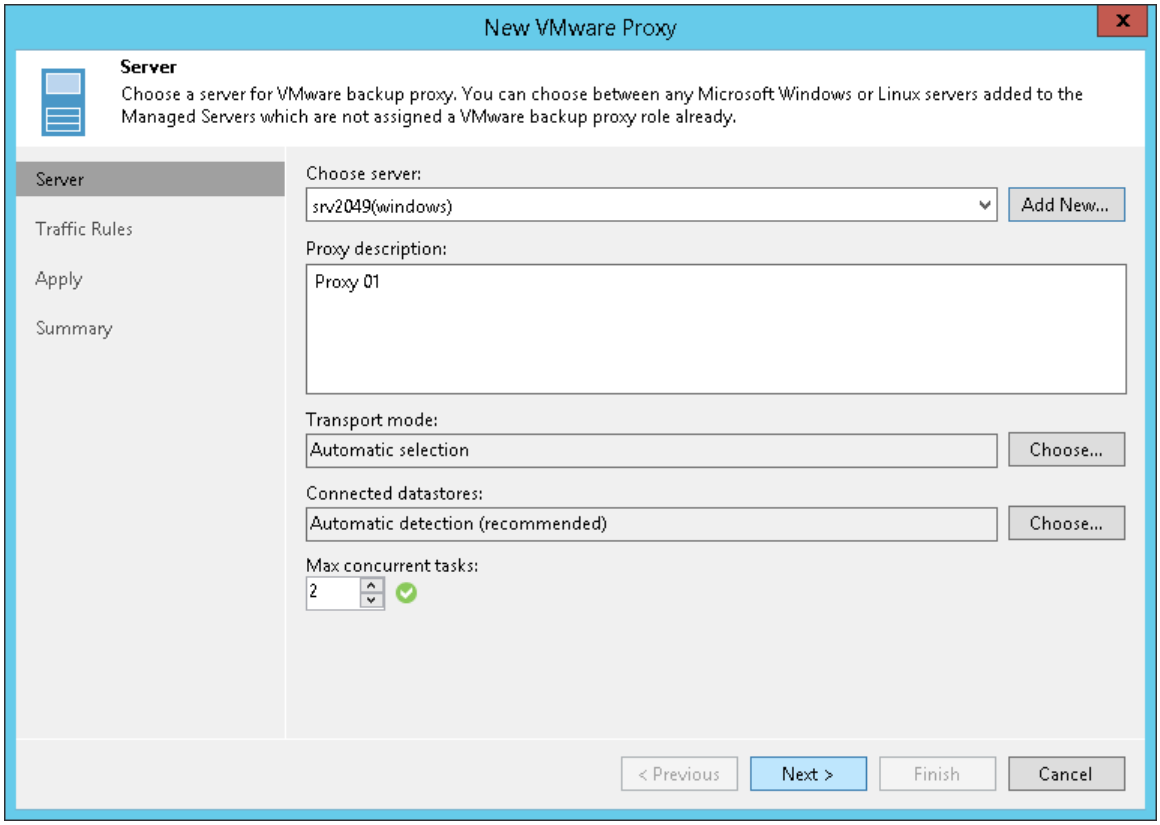
   You can set up the list of datastores if you want the backup proxy to work with specific datastores. Click **Choose** on the right of the **Connected datastores** field, choose **Manual selection** and add datastores with which the backup proxy must work in the Direct storage access mode.

6. In the **Max concurrent** tasks field, specify the number of tasks that the backup proxy must handle in parallel. If this value is exceeded, the backup proxy will not start a new task until one of current tasks finishes.

   Veeam Backup & Replication creates one task per every VM disk. The recommended number of concurrent tasks is calculated automatically based on available resources. Backup proxies with multi-core CPUs can handle more concurrent tasks. For example, for a 4-core CPU, it is recommended that you specify maximum 4 concurrent tasks, for an 8-core CPU — 8 concurrent tasks. When defining the number of concurrent tasks, keep in mind network traffic throughput in the virtual infrastructure.

   > **IMPORTANT**
   >
   > Limitation of concurrent tasks is not supported for vCD backup jobs.

7. Click **Next**.



Step 3. Select Proxy VM

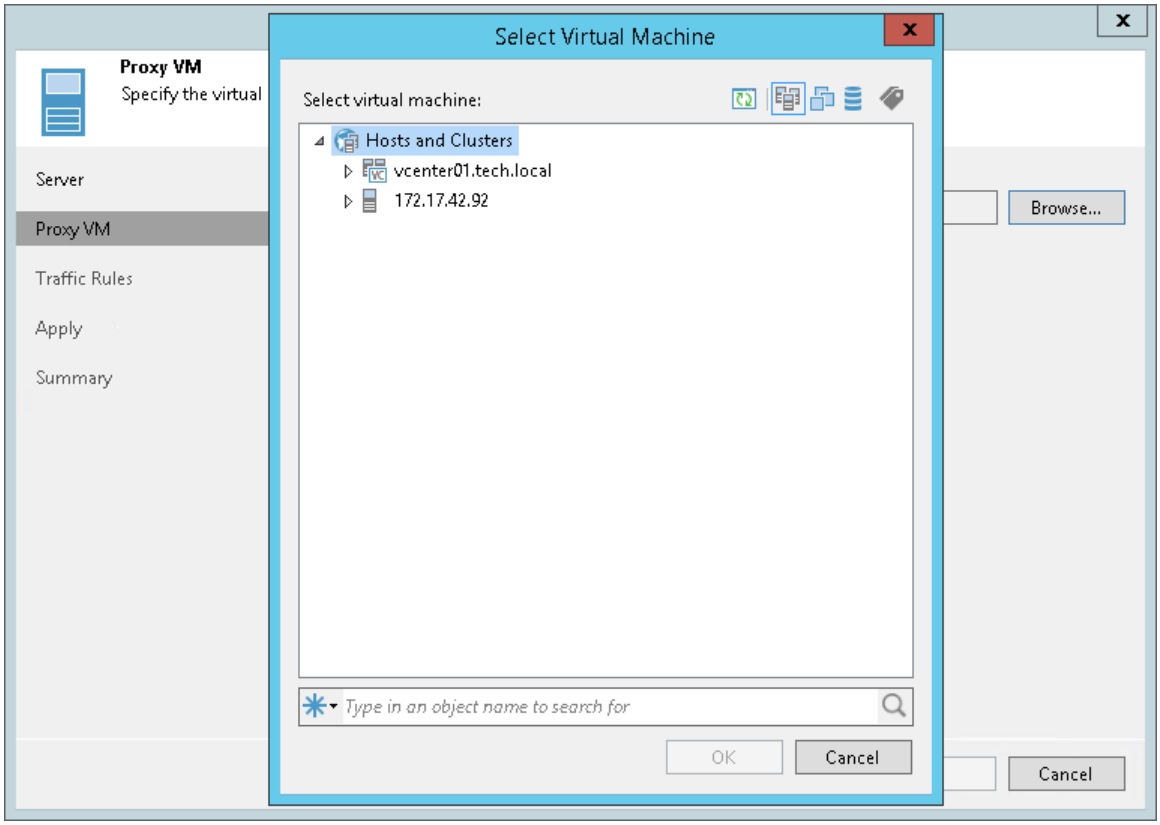
The additional **Proxy VM** step will appear in the following cases:

- If the `disk.EnableUUID` parameter is not set for the selected Linux server. In this case, select this Linux server once again at the **Proxy VM** step. After that, Veeam Backup & Replication will need to reboot the server to set the `disk.EnableUUID` parameter to `TRUE`.
- If Veeam Backup & Replication cannot identify the selected VM: for example, there are two VMs with the same BIOS UUID or BIOS UUID is not specified. In this case, choose the VM from the backup infrastructure at the **Proxy VM** step.

  To select a VM:

  a. At the **Proxy VM** step of the wizard, click the **Browse** button.

  b. In the **Select Virtual Machine** window, select a VM.



Step 4. Configure Traffic Rules

At the **Traffic Rules** step of the wizard, configure network traffic rules. These rules help you throttle and encrypt traffic transferred between backup infrastructure components. For more information, see Configuring Network Traffic Rules.

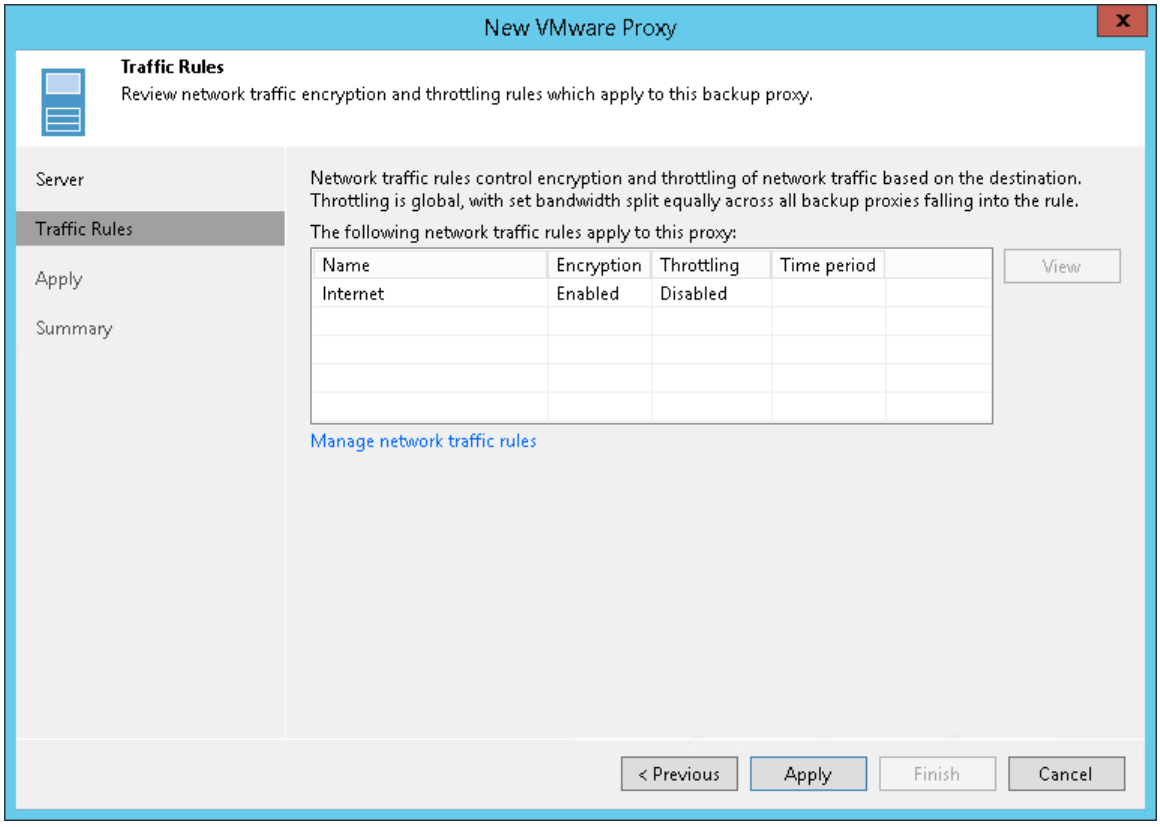
The list of network traffic rules contains only the rules that are applicable to the backup proxy: its IP address falls into the IP range of the rule.

To view rule settings:

1. Select a rule in the list.
2. Click **View** on the right of the rule list.

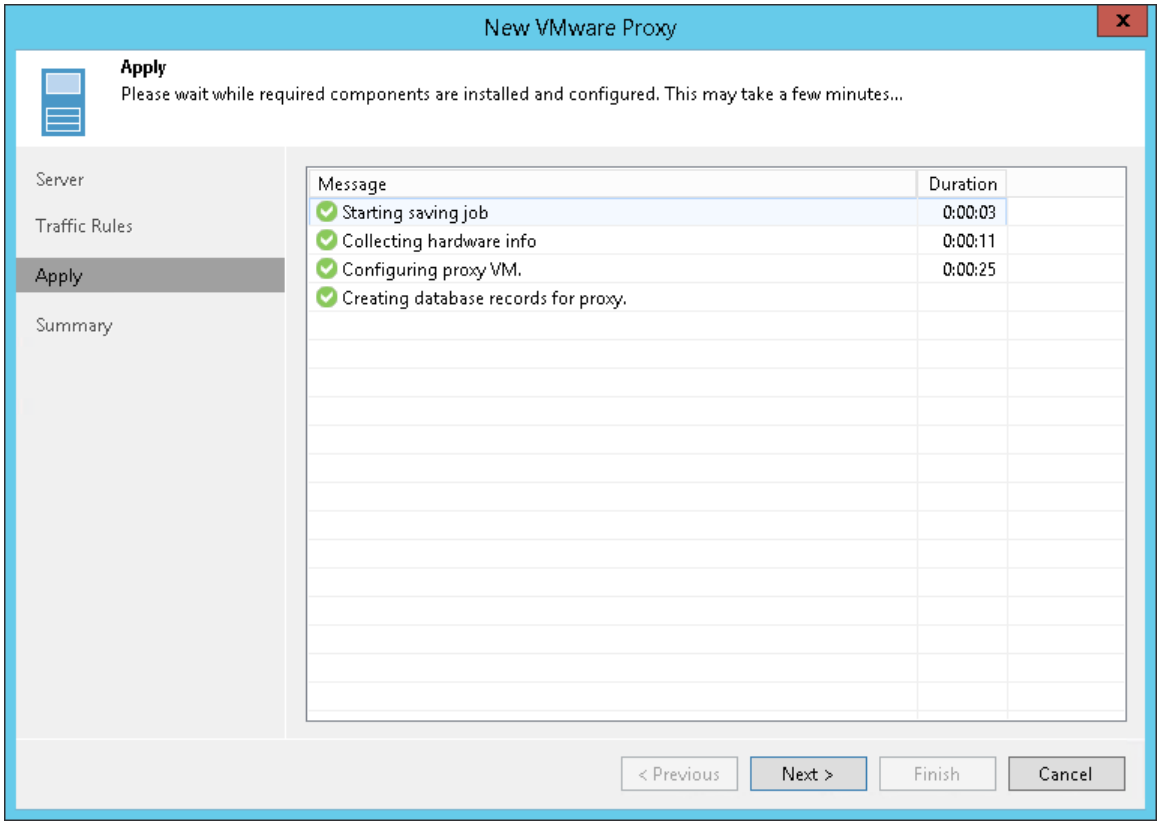
You can also modify network traffic settings:

1. Click **Manage network traffic rules** link at the bottom of the wizard.
2. In the opened window, you will see all global network traffic rules.
3. Select the rule that you want to modify and click **Edit**.



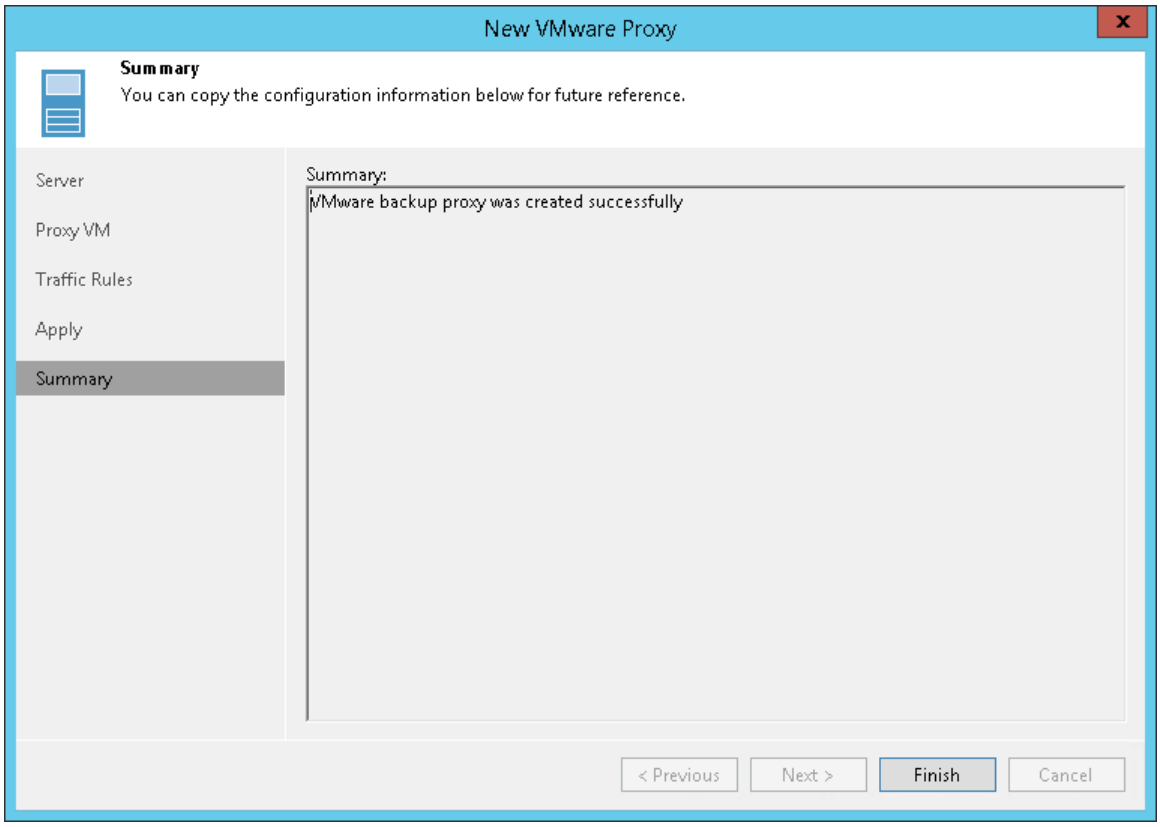
Step 5. Apply Backup Proxy Settings

At the **Apply** step of the wizard, wait for the backup proxy to be added to the backup infrastructure. Then click **Next**.



Step 6. Finish Working with Wizard

At the **Summary** step of the wizard, check that the backup proxy is added. Then click **Finish** to exit the wizard.



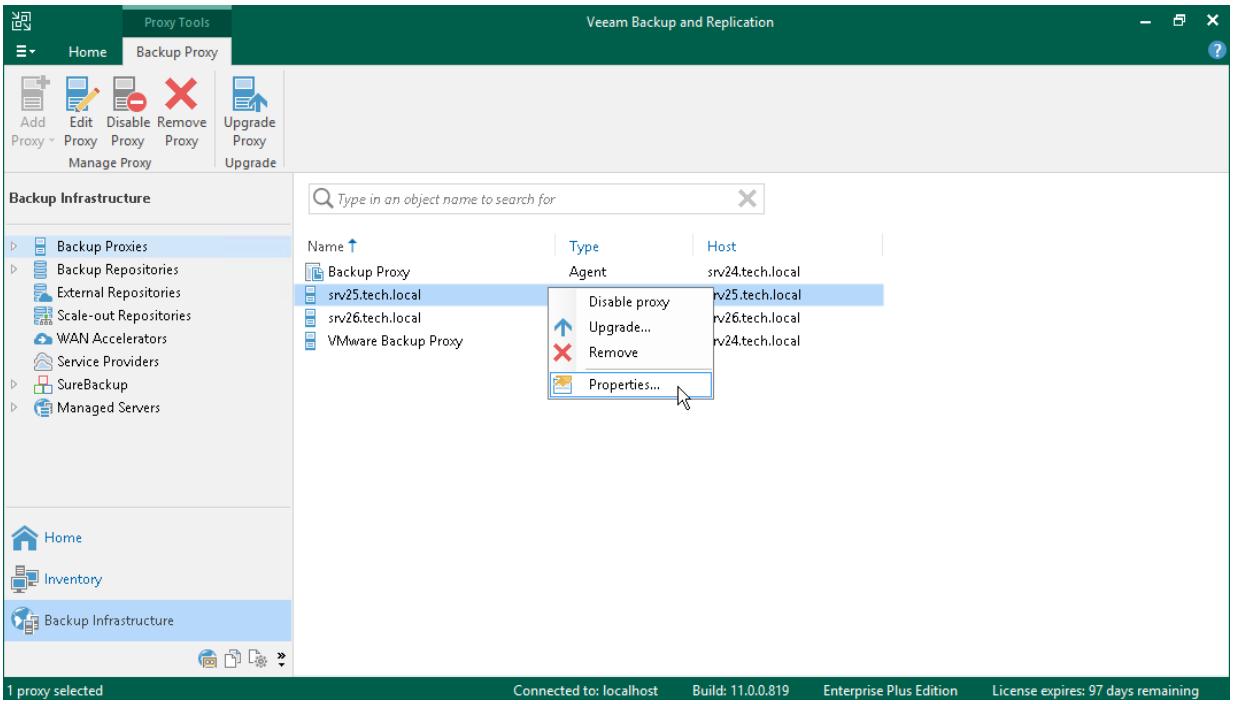
Editing Backup Proxy Settings

You can edit settings of backup proxies you have configured.

To edit backup proxy settings:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select the **Backup Proxies** node.
3. In the working area, select the backup proxy and click **Edit Proxy** on the ribbon or right-click the backup proxy and select **Properties**.

4. Edit backup proxy settings as required.



Disabling and Removing Backup Proxies

You can temporarily disable a backup proxy or remove it from the backup infrastructure.

Disabling Backup Proxies

When you disable a backup proxy, Veeam Backup & Replication does not use this backup proxy for any jobs configured on the backup server. Backup proxy disabling can be helpful if you instruct Veeam Backup & Replication to automatically select backup proxies for jobs and do not want Veeam Backup & Replication to use specific backup proxies.

You can disable all backup proxies, including the default backup proxy installed on the backup server. Do not disable all backup proxies at once. Otherwise, Veeam Backup & Replication will not be able to perform backup, replication and restore operations that use backup proxies.

To disable a backup proxy:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select the **Backup Proxies** node.
3. In the working area, select the backup proxy and click **Disable Proxy** on the ribbon or right-click the backup proxy and select **Disable proxy**.

You can enable a disabled backup proxy at any time:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select the **Backup Proxies** node.
3. In the working area, select the backup proxy and click **Disable Proxy** on the ribbon once again or right-click the backup proxy and select **Disable proxy**.

Removing Backup Proxies

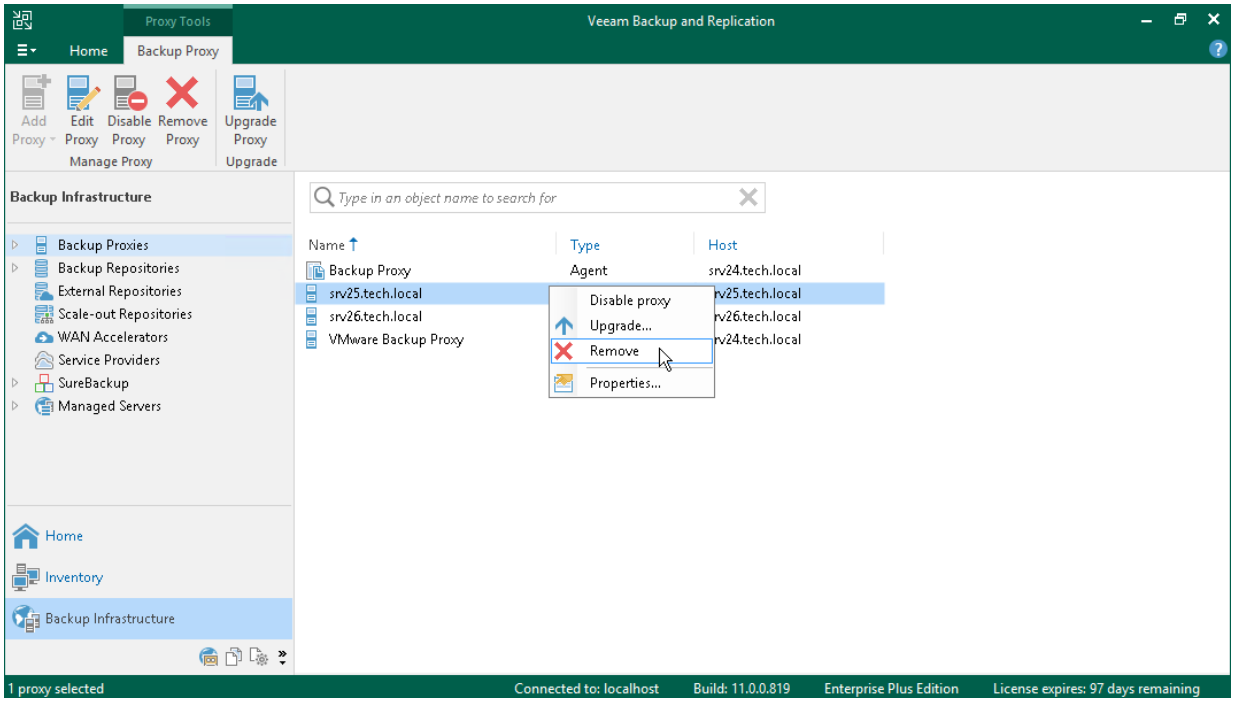
You can permanently remove a backup proxy from the backup infrastructure. When you remove a backup proxy, Veeam Backup & Replication unassigns the backup proxy role from the server, and this server is no longer used as a backup proxy. The actual server remains in the backup infrastructure.

You can remove all backup proxies, including the default backup proxy installed on the backup server. Do not remove all backup proxies at once. Otherwise, Veeam Backup & Replication will not be able to perform backup, replication and restore operations that use backup proxies.

You cannot remove a backup proxy that is explicitly selected in any backup, replication or VM copy job. To remove such backup proxy, you first need to delete a reference to this backup proxy in the job settings.

To remove a backup proxy:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select the **Backup Proxies** node.
3. In the working area, select the backup proxy and click **Remove Proxy** on the ribbon or right-click the backup proxy and select **Remove**.



VMware CDP Proxy

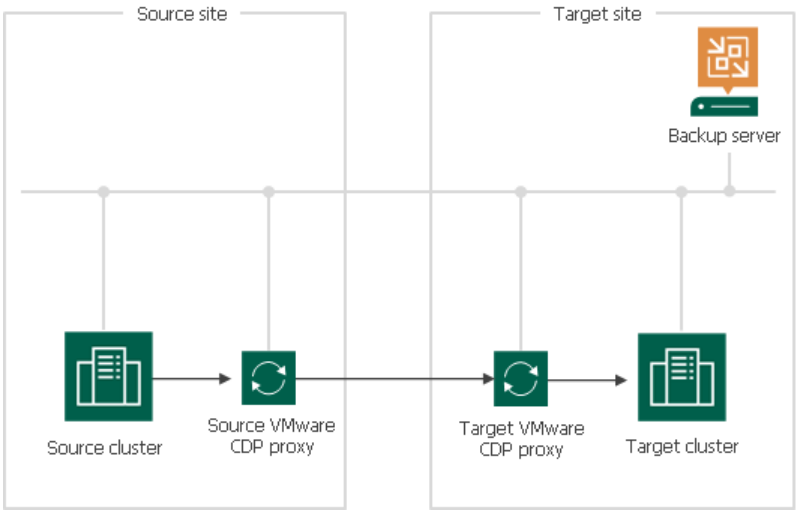
A VMware CDP proxy is a component that operates as a data mover and transfers data between the source and target hosts. Basically, VMware CDP proxy performs the following tasks:

- Receives VM data from the production storage
- Aggregates changed data
- Prepares data for short-term restore points
- Compresses and deduplicates data
- Encrypts and decrypts data
- Sends data to the storage in the disaster recovery site or another VMware CDP proxy

VMware CDP Proxy Deployment

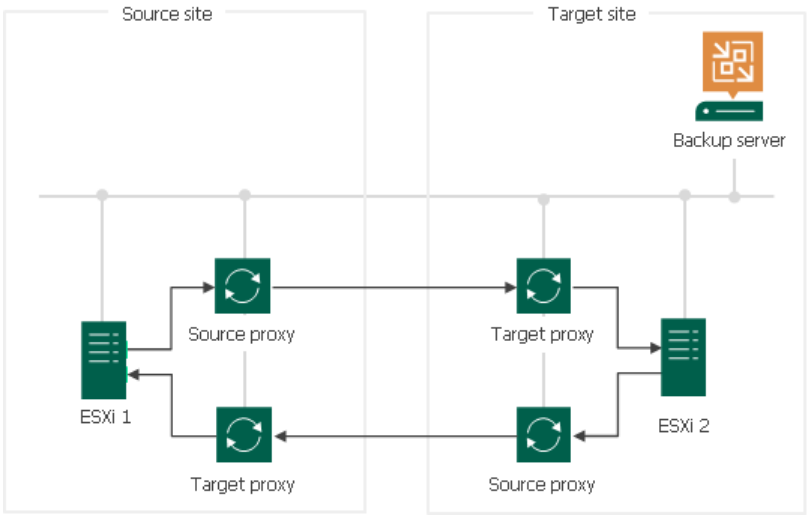
You can assign the role of a VMware CDP proxy to any Windows-managed virtual or physical server added to your Veeam Backup & Replication infrastructure. For information on how to add a server, see Adding Microsoft Windows Servers. For information on how to assign the VMware CDP proxy role, see Adding VMware CDP Proxies.

We recommend you to configure at least two VMware CDP proxies: one (source proxy) in the production site and one (target proxy) in the disaster recovery site. To optimize performance of several concurrent tasks, you can use several VMware CDP proxies in each site. In this case, Veeam Backup & Replication will distribute the restore workload between available proxies on per-task basis, taking into account proxy connectivity and their current load. For more information on which proxies are considered the most appropriate for continuous data protection, see How CDP Works.



Design your system in such a way that one VMware CDP proxy is used only as a source or only as target proxy.

For example, if you have cross cluster or cross host replication (from ESXi 1 to ESXi 2, and from ESXi 2 to ESXi 1), it is better to have four VMware CDP proxies: one source proxy and one target proxy for data flow from ESXi 1 to ESXi 2, and one source proxy and one target proxy for data flow from ESXi 2 to ESXi 1.



NOTE

If you deploy VMware CDP proxies on virtual machines, we recommend you to locate source proxies on the source host, target proxies on the target host.

VMware CDP Proxy Services and Components

VMware CDP proxies run light-weight services that take a few seconds to deploy. Deployment is fully automated. Veeam Backup & Replication installs the following components and services:

- **Veeam CDP Proxy Service** manages all CDP activities such as data aggregation, data compression and decompression, data transfer and other.
- **Veeam Installer Service** is an auxiliary service that is installed and started on any Windows server once it is added to the list of managed servers in the Veeam Backup & Replication console. This service analyzes the system, installs and upgrades necessary components and services depending on the role selected for the server.
- **Veeam Data Mover** handles traffic sent during failback.

VMware CDP Proxy Cache

By default, VMware CDP proxies store the received data into RAM. If a VMware CDP proxy gets overloaded and the memory runs out, the proxy starts storing data into the cache. Data is deleted from the cache or memory only after the proxy gets a notification that the target host has successfully saved data sent by the proxy.

If the cache gets full and there is no "free" VMware CDP proxy, Veeam Backup & Replication will stop the CDP process. CDP will be resumed only after a new or "free" proxy appears, or after the proxy cache gets free space.

From the technical point of view, the cache is a folder on a VMware CDP proxy. We recommend you to allocate at least 10 GB or 1 GB per each VM disk if you plan to protect more than 10 disks in one CDP policy. Also, the longer RPO you set in a CDP policy, the larger cache you may need.

Requirements for VMware CDP Proxy

Before you assign the role of a backup proxy, check the following requirements:

- For system requirements, see System Requirements for VMware CDP Proxy Server.
- A VMware CDP proxy must be a Windows-managed virtual or physical server.
- [For VMware CDP proxies deployed on physical servers] Fast network between hosts and VMware CDP proxies is required.

Adding VMware CDP Proxies

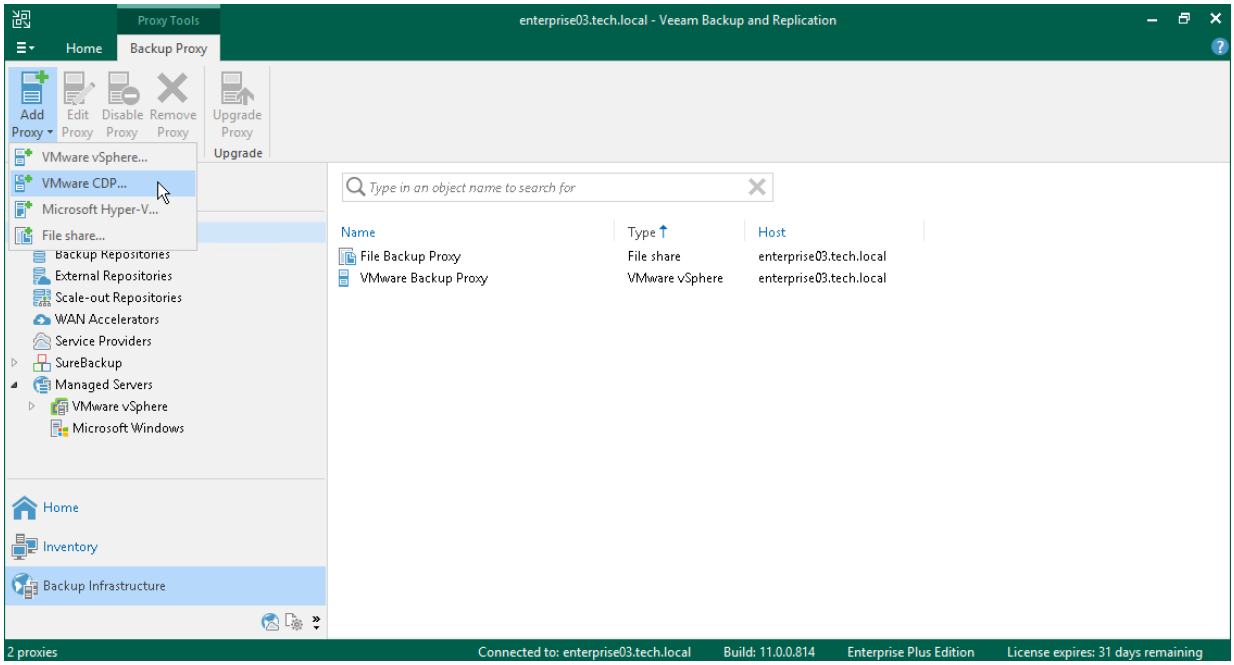
You must add to the backup infrastructure VMware CDP proxies that you plan to use for continuous data protection of VMs.

To add a VMware CDP proxy, use the **New VMware CDP Proxy** wizard.

Step 1. Launch New VMware Proxy Wizard

To launch the **New VMware CDP Proxy** wizard, do one of the following:

- Open the **Backup Infrastructure** view. In the inventory pane, right-click the **Backup Proxies** node and select **Add Veeam CDP Proxy**.
- Open the **Backup Infrastructure** view. In the inventory pane, select the **Backup Proxies** node. On the ribbon, click **Add Proxy** > **VMware CDP**.



Step 2. Select Server and Traffic Ports

At the **Server** step of the wizard, select a server which you want to use as the VMware CDP proxy, specify description and ports that will be used for communication:

1. From the **Choose server** drop-down list, select a physical or virtual Microsoft Windows server to which you want to assign the VMware CDP proxy role.

   If you have not added a server to the backup infrastructure, click **Add New** to open the **Add Server** wizard, and follow the instructions from the Adding Microsoft Windows Servers section.

2. In the **Proxy** description field, provide a description for future reference.

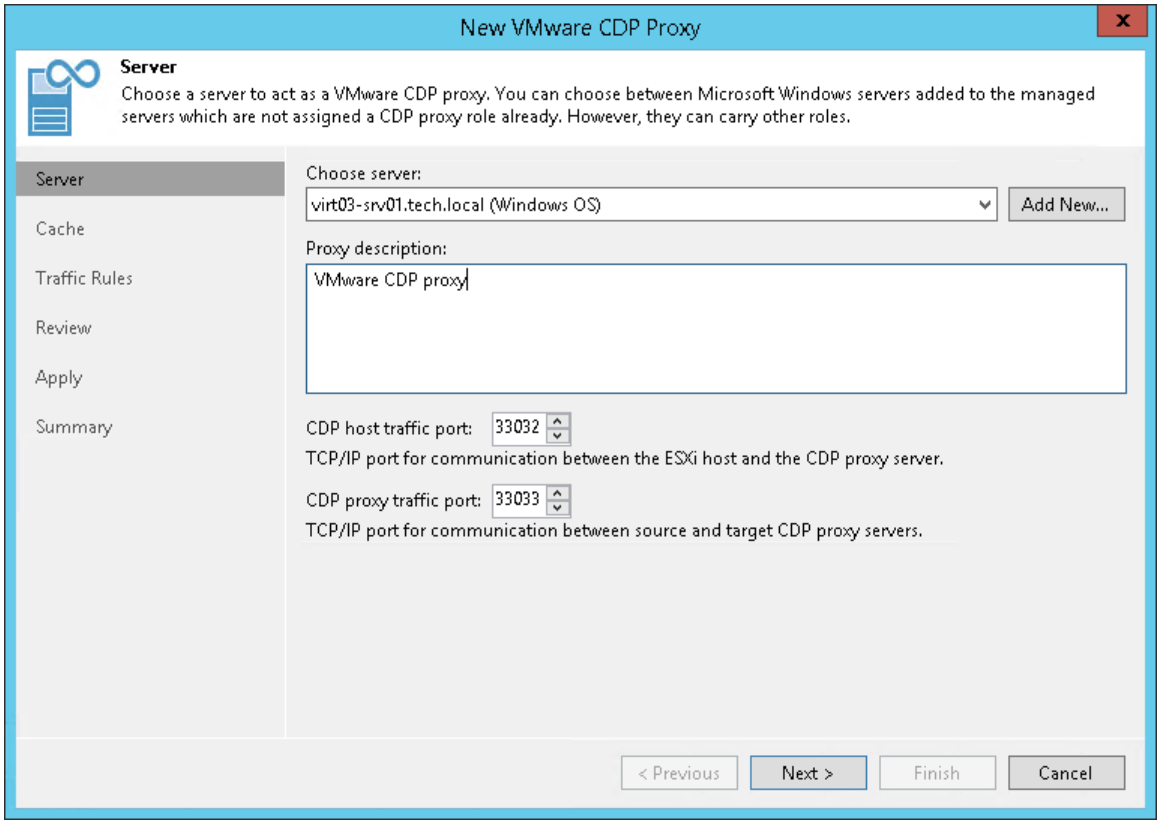
3. In the **CDP host traffic port** field, specify a port that the source and target host will use to communicate with the VMware CDP proxy.

   The range of available ports is 33032–33039.

4. In the **CDP proxy traffic port** field, specify a port that the proxy will use to communicate with other VMware CDP proxies.

   The range of available ports is 33032–33039.

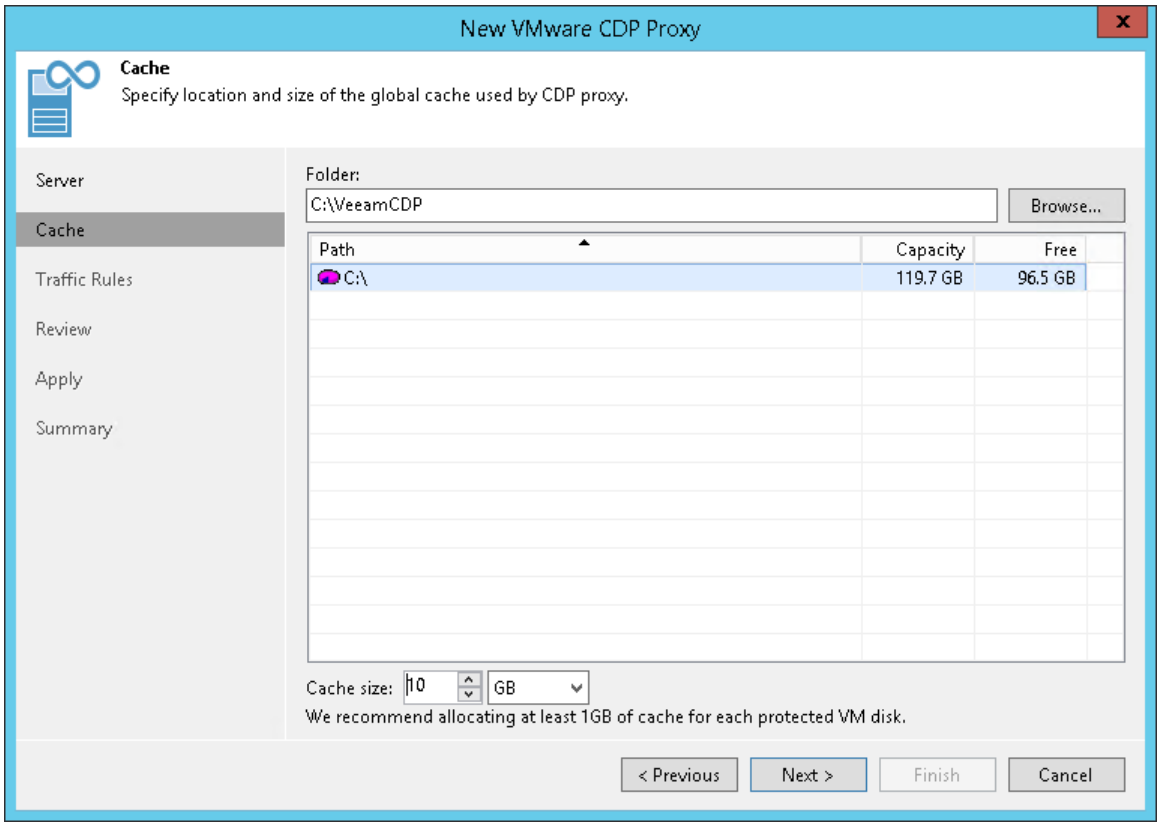
> **IMPORTANT**
>
> You must specify different values in the **CDP host traffic port** and **CDP proxy traffic port** fields.



Step 3. Configure Cache

At the **Cache** step of the wizard, specify the path to a folder where cached data will be stored and how much space can be used for storing the cache. We recommend you to allocate at least 10 GB or 1 GB per each replicated disk if you plan to replicate more than 10 disks in one CDP policy.

For more information on the cache, see VMware CDP Proxy Cache.



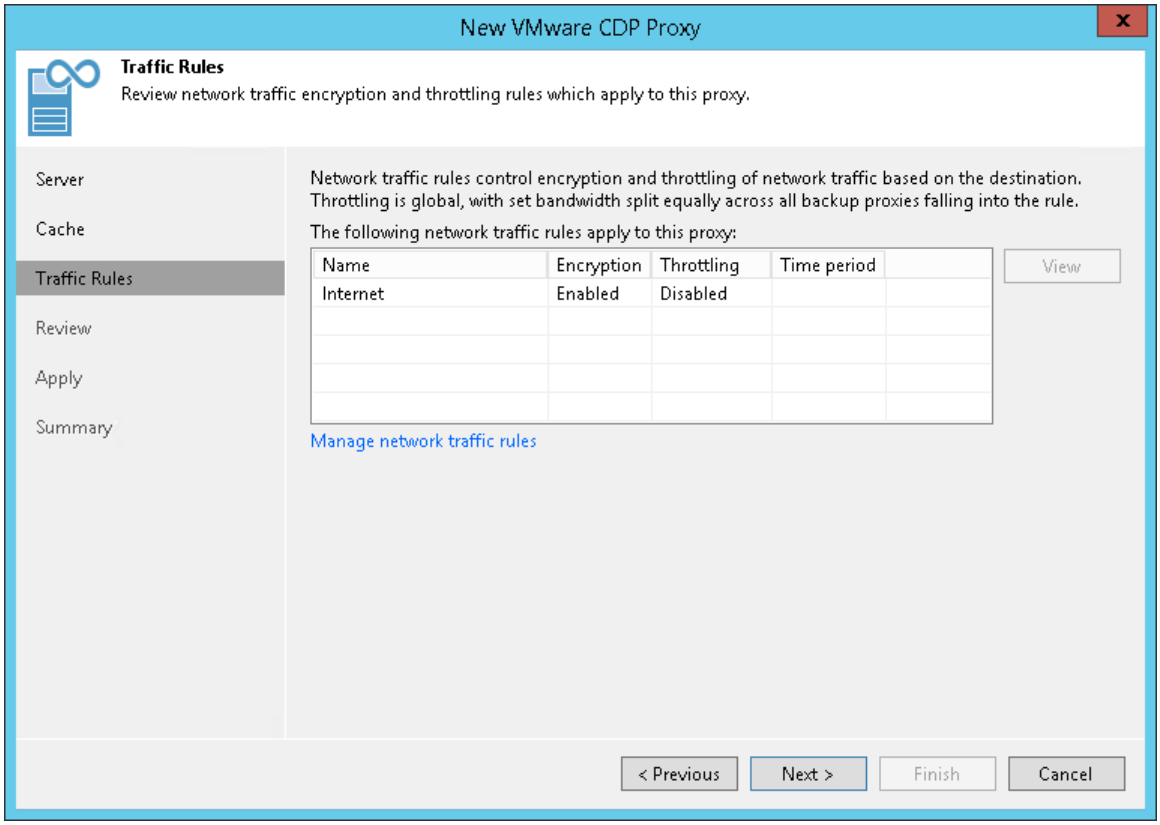
Step 4. Configure Network Traffic Rules

At the **Traffic Rules** step of the wizard, configure network traffic rules. These rules help you reduce, throttle and encrypt traffic sent between backup infrastructure components. For more information, see Network Traffic Management.

The list of network traffic rules contains only the rules that are applicable to the VMware CDP proxy — this means that the proxy IP address falls into the IP range of a rule.

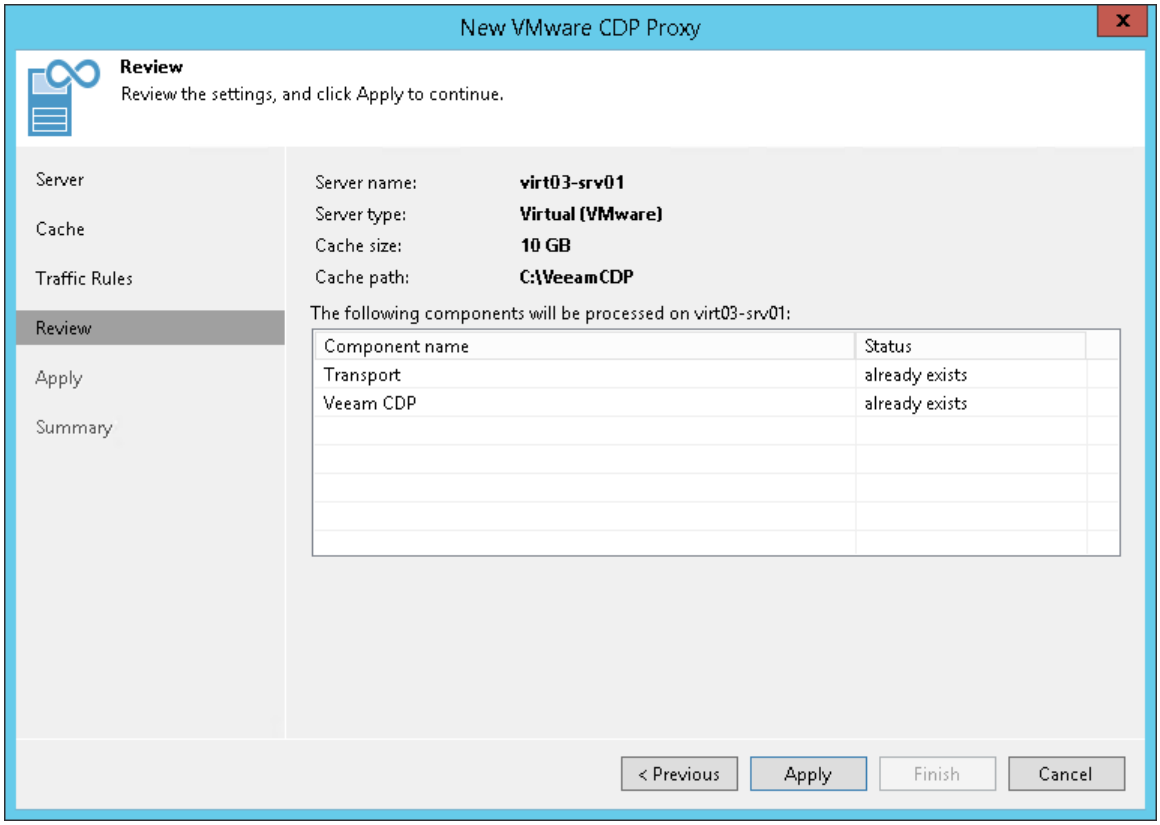
To change network traffic settings:

1. Click the **Manage network traffic rules** link.
2. To edit or remove a rule, select the required rule and click **Edit** or **Remove**.
3. To add a rule, click **Add** and follow the instructions from the Configuring Network Traffic Rules section.
4. To choose networks over which Veeam Backup & Replication will send data between backup infrastructure components, click **Networks**. Then follow the instructions in steps 3–6 in the Specifying Preferred Networks section.



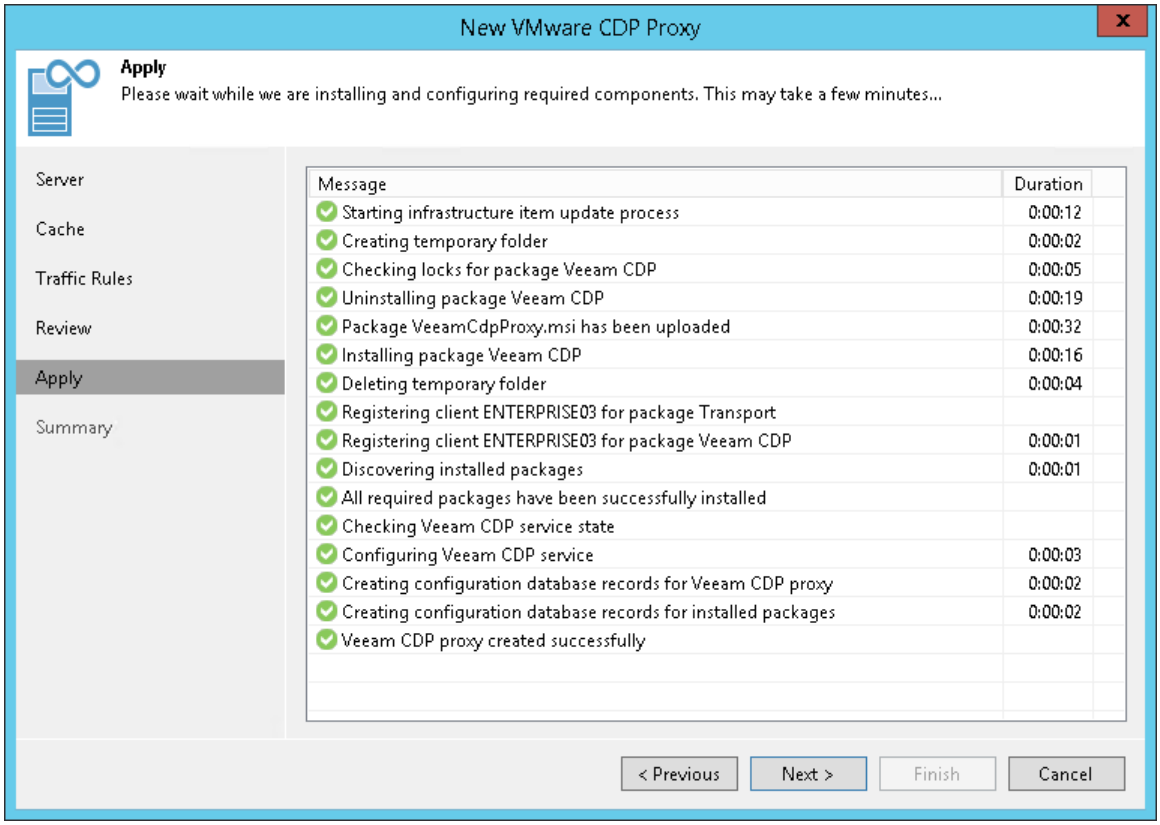
Step 5. Review Settings and Install Components

At the **Review** step of the wizard, review components which are already installed on the server and which will be installed. Click **Apply** to start installation of missing components.



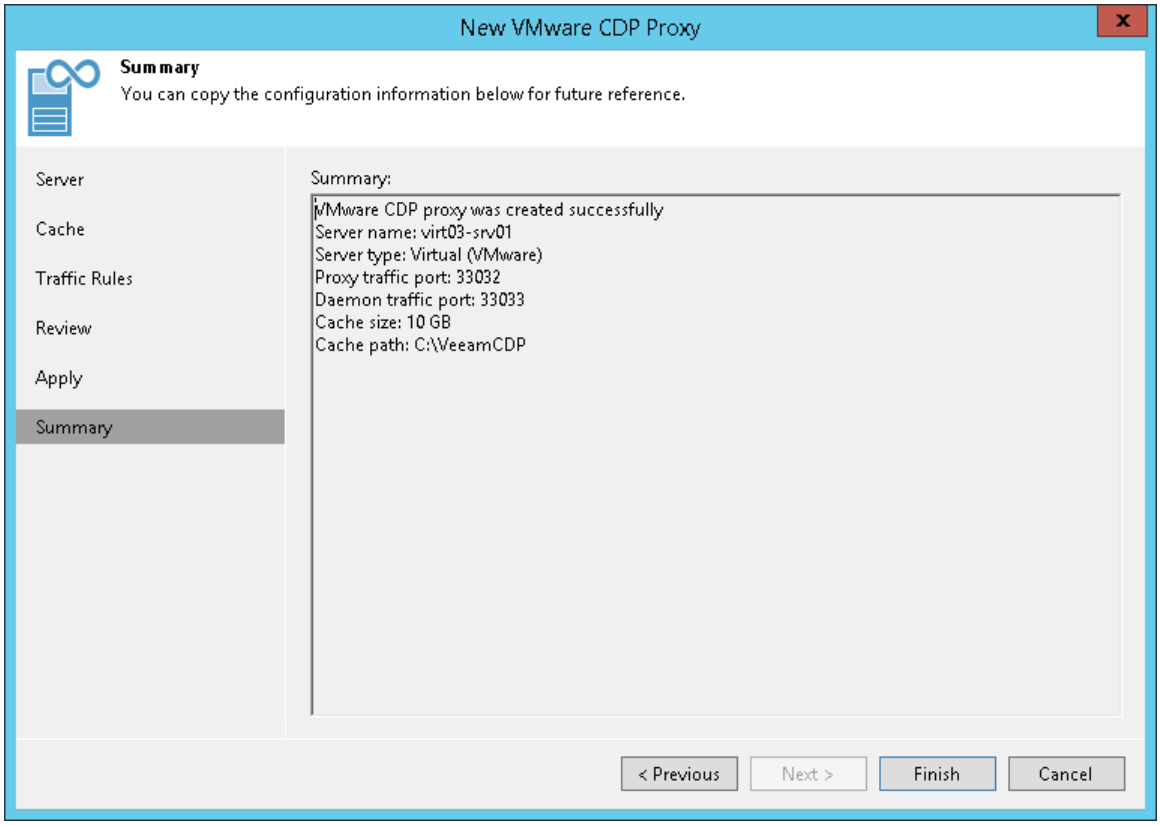
Step 6. Apply Proxy Settings

At the **Apply** step of the wizard, wait till Veeam Backup & Replication installs and configures all the required components. Click **Next** to complete the procedure of the VMware CDP proxy role assignment.



Step 7. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the server added as the VMware CDP proxy and click **Finish** to exit the wizard.



Cache Repository

A cache repository is a storage location where Veeam Backup & Replication keeps temporary cached metadata for the data backed up by the file share backup jobs. For more information about cache repository, see NAS Backup Support.

Backup Repository

A backup repository is a storage location where Veeam keeps backup files, VM copies and metadata for replicated VMs. To configure a backup repository, you can use the following storage types:

- **Direct attached storage**. You can add virtual and physical servers as backup repositories:
  - Microsoft Windows server
  - Linux server
- **Network attached storage**. You can add the following network shares as backup repositories:
  - SMB (CIFS) share
  - NFS share
- **Deduplicating storage appliances**. You can add the following deduplicating storage appliances as backup repositories:
  - Dell EMC Data Domain
  - ExaGrid
  - HPE StoreOnce
  - Quantum DXi
- **Object storage**. You can use cloud storage services as backup repositories. For details, see Object Storage Repository.

> **NOTE**
>
> Do not configure multiple backup repositories pointing to the same location or using the same path.

Related Topics

- Scale-Out Backup Repository
- External Repository

Microsoft Windows Server

You can use a Microsoft Windows server with local or directly attached storage as a backup repository. The storage can be a local disk, directly attached disk-based storage (such as a USB hard drive), or iSCSI/FC SAN LUN in case the server is connected into the SAN fabric.

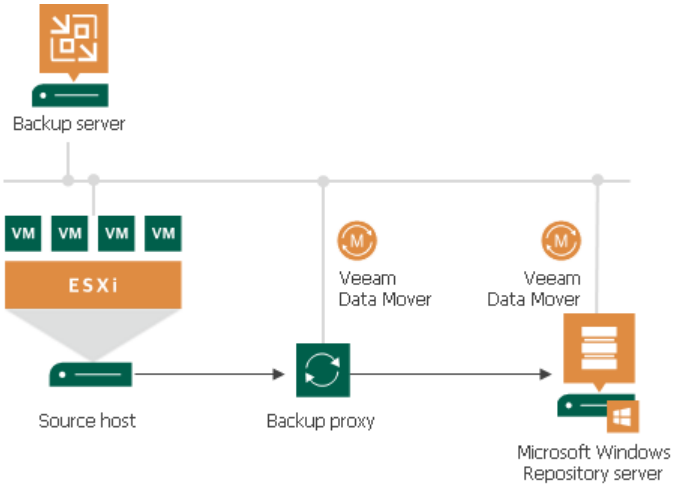
Microsoft Windows Repository Deployment

To communicate with a Microsoft Windows-based repository, Veeam Backup & Replication uses two Data Movers that are responsible for data processing and transfer:

- Veeam Data Mover on a backup proxy
- Veeam Data Mover on the Microsoft Windows repository

When any job addresses the backup repository, Veeam Data Mover on the backup proxy establishes a connection with Veeam Data Mover on the backup repository, enabling efficient data transfer over LAN or WAN.

The Data Mover is installed automatically when you add a server to Veeam Backup & Replication as a managed server.



vPower NFS Server

Windows repositories can be configured to function as vPower NFS Servers. In this case, Veeam Backup & Replication will run the Veeam vPower NFS Service directly in the backup repository (namely, on the managing Windows server to which storage is attached) and provide ESXi hosts with transparent access to backed-up VM images stored on the backup repository. For more information, see Veeam vPower NFS Service.

Requirements for Microsoft Windows Server Based Repositories

A machine performing the role of a repository must meet the following requirements:

- The machine must meet the system requirements. For more information, see System Requirements.
- The role of the repository can be assigned to a Microsoft Windows machine (physical or virtual).
- You must add the machine to the Veeam Backup & Replication console as a managed server.
- If you want to use Fast Clone in the Microsoft Windows backup repository, the machine must also meet requirements listed in the Fast Clone section.

Linux Server

You can add a Linux server with local, directly attached storage or mounted NFS as a backup repository. The storage can be a local disk, directly attached disk-based storage (such as a USB hard drive), NFS share, or iSCSI/FC SAN LUN in case the server is connected into the SAN fabric.

A Linux repository supports the immutability feature that allows you to make your backup files immutable for the specified time period. For more information, see Hardened Repository.

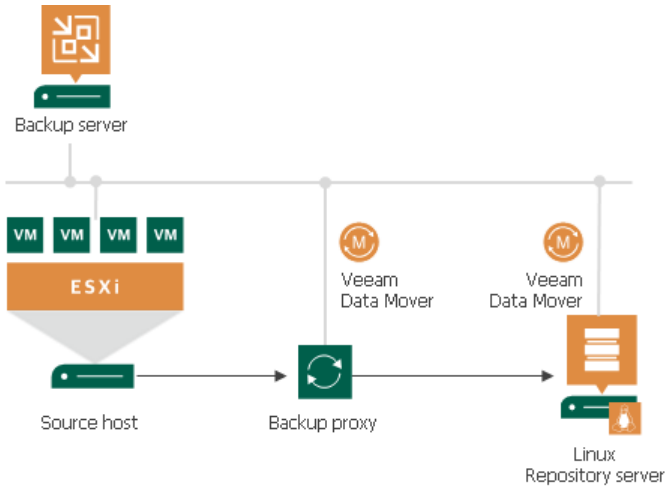
Linux Backup Repository Deployment

To communicate with a Linux-based repository, Veeam Backup & Replication uses two Veeam Data Movers that are responsible for data processing and transfer:

- Veeam Data Mover on the backup proxy
- Veeam Data Mover on the Linux backup repository

For more information about Veeam Data Movers communication with a Linux-based server, see Veeam Data Movers.

Veeam Data Mover establishes a connection with the source-side Data Mover on the backup proxy, enabling efficient data transfer over LAN or WAN.



Requirements for Linux Backup Repositories

A machine performing the role of a repository must meet the following requirements:

- System Requirements.
- The role of the repository can be assigned to a Linux machine (physical or virtual).
- You must add the machine to the Veeam Backup & Replication console as a managed server.
- If Veeam Data Mover is non-persistent, Veeam Backup & Replication uses the SSH protocol to communicate with Linux backup repositories and requires the SCP utility in Linux repositories. Make sure that the SSH daemon is properly configured and SCP utility is available on the Linux host.
- If you want to use Fast Clone in the Linux backup repository, the machine must also meet requirements listed in the Fast Clone section.

For hardened repository limitations, see Limitations and Considerations.

Hardened Repository

Veeam Backup & Replication allows you to make a hardened repository using a Linux server as a storage. A hardened repository protects your backup files from loss as a result of malware activity or unplanned actions with the help of the following features:

- **Single-use credentials**: credentials that are used only once to deploy Veeam Data Mover while adding the Linux server to the backup infrastructure. These credentials are not stored in the backup infrastructure, so that backups files will be safe even if the Veeam Backup & Replication server is compromised.
- **Immutability**: when you add a Linux repository, you can select the **Make recent backups immutable for** check box and specify the time period while backup files must be immutable. During this period, backup files stored in this repository cannot be modified or deleted.

How Immutability Works

After you add a hardened repository with immutability:

1. Veeam Backup & Replication creates a `.veeam.N.lock` file with the information about immutability time period of each backup file in the active chain. The `.veeam.N.lock` files are stored on a Linux host.
2. Backup files become immutable for the configured time period (minimum 7 days, maximum — 9999). The immutability period is extended only for the active backup chain. If there are several chains in the backup, then Veeam Backup & Replication does not extend the immutability for old backups in the chain.
3. After the time period expiration, Veeam Backup & Replication makes backup files non-immutable again so they can be deleted or modified.

The count of the immutability period indicated in the backup repository settings starts from the moment the last restore point in the active chain is created. For example:

- The full backup file of the active backup chain was created on January 12. The first increment was created on January 13. The second and last increment was created on January 14.
- The immutability period indicated at the backup repository settings is 10 days.
- The backup files will be immutable until January 24: the date of the last restore point creation (January 14) + 10 days.

Supported Job Types

The hardened repository supports backups created with the following types of jobs:

- VMware, Hyper-V VM backup jobs and backup copy jobs created by Veeam Backup & Replication
- Backup copy jobs for backups created by Veeam Backup for Azure, Veeam Backup for AWS and Veeam Backup for Google Cloud Platform
- Physical machines backup jobs created by Veeam Agents (Windows, Linux, MAC, AIX, Solaris)
- vCD VM backup jobs
- VeeamZIP backup jobs
- Nutanix AHV VM backup jobs created by Veeam Backup for Nutanix AHV
- RHV VM backup jobs created by Veeam Backup for Red Hat Virtualization

IMPORTANT

You can store backup files and backup copy files of NAS backup jobs, transaction log backup jobs, RMAN/SAP HANA/SAP on Oracle backups jobs in a hardened repository with immutability, but these files will not be immutable.

In This Section

- Limitations and Considerations
- Tips for Enhanced Security of Hardened Repository
- Deploying Hardened Repository

Limitations and Considerations

All requirements and limitations for a Linux backup repository apply to a hardened repository. In addition, mind the following limitations and considerations.

Hosts and File Systems Requirements

- Due to Veeam Data Mover requirements, the Linux host version must be 64-bit.
- Linux host file system must support extended attributes modified by the `chattr` and `setfattr` commands. For more information, see these Linux articles: lsattr, xattr.

  Many common Linux file systems support these attributes (XFS, EXT3 / 4, RFS). We recommend XFS for the hardened repository for performance and space efficiency reasons (block cloning support).
- ReiserFS file system does not support immutable files. You cannot select a Linux server running on ReiserFS as a hardened repository with immutability.
- To support Linux server with the NIST 800-171 security profile, make sure the requirements listed in this Veeam KB article are met.

Limitations for Repositories

- The hardened repository cannot be shared between different Veeam Backup & Replication servers.
- If you want to store backup files in a repository with immutability, you cannot select a reverse or a forever forward incremental backup mode. Once a backup file becomes immutable, it can be merged or deleted only when the immutability time period expires. For this reason, you must enable active full backup or synthetic full backup in the backup job settings.

  The requirement for periodic fulls also means that XFS-based repository is more preferable from a disk space usage perspective because thanks to fast cloning, synthetic full backups take no physical disk space (except for metadata).
- You cannot delete immutable backup files manually from a hardened repository.
- When configuring a repository with immutability, you can use either persistent or single-use credentials. For more information, see Specify Credentials and SSH Settings. But the corresponding Linux host should not be used twice in the database:
  - If you use single-use credentials, the host where the repository resides cannot have any other role: you cannot add it as a proxy or as a file server.

- If you use persistent credentials, the host where the repository resides cannot have the proxy role, and the file server role is not recommended.

- NFS does not accept immutability commands from Linux. Due to this, mind the following:
  - You cannot use an NFS Share as a repository with the immutability.
  - You cannot use an NFS volume mounted on a Linux server as a hardened repository with immutability (needs to be Local or DAS). You can use an NFS volume mounted on a Linux server as a hardened repository without immutability.
- You can place both repositories (hardened and standard) on one Linux server only if you used single-use credentials when adding the host. Standard repository is a repository added with persistent credentials and disabled immutability.
- For a scale-out backup repository:
  - You can add a hardened repository to your scale-out backup repository as a performance extent. For more information, see Scale-Out Backup Repository and Add Performance Extents.
  - If you use the Capacity Tier option, keep in mind that having a hardened repository with immutability as a performance extent will affect the Capacity Tier behavior. You will not be able to move immutable backup files, because they cannot be deleted from the performance extent.
    Veeam Backup & Replication will copy such backup files to the Capacity Tier. When the immutability time period is over, Veeam Backup & Replication will delete these files from the performance extent. For more information on copy and move policies, see Copying Backups to Capacity Tier and Moving Backups to Capacity Tier.
  - If you evacuate your backups from an immutable performance extent, Veeam Backup & Replication will copy them instead of moving. If the target extent is also immutable, then the immutability of the target extent will apply to copied backup files. For more information on evacuating backups, see Evacuating Backups from Performance Extents.
  - We recommend to avoid mixing mutable and immutable extents within one scale-out backup repository. You can mix them only during migration scenarios when you want to make a hardened repository from an existing Linux extent.

Limitations and Considerations for Immutability Feature

- An immutability retention overrides a job retention: if the job retention period is shorter than the immutability period, Veeam Backup & Replication does not delete backup files when the retention period is over, but only when the immutability period expires.
- For importing a backup, we recommend to use VBK backup files. Metadata files of a backup chain (.VBM) cannot be immutable because they are updated on every job pass.
- For backup copy jobs, enable the GFS retention policy. Otherwise, you will not be able to use the immutability feature. For more information, see Long-Term Retention Policy (GFS).
- If a hardened repository with immutability is a part of a scale-out backup repository with the capacity tier added, the immutability time period for full backup files with GFS retention policy is set according to the following:
  - [For capacity tier with disabled move policy] Veeam Backup & Replication compares the immutability period of the backup repository and the GFS backup file lifetime, and sets an immutability period for full backup files with GFS retention policy as equal to the longest of these periods.

    For example: the backup repository immutability period is 10 days; the GFS backup file lifetime is 3 years; the backup file will be immutable for 3 years; the increments from this full backup file will be immutable for 10 days from the moment of the last increment creation.

- [For capacity tier with enabled move policy] Veeam Backup & Replication ignores the GFS retention policy. The immutability time period for full backup files equals the period specified in the setting of a hardened repository.

- If a hardened repository with immutability is a part of a scale-out backup repository (with the capacity tier added and enabled move policy) and is used as a target for VeeamZIP jobs — the immutability time period for backup files equals the period specified in the setting of a hardened repository.
- The immutability feature is supported for image-level backups only. You can use a hardened repository to store NAS backups, transaction backups, RMAN/SAP HANA/SAP on Oracle backups, but you cannot use the immutability feature to protect these backups.
- We do not recommend to use the immutability feature for a Nutanix Mine infrastructure. As Mine repositories contain thin-provisioned disks, there may be the case when Veeam Backup & Replication uses full storage capacity of a repository and is not able to delete backup files from the file system.
- The immutability time period for backup files produced with VeeamZIP or Export Backup jobs is set according to the following:
  - [With enabled retention period] Veeam Backup & Replication compares the immutability period of the backup repository and the retention period lifetime, and sets an immutability period for backup files with retention period as equal to the longest of these periods.

    For example: the backup repository immutability period is 1 month; the VeeamZIP or Export Backup backup file lifetime is 7 years; the backup file will be immutable for 7 years.
  - [With disabled retention period] Veeam Backup & Replication ignores the VeeamZIP or Export Backup retention period. The immutability time period for backup files equals the period specified in the setting of a hardened repository.

Hardened Repository Availability

Veeam Backup & Replication does not store singe-use credentials in the configuration database. Thus, if you backup Veeam Backup & Replication configuration database and restore settings of Veeam Backup & Replication, a hardened repository becomes unavailable in the Veeam Backup & Replication console. To make a repository available again, do the following after restore:

1. In the Veeam Backup & Replication console, open the **Backup Infrastructure** view.
2. In the navigation pane, select **Managed Servers.**
3. In the working area, right-click the Linux server used as a hardened repository and select **Properties**.
4. At the Step 3. Specify Credentials and SSH Settings of the **Edit Backup Server** wizard, use single-use credentials and click **Finish** to update settings.

Related Topics

- Tips for Enhanced Hardened Repository Security

Tips for Enhanced Security of Hardened Repository

We recommend to follow the tips below to maximize the repository security and protect your data from different attacks:

- Change file permissions for authentication certificates on the Linux server, and unwelcome non-root users cannot connect to Veeam Data Mover.

  Use the following commands:

  - To create the folder:

    ```
    mkdir -p /opt/veeam/transport/certs
    ```

  - To change the folder's owner:

    ```
    chown owner:group /opt/veeam/transport/certs
    ```

    Both `owner` and `group` can be the account that you plan to use to connect to the Linux server.

  - To allow access to the folder for root account and account for Veeam Data Mover functioning:

    ```
    chmod 700 /opt/veeam/transport/certs
    ```

    You can also use `chmod 770` to add same permissions to the group.

  Keep in mind that a hardened repository requires persistent Veeam Data Movers. For Veeam Data Mover to be persistent, you must specify an account with equivalent to root permissions when adding the Linux server to the backup infrastructure.

  > **IMPORTANT**
  > During the deployment of Veeam Data Mover, the account that you plan to use to connect to the Linux server requires read and write permissions on the folder where authentication certificates are stored. Make sure the `umask` command value is no more restrictive than `022`. For more information, see this Veeam KB article.

  For security purposes, the rights of Veeam Data Mover are reduced: SSH connection is necessary only for a deployment of Veeam Data Mover to the Linux server. After Veeam Data Mover is deployed, you can disable SSH, so that backup infrastructure components use server and client certificates for authentication.

- Deploy Veeam Backup Enterprise Manager on a server different from the Veeam Backup & Replication server to prevent a key change attack. Even if passwords are lost due to unauthorized access, you can restore lost data with the help of Enterprise Manager. For more information, see Decrypting Data Without Password.

Deploying Hardened Repository

Deploying a hardened repository is similar to setting up a Linux backup repository but the hardened repository is an exclusive role on a server with single-use credentials and with enabled **Make recent backups immutable for** check box.

> **TIP**
>
> We recommend to build a hardened repository with both single-use credentials and immutability features to maximize data security. You can still add a repository with single-use credentials but without immutability.

Before you deploy a hardened repository, check limitations and considerations and security recommendations.

If you want to deploy a hardened repository, perform the following steps:

1. Prepare the directory on the Linux server for backups.
2. Add the Linux server to the backup infrastructure.
3. Add the backup repository role to the Linux server and enable the immutability feature.

If you want to use the Linux repository existing before Veeam Backup & Replication 11a, you can upgrade it to the hardened repository.

Step 1. Prepare Directory on Linux Server for Backups

Create a separate folder where immutable backups will be stored. Allow access to this folder only for the account that you plan to use to connect to the Linux server. Use the following commands:

- To create the folder:

  ```
  mkdir <folder_path>
  ```

  where `<folder_path>` — path to the folder you are creating.

- To assign the folder's owner:

  ```
  chown -R owner:group <folder_path>
  ```

  Both `owner` and `group` can be the account that you plan to use to connect to the Linux server.

- To allow access to the folder only for its owner and root account:

  ```
  chmod 700 <folder_path>
  ```

Step 2. Add Linux Server to Backup Infrastructure

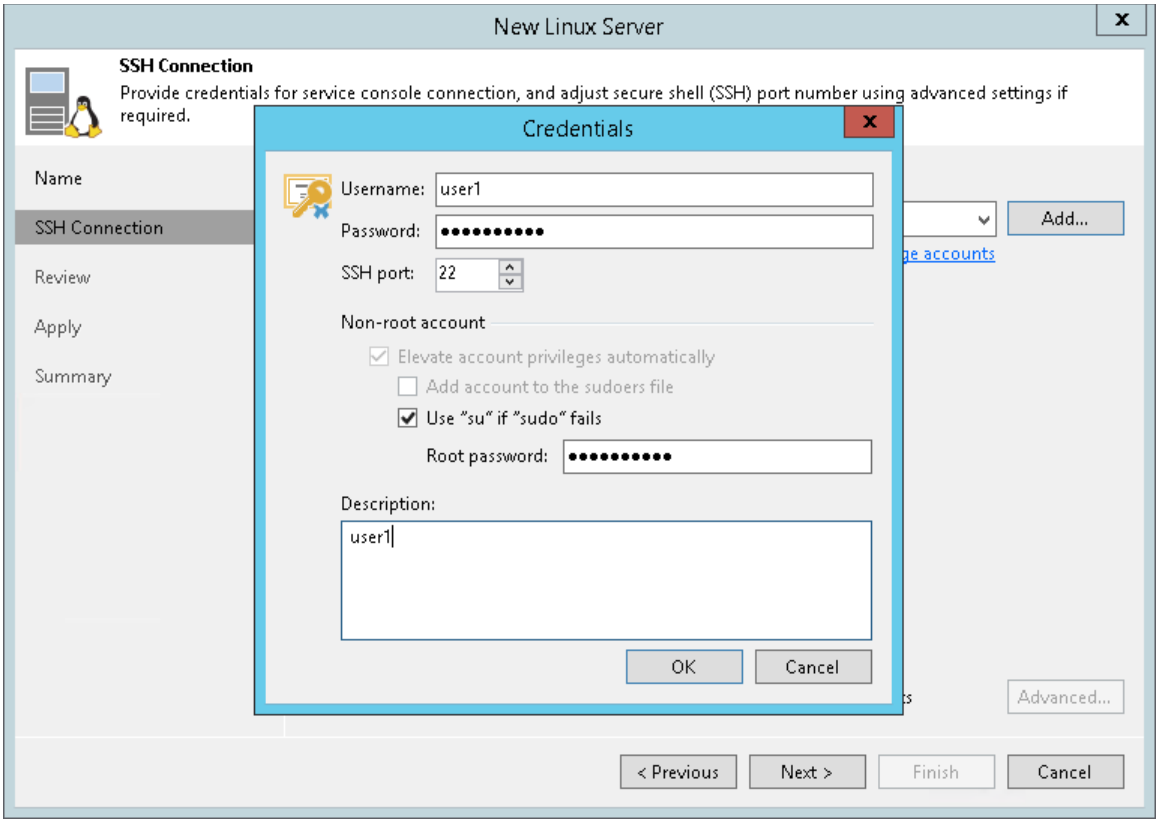
Alternatively, you can use a Linux server that is already added to a backup infrastructure. To add new server, use the **New Linux Server** wizard. For more information, see Adding Linux Servers.

Pay attention to the following settings at the Step 3. Specify Credentials and SSH Settings:

- Use temporary credentials to avoid storing the credentials in the Veeam Backup & Replication configuration database. To do that, click **Add** and select **Single-use credentials for hardened repository**.

- In the **Credentials** window, within the user account that you plan to use to connect to the Linux server select the **Use "su" if "sudo" fails** check box. The **Elevate account privileges automatically** check box is used by default. Both selected check boxes mean if the user is not in the `sudoers` file you can use `su` command instead of `sudo`.

  After the user will have temporary root- or sudo-permissions you can remove the user from the sudo group after the server is added. Further, you can use the immutability feature with an existing repository if you have enough rights to use this repository as a user without root credentials. For more information on these check boxes, see Linux Accounts (User Name and Password).



Step 3. Add Backup Repository Role to Linux Server and Enable Immutability Feature

Use the **New Backup Repository** wizard to add new backup repository. For more information, see Adding Backup Repositories. Pay attention to the following steps:

- In the **Add Backup Repository** window, select the **Direct Attached Storage > Linux** type of the backup repository.
- At the Step 4. Configure Backup Repository Settings, select the **Make recent backups immutable for** check box and specify the immutability time period.

After you added the host (for single-use credentials) or the repository (for persistent credentials), disable SSH connection for the account that you plan to use to connect to the Linux server. If you can work with the server from the console, disable SSH connection for the server itself.

> **IMPORTANT**
>
> Once a backup file becomes immutable, it can be merged or deleted only when the immutability time period expires. For this reason, if you want to store backup files in a hardened repository with immutability, you must enable active full backup or synthetic full backup in the backup job settings.

Upgrading Existing Linux Repository to Hardened Repository

To upgrade the Linux repository existing before Veeam Backup & Replication 11a to the hardened repository, perform the following steps:

1. Change access to the folder where immutable backups are stored. Allow access to this folder for the account that you plan to use to connect to the Linux server. Use the following command:

   ```
   chown -R username:groupname <folder_path>
   ```

   where `<folder_path>` — path to the folder.

2. Edit server settings and use **Single-use credentials for hardened repository** at the Step 3. Specify Credentials and SSH Settings.

3. Editing settings of the backup repositories and select the **Make recent backups immutable for** check box and specify the immutability time period at the Step 4. Configure Backup Repository Settings.

SMB (CIFS) Share

You can use SMB (CIFS) shares as backup repositories.

SMB Backup Repository Deployment

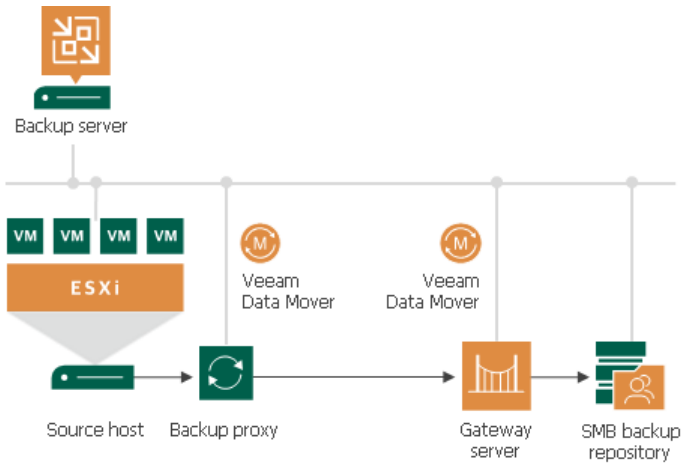
To communicate with an SMB backup repository, Veeam Backup & Replication uses two Veeam Data Movers that are responsible for data processing and transfer:

- Veeam Data Mover on the backup proxy
- Veeam Data Mover on the gateway server

An SMB share cannot host Veeam Data Movers. For this reason, to communicate with the SMB share, you need to deploy a gateway server. Veeam Backup & Replication will automatically deploy a Veeam Data Mover on this gateway server. For more information, see Gateway Server.

When any job addresses the SMB backup repository, Veeam Data Mover on the gateway server establishes a connection with Veeam Data Mover on the backup proxy, enabling efficient data transfer over LAN or WAN.

If you plan to move VM data to an off-site SMB repository over a WAN link, it is recommended that you deploy an additional gateway server in the remote site, closer to the SMB repository.



Requirements for SMB Backup Repositories

A machine performing the role of an SMB repository must meet the following requirements:

- The machine must meet the system requirements. For more information, see System Requirements.
- The role of the repository can be assigned to a Microsoft Windows machine (physical or virtual).

NFS Share

You can use NFS shares as backup repositories.

NFS Backup Repository Deployment

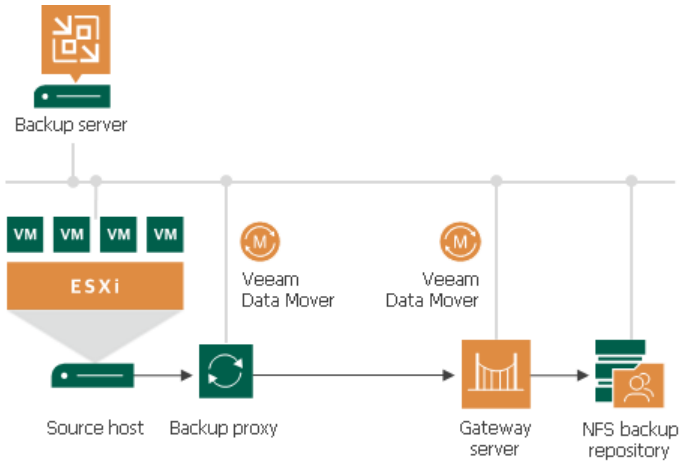
To communicate with an NFS backup repository, Veeam Backup & Replication uses two Veeam Data Movers that are responsible for data processing and transfer:

- Veeam Data Mover on the backup proxy
- Veeam Data Mover on the gateway server

An NFS share cannot host Veeam Data Movers. For this reason, to communicate with the NFS share, you need to deploy a gateway server. Veeam Backup & Replication will automatically deploy a Veeam Data Mover on this gateway server. For more information, see Gateway Server.

When any job addresses the NFS backup repository, Veeam Data Mover on the gateway server establishes a connection with Veeam Data Mover on the backup proxy, enabling efficient data transfer over LAN or WAN.

If you plan to move VM data to an off-site NFS repository over a WAN link, it is recommended that you deploy an additional gateway server in the remote site, closer to the NFS repository.



Requirements and Limitations for NFS Backup Repositories

A machine performing the role of an NFS repository must meet the following requirements:

- The machine must meet the system requirements. For more information, see System Requirements.
- The role of the NFS repository can be assigned to a Microsoft Windows or Linux machine (physical or virtual) or to NAS storage supporting NFS protocol.
- The NFS repository must present read and write access rights to the gateway.

Mind that Veeam Backup & Replication does not support multipathing for NFS repository.

Requirements for Gateway Server

A machine performing the role of a gateway server for communication with the NFS backup repository must meet the following requirements:

- The machine must meet the system requirements. For more information, see System Requirements.
- The role of the gateway server can be assigned to a Microsoft Windows or Linux machine (physical or virtual).
- You must add the machine to the Veeam Backup & Replication console as a managed server.
- [For automatic gateway selection] The backup server must have read and write access in the NFS repository.
- [For automatic gateway selection] If you configure automatic gateway selection for NFS repository, Veeam Backup & Replication may use the same machines as gateways for the repository and as proxies for backup jobs. Make sure that the backup proxies meet the following requirements:
  - If you explicitly choose backup proxies for backup jobs, provide read and write access rights to all proxies chosen for backup jobs that are targeted to the NFS repository.
  - If you configure automatic proxy selection for backup jobs, provide read and write access rights to all proxies in the Veeam Backup & Replication infrastructure.
  - If backup jobs that are targeted to the NFS repository use Linux proxies, check that the NFS client package is installed on the Linux proxy server.

Requirements and Limitations for Linux Gateway Server

In addition to the general requirements, the Linux gateway server must meet the following requirements:

- The Linux gateway server must have NFS client package installed.
- The credentials to authenticate with the Linux gateway server must have root or elevated to root permission.
- Veeam Backup & Replication uses the highest NFS protocol version supported by the gateway and the repository.

  Note that the suffix indicating the NFS version in the NFS share properties may not be displayed correctly, this is a known issue.

> **TIP**
>
> If the NFS protocol version has changed (for example, if you updated the NFS share), click through the **Edit Backup Repository** wizard to update the information in Veeam Backup & Replication.

Deduplicating Storage Appliances

You can add deduplicating storage appliances as backup repositories.

In This Section

- Dell EMC Data Domain

- ExaGrid
- HPE StoreOnce
- Quantum DXi

Dell EMC Data Domain

You can use Dell EMC Data Domain storage systems with Data Domain Boost (DD Boost) as backup repositories.

To support the DD Boost technology, Veeam Backup & Replication leverages the following Dell EMC Data Domain components:

- **DD Boost library**. The DD Boost library is a component of the Dell EMC Data Domain system. The DD Boost library is embedded into the Veeam Data Mover setup. When you add a Microsoft Windows server to the backup infrastructure, the DD Boost Library is automatically installed on the added server together with Veeam Data Mover.
- **DD Boost server**. The DD Boost server is a target-side component. The DD Boost server runs on the OS of the Dell EMC Data Domain storage system.

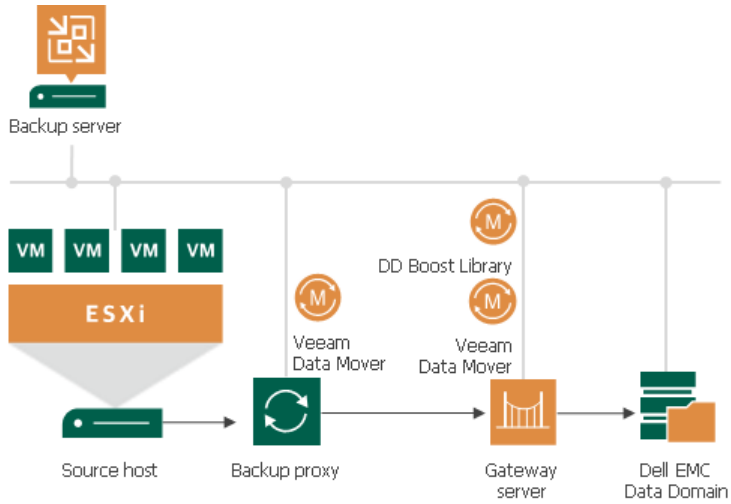
Dell EMC Data Domain Deployment

To communicate with Dell EMC Data Domain, Veeam Backup & Replication uses two Veeam Data Movers that are responsible for data processing and transfer:

- Veeam Data Mover on the backup proxy
- Veeam Data Mover on the gateway server

The Dell EMC Data Domain storage cannot host Veeam Data Mover. For this reason, to communicate with the Dell EMC Data Domain storage, you need to deploy a gateway server. Veeam Backup & Replication will automatically deploy Veeam Data Mover on this gateway server. For more information, see Gateway Server.

When any job addresses the backup repository, Veeam Data Mover on the gateway server establishes a connection with Veeam Data Mover on the backup proxy, enabling efficient data transfer over LAN or WAN.



You define what gateway server to use when you assign a backup repository role to Dell EMC Data Domain. You can define the gateway server explicitly or instruct Veeam Backup & Replication to select it automatically.

IMPORTANT

For Dell EMC Data Domain storage systems working over Fibre Channel, you must explicitly define the gateway server that will communicate with Dell EMC Data Domain. As a gateway server, you must use a Microsoft Windows server that is added to the backup infrastructure and has access to Dell EMC Data Domain over Fibre Channel.

Supported Protocols

Veeam Backup & Replication supports Dell EMC Data Domain storage systems working over the following protocols:

- TCP/IP protocol: Veeam Backup & Replication communicates with the Dell EMC Data Domain server by sending commands over the network.
- Fibre Channel protocol: Veeam Backup & Replication communicates with the Dell EMC Data Domain Fibre Channel server by sending SCSI commands over Fibre Channel.

Limitations for Dell EMC Data Domain

If you plan to use Dell EMC Data Domain as a backup repository, mind the following limitations and recommendations:

- Use of Dell EMC Data Domain with DD Boost does not guarantee improvement of job performance. It reduces the load on the network and improves the network throughput.
- NFS services must be enabled on Dell EMC Data Domain. Otherwise, Veeam Backup & Replication will not be able to access the storage system.
- Dell EMC Data Domain does not support the reverse incremental backup method.
- We recommend against enabling encryption for the jobs targeted at the deduplication storage appliance. Encryption has a negative effect on the deduplication ratio. For more information, see Data Encryption.
- You cannot use Dell EMC Data Domain backup repositories as sources or targets for file copy jobs.
- When you create a backup job targeted at an Dell EMC Data Domain backup repository, Veeam Backup & Replication will offer you to switch to optimized job settings and use the 4 MB size of data block for VM data processing. It is recommended that you use optimized job settings. Large data blocks produce a smaller metadata table that requires less memory and CPU resources to process.
- The length of forward incremental and forever forward incremental backup chains (chains that contain one full backup and a set of subsequent incremental backups) cannot be greater than 60 restore points. To overcome this limitation, schedule full backups (active or synthetic) to split the backup chain into shorter series. For example, to perform backups at 30-minute intervals 24 hours a day, you must schedule synthetic fulls every day. In this scenario, intervals immediately after midnight may be skipped due to duration of synthetic processing. For more information, see How Synthetic Full Backup Works.
- If you connect to an Dell EMC Data Domain backup repository over Fibre Channel, you must explicitly define a gateway server to communicate with Dell EMC Data Domain. As a gateway server, you must use a Microsoft Windows server that is added to the backup infrastructure and has access to the Dell EMC Data Domain backup repository over Fibre Channel.
- During backup repository rescan, Veeam Backup & Replication detects if the hard stream limit is set for a storage unit, and displays this information in backup repository rescan statistics. If the hard stream limit is exceeded when Veeam Backup & Replication runs tasks against the backup repository, Veeam Backup & Replication will fail to create new I/O streams.

For more information and recommendations on working with Dell EMC Data Domain, see this Veeam KB article.

Dell EMC Data Domain Supported Features

The DD Boost technology offers a set of features for advanced data processing. Veeam Backup & Replication supports the following features:

- Distributed Segment Processing
- Advanced Load Balancing and Link Failover
- Virtual Synthetics

In addition to these technologies, Veeam Backup & Replication supports in-flight data encryption and per storage unit streams.

> **NOTE**
>
> You cannot configure Managed File Replication using Veeam Backup & Replication. However, you can import and map backups replicated between Data Domain storage systems to backup, backup copy or replication jobs, or perform restore operations from such backups.

Distributed Segment Processing

Distributed Segment Processing lets Dell EMC Data Domain 'distribute' the deduplication process and perform a part of data deduplication operations on the backup proxy side.

Without Distributed Segment Processing, Dell EMC Data Domain performs deduplication on the Dell EMC Data Domain storage system. The backup proxy sends unfiltered data blocks to Dell EMC Data Domain over the network. Data segmentation, filtering and compression operations are performed on the target side, before data is written to disk.

With Distributed Segment Processing, operations on data segmentation, filtering and compression are performed on the backup proxy side. The backup proxy sends only unique data blocks to Dell EMC Data Domain. As a result, the load on the network reduces and the network throughput improves.

Advanced Load Balancing and Link Failover

Advanced Load Balancing and Link Failover allow you to balance data transfer load and route VM data traffic to a working link in case of network outage problems.

Without Advanced Load Balancing, every backup server connects to Data Domain on a dedicated Ethernet link. Such configuration does not provide an ability to balance the data transfer load across the links. If a network error occurs during the data transfer process, the backup job fails and needs to be restarted.

Advanced Load Balancing allows you to aggregate several Ethernet links into one interface group. As a result, Dell EMC Data Domain automatically balances the traffic load coming from several backup servers united in one group. If some link in the group goes down, Dell EMC Data Domain automatically performs link failover, and the backup traffic is routed to a working link.

Virtual Synthetics

Veeam Backup & Replication supports Virtual Synthetic Fulls by Dell EMC Data Domain. Virtual Synthetic Fulls let you synthesize a full backup on the target backup storage without physically copying data from source datastores. To construct a full backup file, Dell EMC Data Domain uses pointers to existing data segments on the target backup storage. Virtual Synthetic Fulls reduce the workload on the network and backup infrastructure components and increase the backup job performance.

In-Flight Data Encryption

Veeam Backup & Replication supports in-flight encryption introduced in Dell EMC Data Domain Boost 3.0. If necessary, you can enable data encryption at the backup repository level. Veeam Backup & Replication will leverage the Dell EMC Data Domain technology to encrypt data transported between the DD Boost library and Data Domain system.

Per Storage Unit Streams

Veeam Backup & Replication supports per storage unit streams on Dell EMC Data Domain. The maximum number of parallel tasks that can be targeted at the backup repository (the **Limit maximum concurrent tasks to N** setting) is applied to the storage unit, not the whole Dell EMC Data Domain system.

Supported Protocols

Veeam Backup & Replication supports Dell EMC Data Domain storage systems working over the following protocols:

- TCP/IP protocol: Veeam Backup & Replication communicates with the Dell EMC Data Domain server by sending commands over the network.
- Fibre Channel protocol: Veeam Backup & Replication communicates with the Dell EMC Data Domain Fibre Channel server by sending SCSI commands over Fibre Channel.

Accelerated Restore of Entire VM

To speed up entire VM restore on Dell EMC Data Domain, Veeam Backup & Replication uses the mechanism of sequential data reading from backups and parallel VM disks restore.

Dell EMC Data Domain storage systems are optimized for sequential I/O operations. However, data blocks of VM disks in backup files are stored not sequentially, but in the random order. If data blocks of VM disks are read at random, the restore performance from backups on Dell EMC Data Domain degrades.

To accelerate the restore process, Veeam Backup & Replication creates a map of data blocks in backup files. It uses the created map to read data blocks of VM disks from backup files sequentially, as they reside on disk. Veeam Backup & Replication writes data blocks to target in the order in which they come from the target Veeam Data Mover, restoring several VM disks in parallel.

This accelerated restore mechanism is enabled by default, and is used for the entire VM restore scenario.

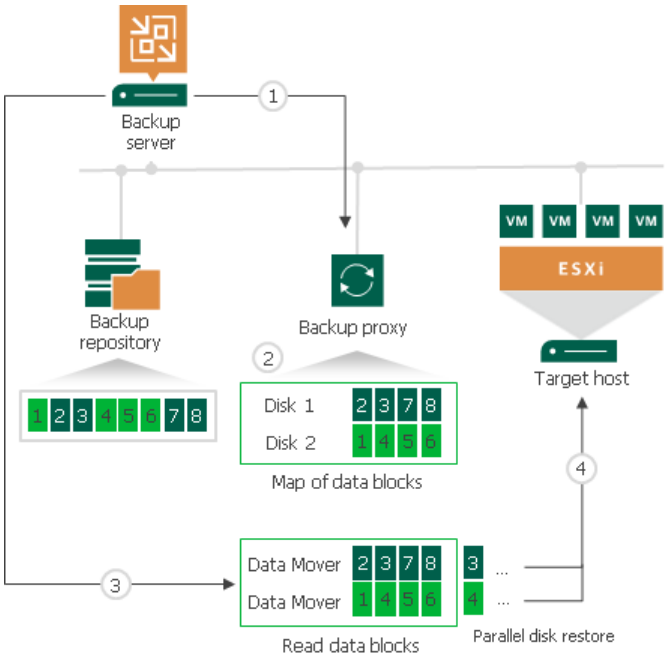
> **NOTE**
>
> To further accelerate the process of entire VM restore, Veeam Backup & Replication reads VM data from Dell EMC Data Domain in multiple threads.

How Accelerated Restore Works

Entire VM restore from backups on Dell EMC Data Domain is performed in the following way:

1. Veeam Backup & Replication opens all backup files in the backup chain, reads metadata from these backup files and caches this metadata on the backup proxy that is assigned for the restore task.
2. Veeam Backup & Replication uses the cached metadata to build a map of data blocks. The map contains references to VM data blocks, sorted by VM disks.
3. Every VM disks is processed in a separate task. For every task, Veeam Backup & Replication starts a separate Veeam Data Mover on the backup proxy.

   Veeam Data Movers read data blocks of VM disks from the backup repository sequentially, as these blocks reside on disk, and put read data blocks to the buffer on the backup proxy.
4. Data blocks are written to target in the order in which they come from the target Veeam Data Mover.



Backup Proxy for Accelerated Restore

Veeam Backup & Replication restores all disks of a VM through one backup proxy. If you instruct Veeam Backup & Replication to select a backup proxy for the restore task automatically, it picks the least loaded backup proxy in the backup infrastructure. If you assign a backup proxy explicitly, Veeam Backup & Replication uses the selected backup proxy.

For every VM disk, Veeam Backup & Replication starts a separate Veeam Data Mover on the backup proxy. For example, if you restore a VM with 10 disks, Veeam Backup & Replication starts 10 Veeam Data Movers on the backup proxy.

The backup proxy assigned for the entire VM restore task must have enough RAM resources to be able to restore VM disks in parallel. For every VM disk, 200 MB of RAM is required. The total amount of required RAM resources is calculated by the following formula:

```
Total amount of RAM = Number of VM disks * 200 MB
```

Before starting the restore process, Veeam Backup & Replication checks the amount of RAM resources on the backup proxy. If the backup proxy does not have enough RAM resources, Veeam Backup & Replication displays a warning in the job session details and automatically fails over to a regular VM disks processing mode (data of VM disks is read at random and VM disks are restored sequentially).

Limitations for Accelerated Restore

The accelerated restore of entire VM has the following limitations:

- Accelerated restore works on Dell EMC Data Domain systems with DD Boost.
- If you restore a VM with dynamically expanding disks, the restore process may be slow.
- If you restore a VM using the Network transport mode, the number of VM disks restored in parallel cannot exceed the number of allowed connections to an ESXi host.
- If Dell EMC Data Domain is added as an extent to a scale-out backup repository, you must set the backup file placement policy to Locality. If the backup file placement policy is set to Performance, parallel VM disk restore will be disabled.

ExaGrid

You can use ExaGrid appliances as backup repositories.

Adaptive Deduplication

ExaGrid uses adaptive deduplication. Data deduplication is performed on the target storage system. After VM data is written to disk, ExaGrid analyses bytes in the newly transferred data portions. ExaGrid compares versions of data over time and stores only the differences to disk.

ExaGrid deduplicates data at the storage level. Identical data is detected throughout the whole storage system, which increases the deduplication ratio.

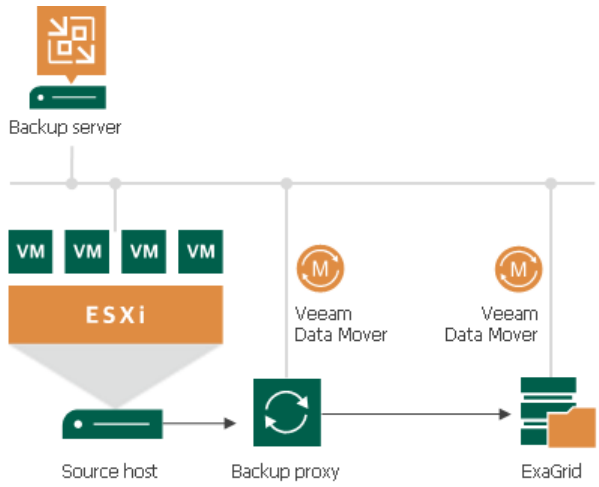
ExaGrid Deployment

To communicate with ExaGrid, Veeam Backup & Replication uses two Veeam Data Movers that are responsible for data processing and transfer:

- Veeam Data Mover on the backup proxy
- Veeam Data Mover on the ExaGrid appliance

ExaGrid does not host Veeam Data Mover permanently. When any task addresses an ExaGrid storage, Veeam Backup & Replication deploys and starts the Veeam Data Mover on the ExaGrid appliance.

Veeam Data Mover establishes a connection with Veeam Data Mover on the backup proxy, enabling efficient data transfer over LAN or WAN.



Requirements and Recommendations for ExaGrid

To perform backup to ExaGrid appliances, it is recommended to configure backup repositories and jobs in the following way:

Backup repositories

Configure ExaGrid backup repositories in the following way:

1. Create at least one share on each ExaGrid appliance. Enable the **ExaGrid-Veeam Accelerated Data Mover transport** option for the created share. Leave default compression and deduplication settings for the share.

   > **NOTE**
   >
   > When adding credentials for ExaGrid, use only lower case symbols for the user name.

2. In Veeam Backup & Replication, perform the following actions:

   a. Configure ExaGrid backup repositories and point them at the created shares on each ExaGrid appliance. Set the **Limit maximum concurrent tasks to N** option to 10 tasks. This limit can be tuned up or down with assistance from ExaGrid Customer Support.

   When you add ExaGrid servers to the Veeam backup infrastructure, and you use the UPN format for an Active Directory account user name (for example, john.doe@domain.local), make sure you enter the user name in lowercase letters only.

   b. Add ExaGrid backup repositories as extents to a scale-out backup repository.

Backup Jobs

Configure backup jobs in the following way:

1. Backup job settings:

   a. Use the forward incremental backup method.

   b. Enable synthetic full backups and schedule them to run on a weekly basis.

   c. Enable active full backups and schedule them to run on a monthly basis.

2. Backup target: Assign backup jobs to the scale-out backup repository with ExaGrid appliances as extents.

NOTE

Mind the following:

- Do not create multiple backup repositories directed at the same folder/path on the same device.
- We recommend against enabling encryption for the jobs targeted at the deduplication storage appliance. Encryption has a negative effect on the deduplication ratio. For more information, see Data Encryption.

For more information and recommendations on working with ExaGrid, see this Veeam KB article.

HPE StoreOnce

You can use HPE StoreOnce storage appliances as backup repositories.

To work with HPE StoreOnce, Veeam Backup & Replication leverages the HPE StoreOnce Catalyst technology and two HPE StoreOnce components:

- **HPE StoreOnce Catalyst agent**. The HPE StoreOnce Catalyst agent is a component of the HPE StoreOnce Catalyst software. The HPE StoreOnce Catalyst agent is embedded into the Veeam Data Mover setup. When you add a Microsoft Windows server to the backup infrastructure, the HPE StoreOnce Catalyst agent is automatically installed on the added server together with Veeam Data Mover.
- **HPE StoreOnce appliance**. The HPE StoreOnce appliance is an HPE StoreOnce storage system on which Catalyst stores are created.

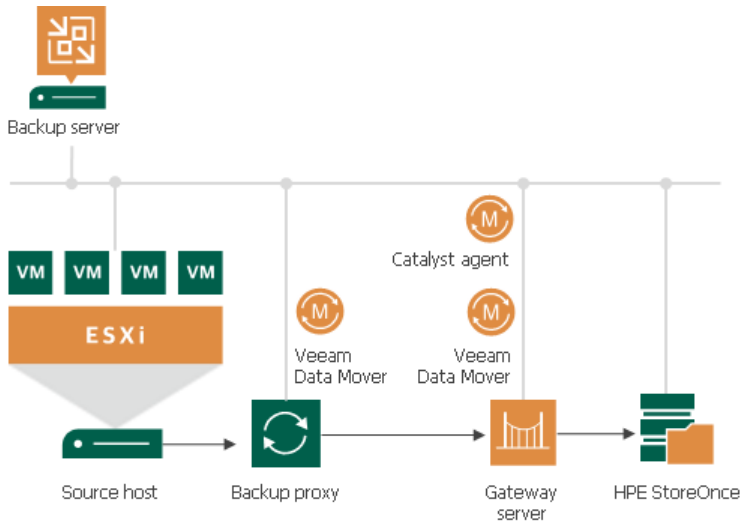
HPE StoreOnce Deployment

To communicate with HPE StoreOnce, Veeam Backup & Replication uses two Data Movers that are responsible for data processing and transfer:

- Veeam Data Mover on the backup proxy
- Veeam Data Mover on the gateway server

The HPE StoreOnce storage cannot host Veeam Data Mover. For this reason, to communicate with the HPE StoreOnce storage, you need to deploy a gateway server. Veeam Backup & Replication will automatically deploy Veeam Data Mover on this gateway server. For more information, see Gateway Server. For communicating with the HPE StoreOnce storage appliances, the gateway server must run a 64-bit Microsoft Windows version.

When any job addresses the backup repository, Veeam Data Mover on the gateway server establishes a connection with Veeam Data Mover on the backup proxy, enabling efficient data transfer over LAN or WAN.



The gateway server is selected when you assign a backup repository role to the HPE StoreOnce appliance. You can define the gateway server explicitly or instruct Veeam Backup & Replication to select it automatically.

> **TIP**
>
> For work with HPE StoreOnce, Veeam Backup & Replication uses the Catalyst agent installed on the gateway server. If you want to reduce the load on the network between the source and target side, assign the gateway server role to a machine on the source side, closer to the backup proxy.

Limitations and Recommendations for HPE StoreOnce

If you plan to use HPE StoreOnce storage appliances for NAS backup, mind the following recommendations for optimal performance:

- A StoreOnce system can have multiple Catalyst stores, and large backup loads (exceeding 1PB) should be spread across more than one Catalyst store on the same StoreOnce system.
- Do not include Catalyst stores in a SOBR intended for NAS backups. This will reduce the global deduplication of the StoreOnce system.

If you plan to use HPE StoreOnce as a backup repository for Veeam Plug-in for Oracle RMAN or Veeam Plug-in for SAP HANA, the total number of stored files (data and metadata) must not exceed 3,000,000 per Catalyst store. If necessary, multiple Catalyst stores may be created on the same StoreOnce system.

If you plan to use HPE StoreOnce as a backup repository for other jobs, mind the following limitations. Limitations apply only if you use HPE StoreOnce in the integration mode, not the shared folder mode.

- When you create a job targeted at HPE StoreOnce, Veeam Backup & Replication will offer you to switch to optimized job settings and use the 4 MB size of data block for VM data processing. It is recommended that you use optimized job settings. Large data blocks produce a smaller metadata table that requires less memory and CPU resources to process.
- The HPE StoreOnce backup repository always works in the **Use per-machine backup files** mode. For more information, see Per-Machine Backup Files.
- We recommend against enabling encryption for the jobs targeted at the deduplication storage appliance. Encryption has a negative effect on the deduplication ratio. For more information, see Data Encryption.

- HPE StoreOnce does not support the reverse incremental backup method.
- For backup jobs, HPE StoreOnce does not support the forever forward incremental backup method. When creating a backup job, you must enable synthetic and/or active full backups. For more information on how to enable such backups, see Backup Settings.
- The HPE StoreOnce backup repository does not support the **Defragment and compact full backup file** option.
- You cannot use HPE StoreOnce backup repositories as sources or targets for file copy jobs.
- You cannot copy backup files (VBK, VIB and VRB) manually to the HPE StoreOnce backup repository. To copy such files, use backup copy jobs or evacuate backups if you use a scale-out backup repository.
- You cannot use the HPE StoreOnce backup repository as a cloud repository hosted behind a Cloud Connect Gateway server.
- Considering possible bandwidth and latency constraints, Veeam recommends inline with the HPE guidelines the usage of HPE Cloud Volumes Backups as a secondary target repository, meaning that it is ideally used as a target for backup copy jobs with enabled HPE Catalyst Copy.

  > **NOTE**
  >
  > Due to the limit on the number of concurrently opened files on HPE Cloud Volumes Backup, the maximum length of backup chains (chains that contain one full backup and a set of subsequent incremental backups) on HPE Cloud Volumes Backup repositories is limited to 32 for backup copy jobs with enabled HPE Catalyst Copy. For other backup copy jobs and backup jobs, the maximum length of backup chains is limited to 7. That is why it is recommended that you use HPE Cloud Volumes Backup only as a target for backup copy jobs with enabled HPE Catalyst Copy.

- To optimize data transfer between two HPE StoreOnce repositories, use backup copy jobs with HPE Catalyst Copy enabled. For more information on how to create jobs and recommendations for them, see Creating Backup Copy Jobs for HPE StoreOnce Repositories.
- HPE StoreOnce has a limit on the number of concurrently opened files. Due to this limit, the maximum length of backup chains (chains that contain one full backup and a set of subsequent incremental backups) on HPE StoreOnce is also limited and depends on the particular storage model (see the following table).

Product	Maximum number of restore points per backup chain
VSA	
VSA Gen3	7
VSA Gen4	7 to 14 (for version 4.1.1 varies depending on the amount of available memory)
Proliant Gen7	
6200	14 (per node)
Proliant Gen8	

Product	Maximum number of restore points per backup chain
2700	7
2900	14
4500	14
4700	14
4900	28
6500	28 (per node)
Proliant Gen9	
3100	7
3500	14
5100	21
5500	35
6600	42 (per node)
Proliant Gen10	
3620	14
3640	14
5200	28
5250	28
5650	42

Several Backup Repositories on HPE StoreOnce

You can configure several backup repositories on one HPE StoreOnce appliance and associate them with different gateway servers.

Mind the following:

- If you configure several backup repositories on HPE StoreOnce and add them as extents to a scale-out backup repository, make sure that all backup files from one backup chain are stored on one extent. If backup files from one backup chain are stored to different extents, the performance of transformation processes will be lower. For more information about transformation performance, see this Veeam blog post.
- HPE StoreOnce has a limit on the number of opened files that applies to the whole appliance. Tasks targeted at different backup repositories on HPE StoreOnce and run in parallel will equally share this limit.
- For HPE StoreOnce working over Fibre Channel, there is a limitation on the number of connections from one host. If you connect several backup repositories to one gateway, backup repositories will compete for connections.
- Deduplication on HPE StoreOnce works within the limits of one object store.
- If your VMs contain similar data, it is recommended to create backup repositories on a single HPE StoreOnce Catalyst Store. This minimizes backup job duration and reduces disk space used for backups. For details, see this Veeam KB article.

For more information and recommendations on working with HPE StoreOnce, see this Veeam KB article.

Operational Modes

Depending on the storage configuration and type of the backup target, HPE StoreOnce can work in the following modes:

- Source-side deduplication
- Target-side deduplication
- Shared folder mode

Source-Side Data Deduplication

HPE StoreOnce performs source-side deduplication if the backup target meets the following requirements:

- You have a Catalyst license installed on HPE StoreOnce.
- You use a Catalyst store as a backup repository.
- The Catalyst store is configured to work in the Low Bandwidth mode (Primary Transfer Policy).
- The HPE StoreOnce Catalyst is added to the backup repository as a deduplicating storage appliance, not as a shared folder.

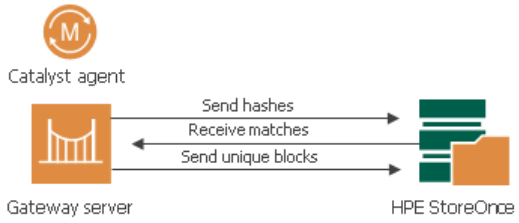
To deduplicate data on the source side, HPE StoreOnce uses the HPE StoreOnce Catalyst agent. The HPE StoreOnce Catalyst agent is a component of the HPE StoreOnce Catalyst software. It is installed on the gateway server communicating with the HPE StoreOnce appliance.

HPE StoreOnce deduplicates data on the source side, before writing it to target:

1. During the backup job session, HPE StoreOnce analyzes data incoming to the HPE StoreOnce appliance in chunks and computes a hash value for every data chunk. Hash values are stored in an index on disk.

2. The HPE StoreOnce Catalyst agent calculates hash values for data chunks in a new data flow and sends these hash values to target.
3. HPE StoreOnce identifies which data blocks are already saved on disk and communicates this information to the HPE StoreOnce Catalyst agent. The HPE StoreOnce Catalyst agent sends only unique data blocks to target.

As a result, the load on the network reduces, the backup job performance improves, and you can save on disk space.



Target-Side Data Deduplication

HPE StoreOnce performs target-side deduplication if the backup target is configured in the following way:

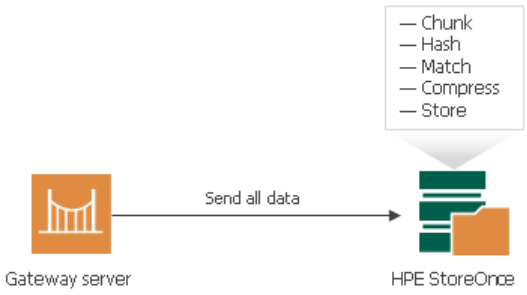
- For a Catalyst store:
  - The Catalyst store works in the High Bandwidth mode (Primary Transfer Policy is set to High Bandwidth).
  - The Catalyst license is installed on the HPE StoreOnce (required).
  - The Catalyst store is added to the backup repository as a deduplicating storage appliance, not as a shared folder.
- For a CIFS store:
  - The Catalyst license is not required.
  - The CIFS store is added as a shared folder backup repository to the backup infrastructure.

  For more information about working with CIFS stores, see Shared Folder Mode.

HPE StoreOnce deduplicates data on the target side, after the data is transported to HPE StoreOnce:

1. HPE StoreOnce analyzes data incoming to the HPE StoreOnce appliance in chunks and creates a hash value for every data chunk. Hash values are stored in an index on the target side.
2. HPE StoreOnce analyzes VM data transported to target and replaces identical data chunks with references to data chunks that are already saved on disk.

As a result, only new data chunks are written to disk, which helps save on disk space.



Shared Folder Mode

If you do not have an HPE StoreOnce Catalyst license, you can add the HPE StoreOnce appliance as a shared folder backup repository. In this mode, HPE StoreOnce will perform target-side deduplication.

If you work with HPE StoreOnce in the shared folder mode, the performance of backup jobs and transformation processes is lower (in comparison with the integration mode, when HPE StoreOnce is added as a deduplicating storage appliance).

HPE StoreOnce Supported Features

The HPE StoreOnce Catalyst technology offers a set of features for advanced data processing. Veeam Backup & Replication supports the following features:

Synthetic Full Backups

HPE StoreOnce Catalyst improves synthetic full backup file creation and transformation performance. When Veeam Backup & Replication creates or transforms a synthetic full backup, HPE StoreOnce does not physically copy data between the existing backup chain and the target full backup file. Instead, it performs a metadata-only operation — updates pointers to existing data blocks on the storage device. As a result, the operation completes much faster. This mechanism helps improve performance of primary backup jobs and backup copy jobs that are scheduled to create periodic archive full backups (GFS).

Accelerated vPower-Enabled Operations

If HPE StoreOnce Catalyst is integrated with Veeam Backup & Replication, it allows to benefit from improved performance of vPower-enabled operations — Instant Recovery, SureBackup and On-Demand Sandbox — from backups residing on HPE StoreOnce. To get the maximum vPower performance, HPE StoreOnce must be running firmware version 3.15.1 or later.

Accelerated Data Recovery

Integration with HPE StoreOnce improves data recovery performance for different restore scenarios: Instant Recovery, file-level recovery and application items recovery with Veeam Explorers.

WAN-based Catalyst Store Support

Veeam Backup & Replication provides advanced support for WAN-based HPE Catalyst stores. If a WAN connection to HPE StoreOnce is weak, you can instruct Veeam Backup & Replication to compress VM data and calculate checksums for data blocks going from the source side to HPE StoreOnce.

HPE StoreOnce Replication

HPE StoreOnce replication improves copying data between two HPE StoreOnce backup repositories. For more information on copying, see Creating Backup Copy Jobs for HPE StoreOnce Repositories.

Supported Protocols

Veeam Backup & Replication supports HPE StoreOnce storage systems working over the following protocols:

- TCP/IP protocol: Veeam Backup & Replication communicates with the HPE StoreOnce appliance by sending commands over the LAN.
- Fibre Channel protocol: Veeam Backup & Replication communicates with the HPE StoreOnce appliance by sending SCSI commands over Fibre Channel.

  Data processing over Fibre Channel (FC) connectivity enables local area network-free backup to HP StoreOnce, eliminates the load from backup activities and increases availability of LAN resources to production workloads.

Quantum DXi

You can use Quantum DXi appliances as backup repositories.

Quantum DXi Deduplication

Quantum DXi appliances use Quantum's patented data deduplication technology. During backup data transfer, Quantum analyses blocks in a data stream. Instead of copying redundant data blocks Quantum uses reference pointers to existing blocks on the storage device.

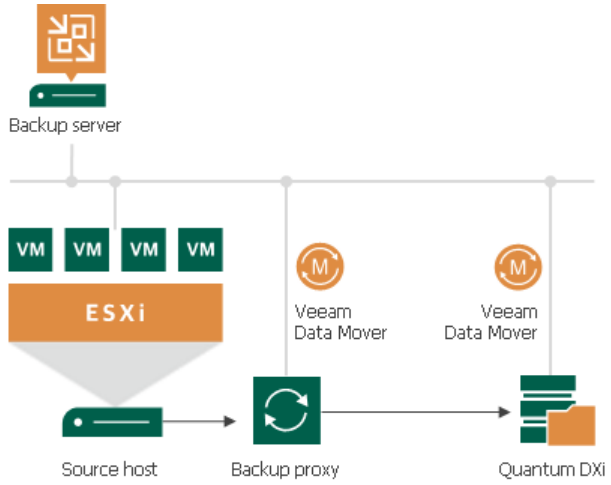
Quantum DXi Deployment

To communicate with Quantum DXi, Veeam Backup & Replication uses two Veeam Data Movers that are responsible for data processing and transfer:

- Veeam Data Mover on the backup proxy
- Veeam Data Mover on the Quantum DXi appliance

Quantum DXi does not host the Veeam Data Mover permanently. When any task addresses a Quantum DXi storage, Veeam Backup & Replication deploys and starts Veeam Data Mover on the Quantum DXi system.

Veeam Data Mover establishes a connection with Veeam Data Mover on the backup proxy, enabling efficient data transfer over LAN or WAN.



> **NOTE**
>
> We recommend against enabling encryption for the jobs targeted at the deduplication storage appliance. Encryption has a negative effect on the deduplication ratio. For more information, see Data Encryption.

For more information and recommendations on working with Quantum DXi, see this Veeam KB article.

Backup Repositories with Rotated Drives

A backup repository can use rotated drives. Rotated drives can be detachable USB or eSATA hard drives. This scenario can be helpful if you want to store backups on several external hard drives that you plan to regularly move between different locations.

To use rotated drives, you must enable the **This repository is backed by rotated hard drives** option in the advanced settings of the backup repository. When this option is enabled, Veeam Backup & Replication recognizes the backup target as a backup repository with rotated drives and uses a specific algorithm to make sure that the backup chain created on these drives is not broken.

Limitations for Backup Repositories with Rotated Drives

Backup repositories with rotated drives have the following limitations:

- On one managed server, you must create only one repository with rotated drives.
- You cannot store archive full backups (GFS backups) created with backup jobs or backup copy jobs in backup repositories with rotated drives.
- You cannot store per-machine backup files in backup repositories with rotated drives.
- You cannot rescan backup repositories with rotated drives.
- NFS backup repositories do not support rotated drives. If you enable the **This repository is backed by rotated hard drives** setting on the repository, Veeam Backup & Replication will ignore this setting.
- Scale-out backup repositories do not support rotated drives. If you enable the **This repository is backed by rotated hard drives** setting on an extent, Veeam Backup & Replication will ignore this setting and will work with such repository as with a standard extent.
- Repositories with rotated drives are not supported as primary backup repositories, archive repositories, and secondary target repositories for NAS backup.

In This Section

- How Repository with Rotated Drives Works
- Adding Backup Repositories with Rotated Drives

How Repository with Rotated Drives Works

You can use Microsoft Windows server or Linux server or SMB (CIFS) share as a backup repository with rotated drives.

Microsoft Windows Backup Repository

Veeam Backup & Replication performs backup jobs and backup copy jobs targeted at a backup repository with rotated drives in different ways.

> **IMPORTANT**
>
> When you specify retention settings for a backup job or a backup copy job targeted at a backup repository with rotated drives, you must define the total number of restore points that you want to retain on all drives in the set. For example, if you set retention to 14, the job will keep the total of 14 restore points across all drives.
>
> Before the drive change, we recommend to check the retention setting in backup chains on all drives in the set. Otherwise, Veeam Backup & Replication will delete all restore points on the drive that you plan to use in the following cases:
>
> - All restore points are outdated.
> - A number of restore points equals the retention value.
>
> To eliminate the risk of deleting a backup chain from a rotated drive, make sure you have increased the retention value. You can also disable a backup copy job before you will use the drive.

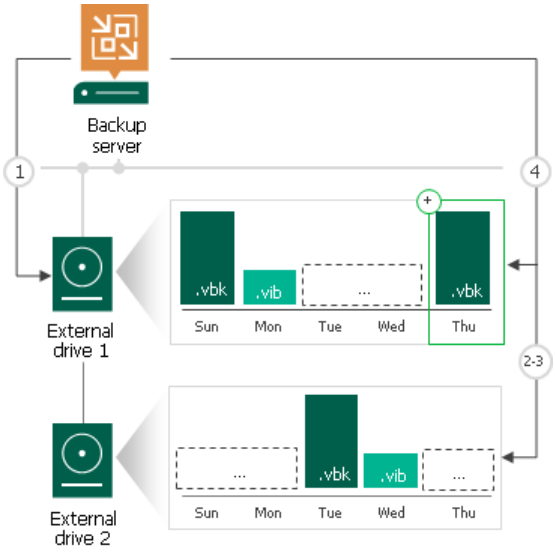
Backup Jobs

Backup jobs are performed in the following way:

1. Veeam Backup & Replication creates a regular backup chain on the currently attached drive.
2. When a new job session starts, Veeam Backup & Replication checks if the backup chain on the currently attached drive is consistent. The consistent backup chain must contain a full backup and all incremental backups that have been produced by the job. This requirement applies to all types of backup chains: forever forward incremental, forward incremental and reverse incremental.

   If external drives have been swapped, and the full backup or any incremental backups are missing from the currently attached drive, Veeam Backup & Replication starts the backup chain anew. It creates a new full backup file on the drive, and this full backup is used as a starting point for subsequent incremental backups.
3. [For external drives attached to Microsoft Windows servers] Veeam Backup & Replication checks the retention policy set for the job. If some backup files in the backup chain are outdated, Veeam Backup & Replication removes them from the backup chain.

4. When you swap drives again, Veeam Backup & Replication checks the backup chain for consistency and creates a new full backup.



Backup Copy Jobs

Backup copy jobs are performed in the following way:

1. Veeam Backup & Replication creates a regular backup chain on the currently attached drive.
2. When you swap drives, and the attached drive is empty, Veeam Backup & Replication creates a full backup on it. If there is a backup chain on the drive, Veeam Backup & Replication creates a new incremental backup and adds it to the backup chain. The latest incremental backup existing in the backup chain is used as a starting point for the new incremental backup.
3. [For external drives attached to Microsoft Windows servers] Veeam Backup & Replication checks the retention policy set for the job. If some backup files in the backup chain are outdated, Veeam Backup & Replication removes them from the backup chain.

Drive Detection

Drive letters for external drives may change when you add new volumes or storage hardware such as CD-ROM on the server. In Microsoft Windows backup repositories, Veeam Backup & Replication can keep track of drives and detect them even if the drive letter changes.

To detect a drive correctly, Veeam Backup & Replication must have a record about it in the configuration database. Consider the following requirements:

- When you insert a drive for the first time, the drive is not registered in the configuration database. Such drive must have the same letter as the one specified in the **Path to folder** field in backup repository settings. For more information, see Configuring Path and Load Control Settings.

  If the drive has some other letter, Veeam Backup & Replication will not be able to detect and use it.
- When you insert a drive that has already been used and has some restore points on it, the drive is already registered in the configuration database. Veeam Backup & Replication will be able to detect and use it, even if the drive letter changes.

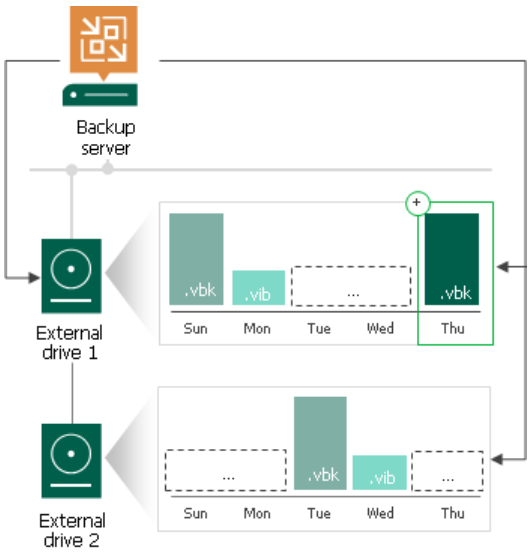
Linux and Shared Folder Backup Repository

If you use a Linux server or CIFS share as a backup repository with rotated drives, Veeam Backup & Replication employs a "cropped" mechanism of retention with rotated drives. Veeam Backup & Replication keeps information only about the latest backup chain in the configuration database. Information about previous backup chains is removed from the database. For this reason, the retention policy set for the job may not work as expected.

A job targeted at a backup repository with rotated drives is performed in the following way:

1. During the first run of the job, Veeam Backup & Replication creates a regular backup full backup on the drive that is attached to the backup repository server.
2. During the next job session, Veeam Backup & Replication checks if the current backup chain on the attached drive is consistent. The consistent backup chain must contain a full backup and all incremental backups subsequent to it. This requirement applies to all types of backup chains: forever forward incremental, forward incremental and reverse incremental.
   - If the current backup chain is consistent, Veeam Backup & Replication adds a new restore point to the backup chain.
   - If external drives have been swapped, and the current backup chain is not consistent, Veeam Backup & Replication always starts a new backup chain (even if restore points from previous backup chains are available on the attached drive). Veeam Backup & Replication creates a new full backup file on the drive, and this full backup is used as a starting point for subsequent incremental backups.

   As soon as Veeam Backup & Replication starts a new backup chain on the drive, it removes information about restore points from previous backup chains from the configuration database. Backup files corresponding to these previous restore points are not deleted, they remain on disk. This happens because Veeam Backup & Replication applies the retention policy only to the current backup chain, not to previous backup chains.



Deploying Backup Repositories with Rotated Drives

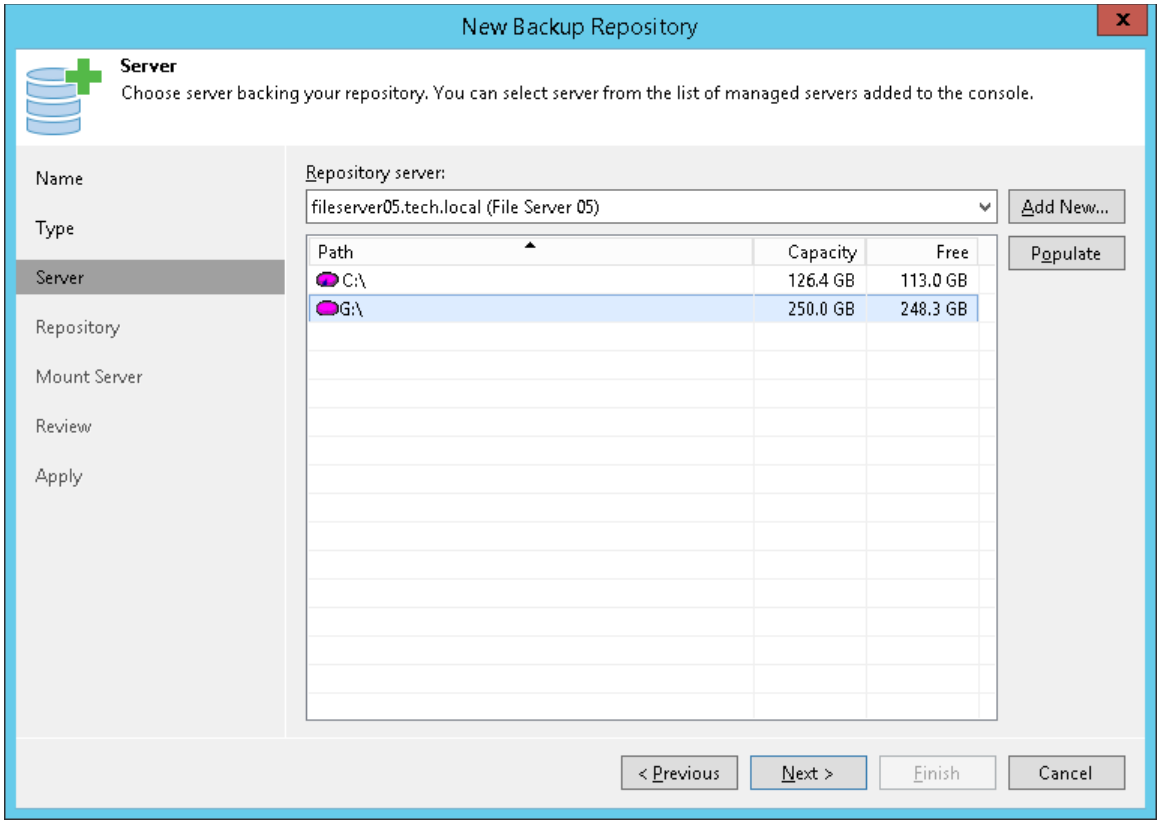
To deploy a backup repository with rotated drives, do the following:

1. Attach one of external drives from the set to a Microsoft Windows or Linux server. The server must be added to the backup infrastructure. For more information on how to add a server, see Virtualization Servers and Hosts.

   You can also attach the external hard drive to the backup server itself. In this case, the VM traffic will path through the backup server, which will produce additional workload on it.

2. Use the **New Backup Repository** wizard to add a direct attached storage. For more information, see Adding Backup Repositories. Pay attention to the following settings:

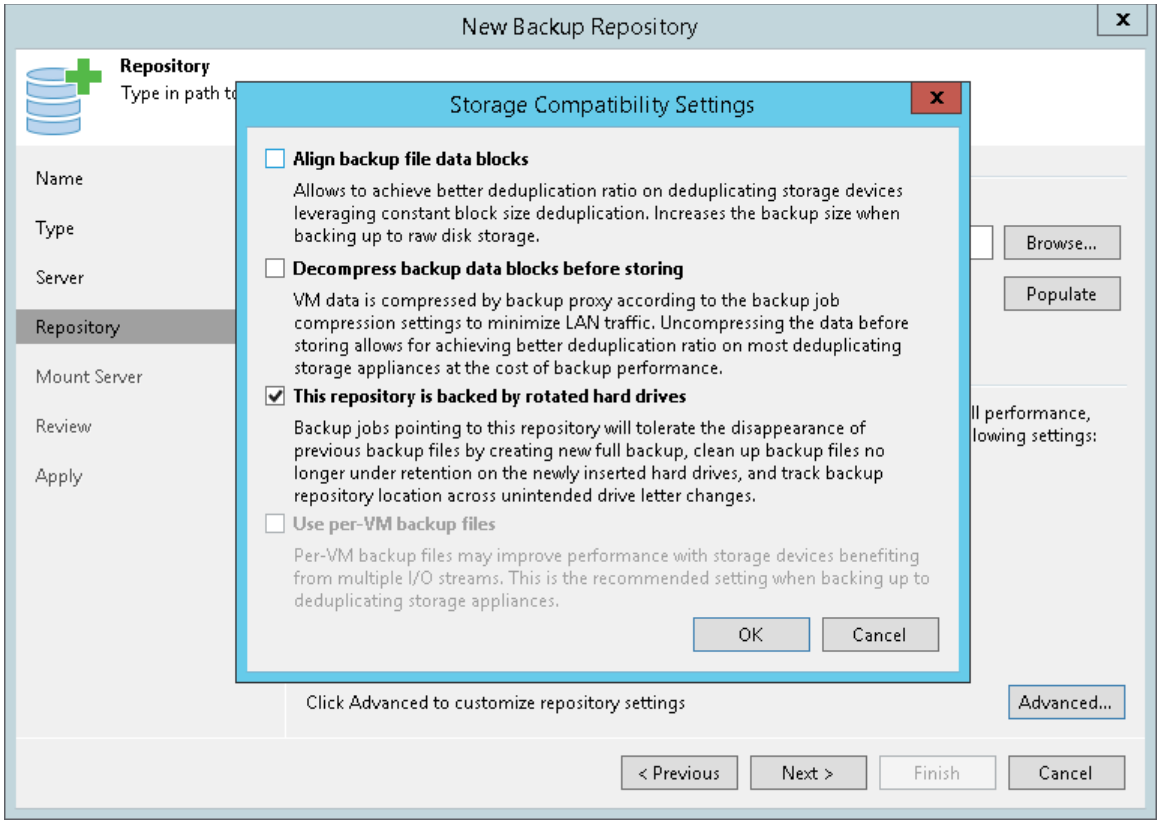
   a. At the **Server** step of the wizard, select the server to which the drive is attached.



   b. At the **Repository** step of the wizard, in the **Path to folder** field, specify a path to the folder where backup files must be stored.

> **IMPORTANT**
>
> Later, when you attach another external hard drive to the server for the first time, this drive must have the same drive letter as specified in the **Path to folder** field. For more information, see Drive Detection.

c. Click the **Advanced** button and select the **This repository is backed by rotated hard drives** check box.



d. Configure other settings of the backup repository as required and finish working with the wizard.

Object Storage Repository

An object storage repository is a repository intended for long-term data storage and based on either a cloud solution or an S3 compatible storage solution.

Veeam Backup & Replication supports the following types of object storage repositories:

- Amazon S3, Amazon S3 Glacier and AWS Snowball Edge
- S3 compatible
- Google Cloud
- IBM Cloud
- Microsoft Azure Blob, Azure Archive Storage and Azure Data Box

Object storage repositories can be used in the following ways:

- In Scale-Out Backup Repository as a part of Capacity Tier. Capacity tier of scale-out backup repository allows to offload existing backup data directly to cloud-based object storage such as Amazon S3, Microsoft Azure Blob storage, Google Cloud Storage, and IBM Cloud Object Storage, or using on-premises S3 compatible devices. For more information, see Capacity Tier.
- In Scale-Out Backup Repository as a part of Archive Tier. Archive tier of scale-out backup repository allows to transport infrequently accessed data from the capacity tier for archive storage. For more information, see Archive Tier.
- As target archive repositories for NAS Backup.

Considerations and Limitations

This section lists considerations and known limitations for object storage repositories.

General Considerations and Limitations

- Make sure to open required ports to communicate with object storage repositories in advance, as described in Used Ports.
- Object storage gateway appliances that are used to store backup data in filer (CIFS/NFS) or block device mode (iSCSI/FC/SAS) are not supported if the backup data is offloaded to object storage and is no longer stored directly on the appliance.

  Such gateway appliances are only supported in the following cases:

  - All of the backup data is stored on the appliance altogether (that is, all of the backup chains are stored on the appliance as a whole and not scattered across multiple devices) and only additional copies of the backup data are transported to object storage.
  - These appliances emulate a tape system (VTL) as an access protocol for Veeam Backup & Replication.
- Data in object storage bucket/container must be managed solely by Veeam Backup & Replication, including retention and data management. Enabling lifecycle rules is not supported, and may result in backup and restore failures.
- Use one bucket per scale-out backup repository to reduce metadata. Creating folders for multiple scale-out backup repositories within a bucket slows down processing, as metadata operations within the object storage are handled per bucket.
- Object storage systems that offload data to tape devices or any other sequential storage devices are supported only if they are added as an archive tier in a scale-out backup repository.
- If a backup chain contains backup files that are marked as corrupted by Health Check, then such corrupted files, as well as all subsequent files that go after the corrupted one are never offloaded. In such a scenario, offload is only possible starting from the full backup file that succeeds the backup chain with corrupted backups.
- Different object storage repositories mapped to the same cloud folder can be used for storing both the capacity tier backups and the NAS backups.

  > **IMPORTANT**
  >
  > The same object storage repository (mapped to the same cloud folder) must not be used across multiple Veeam Backup & Replication servers for the same purposes as it leads to unpredictable system behavior and data loss.
  >
  > For the same reason, two object storage repositories mapped to the same cloud folder must not be added to different scale-out backup repositories within one Veeam Backup & Replication server.
- Within a scale-out backup repository, the mount server of a performance extent will act as a gateway server of the capacity extent if all of the following is true:

  a. You use SMB share/NFS share/deduplicating storage appliances as performance extents of your scale-out backup repository.

  b. You have chosen **Automatic selection** for the gateway server at the Specify Shared Folder Settings step of the **New backup repository** wizard.

c. For the object storage that you use as the capacity extent, you have not selected the **Use the following gateway server** check box at the **Specify Object Storage Account** step of the **New object storage repository** wizard.

Limitations for Amazon and S3 Compatible Object Storage

- Make sure that you add an S3-compatible object storage device fully compatible with the AWS S3 operations and AWS S3 Signature Version 4 standard.
- [For Amazon S3] Only the Standard, Standard-IA and One Zone-IA storage classes are supported. For more information about Amazon S3 storage classes, see this Amazon article.

Limitations for Microsoft Azure Object Storage

- Azure object-level immutability is not supported by Veeam Backup & Replication. In Azure, it is possible to set account-level or container-level immutability policy, these can be overridden by object-level immutability. Veeam Backup & Replication currently does not support these features. Do not enable this feature as this might lead to a significant data loss.
- [For Microsoft Azure Blob storage] Veeam Backup & Replication supports specific types of storage accounts and tiers. For more information, see Microsoft Azure Storage Accounts.
- Currently, Veeam Backup & Replication does not support the *Versioning* feature for Microsoft Azure object storage. If you plan to use an account with blob versioning enabled, keep in mind this may result in extra costs for storing objects that have been removed by the retention policy.
- [For Microsoft Azure Archive storage] Microsoft Azure has certain limits (quotas) on maximum amount of resources used. The quotas depend on the type of proxies you have selected. If you exhaust a quota, you will be unable to use Microsoft Azure Archive storage. For more information about Microsoft Azure quotas, see this Microsoft article.
- [For Microsoft Azure Archive storage and Azure Blob storage] Soft delete is not supported by Veeam Backup & Replication.

Limitations for Google Cloud Object Storage

Currently, Veeam Backup & Replication does not support the Object Versioning and Bucket Lock features for Google Cloud object storage.

> **IMPORTANT**
>
> Enabling either any or both of these features on the bucket may result in unpredictable system behavior and data loss, as well as in extra costs for storing objects that have been removed by the retention policy.
>
> For more information, see Object Versioning and Bucket Lock.

Limitations for IBM Cloud Object Storage

- For IBM Cloud Object Storage on-premise, Veeam Backup & Replication supports versions starting from 3.15.0.44.
- Currently, only the *Standard* storage class is supported for IBM Cloud object storage.

Considerations and Limitations for Immutability

- After you have created an S3 bucket with *Object Lock* enabled, check that the default retention is disabled. For this, edit Object Lock retention settings as described in AWS documentation.

  The default retention may result in an unpredictable system behavior and data loss. However, note that Veeam Backup & Replication will use *Compliance* object lock mode for each uploaded object. For more information on the retention modes, see AWS documentation.

- Versioning and Object Lock must not be enabled or disabled on buckets that have been added to Veeam Backup & Replication as it may lead to unpredictable system behavior and data loss.

- If you plan to use the immutability feature with the existing S3 bucket containing backups created by 9.5 Update 4, keep in mind that both Versioning and Object Lock must be enabled on the bucket simultaneously and immediately before enabling the immutability feature. Any other approach will lead to backup offload failures and inability to correctly interact with backups in the bucket.

- The immutability feature is applicable to the Capacity Tier and Archive Tier backups. It does not support the NAS backups.

- Immutable data is preserved as described in Block Generation (applicable only to Capacity Tier).

Adding Object Storage Repositories

You can add the following types of object storage repositories:

- S3 Compatible Object Storage
- Amazon S3 Object Storage, Amazon S3 Glacier Storage and AWS Snowball Edge
- Google Cloud Object Storage
- IBM Cloud Object Storage
- Microsoft Azure Object Storage, Microsoft Azure Archive Storage and Data Box

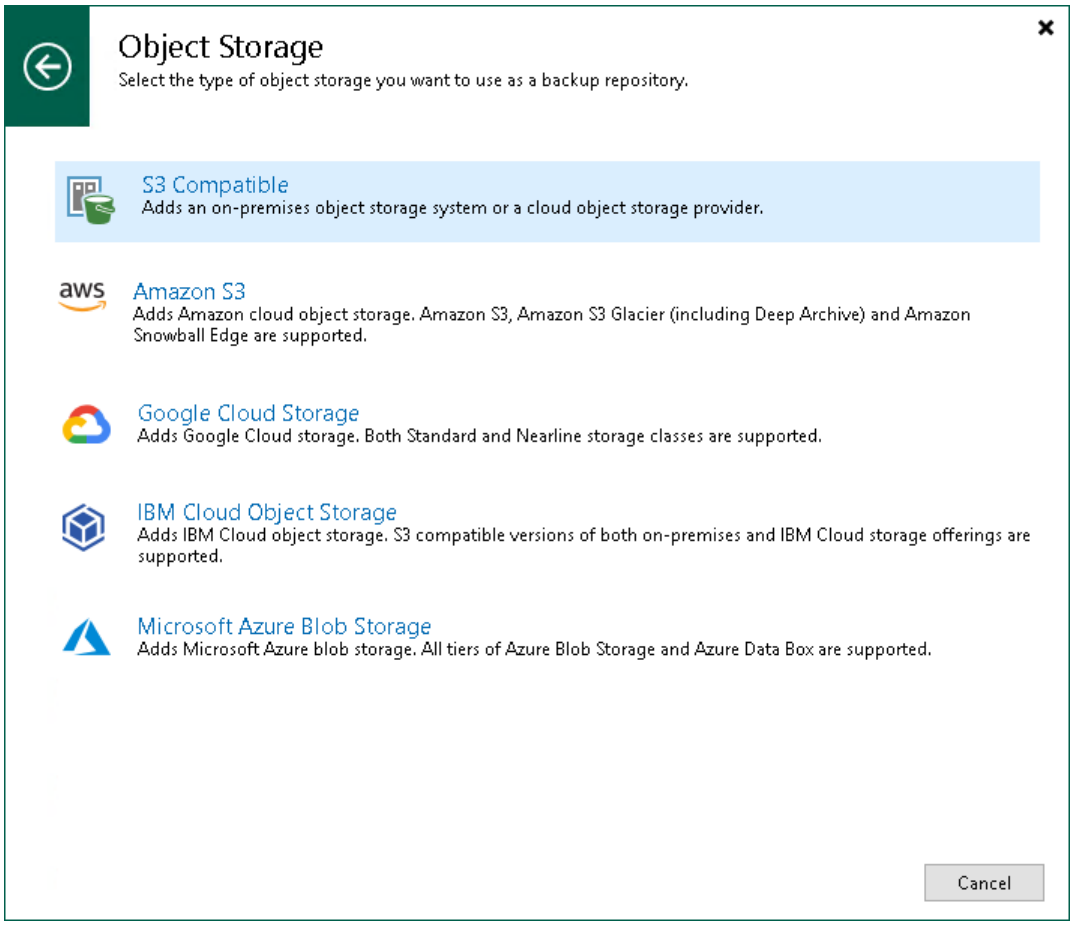
Adding S3 Compatible Object Storage

To add an S3 Compatible object storage to the backup infrastructure, use the **New Object Repository wizard**.

Step 1. Launch New Object Repository Wizard

To launch the **New Object Repository** wizard, do one of the following:

- Open the **Backup Infrastructure** view. In the inventory pane select the **Backup Repositories** node and click **Add Repository** on the ribbon. In the **Add Backup Repository** dialog, select **Object Storage > S3 Compatible**.
- Open the **Backup Infrastructure** view. In the inventory pane right-click the **Backup Repositories** node and select **Add Backup Repository**. In the **Add Backup Repository** dialog, select **Object Storage > S3 Compatible**.

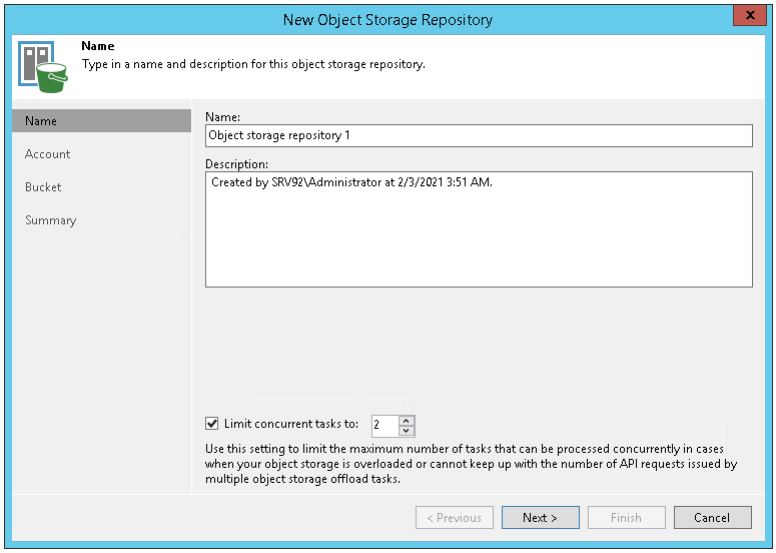


Step 2. Specify Object Storage Name

At the **Name** step of the wizard, specify a name and description for the object storage repository:

- In the **Name** field, specify a name for the new object storage repository.
- In the **Description** field, enter an optional description. The default description contains information about the user who added the object storage repository, date and time when the object storage repository was added.

If you wish to limit the maximum number of tasks that can be processed at once, select the **Limit concurrent tasks to N** check box.



Step 3. Specify Object Storage Account

At the **Account** step of the wizard, specify the connection settings:

1. In the **Service point** field, specify an endpoint address of your S3 Compatible object storage.
2. In the **Region** field, specify a region.
3. From the **Credentials** drop-down list, select user credentials to access your S3 Compatible object storage.

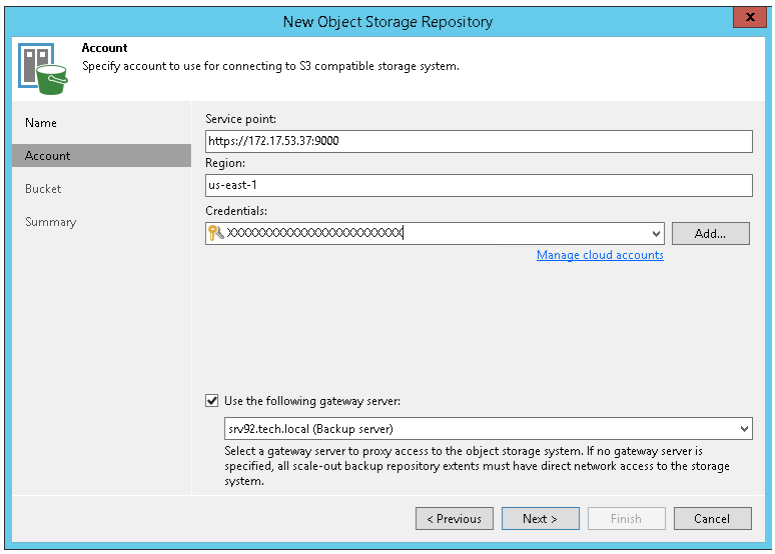
   If you already have a credentials record that was configured in advance, select it from the drop-down list. Otherwise, click **Add** and provide your access and secret keys, as described in Access Keys for AWS Users. You can also click the **Manage cloud accounts** link to add, edit or remove a credentials record.

To use a gateway server, select the **Use the following gateway server** check box and choose an appropriate server from the list. You may want to use a gateway server, for example, if your organization has NAT or different types of firewalls and your access to the internet is limited.

You can select any Microsoft Windows or Linux server that is added to your backup infrastructure and has internet connection. For more information on how to add such a server to your environment, see Adding Microsoft Windows Servers and Adding Linux Servers respectively. By default, the role of a gateway server is assigned to the machine where Veeam Backup & Replication is installed.

If you choose not to use a gateway server, make sure that all scale-out repository extents have direct internet access.

For information on considerations and limitations for using gateway server, see Considerations and Limitations.



Step 4. Specify Object Storage Settings

At the **Bucket** step of the wizard, specify the bucket and folder that will be used to store data:

1. From the **Bucket** drop-down list, select a bucket.

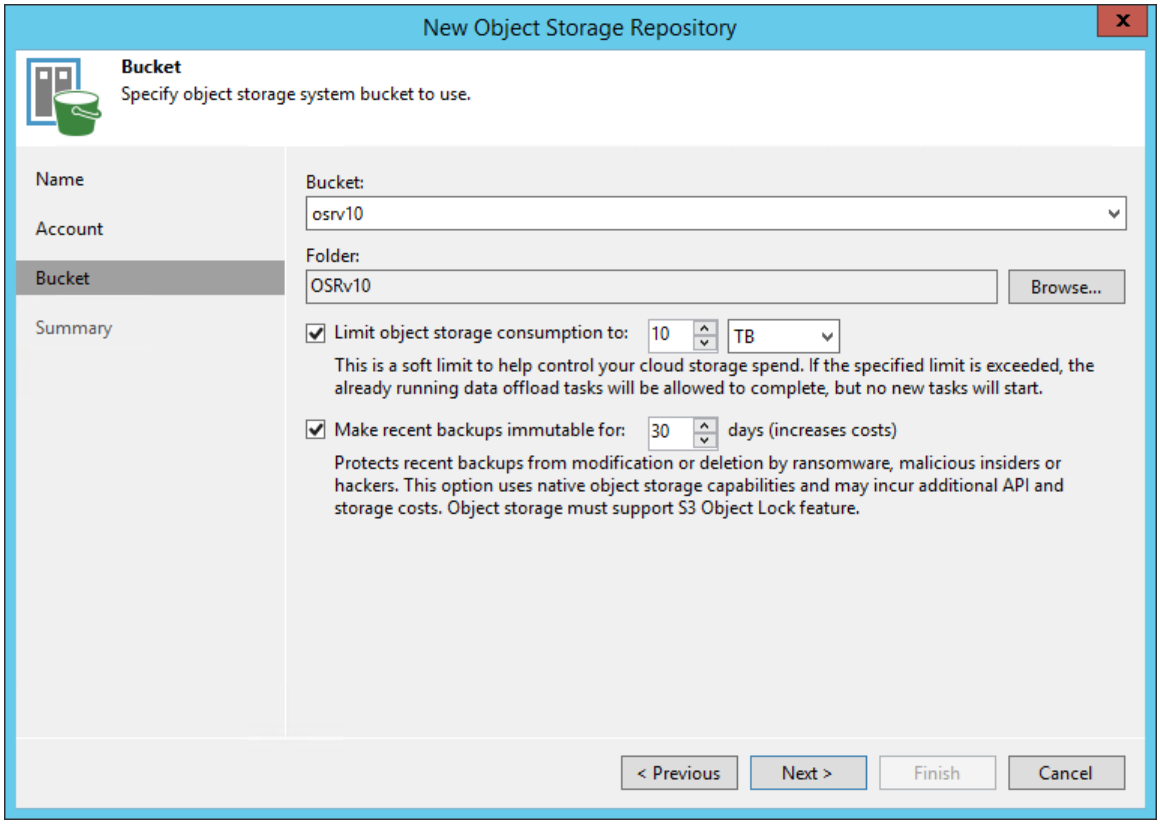
   Make sure that the bucket where you want to store your backup data was created in advance.

2. In the **Select Folder** field, select a cloud folder to which you want to map your object storage repository.

   To select a folder, click **Browse** and either select an existing folder or create a new one by clicking **New Folder**.

To define a soft limit that can be exceeded temporarily for your object storage consumption, select the **Limit object storage consumption to** check box and provide the value in TB or PB.

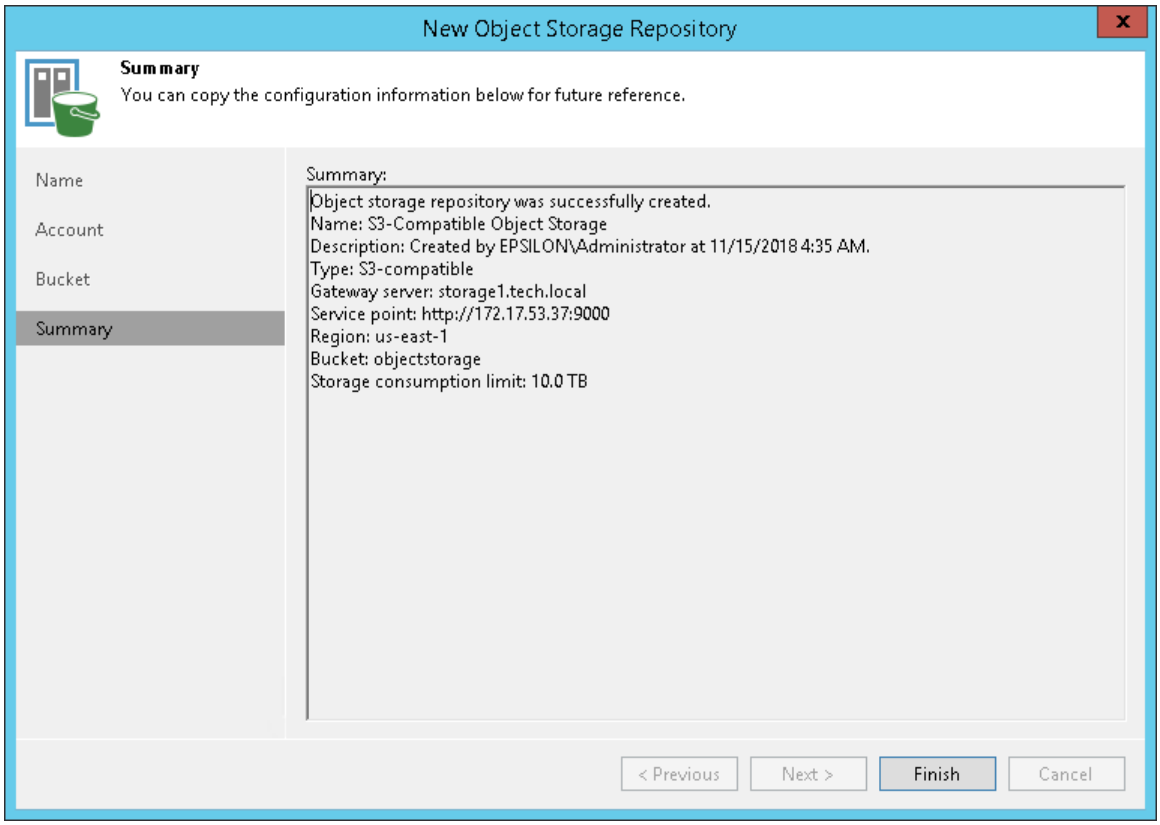
To prohibit deletion of blocks of data from object storage, select the **Make recent backups immutable for** check box and specify the immutability period. For more information about immutability, see Immutability.



Step 5. Finish Working with Wizard

At the **Summary** step of the wizard, complete the object storage repository configuration:

1. Review details of the object storage repository.
2. Click **Finish** to exit the wizard.



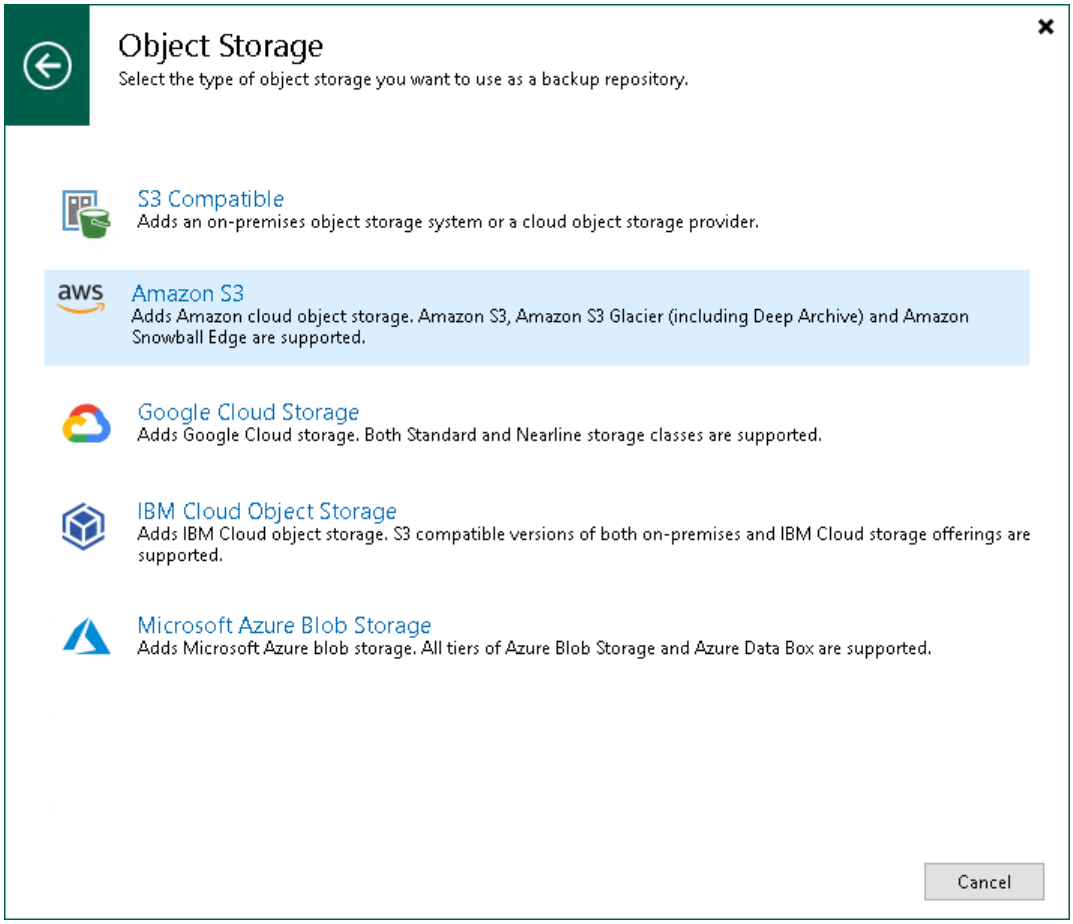
Adding Amazon S3 Object Storage, Amazon S3 Glacier Storage and AWS Snowball Edge

Before you add an Amazon S3 storage, Amazon Glacier or Amazon Snowball Edge to the backup infrastructure, check prerequisites. After that, use the **New Object Repository wizard**.

Step 1. Launch New Object Repository Wizard

To launch the **New Object Repository** wizard, do one of the following:

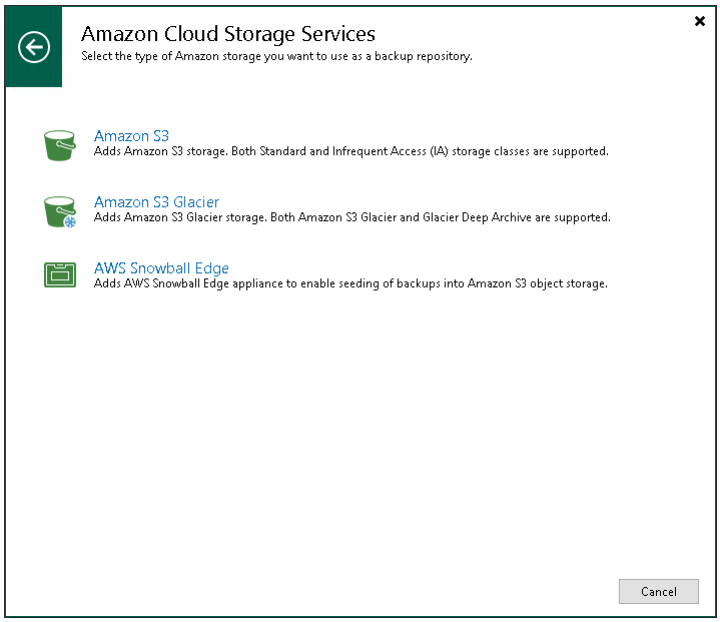
- Open the **Backup Infrastructure** view. In the inventory pane select the **Backup Repositories** node and click **Add Repository** on the ribbon. In the **Add Backup Repository** dialog, select **Object Storage > Amazon S3**.
- Open the **Backup Infrastructure** view. In the inventory pane right-click the **Backup Repositories** node and select **Add Backup Repository**. In the **Add Backup Repository** dialog, select **Object Storage > Amazon S3**.



Step 2. Select Amazon S3 Storage Type

At the **Amazon Cloud Storage Services** step of the wizard, select one of the following storage types:

- Amazon S3 Storage
- Amazon S3 Glacier Storage
- AWS Snowball Edge Storage



Adding Amazon S3 Storage

For more information about Amazon S3 object storage, see this Amazon article.

To add Amazon S3 object storage, do the following:

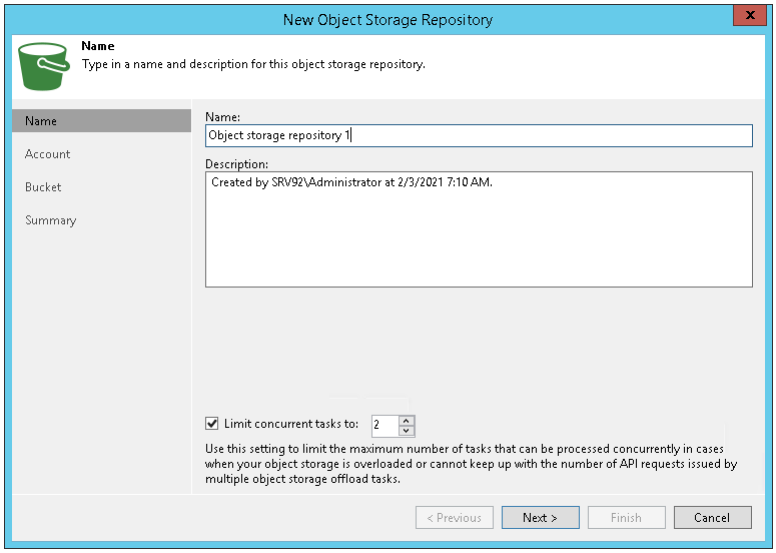
1. Specify object storage name
2. Specify object storage account
3. Specify object storage settings
4. Finish working with the wizard

Step 1. Specify Object Storage Name

At the **Name** step of the wizard, specify a name and description for the object storage repository:

- In the **Name** field, specify a name for the new object storage repository.
- In the **Description** field, enter an optional description. The default description contains information about the user who added the object storage repository, date and time when the object storage repository was added.

If you wish to limit the maximum number of tasks that can be processed at once, select the **Limit concurrent tasks to N** check box.



Step 2. Specify Object Storage Account

At the **Account** step of the wizard, specify the connection settings:

1. From the **Credentials** drop-down list, select user credentials to access your Amazon S3 object storage.

   If you already have a credentials record that was configured in advance, select it from the drop-down list. Otherwise, click **Add** and provide your access and secret keys, as described in Access Keys for AWS Users. You can also click the **Manage cloud accounts** link to add, edit or remove a credentials record.

   The user account must have permissions listed in Required Permissions.

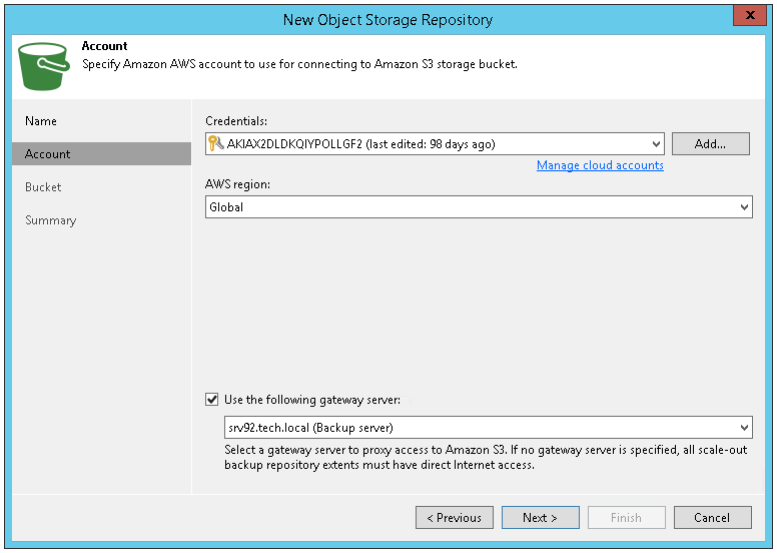
2. From the **AWS region** drop-down list, select the AWS region where the Amazon S3 bucket is located.

To use a gateway server, select the **Use the following gateway server** check box and choose an appropriate server from the list. You may want to use a gateway server, for example, if your organization has NAT or different types of firewalls and your access to the internet is limited.

You can select any Microsoft Windows or Linux server that is added to your backup infrastructure and has internet connection. For more information on how to add such a server to your environment, see Adding Microsoft Windows Servers and Adding Linux Servers respectively. By default, the role of a gateway server is assigned to the machine where Veeam Backup & Replication is installed.

If you choose not to use a gateway server, make sure that all scale-out repository extents have direct internet access.

For information on considerations and limitations for using gateway server, see Considerations and Limitations.



Step 3. Specify Object Storage Settings

At the **Bucket** step of the wizard, specify Amazon S3 bucket and folder that will be used to store data:

1. From the **Data center region** drop-down list, select a region.
2. From the **Bucket** drop-down list, select a bucket.

   Make sure that the bucket where you want to store your backup data was created in advance.
3. In the **Folder** field, select a cloud folder to which you want to map your object storage repository.

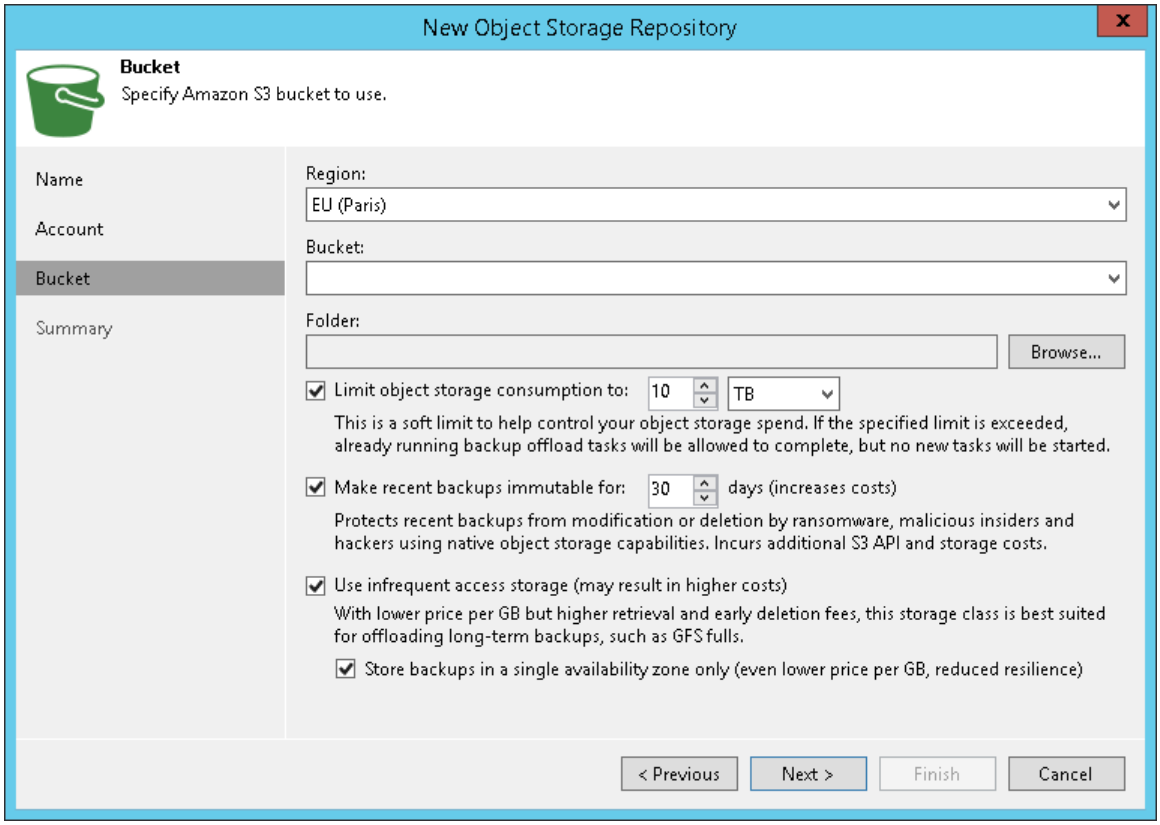
   To select a folder, click **Browse** and either select an existing folder or create a new one by clicking **New Folder**.

> **IMPORTANT**
>
> Veeam Backup & Replication supports specific storage classes. For more information, see Considerations and Limitations.

To define a soft limit that can be exceeded temporarily for your object storage consumption, select the **Limit object storage consumption to** check box and provide the value in TB or PB.

To prohibit deletion of blocks of data from object storage, select the **Make recent backups immutable for** check box and specify the immutability period. For more information about immutability, see Immutability.

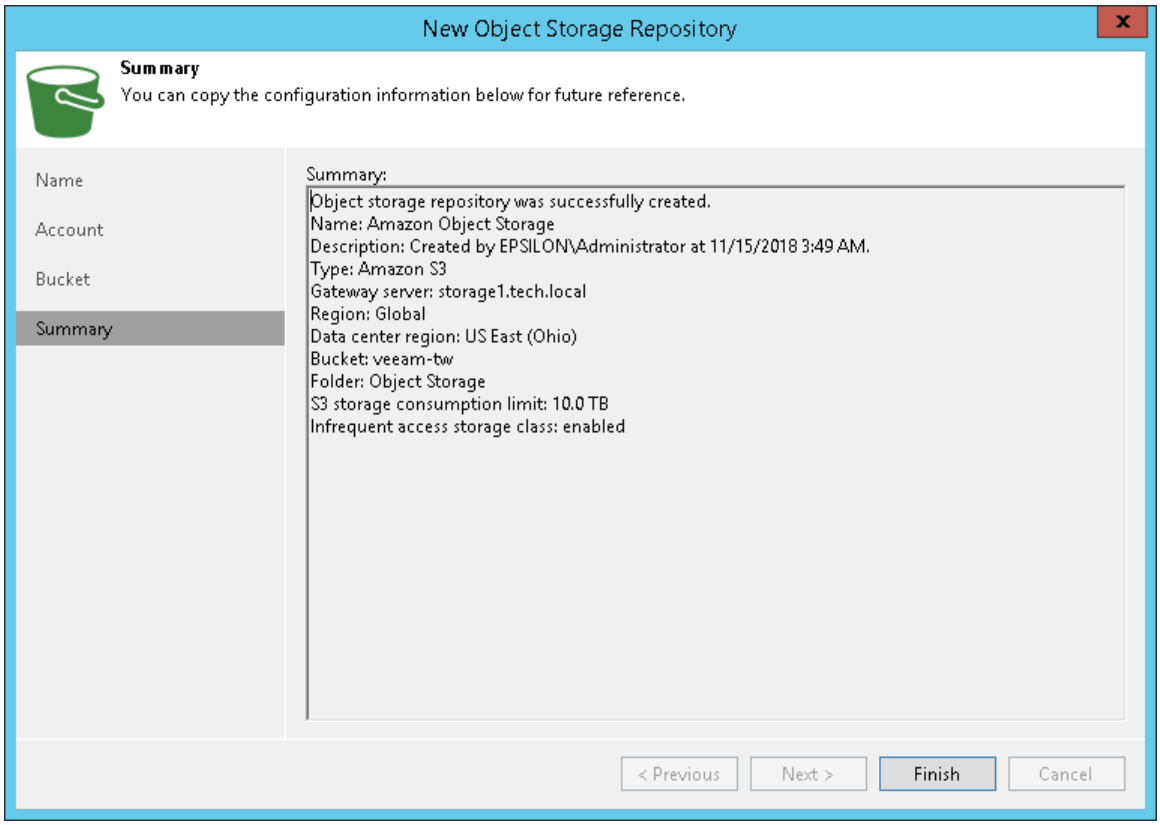
If you plan to access your backup data in an infrequent manner, select the **Use infrequent access storage class** check box to mark each block as *Standard IA* (*Standard Infrequent Access*). To enable *Amazon S3 One Zone-Infrequent Access*, select the **Store backups in a single availability zone only** check box. For more information, see this Amazon article.



Step 4. Finish Working with Wizard

At the **Summary** step of the wizard, complete the object storage repository configuration:

1. Review details of the object storage repository.
2. Click **Finish** to exit the wizard.



Adding Amazon S3 Glacier Storage

Veeam Backup & Replication uses Amazon S3 Glacier object storage with Amazon S3 service. For more information about Amazon S3 storage classes that you can use for this object storage repository, see this Amazon article.

> **NOTE**
>
> Veeam Backup & Replication does not create or use any S3 Glacier vaults in your AWS environment. Glacier vaults is an archive storage solution independent from AWS. It uses storage containers named vaults (opposed to S3 buckets) and its own set of APIs to upload and retrieve data.
>
> Amazon S3 Glacier storage uses S3 APIs to manage data. It also uses S3 storage as a repository for metadata of the Glacier-stored objects. Veeam Backup & Replication assigns the added storage class to backups stored in the repository. That is why the archived backups remain in Amazon S3 and cannot be accessed directly through the Amazon S3 service.

You can only use this repository as an archive extent of the scale-out backup repository. For more information, see Archive Tier.

To add Amazon S3 Glacier object storage, do the following:

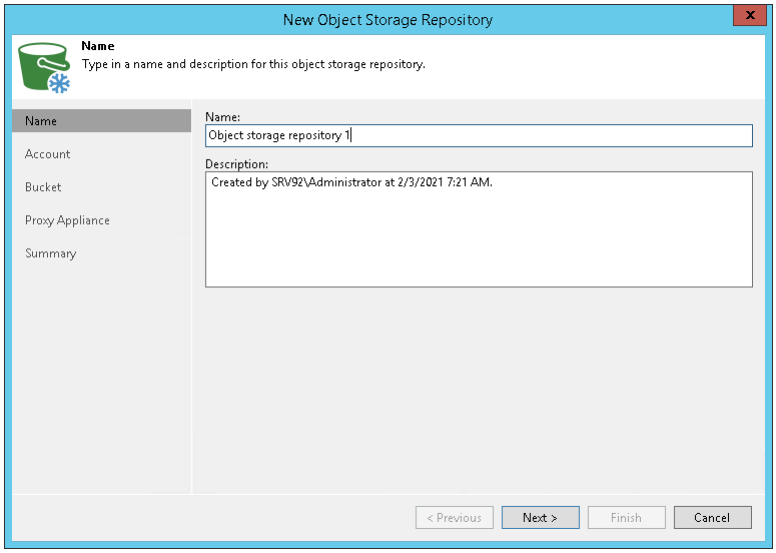
1. Specify object storage name

2. Specify object storage account
3. Specify object storage settings
4. Specify proxy appliance
5. Finish working with the wizard

Step 1. Specify Object Storage Name

At the **Name** step of the wizard, specify a name and description for the object storage repository:

- In the **Name** field, specify a name for the new object storage repository.
- In the **Description** field, enter an optional description. The default description contains information about the user who added the object storage repository, date and time when the object storage repository was added.



Step 2. Specify Object Storage Account

At the **Account** step of the wizard, specify the connection settings:

1. From the **Credentials** drop-down list, select user credentials to access your Amazon S3 object storage.

   If you already have a credentials record that was configured in advance, select it from the drop-down list. Otherwise, click **Add** and provide your access and secret keys, as described in Access Keys for AWS Users. You can also click the **Manage cloud accounts** link to add, edit or remove a credentials record.

   The user account must have permissions listed in Required Permissions.

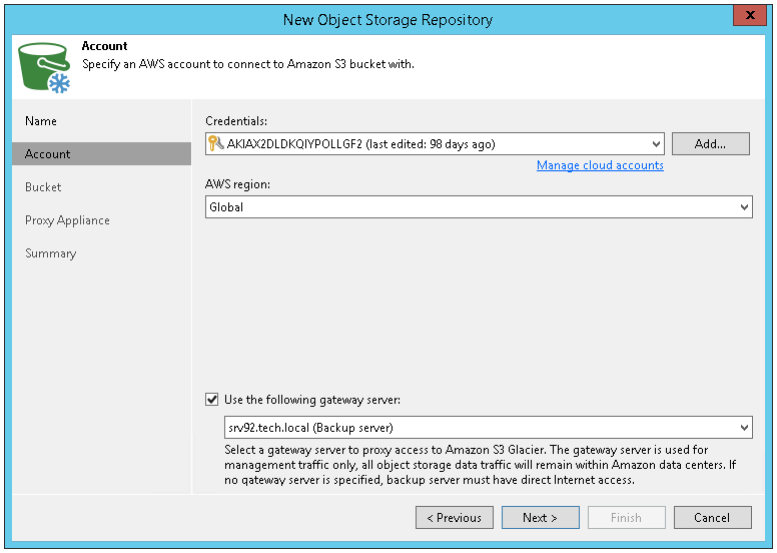
2. From the **AWS region** drop-down list, select the geographical location of the Amazon datacenter.

To use a gateway server, select the **Use the following gateway server** check box and choose an appropriate server from the list. You may want to use a gateway server, for example, if your organization has NAT or different types of firewalls and your access to the internet is limited.

You can select any Microsoft Windows or Linux server that is added to your backup infrastructure and has internet connection. For more information on how to add such a server to your environment, see Adding Microsoft Windows Servers and Adding Linux Servers respectively. By default, the role of a gateway server is assigned to the machine where Veeam Backup & Replication is installed.

If you choose not to use a gateway server, make sure that all scale-out repository extents have direct internet access.

For information on considerations and limitations for using gateway server, see Considerations and Limitations.



Step 3. Specify Object Storage Settings

At the **Bucket** step of the wizard, specify Amazon S3 bucket and folder that will be used to store data:

1. From the **Data center** drop-down list, the AWS region where the Amazon S3 bucket is located.
2. From the **Bucket** drop-down list, select a bucket where you want to store your backup data.

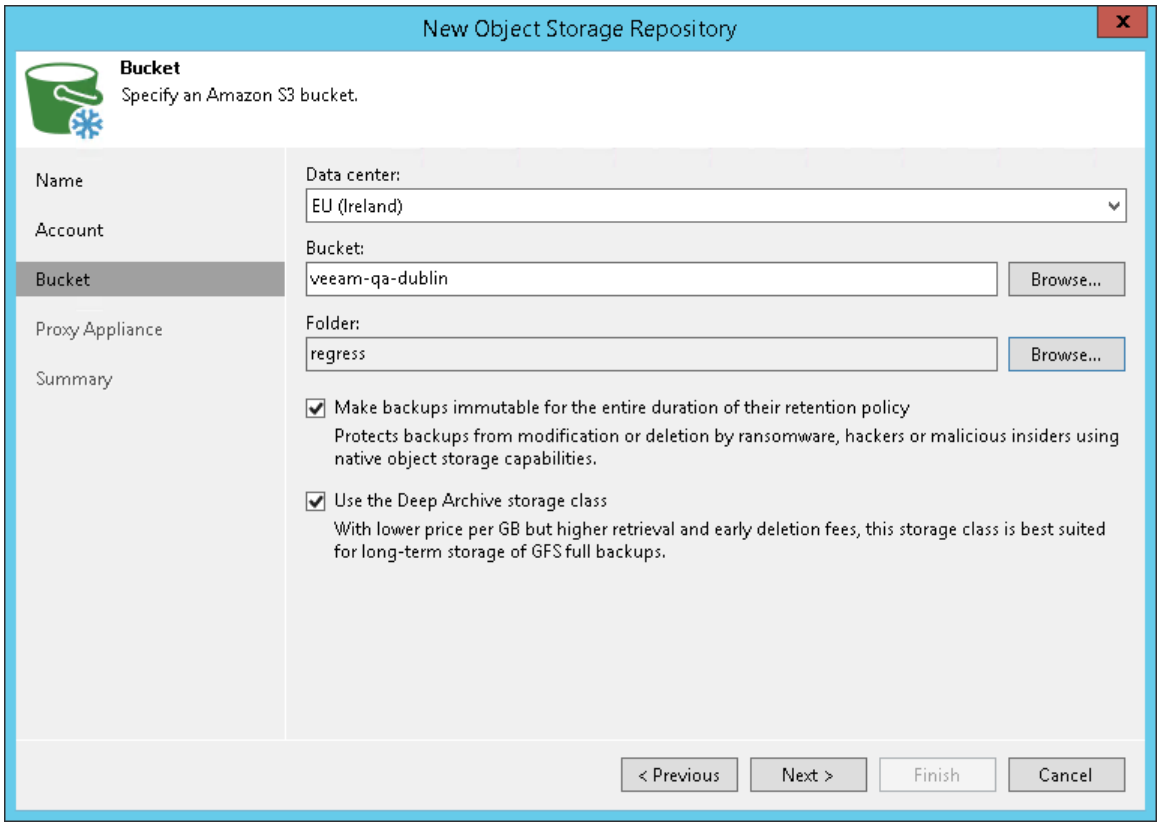
   Make sure that the bucket was created in advance.
3. In the **Folder** field, select a cloud folder to which you want to map your object storage repository and which will be used to store data.

   To select a folder, click **Browse** and either select an existing folder or create a new one by clicking **New Folder**.

To prohibit deletion of blocks of data from object storage, select the **Make backups immutable for the entire duration of their retention policy** check box. The immutability period will be equal to the retention period (if any) of the data blocks. All the types of files that are eligible for archive storage can be made immutable. For more information on the immutability feature and the retention policy for each file type, see Immutability for Archive Tier.

Keep in mind that in order to use immutability, you need to enable the *Object Lock* and *Versioning* features on your S3 bucket at the time when you create the bucket. For more information, see Preparing to Use Immutability.

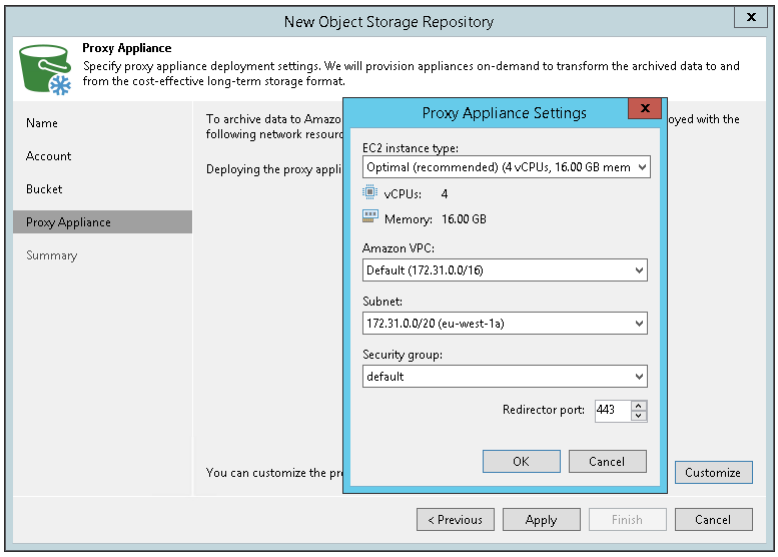
If you plan to access your backup data in an infrequent manner, on average once or twice in a year, select the **Use Glacier Deep Archive for backups with the retention policy over 180 days** check box. For more information on Amazon storage classes, see this Amazon article. For more information on Amazon Glacier Deep Archive, see this Amazon article.



Step 4. Specify Proxy Appliance

At the **Proxy Appliance** step of the wizard, you can specify the proxy appliance settings. A proxy appliance is an auxiliary instance that is necessary for transferring the data from Amazon S3 to Amazon S3 Glacier. You can specify proxy appliance settings in one of the following ways:

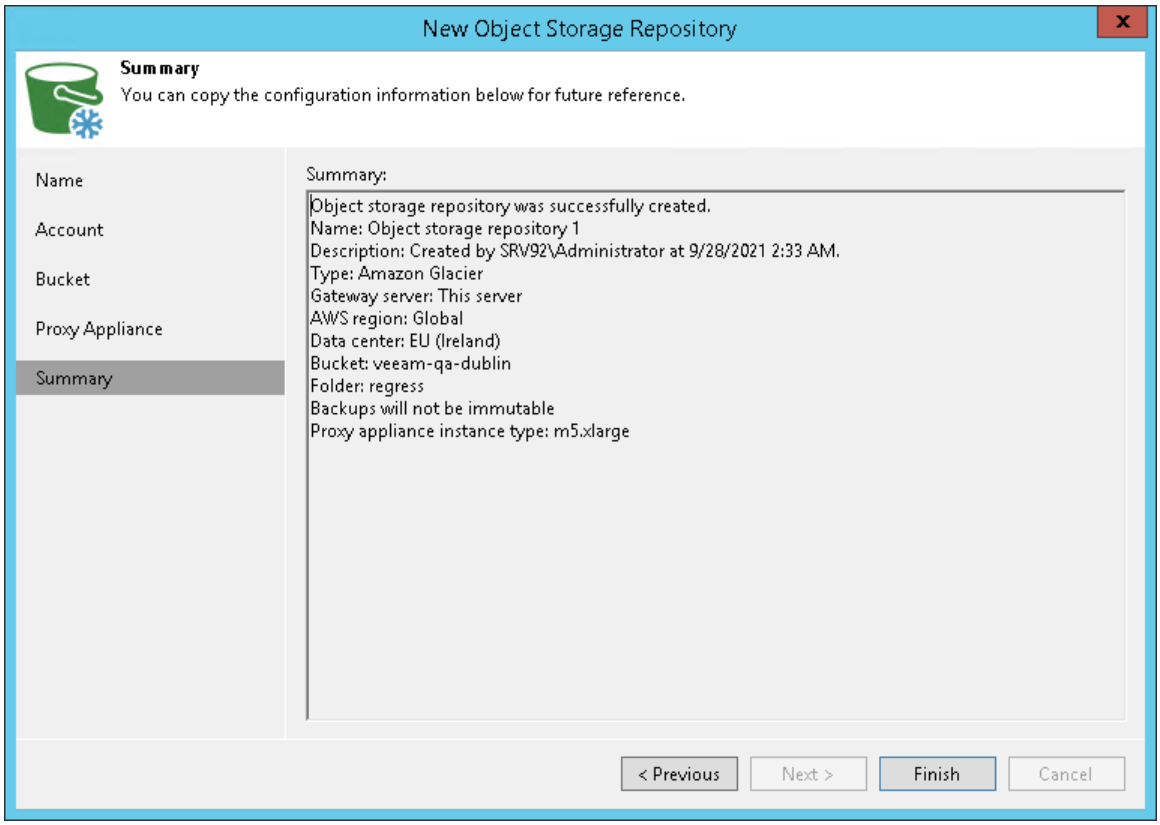
- Use the default settings. In that case, Veeam Backup & Replication will select the necessary settings from those available in your account, or will create for you new settings for the EC2 instance type, Amazon VPC, subnet, security group and redirector port.
- Specify the settings manually:
  a. Click **Customize**.
  b. From the **EC2 instance type** drop-down list, select the instance type for the proxy appliance. The EC2 instance type affects the speed and the cost of transferring the backup files to the Archive Tier of a scale-out backup repository. For information on instance types, see this Amazon article.
  c. From the **Amazon VPC** drop-down list, select the Amazon VPC where Veeam Backup & Replication will launch the target instance. For information on the Amazon VPC, see this Amazon article.
  d. From the **Subnet** drop-down list, select the subnet for the proxy appliance.
  e. From the **Security group** drop-down list, select a security group that will be associated with the proxy appliance. For information on security groups for Amazon VPC, see this Amazon article.
  f. In the **Redirector port** field, specify the TCP port that Veeam Backup & Replication will use to route requests between the proxy appliance and backup infrastructure components.
  g. Click **OK** to finish configuring the settings.



Step 5. Finish Working with Wizard

At the **Summary** step of the wizard, complete the procedure of object storage repository configuration:

1. Review details of the object storage repository.
2. Click **Finish** to exit the wizard.



Adding AWS Snowball Edge Storage

AWS Snowball Edge is a physical device that is requested for a short period of time from Amazon. It can be temporarily attached to your Veeam Backup & Replication backup infrastructure and used as an object storage. For more information about ordering AWS Snowball Edge and preparing to use it, see this Amazon article.

This device might become useful when you need to offload a large number of backup files that are occupying a lot of storage space on your extents, as offloading data to the AWS Snowball Edge device might occur a lot faster that transferring the same amount of data directly to Amazon object storage.

Once you have offloaded backups to AWS Snowball Edge, you need to ship the device back to Amazon for further data synchronization with your storage account, as described in Seeding Backups to AWS Snowball Edge Storage.

IMPORTANT

To eliminate any risk of data loss in case the device is damaged during the shipping, we recommend to copy your backup files to AWS Snowball Edge instead of moving them, and to keep them in your local infrastructure until the device is delivered and your data is safely transferred into your storage account.

NOTE

For information on FIPS status of AWS Snowball Edge Storage, see Amazon official updates.

To add AWS Snowball Edge storage, do the following:

1. Specify object storage name
2. Specify object storage account
3. Specify object storage settings
4. Finish working with the wizard

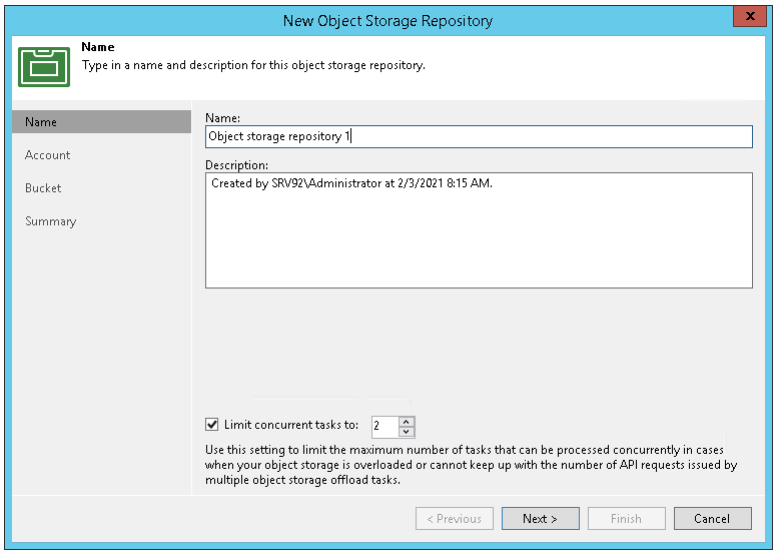
Step 1. Specify Object Storage Name

At the **Name** step of the wizard, specify a name and description for the object storage repository:

- In the **Name** field, specify a name for the new object storage repository.
- In the **Description** field, enter an optional description. The default description contains information about the user who added the object storage repository, date and time when the object storage repository was added.

If you wish to limit the maximum number of tasks that can be processed at once, select the **Limit concurrent tasks to N** check box.

> **TIP**
>
> Set the maximum number of concurrent task to a reasonable number to avoid overloading if you plan to upload significant amount of backup chains to the device.



Step 2. Specify Object Storage Account

At the **Account** step of the wizard, specify the connection settings:

1. In the **Service point** field, specify a service point address of your AWS Snowball Edge device and define a port. To learn what port you need to define, see AWS Snowball Edge Developer Guide, section Ports Required to Use AWS Services on an AWS Snowball Edge Device. Otherwise, the default 443 port is used.
2. In the **Region** field, type "snow".
3. From the **Credentials** drop-down list, select user credentials to access your AWS Snowball Edge storage.

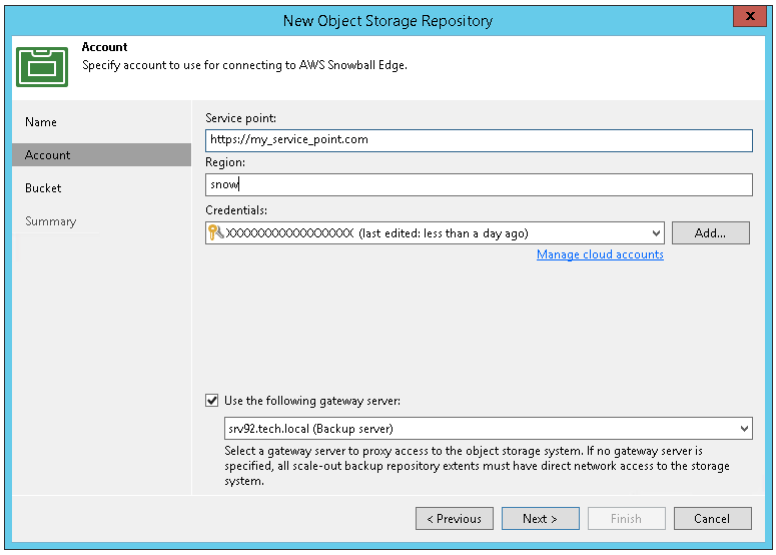
   If you already have a credentials record that was configured in advance, select it from the drop-down list. Otherwise, click **Add** and provide your access and secret keys, as described in Access Keys for AWS Users. You can also click the **Manage cloud accounts** link to add, edit or remove a credentials record.

To use a gateway server, select the **Use the following gateway server** check box and choose an appropriate server from the list. You may want to use a gateway server, for example, if your organization has NAT or different types of firewalls and your access to the internet is limited.

You can select any Microsoft Windows or Linux server that is added to your backup infrastructure and has internet connection. For more information on how to add such a server to your environment, see Adding Microsoft Windows Servers and Adding Linux Servers respectively. By default, the role of a gateway server is assigned to the machine where Veeam Backup & Replication is installed.

If you choose not to use a gateway server, make sure that all scale-out repository extents have direct internet access.

For information on considerations and limitations for using gateway server, see Considerations and Limitations.



Step 3. Specify Object Storage Settings

At the **Bucket** step of the wizard, specify Amazon S3 bucket and folder that will be used to store data:

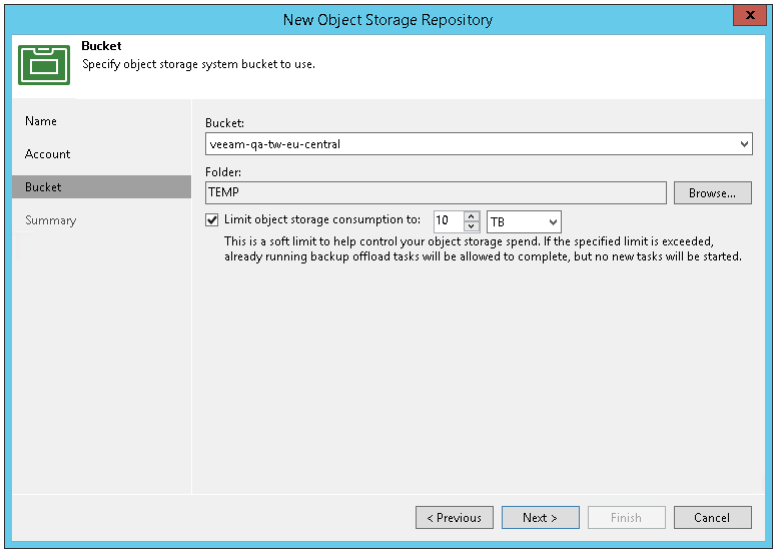
1. From the **Bucket** drop-down list, select a bucket.

   Make sure that the bucket where you want to store your backup data was created in advance.

2. In the **Folder** field, select a cloud folder to which you want to map your object storage repository.

   To select a folder, click **Browse** and either select an existing folder or create a new one by clicking **New Folder**.

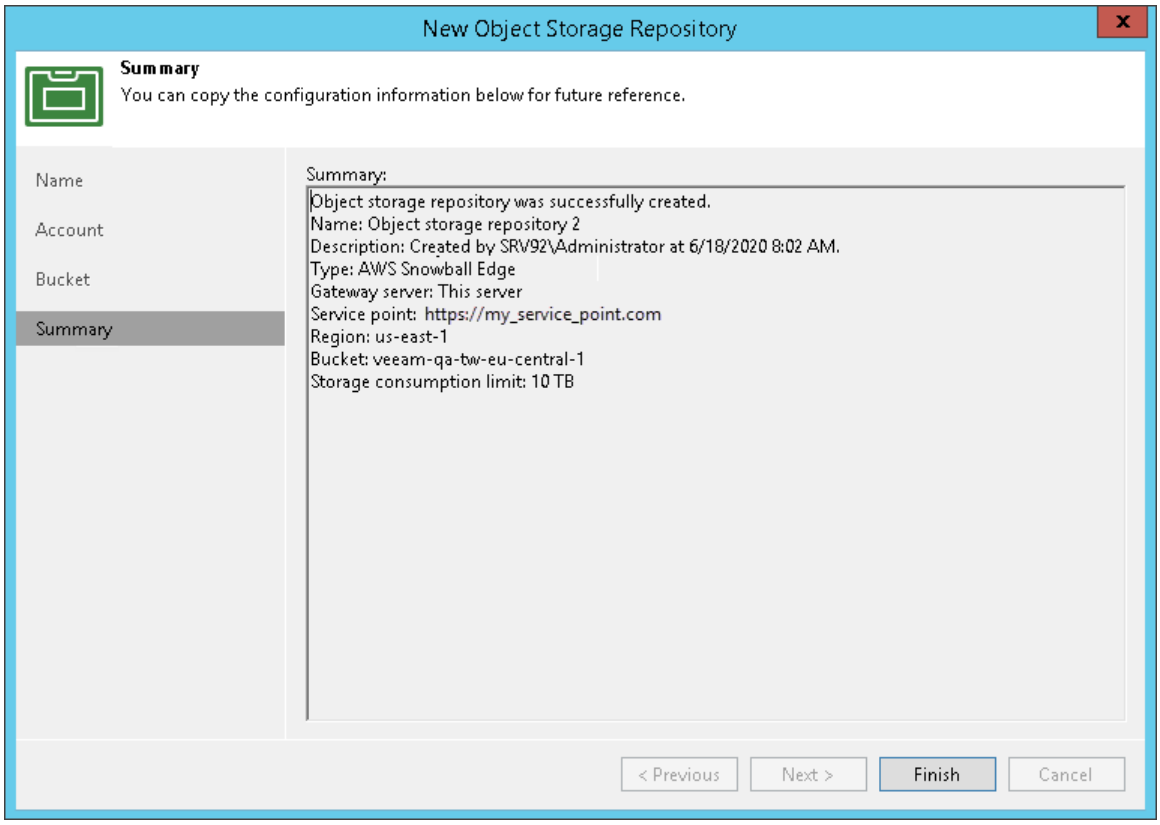
To define a soft limit that can be exceeded temporarily for your object storage consumption, select the **Limit object storage consumption to** check box and provide the value in TB or PB.



Step 4. Finish Working with Wizard

At the **Summary** step of the wizard, complete the object storage repository configuration:

1. Review details of the object storage repository.
2. Click **Finish** to exit the wizard.



Adding Google Cloud Object Storage

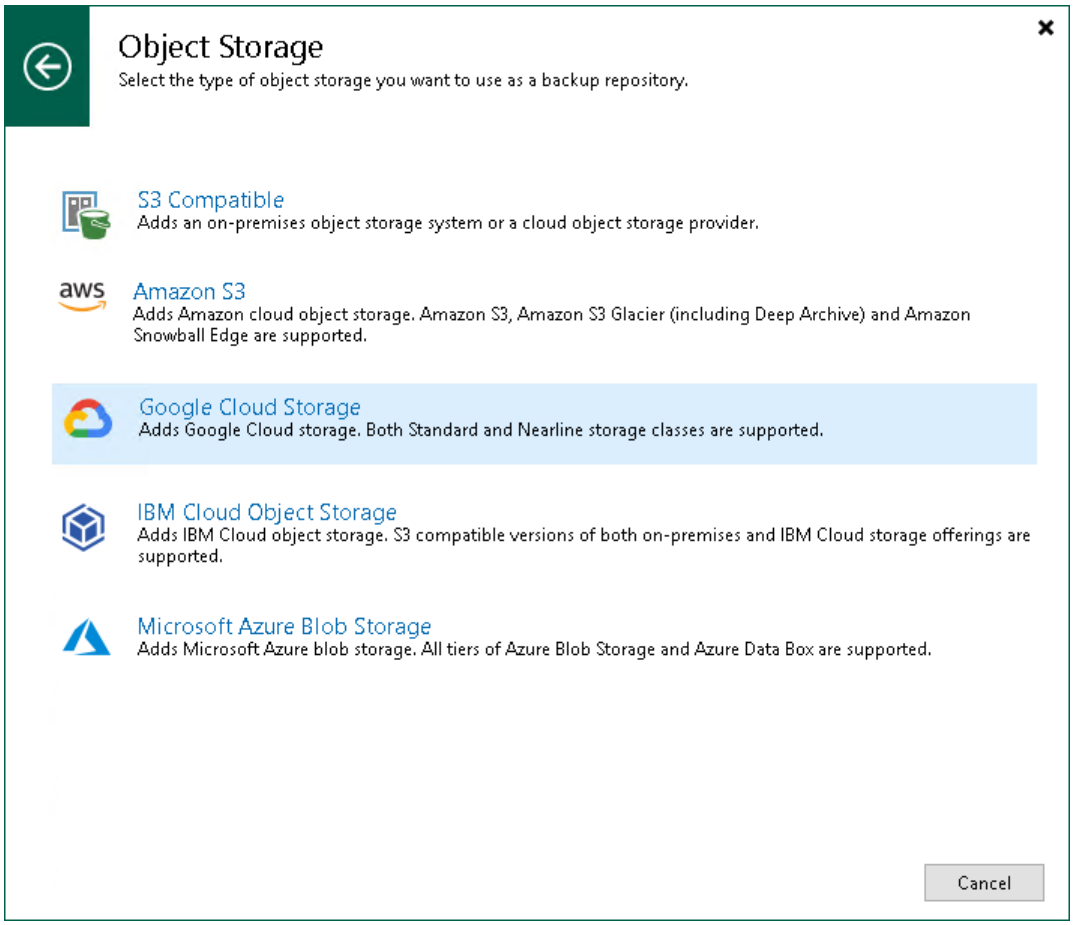
For more information about Google Cloud Object Storage, see this Google article.

Before you add a Google Cloud object storage to the backup infrastructure, check prerequisites. After that, use the **New Object Repository** wizard.

Step 1. Launch New Object Repository Wizard

To launch the **New Object Repository** wizard, do one of the following:

- Open the **Backup Infrastructure** view. In the inventory pane, select the **Backup Repositories** node and click **Add Repository** on the ribbon. In the **Add Backup Repository** dialog, select **Object Storage > Google Cloud Storage**.
- Open the **Backup Infrastructure** view. In the inventory pane, right-click the **Backup Repositories** node and select **Add Backup Repository**. In the **Add Backup Repository** dialog, select **Object Storage > Google Cloud Storage**.

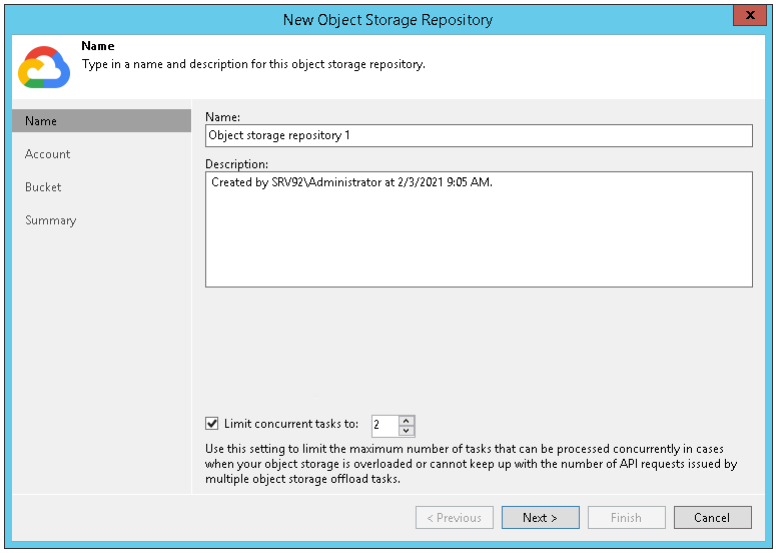


Step 2. Specify Object Storage Name

At the **Name** step of the wizard, specify a name and description for the object storage repository:

- In the **Name** field, specify a name for the new object storage repository.
- In the **Description** field, enter an optional description. The default description contains information about the user who added the object storage repository, date and time when the object storage repository was added.

If you wish to limit the maximum number of tasks that can be processed at once, select the **Limit concurrent tasks to N** check box.



Step 3. Specify Object Storage Account

At the **Account** step of the wizard, specify credentials: from the **Credentials** drop-down list, select user credentials to access your Google Cloud object storage.

If you already have a credentials record that was configured in advance, select it from the drop-down list. Otherwise, click **Add** and provide your access and secret keys, as described in Cloud Credentials Manager. You can also click the **Manage cloud accounts** link to add, edit or remove a credentials record.

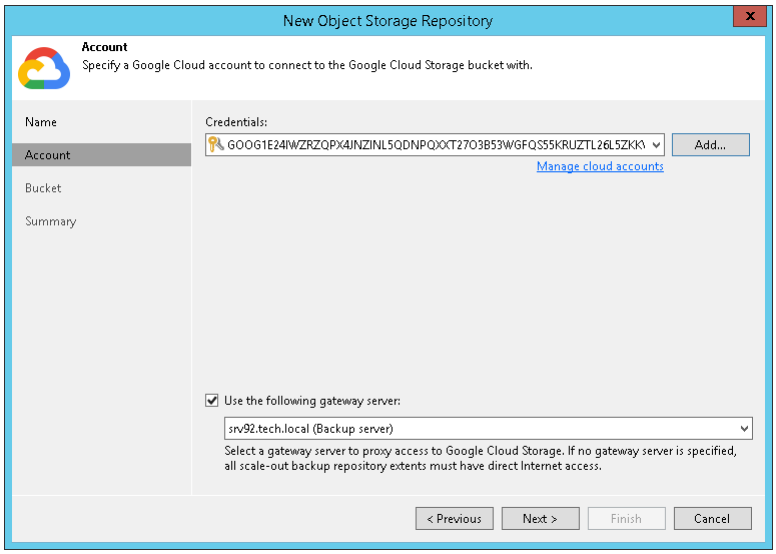
The user account must have permissions listed in Required Permissions.

To use a gateway server, select the **Use the following gateway server** check box and choose an appropriate server from the list. You may want to use a gateway server, for example, if your organization has NAT or different types of firewalls and your access to the internet is limited.

You can select any Microsoft Windows or Linux server that is added to your backup infrastructure and has internet connection. For more information on how to add such a server to your environment, see Adding Microsoft Windows Servers and Adding Linux Servers respectively. By default, the role of a gateway server is assigned to the machine where Veeam Backup & Replication is installed.

If you choose not to use a gateway server, make sure that all scale-out repository extents have direct internet access.

For information on considerations and limitations for using gateway server, see Considerations and Limitations.



Step 4. Specify Object Storage Settings

At the **Bucket** step of the wizard, specify the bucket and folder that will be used to store data:

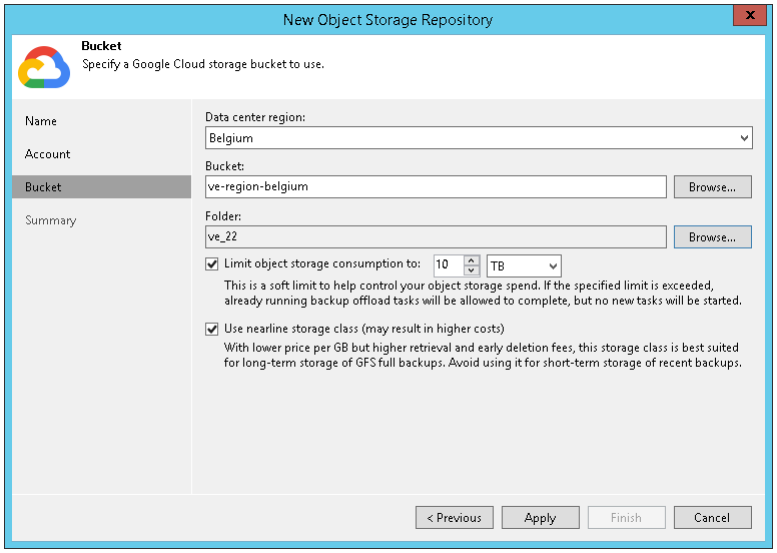
1. In the **Data center region** drop-down list, select a region.
2. In the **Bucket** drop-down list, select a bucket.

   Make sure that the bucket where you want to store your backup data was created in advance.
3. In the **Select Folder** field, select a cloud folder to which you want to map your object storage repository.

   To select a folder, click **Browse** and either select an existing folder or create a new one by clicking **New Folder**.

To define a soft limit that can be exceeded temporarily for your object storage consumption, select the **Limit object storage consumption to** check box and provide the value in TB or PB.

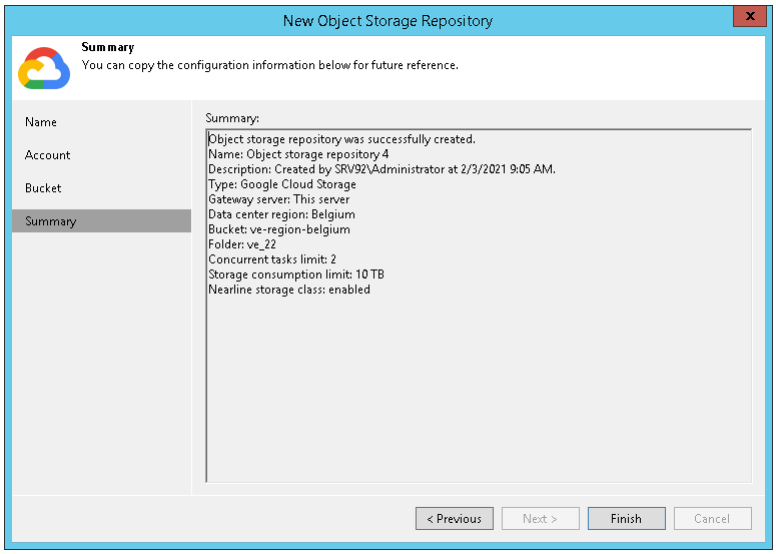
If you plan to access your backup data rarely, on average once per month or less, select the **Use nearline storage class** check box. For more information about Google Cloud storage classes, see this Google article.



Step 5. Finish Working with Wizard

At the **Summary** step of the wizard, complete the object storage repository configuration:

1. Review details of the object storage repository.
2. Click **Finish** to exit the wizard.



Adding IBM Cloud Object Storage

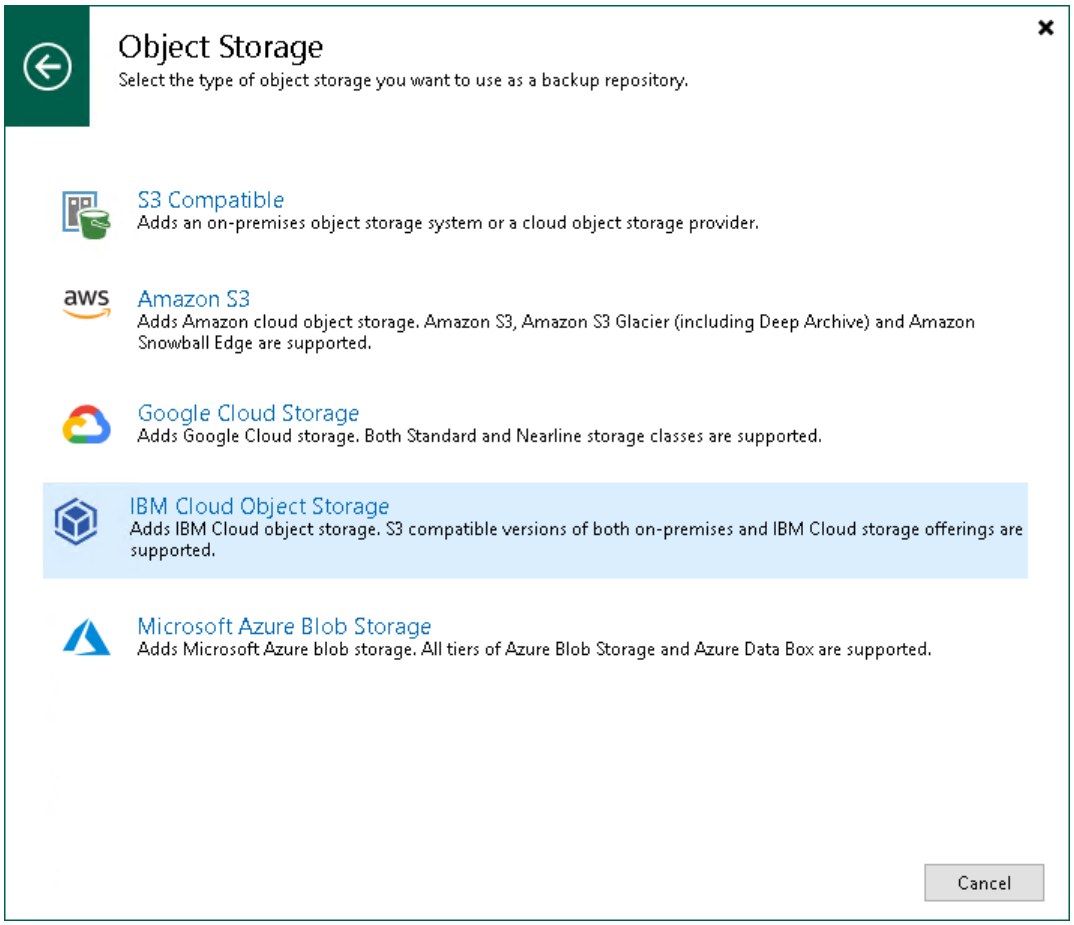
For more information about IBM Cloud object storage, see this IBM article.

Before you add an IBM Cloud object storage to the backup infrastructure, check prerequisites. After that, use the **New Object Repository wizard**.

Step 1. Launch New Object Repository Wizard

To launch the **New Object Repository** wizard, do one of the following:

- Open the **Backup Infrastructure** view. In the inventory pane select the **Backup Repositories** node and click **Add Repository** on the ribbon. In the **Add Backup Repository** dialog, select **Object Storage > IBM Cloud Object Storage**.
- Open the **Backup Infrastructure** view. In the inventory pane right-click the **Backup Repositories** node and select **Add Backup Repository**. In the **Add Backup Repository** dialog, select **Object Storage > IBM Cloud Object Storage**.

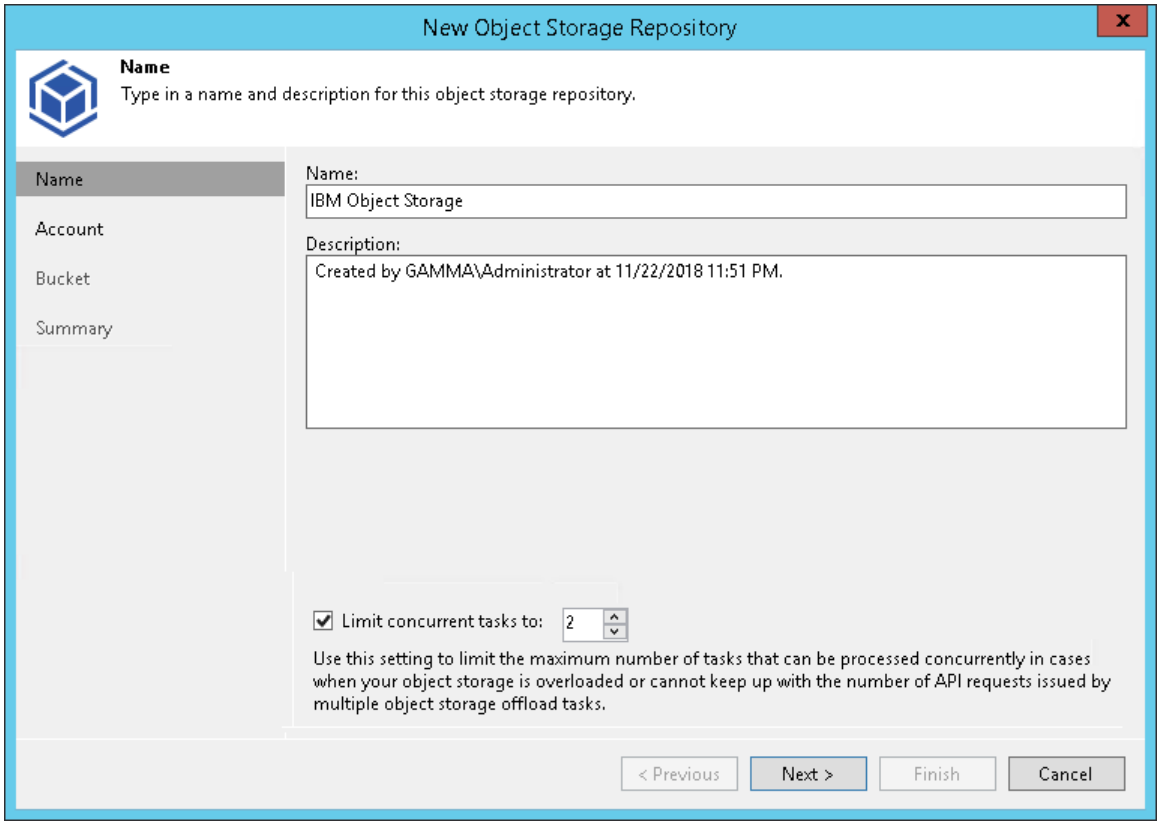


Step 2. Specify Object Storage Name

At the **Name** step of the wizard, specify a name and description for the object storage repository:

- In the **Name** field, specify a name for the new object storage repository.
- In the **Description** field, enter an optional description. The default description contains information about the user who added the object storage repository, date and time when the object storage repository was added.

If you wish to limit the maximum number of tasks that can be processed at once, select the **Limit concurrent tasks to N** check box.



Step 3. Specify Object Storage Account

At the **Account** step of the wizard, specify the connection settings:

1. In the **Service point** field, specify a service point address of your IBM cloud object storage.
2. In the **Region** field, specify a region.
3. From the **Credentials** drop-down list, select user credentials to access your IBM cloud object storage.

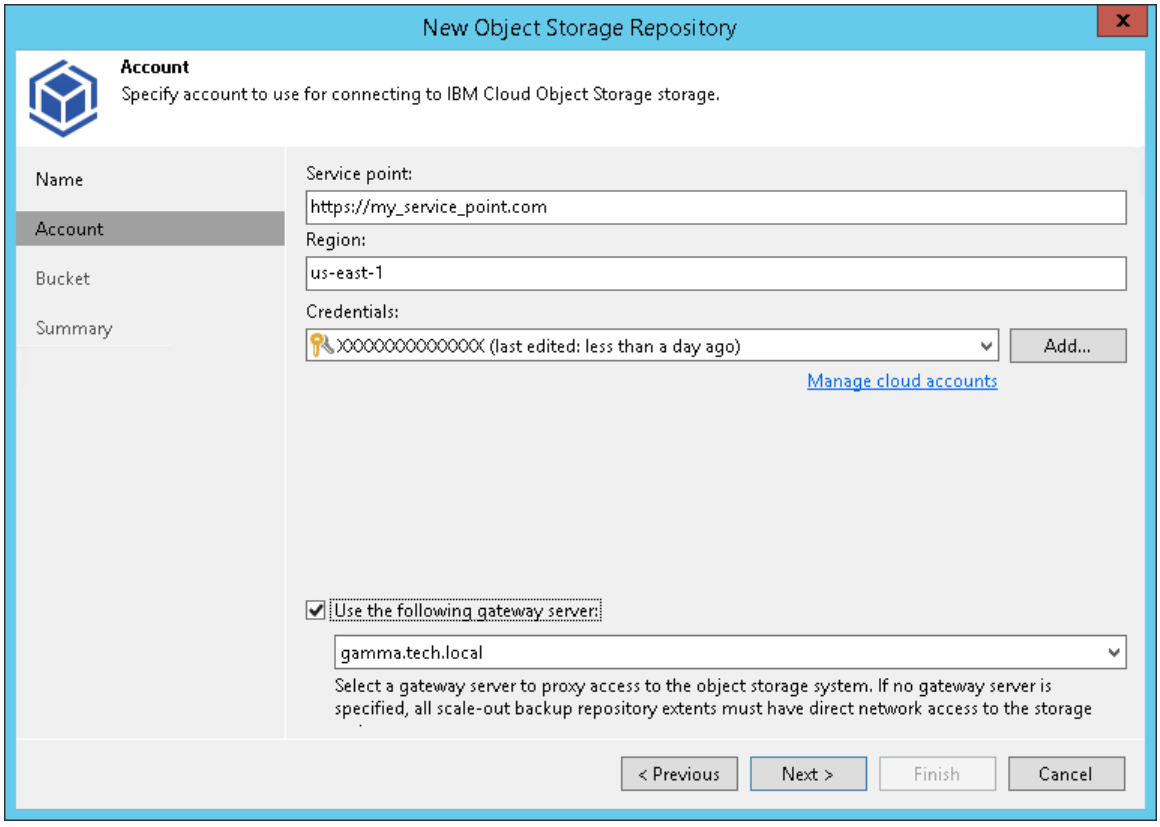
   If you already have a credentials record that was configured in advance, select such a record from the drop-down list. Otherwise, click **Add** and provide your access and secret keys, as described in Cloud Credentials Manager. You can also click the **Manage cloud accounts** link to add, edit or remove a credentials record.

To use a gateway server, select the **Use the following gateway server** check box and choose an appropriate server from the list. You may want to use a gateway server, for example, if your organization has NAT or different types of firewalls and your access to the internet is limited.

You can select any Microsoft Windows or Linux server that is added to your backup infrastructure and has internet connection. For more information on how to add such a server to your environment, see Adding Microsoft Windows Servers and Adding Linux Servers respectively. By default, the role of a gateway server is assigned to the machine where Veeam Backup & Replication is installed.

If you choose not to use a gateway server, make sure that all scale-out repository extents have direct internet access.

For information on considerations and limitations for using gateway server, see Considerations and Limitations.



Step 4. Specify Object Storage Settings

At the **Bucket** step of the wizard, specify the bucket and folder that will be used to store data:

1. From the **Bucket** drop-down list, select a bucket.

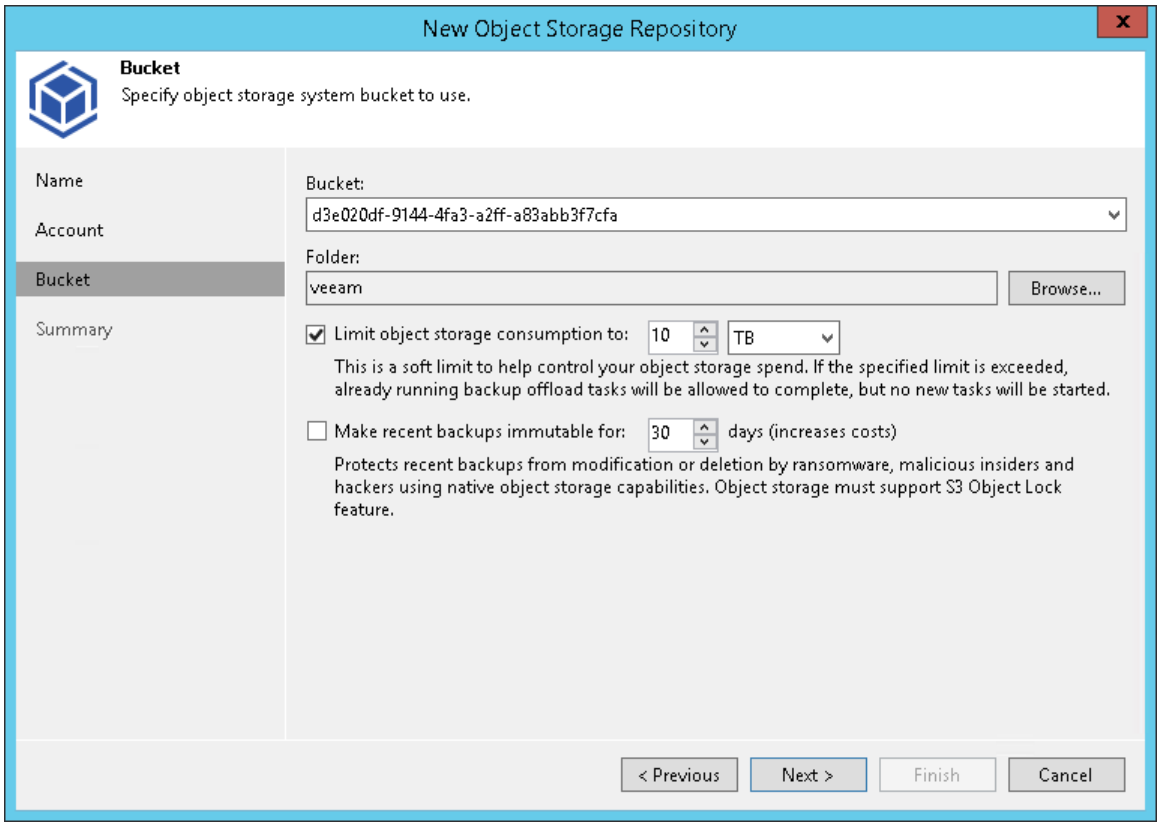
   Make sure that the bucket where you want to store your backup data was created in advance.

2. In the **Select Folder** field, select a cloud folder to which you want to map your object storage repository.

   To select a folder, click **Browse** and either select an existing folder or create a new one by clicking **New Folder**.

To define a soft limit that can be exceeded temporarily for your object storage consumption, select the **Limit object storage consumption to** check box and provide the value in TB or PB.

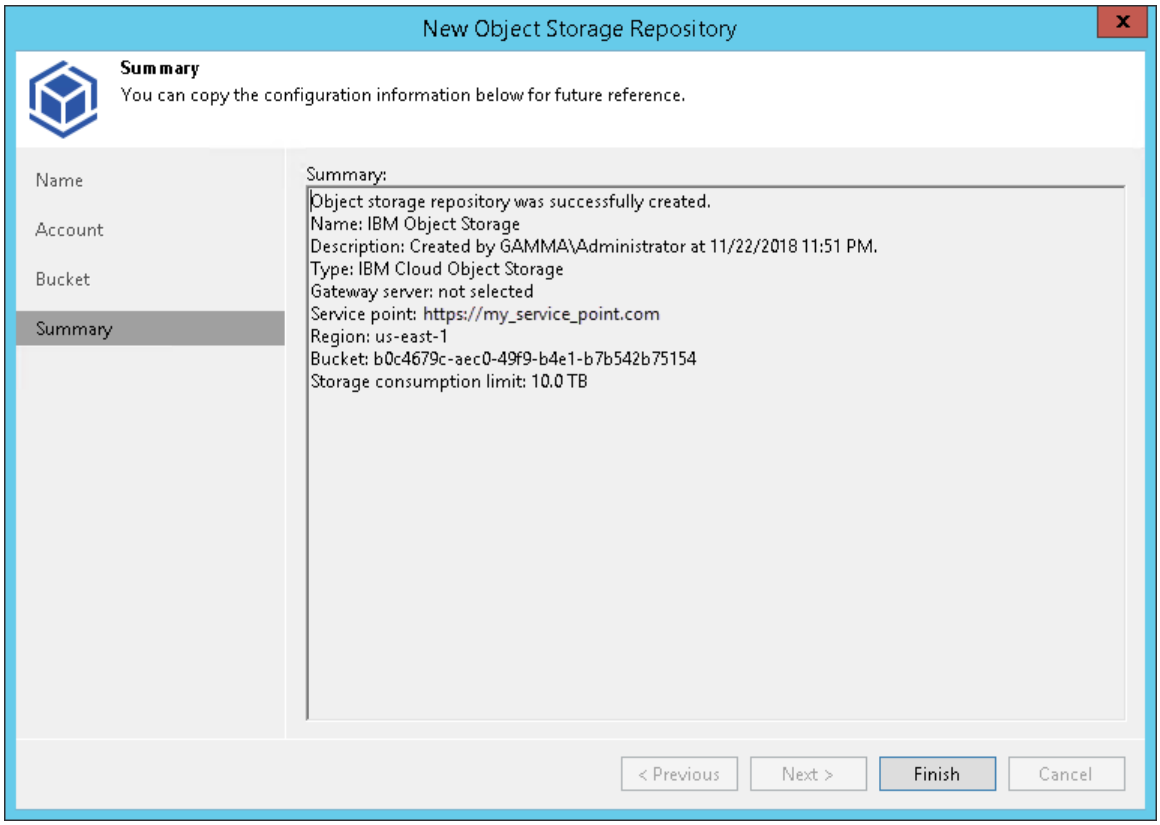
> **NOTE**
>
> The **Make recent backups immutable for** check box is currently unavailable for IBM Cloud Object Storage, since IBM does not currently support S3 Object Lock feature required for the immutability.



Step 5. Finish Working with Wizard

At the **Summary** step of the wizard, complete the object storage repository configuration:

1. Review details of the object storage repository.
2. Click **Finish** to exit the wizard.



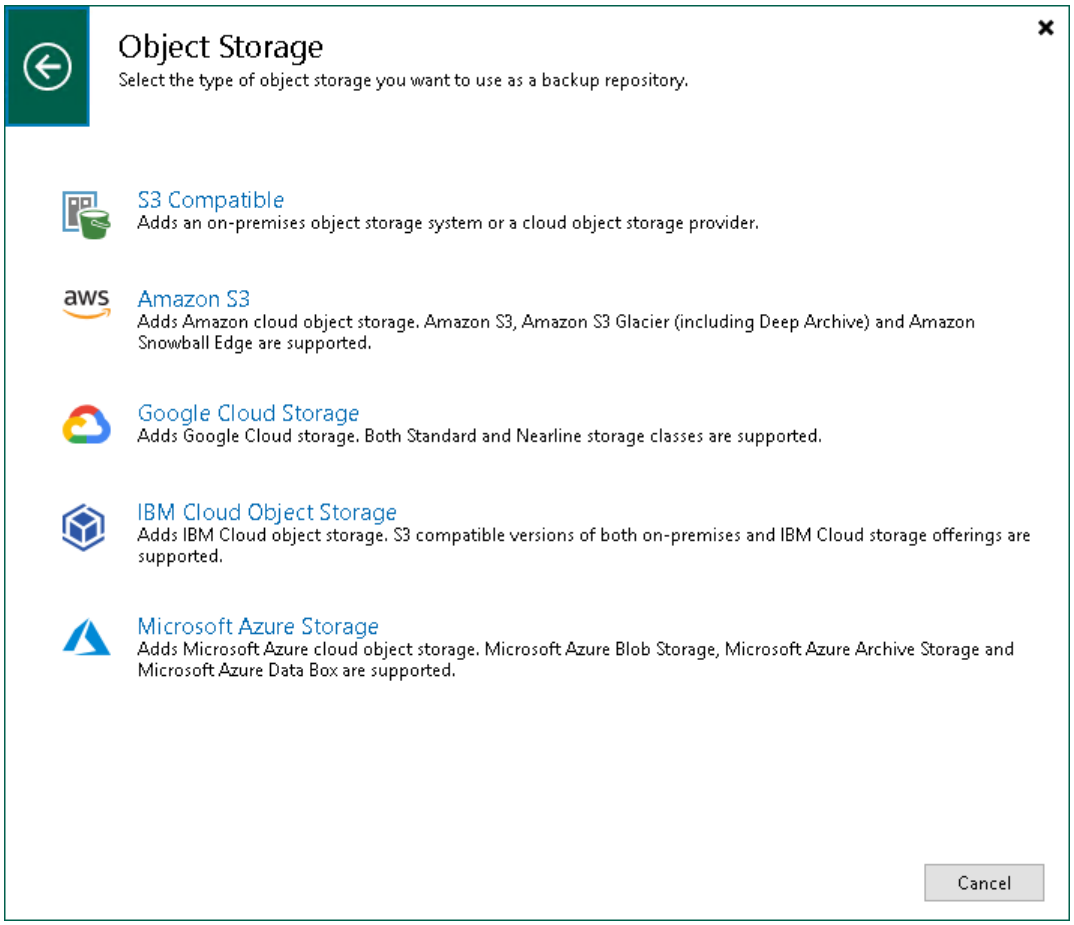
Adding Microsoft Azure Blob Storage, Microsoft Azure Archive Storage and Microsoft Azure Data Box

Before you add a Microsoft Azure Blob storage, Azure Archive storage or Azure Data Box to the backup infrastructure, check prerequisites. After that, use the **New Object Repository wizard**.

Step 1. Launch New Object Repository Wizard

To launch the **New Object Repository** wizard, do one of the following:

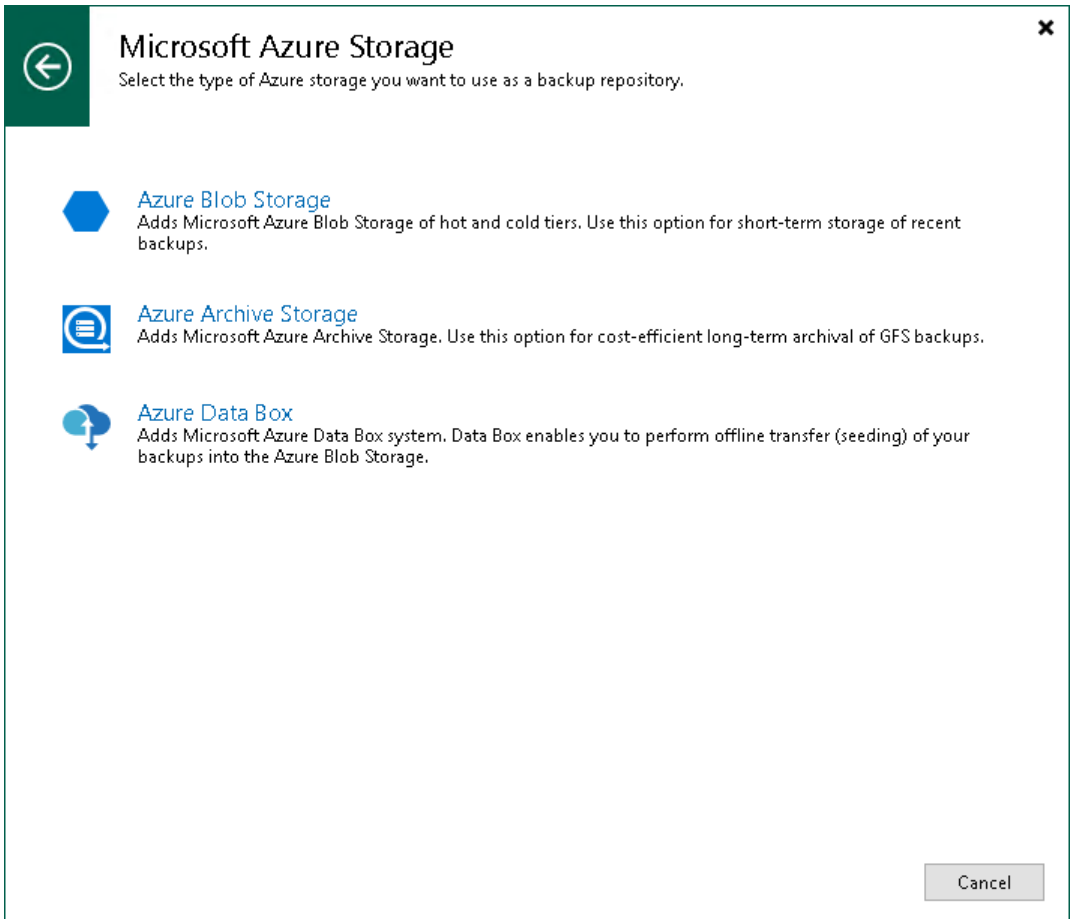
- Open the **Backup Infrastructure** view. In the inventory pane select the **Backup Repositories** node and click **Add Repository** on the ribbon. In the **Add Backup Repository** dialog, select **Object Storage > Microsoft Azure Storage**.
- Open the **Backup Infrastructure** view. In the inventory pane right-click the **Backup Repositories** node and select **Add Backup Repository**. In the **Add Backup Repository** dialog, select **Object Storage > Microsoft Azure Storage**.



Step 2. Select Azure Storage Type

At the **Microsoft Azure Storage** step of the wizard, select one of the following storage types:

- Azure Blob Storage
- Azure Archive Storage
- Azure Data Box



Adding Azure Blob Storage

For more information about Azure Blob storage, see this Microsoft article.

To add Microsoft Azure Blob storage, do the following:

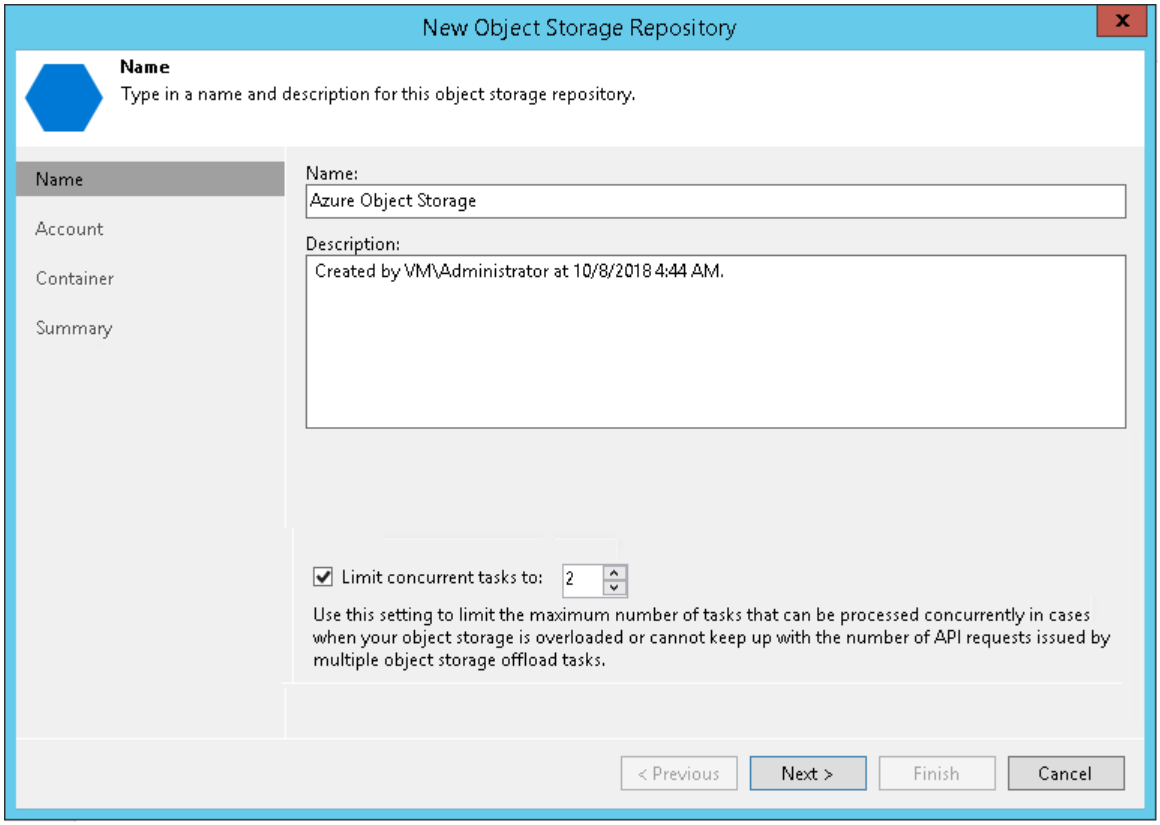
1. Specify object storage name
2. Specify object storage account
3. Specify object storage settings
4. Finish working with the wizard

Step 1. Specify Object Storage Name

At the **Name** step of the wizard, specify a name and description for the object storage repository:

- In the **Name** field, specify a name for the new object storage repository.
- In the **Description** field, enter an optional description. The default description contains information about the user who added the object storage repository, date and time when the object storage repository was added.

If you wish to limit the maximum number of tasks that can be processed at once, select the **Limit concurrent tasks to N** check box.



Step 2. Specify Object Storage Account

At the **Account** step of the wizard, specify the connection settings:

1. From the **Credentials** drop-down list, select user credentials to access your Azure Blob storage.

   If you already have a credentials record that was configured in advance, select such a record from the drop-down list. Otherwise, click **Add** and provide your account name and a shared key, as described in Microsoft Azure Storage Accounts. You can also click the **Manage cloud accounts** link to add, edit or remove a credentials record.

2. From the **Region** drop-down list, select an Azure region.

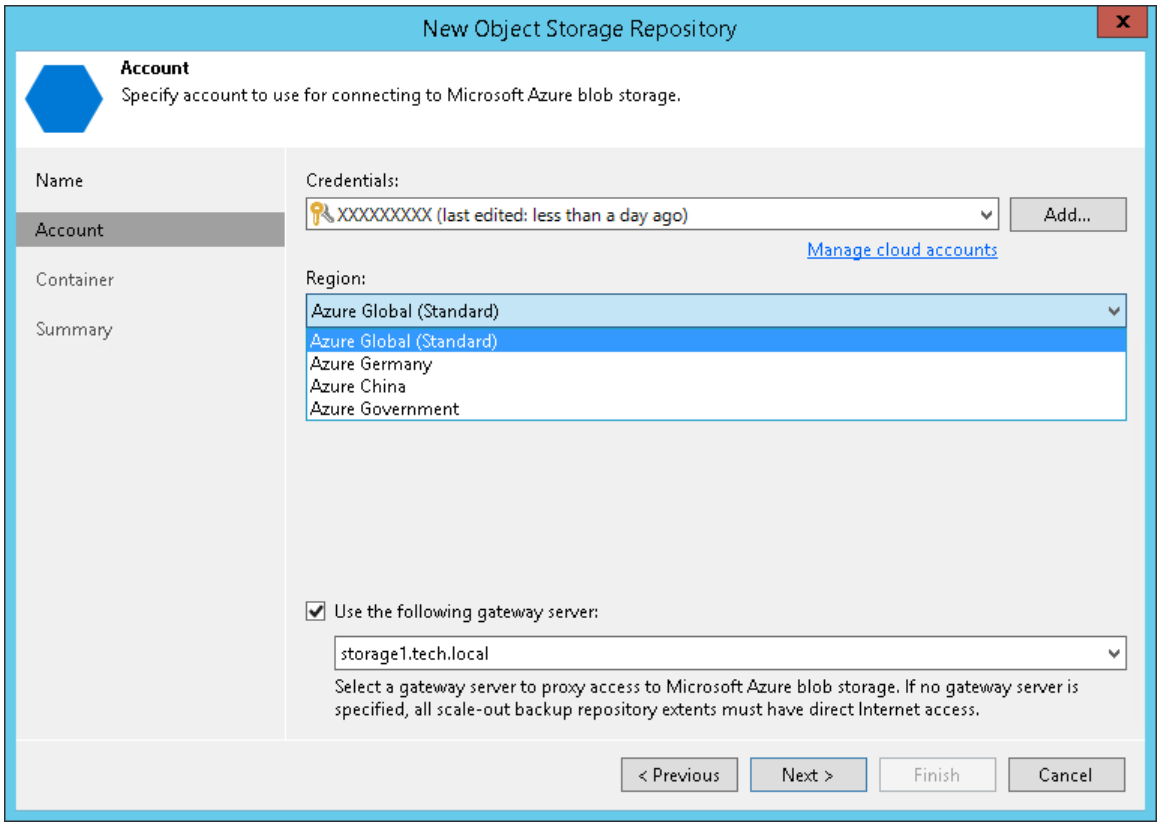
> **IMPORTANT**
>
> Veeam Backup & Replication can no longer support Azure Germany since Microsoft closed Microsoft Cloud Deutschland.

To use a gateway server, select the **Use the following gateway server** check box and choose an appropriate server from the list. You may want to use a gateway server, for example, if your organization has NAT or different types of firewalls and your access to the internet is limited.

You can select any Microsoft Windows or Linux server that is added to your backup infrastructure and has internet connection. For more information on how to add such a server to your environment, see Adding Microsoft Windows Servers and Adding Linux Servers respectively. By default, the role of a gateway server is assigned to the machine where Veeam Backup & Replication is installed.

If you choose not to use a gateway server, make sure that all scale-out repository extents have direct internet access.

For information on considerations and limitations for using gateway server, see Considerations and Limitations.



Step 3. Specify Object Storage Settings

At the **Container** step of the wizard, specify the container and folder that will be used to store data:

1. From the **Container** drop-down list, select a container.

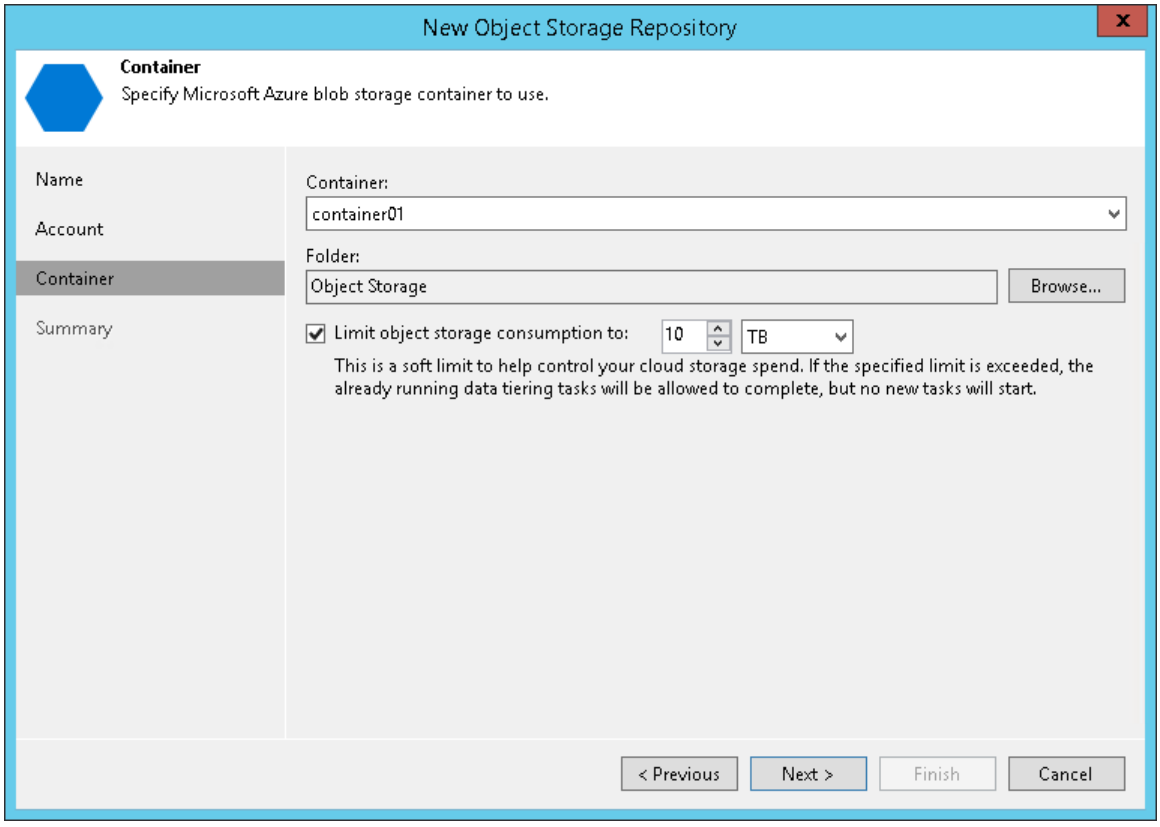
   Make sure that the container where you want to store your backup data was created in advance.

> **NOTE**
>
> The default *Root* container is not supported. For more information about this container, see this Microsoft article.

2. In the **Select Folder** field, select a cloud folder to which you want to map your object storage repository.

   To select a folder, click **Browse** and either select an existing folder or create a new one by clicking **New Folder**.

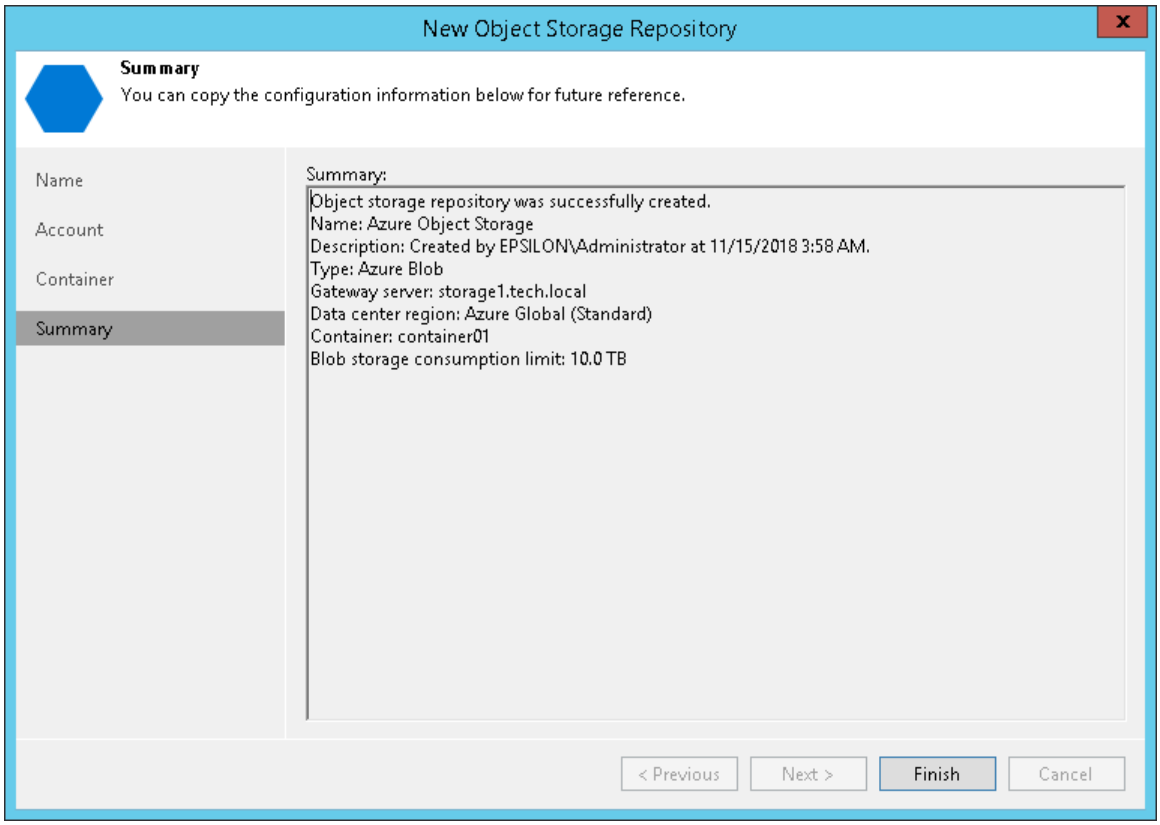
To define a soft limit that can be exceeded temporarily for your object storage consumption, select the **Limit object storage consumption to** check box and provide the value in TB or PB.



Step 4. Finish Working with Wizard

At the **Summary** step of the wizard, complete the object storage repository configuration:

1. Review details of the object storage repository.
2. Click **Finish** to exit the wizard.



Adding Azure Archive Storage

This repository can be used only as an archive extent of the scale-out backup repository.

For more information about Azure Archive Storage, see this Microsoft article.

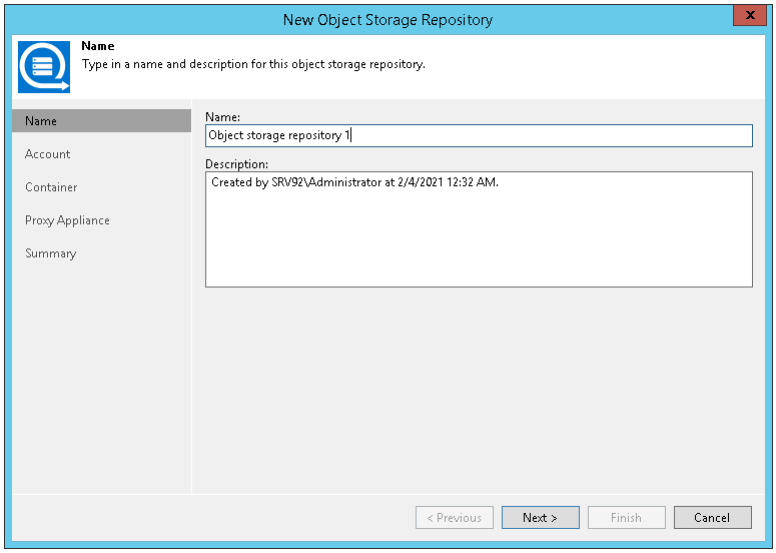
To add Azure Archive Storage, do the following:

1. Specify object storage name
2. Specify object storage account
3. Specify object storage settings
4. Specify proxy appliance
5. Finish working with the wizard

Step 1. Specify Object Storage Name

At the **Name** step of the wizard, specify a name and description for the object storage repository:

- In the **Name** field, specify a name for the new object storage repository.
- In the **Description** field, enter an optional description. The default description contains information about the user who added the object storage repository, date and time when the object storage repository was added.



Step 2. Specify Object Storage Account

At the **Account** step of the wizard, specify the connection settings:

1. From the **Credentials** drop-down list, select user credentials to access your Azure Archive Storage.

   If you already have a credentials record that was configured in advance, select it from the drop-down list. Otherwise, click **Add** and provide your account name and a shared key, as described in Microsoft Azure Storage Accounts. You can also click the **Manage cloud accounts** link to add, edit or remove a credentials record.

   The user account must have permissions listed in Required Permissions.

   > **NOTE**
   >
   > Veeam Backup & Replication archive tier feature does not support Microsoft Azure Stack Hub compute accounts.

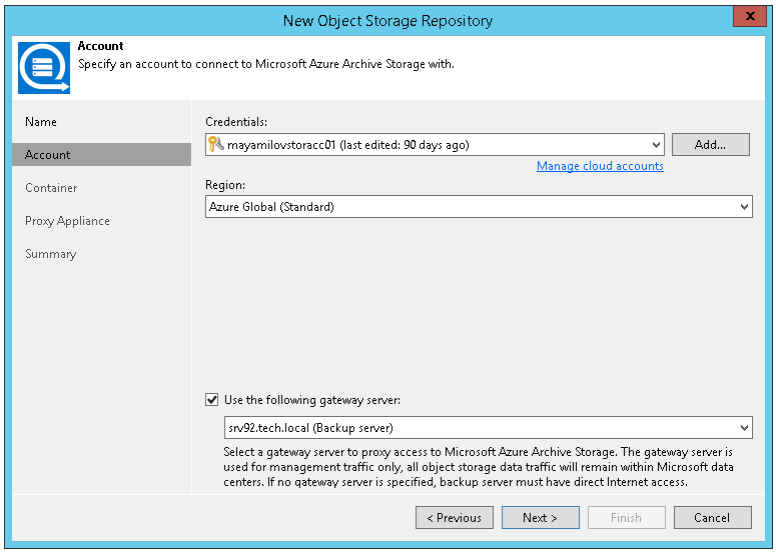
2. From the **Region** drop-down list, select an Azure region.

To use a gateway server, select the **Use the following gateway server** check box and choose an appropriate server from the list. You may want to use a gateway server, for example, if your organization has NAT or different types of firewalls and your access to the internet is limited.

You can select any Microsoft Windows or Linux server that is added to your backup infrastructure and has internet connection, except Linux servers that have the hardened repository role. For more information on how to add such a server to your environment, see Adding Microsoft Windows Servers and Adding Linux Servers respectively. By default, the role of a gateway server is assigned to the machine where Veeam Backup & Replication is installed.

If you choose not to use a gateway server, make sure that all scale-out repository extents have direct internet access.

For information on considerations and limitations for using gateway server, see Considerations and Limitations.



Step 3. Specify Object Storage Settings

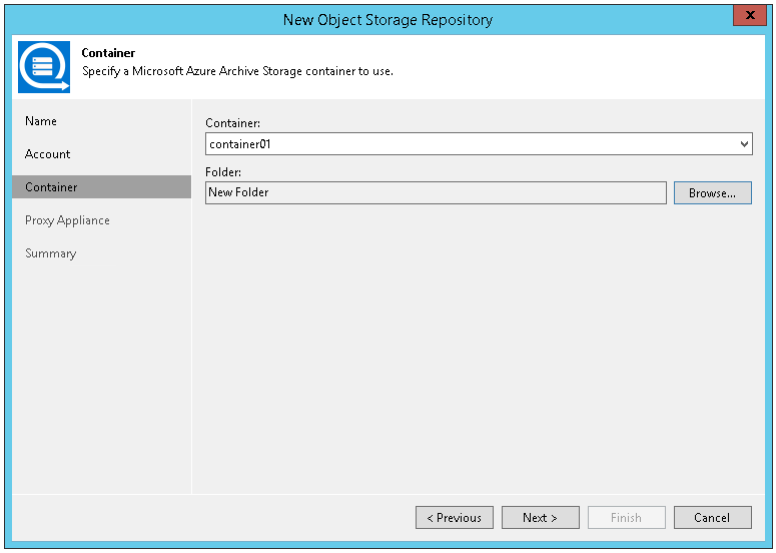
At the **Container** step of the wizard, specify the container and folder that will be used to store data:

1. From the **Container** drop-down list, select a container.

   Make sure that the container where you want to store your backup data was created in advance.

2. In the **Folder** field, select a cloud folder to which you want to map your object storage repository and which will be used to store data.

   To select a folder, click **Browse** and either select an existing folder or create a new one by clicking **New Folder**.

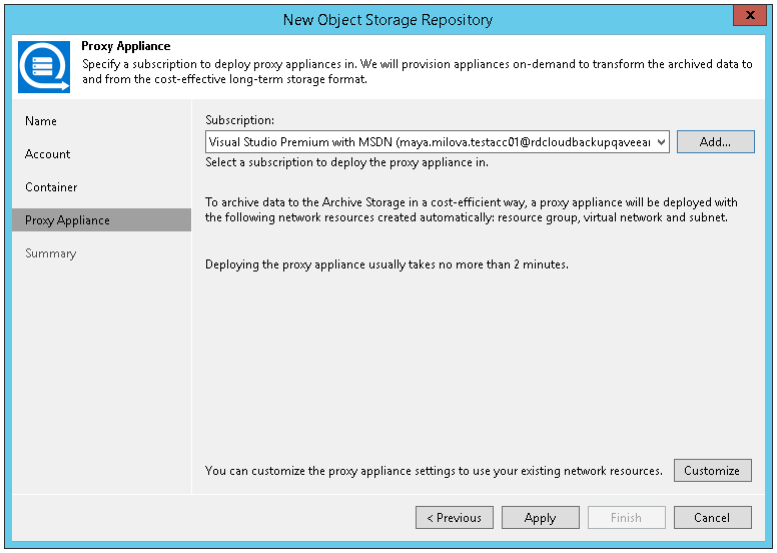


Step 4. Specify Proxy Appliance

At the **Proxy Appliance** step of the wizard, you can specify proxy appliance settings. A proxy appliance is an auxiliary instance that is necessary for transferring data from Azure Blob storage to Azure Archive Storage.

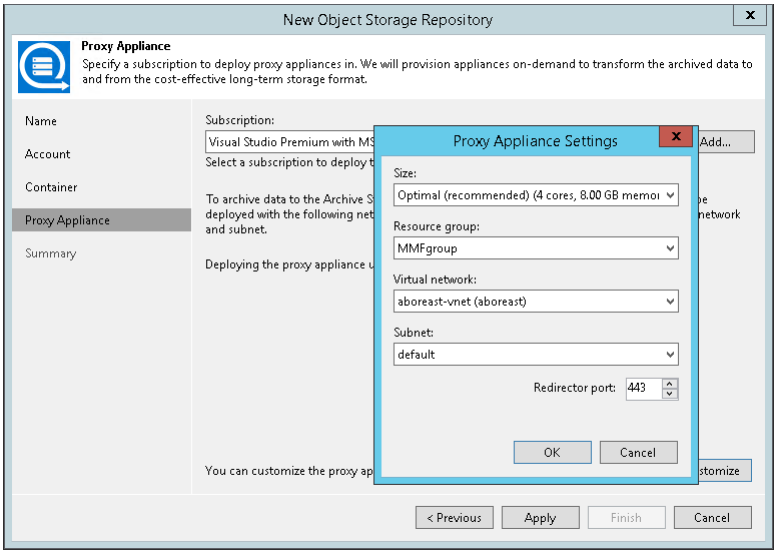
To configure and set up a proxy appliance, you must add a Microsoft Azure account to Veeam Backup & Replication. For that, do the following:

1. From the **Subscription** drop-down list, select your Microsoft Azure subscription credentials. If you have not set up credentials beforehand, click **Add**. You will be prompted to the Adding Microsoft Azure Compute Accounts wizard. Follow the wizard to add your account. Before adding your Microsoft Azure account, check the prerequisites.



2. You can use the default settings or customize the proxy appliance. To specify proxy appliance settings, do the following:
   a. Click **Customize**.
   b. From the **Size** drop-down list, select the size of the appliance.
   c. From the **Resource group** drop-down list, select a resource group that will be associated with the proxy appliance.
   d. From the **Virtual network** drop-down list, select a network to which the helper appliance must be connected.
   e. From the **Subnet** drop-down list, select the subnet for the proxy appliance.
   f. In the **Redirector port** field, specify the port that Veeam Backup & Replication will use to route requests between the proxy appliance and backup infrastructure components.

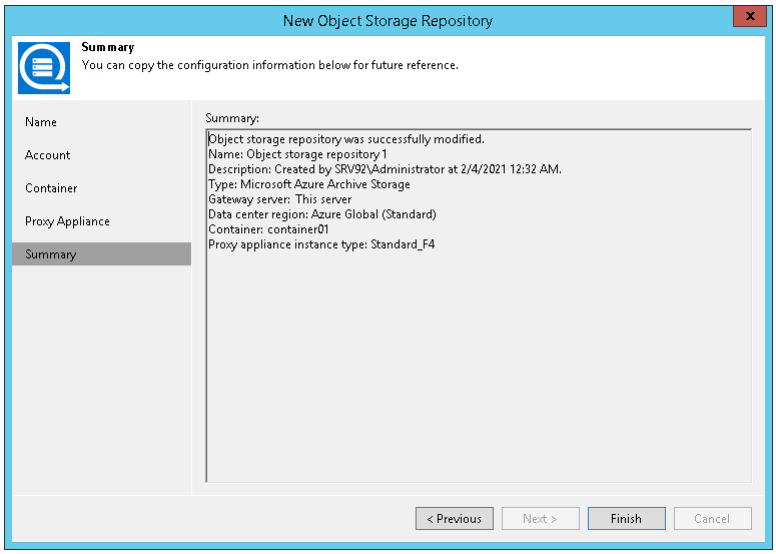
g. Click **OK**.


New Object Storage Repository
Proxy Appliance
Specify a subscription to deploy proxy appliances in. We will provision appliances on-demand to transform the archived data to and from the cost-effective long-term storage format.
Name
Account
Container
Proxy Appliance
Summary
Subscription:
Visual Studio Premium with MS
Add...
Select a subscription to deploy t
To archive data to the Archive S
deployed with the following net
and subnet.
network
Deploying the proxy appliance u
You can customize the proxy ap
stomize
< Previous
Apply
Finish
Cancel
Proxy Appliance Settings
Size:
Optimal (recommended) (4 cores, 8.00 GB memor
Resource group:
MMFgroup
Virtual network:
aboreast-vnet (aboreast)
Subnet:
default
Redirector port:
443
OK
Cancel

Step 5. Finish Working with Wizard

At the **Summary** step of the wizard, complete the object storage repository configuration:

1. Review the details of the object storage repository.
2. Click **Finish** to exit the wizard.



Adding Azure Data Box Storage

Azure Data Box is a physical device that is requested for a short period of time from Microsoft and which can be used as object storage that can be temporarily attached to your Veeam Backup & Replication backup infrastructure. For more information about ordering Azure Data Box and preparing to use it, see this Microsoft article.

This device might become useful when you need to offload a large number of backup files that are occupying a lot of storage space on your extents, as offloading data to the Azure Data Box device occurs a lot faster that transferring the same amount of data directly to Azure object storage.

Once you have offloaded backups to Azure Data Box, you need to ship the device back to Microsoft for further data synchronization with your Azure storage account, as described in Seeding Backups to Azure Blob Storage.

IMPORTANT

To eliminate any risk of data loss in case the device is damaged during the shipping, we recommend to copy your backup files to Azure Data Box Storage instead of moving them, and to keep them in your local infrastructure until the device is delivered and your data is safely transferred into your storage account.

NOTE

Consider that Veeam Backup & Replication supports Azure Data Box devices that are capable of reading/writing data using REST API only; the *Azure Data Box Disk* type is not supported.

For information about other limitations for Microsoft Azure Data Box storage, see this Microsoft article.

To add Microsoft Azure Data Box storage, do the following:

1. Perform the steps described in the Before You Begin section.
2. Specify object storage name

3. Specify object storage account
4. Specify object storage settings
5. Finish working with the wizard

Before You Begin

Before you add Microsoft Azure Data Box to the Veeam Backup & Replication backup infrastructure, complete the following steps:

- Configure Name Resolutions
- Download and Install SSL Certificate

Configuring Name Resolutions

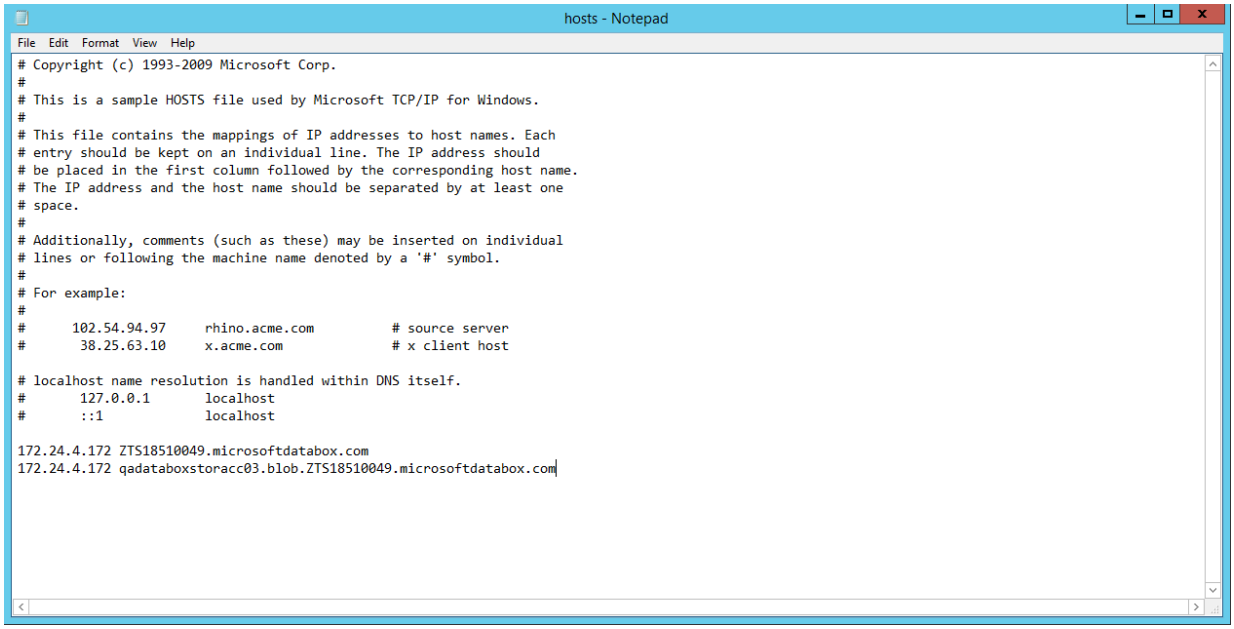
To resolve a DNS name of your Azure Data Box device, make sure to add the following DNS addresses to the *HOST* file that is located on both the Veeam Backup & Replication server and the gateway server (if any additional gateway server is used):

- *<ip_address> <mydataboxno>.microsoftdatabox.com*
- *<ip_address> <storageaccountname>.blob.<mydataboxno>.microsoftdatabox.com*

For more information on how to learn the values for *<ip_address>*, *<mydataboxno>* and *<storageaccountname>*, see this Microsoft article.

As an *<ip_address>*, use any of the IP-addresses listed under the **Data N** section. You can also use the address specified under the **MGMT** section, but due to its slow connection rate (limited to 1 GbE), using such an address is not recommended.

The following is an example of the *HOST* file.



NOTE

Consider the following:

- Make sure to configure name resolutions on each server that might be used as a gateway.
- Alternatively, you can create a *microsoftdatabox.com* DNS zone with corresponding records on your DNS server if you prefer not to modify the *HOST* file.

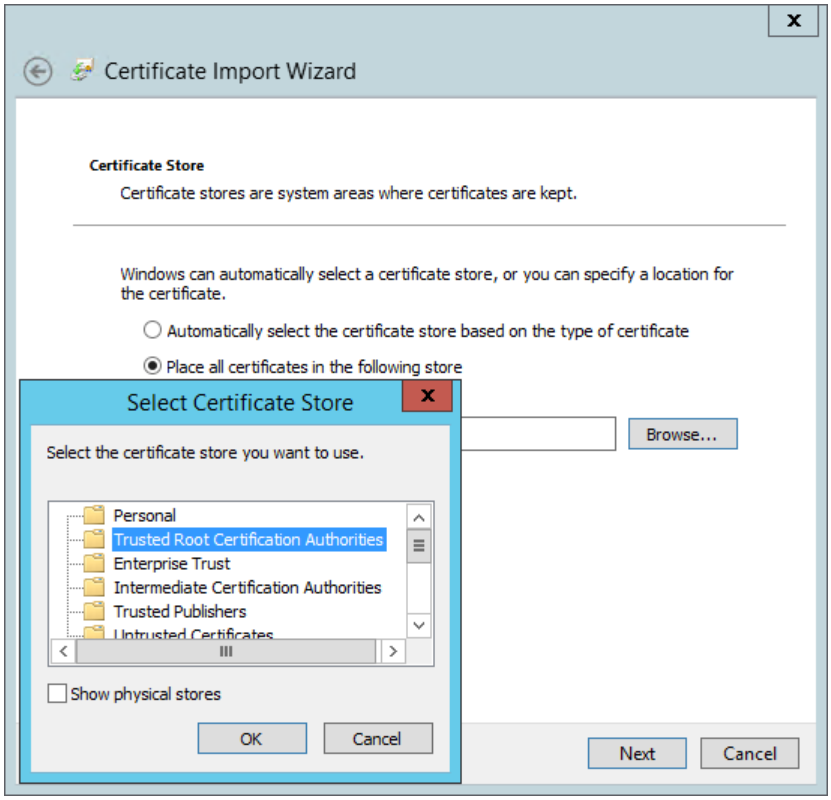
Downloading and Installing SSL Certificate

To establish a secure connection to the Azure Data Box device, make sure to download and import an SSL certificate and install it on both the Veeam Backup & Replication server and the gateway server (if any additional gateway server is used).

For more information on how to download an SSL certificate, see this Microsoft article.

When installing a certificate, do the following:

1. In the **Certificate Import Wizard** dialog, select **Local Machine**.
2. In the **Certificate Store** step, select **Place all certificates in the following store** and click **Browse**.
3. Select **Trusted Root Certification Authorities**.



Configuring registry settings

For information how to configure registry settings for Microsoft Azure Data Box device, see this Veeam KB article.

Sizing Gateway Server

Consider that Veeam Backup & Replication supports Azure Data Box devices that are capable of reading/writing data using REST API only; the Azure Data Box Disk type is not supported.

As with any other object storage, REST API performance depends on scale. As Azure Data Box is a single endpoint, the individual throughput of this REST API may be limited. The block size used in Veeam Backup & Replication capacity tier for object storage offload matches that of the source job. The default object size would be a compressed 1MB block, resulting in objects of around 512KB in size.

The speed of data offload to Azure Data Box devices may reach about 300 MB/s. To achieve this speed, we recommend using a separate gateway server with 8 CPU cores.

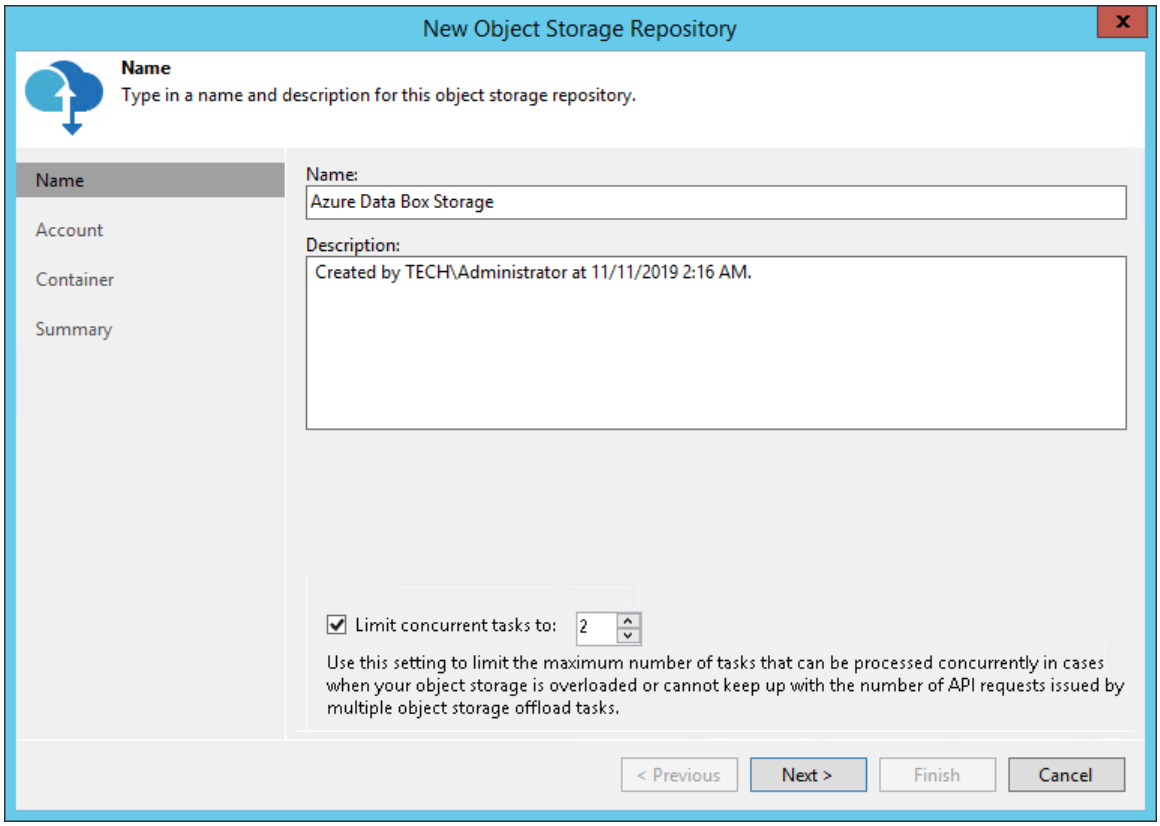
Step 1. Specify Object Storage Name

At the **Name** step of the wizard, specify a name and description for the object storage repository:

- In the **Name** field, specify a name for the new object storage repository.
- In the **Description** field, enter an optional description. The default description contains information about the user who added the object storage repository, date and time when the object storage repository was added.

If you wish to limit the maximum number of tasks that can be processed at once, select the **Limit concurrent tasks to N** check box.

> **TIP**
>
> Set the maximum number of concurrent task to a reasonable number to avoid overloading if you plan to upload significant amount of backup chains to the device.



Step 2. Specify Object Storage Account

At the **Account** step of the wizard, specify the connection settings:

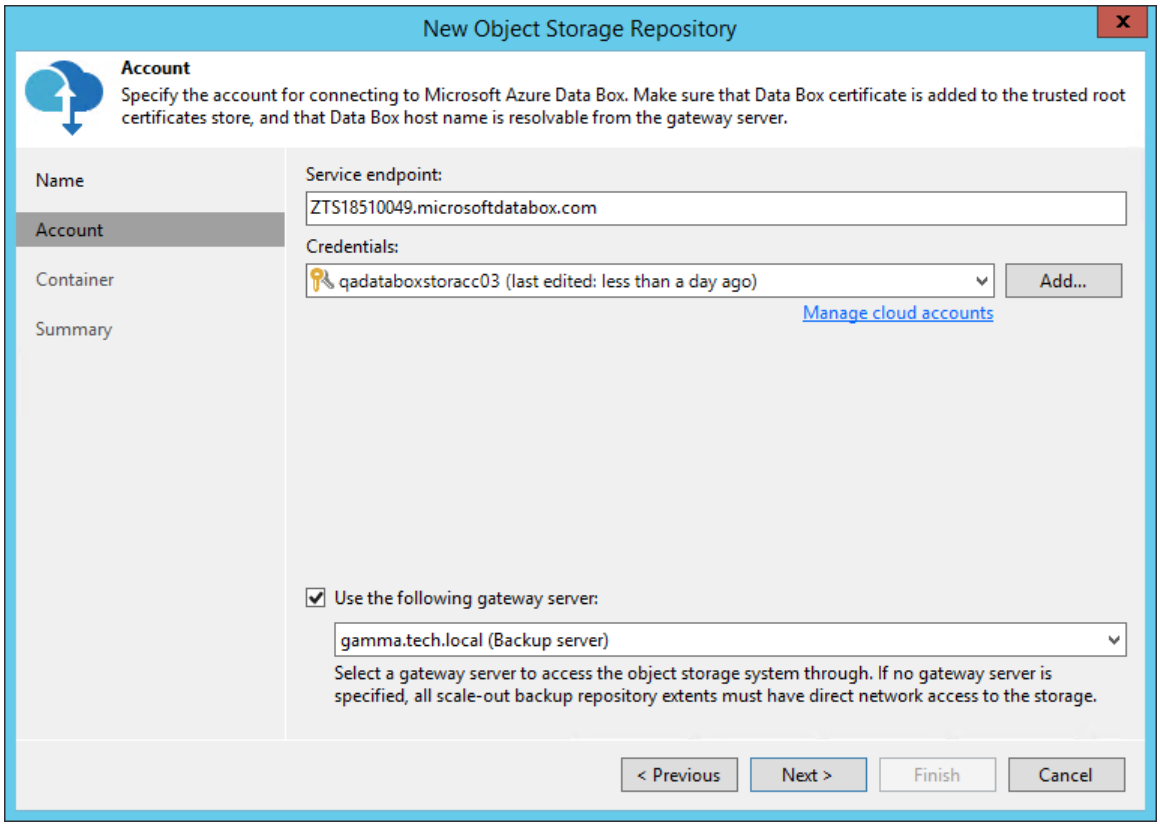
1. In the **Service endpoint** field, specify a service endpoint address of your Azure Data Box device.
2. From the **Credentials** drop-down list, select user credentials to access your Azure Data Box storage.

   If you already have a credentials record that was configured in advance, select such a record from the drop-down list. Otherwise, click **Add** and provide your access and secret keys, as described in Cloud Credentials Manager. You can also click the **Manage cloud accounts** link to add, edit or remove a credentials record.

For more information on where to find connection parameters of your Azure Data Box device, see Getting Data Box Connection Parameters.

To use a gateway server, select the **Use the following gateway server** check box and choose an appropriate server from the list. Consider using Windows-based server as a gateway. For more information on how to add such a server to your environment, see Adding Microsoft Windows Servers. By default, the role of a gateway server is assigned to the machine where Veeam Backup & Replication is installed.

For information on considerations and limitations for using gateway server, see Considerations and Limitations.



Getting Data Box Connection Parameters

To find connection parameters of your Azure Data Box device, do the following:

1. Open the Microsoft Azure Data Box portal.
2. In the navigation pane, click **Connect and Copy**.

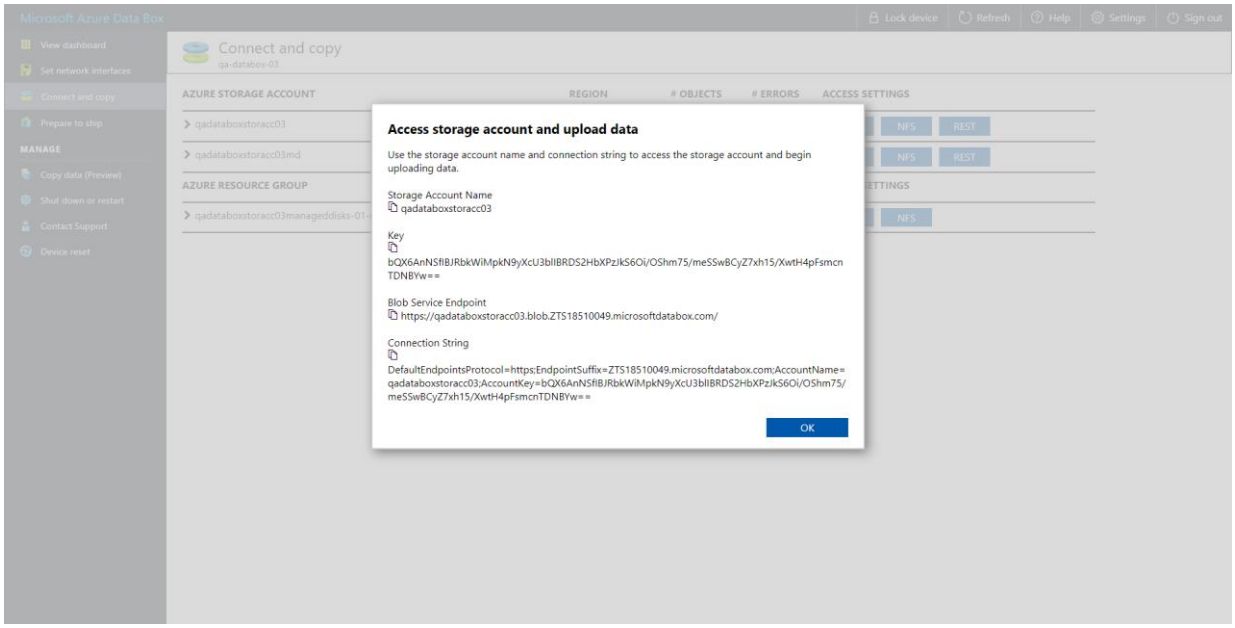
3. Under the **Access Settings** column of the storage account that you want to use, click **REST** and in the **Access storage account and upload data** dialog, copy the following:

   a. Under **Storage Account Name**, copy the Azure storage account name.

   b. Under **Key**, copy the storage account key.

   c. Under **Blob Service Endpoint**, copy the service endpoint address that starts exactly after the *blob* word.

      For example, if the complete service endpoint address is *https://qadataboxstoracc03.blob.ZTS18510049.microsoftdatabox.com*, then you will need to copy everything that starts from *ZTS* only. That is, *ZTS18510049.microsoftdatabox.com*. Make sure not to copy the last slash ("/") symbol.



Step 3. Specify Object Storage Settings

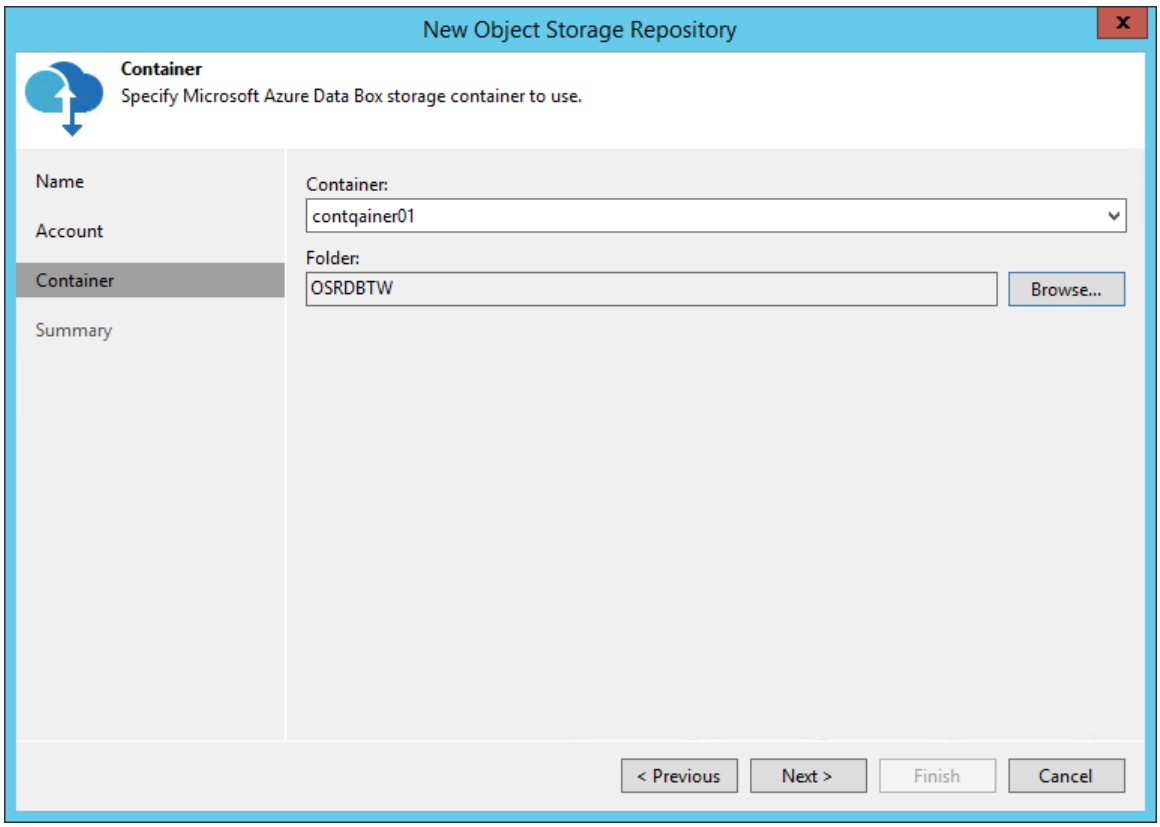
At the **Container** step of the wizard, specify the container and folder that will be used to store data:

1. From the **Container** drop-down list, select a container.

   To create a container, use Microsoft Azure Storage Explorer. For more information on how to connect to the Azure Data Box device using Microsoft Azure Storage Explorer, see this Microsoft article.

2. In the **Select Folder** field, select a cloud folder to which you want to map your object storage repository.

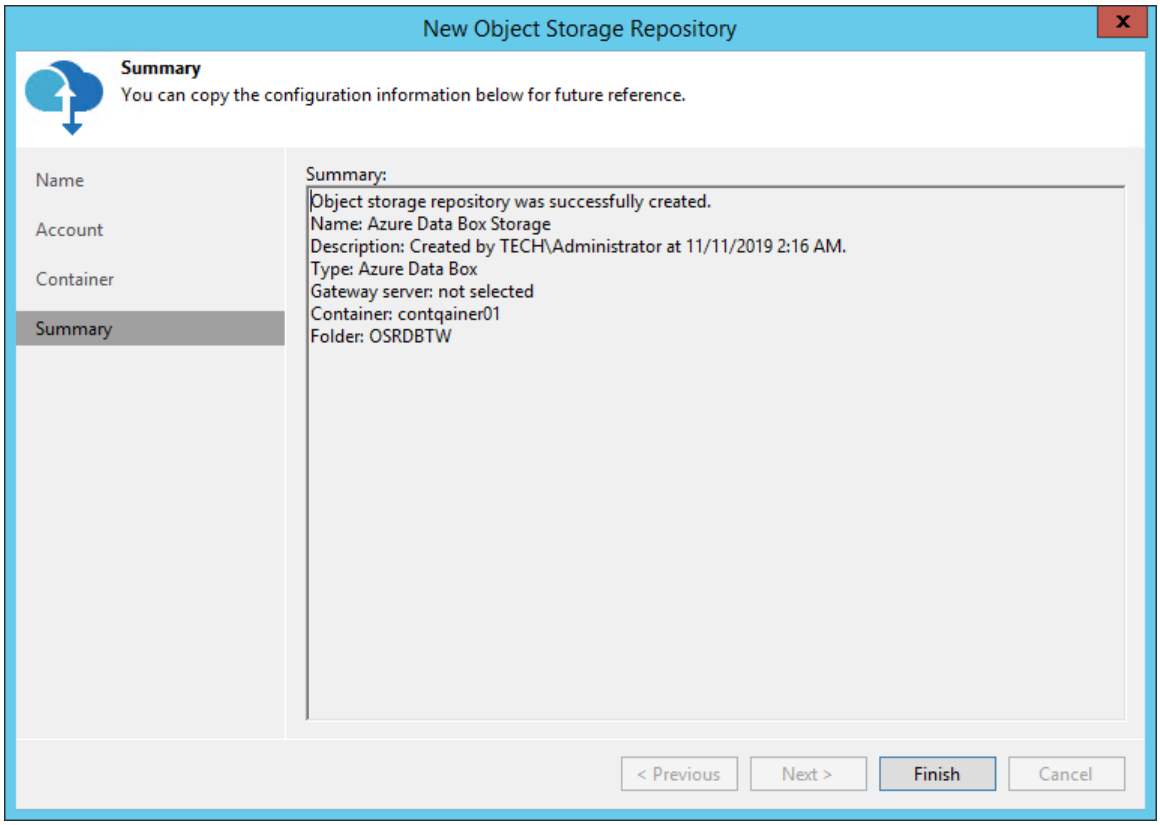
   To select a folder, click **Browse** and either select an existing folder or create a new one by clicking **New Folder**.



Step 4. Finish Working with Wizard

At the **Summary** step of the wizard, complete the object storage repository configuration:

1. Review details of the object storage repository.
2. Click **Finish** to exit the wizard.



Managing Object Storage Repositories

You can manage your object storage repositories in various ways: edit settings of an object storage repository, switch it to the maintenance or the seal mode or remove an object storage repository.

Editing Settings of Object Storage Repository

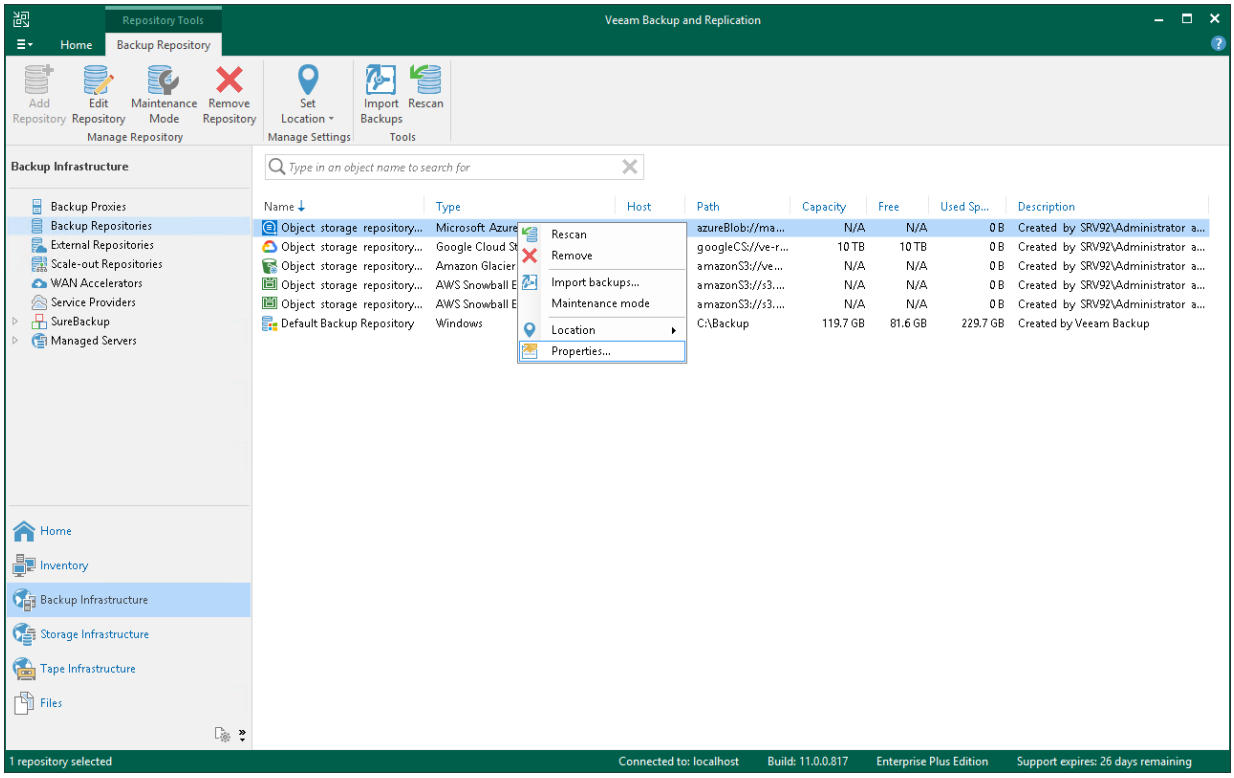
After you have added an object storage repository, you may want to edit its settings.

To edit object storage settings, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Backup Repositories**.
3. In the working area, select an object storage repository and click **Edit Repository** on the ribbon or right-click an object storage repository and select **Properties**.

4. Follow the steps of the **Edit Object Storage Repository** wizard and edit settings as required.

   Mind that some settings cannot be modified and will remain disabled while being edited.



Switching to Maintenance Mode

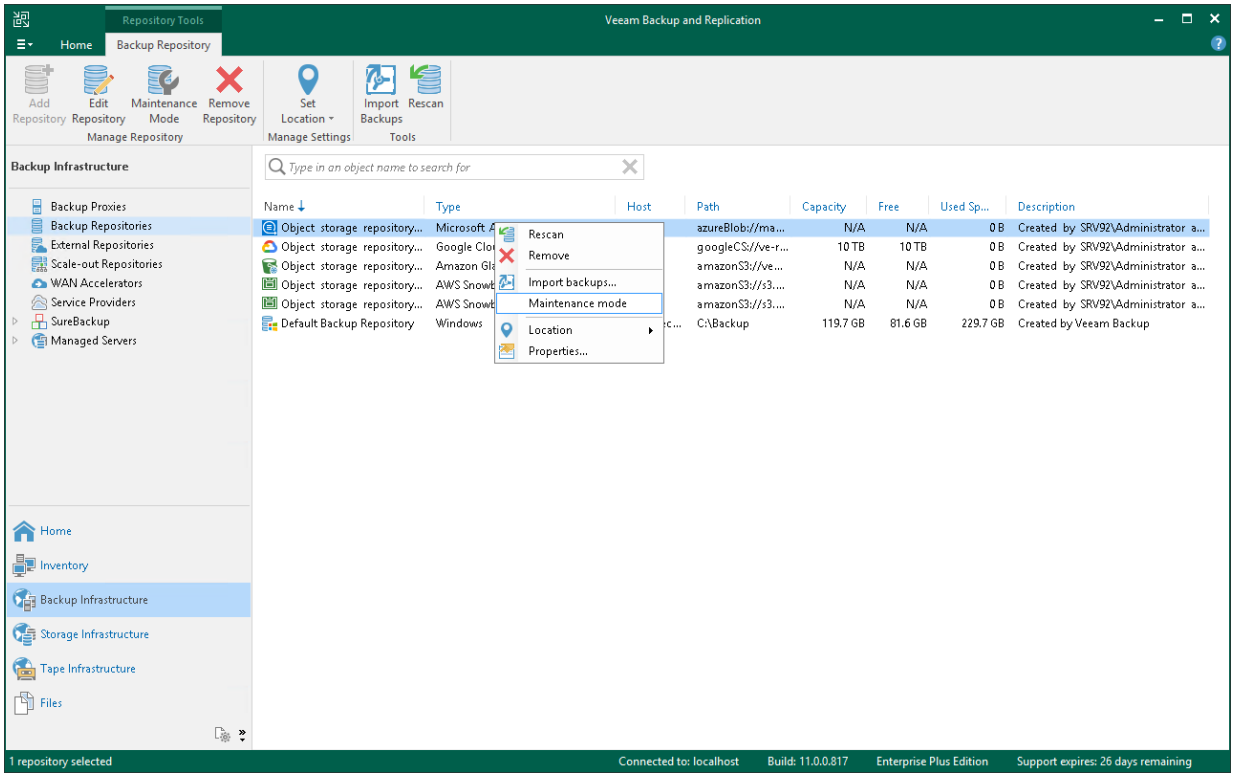
Maintenance mode imposes limitations on usage of object storage repositories, as described in Maintenance Mode Limitations.

An object storage repository can be put into both the maintenance and the seal modes at the same time. When both modes are applied, limitations of maintenance mode override limitations of seal mode.

To put an object storage repository into the maintenance mode, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select the **Backup Repositories** node.
3. In the working area, select an object storage repository and click **Maintenance mode** on the ribbon or right-click an object storage repository and select **Maintenance mode**.

To remove the extent from the maintenance mode, select the extent and click **Maintenance Mode** on the ribbon or right-click the extent and select **Maintenance mode** once again.



Maintenance Mode Limitations

The following table lists limitations that are imposed right after the object storage repository was put into the maintenance mode.

Activity	Restriction Level	Related Topic
Moving to object storage.	○	Moving Inactive Backup Chains to Object Storage
Copying to object storage.	○	Copying Backups to Object Storage
Downloading from object storage.	○	Downloading Data from Object Storage
Moving to capacity tier (manual operation).	○	Moving to Capacity Tier
Copy to performance tier (manual operation).	○	Copying to Performance Tier

Activity	Restriction Level	Related Topic
Restore from offloaded backups.	○	Data Restore
Export as `.vbk` from offloaded backups.	○	Exporting Backups
Removal of backups from configuration.	●	Removing from Configuration
Retention policies.	Synchronization is skipped for backup chains located in object storage repositories. Obsolete restore points will only be removed from backup chains on the extents.	Retention Policy
Removal of backups or VMs created with the per-machine method.	○	Removing Backups from Object Storage Repository
Scale-out backup repository rescan.	Synchronization is skipped for object storage repositories.	Rescanning Scale-Out Repositories
Evacuation of storage along with indexes from on-premise extents.	●	Evacuating Backups from Performance Tiers

Switching to Seal Mode

Veeam Backup & Replication allows you to put an object storage repository into *Seal* mode.

Sealing up object storage gives you the ability to gradually remove data located in such storage by applying a retention policy. You can use this feature to gracefully stop using some of your object storage as backup repositories and exclude them from the scale-out backup repository configuration altogether.

When sealing up object storage, Veeam Backup & Replication restricts any further data transfer to such a sealed storage and allows only operations listed in Seal Mode Limitations.

All restore points that exceed the specified retention period will be continuously removed from the sealed object storage, as described in Retention Policy.

Consider the following:

- An object storage repository can be put into seal mode only if it is a member of the scale-out backup repository.

  If an object storage repository was not added as part of any of your scale-out backup repositories, the **Seal Extent** option will not be available.

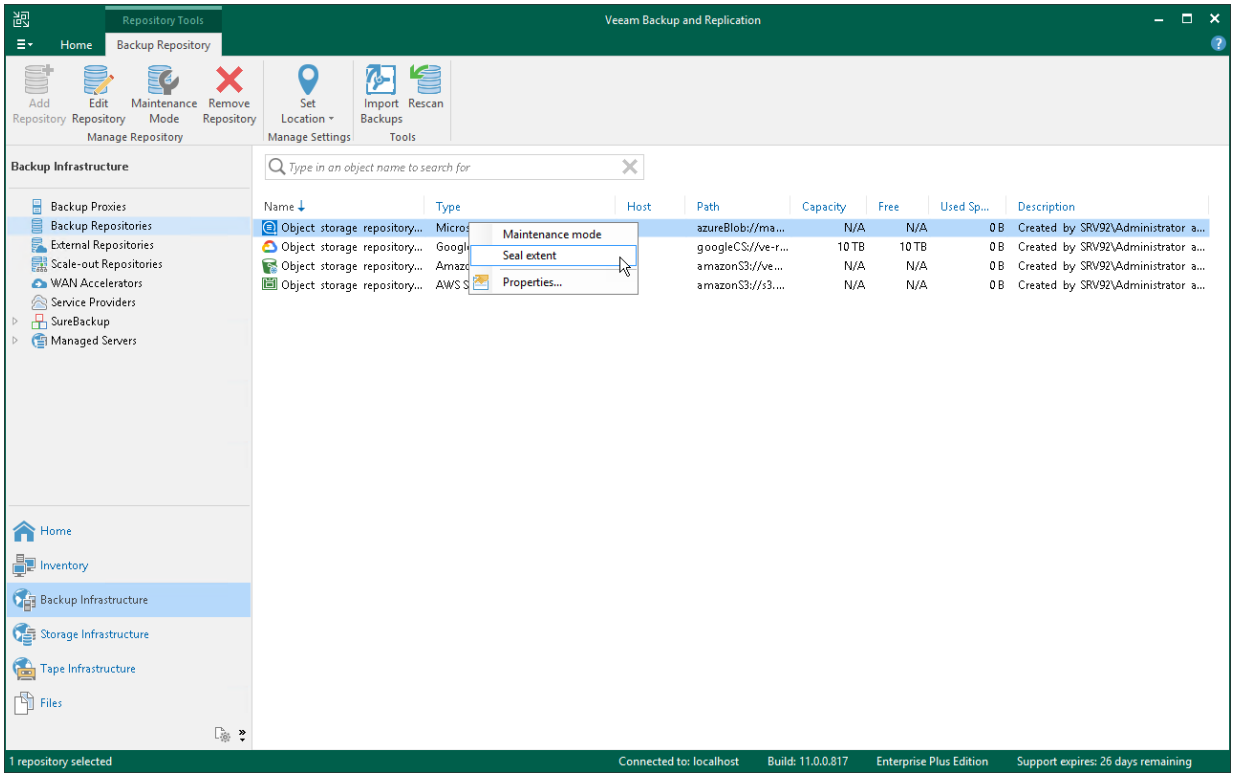
- An object storage repository can be put into both maintenance and seal modes at the same time.

  When both modes are applied, limitations of maintenance mode override limitations of seal mode.

To put an object storage repository into seal mode, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select the **Backup Repositories** node.
3. In the working area, select an object storage repository and click **Seal Extent** on the ribbon or right-click an object storage repository and select **Seal extent**.

To remove the extent from seal mode, select the extent and click **Sealed** on the ribbon or right-click the extent and select **Sealed**.



Seal Mode Limitations

The following table lists limitations that are imposed right after the object storage repository was put into seal mode.

Activity	Restriction Level	Related Topic
Moving to object storage.	○	Moving Inactive Backup Chains to Object Storage
Copying to object storage.	○	Copying Backups to Object Storage
Downloading from object storage.	●	Downloading Data from Object Storage
Moving to capacity tier (manual operation).	○	Moving to Capacity Tier

Activity	Restriction Level	Related Topic
Copy to performance tier (manual operation).		Copying to Performance Tier
Restore from offloaded backups.		Data Restore
Export as `.vbk` from offloaded backups.		Exporting Backups
Removal of backups from configuration.		Removing from Configuration
Retention policies.		Retention Policy
Removal of backups or VMs created with the per-machine method.		Removing Backups from Object Storage Repository
Removal of a VM from a single storage.		Deleting from Disk
Scale-out backup repository rescan.		Rescanning Scale-Out Repositories
Evacuation of storage along with indexes from on-premise extents.		Evacuating Backups from Performance Tiers

Removing Object Storage Repository

You can remove any object storage repository from the application scope if you no longer need it.

Consider the following:

- An object storage repository cannot be removed if it is part of a scale-out backup repository.

  To remove such a repository, you must first exclude an object storage repository from the scale-out backup repository configuration, as described in Excluding Capacity Tier from Scale-Out Repositories.

- An object storage repository cannot be removed if backups located in this repository was imported, as described in Importing Backups.

  To remove such a repository, you must first detach object storage, as described in Detaching Object Storage.

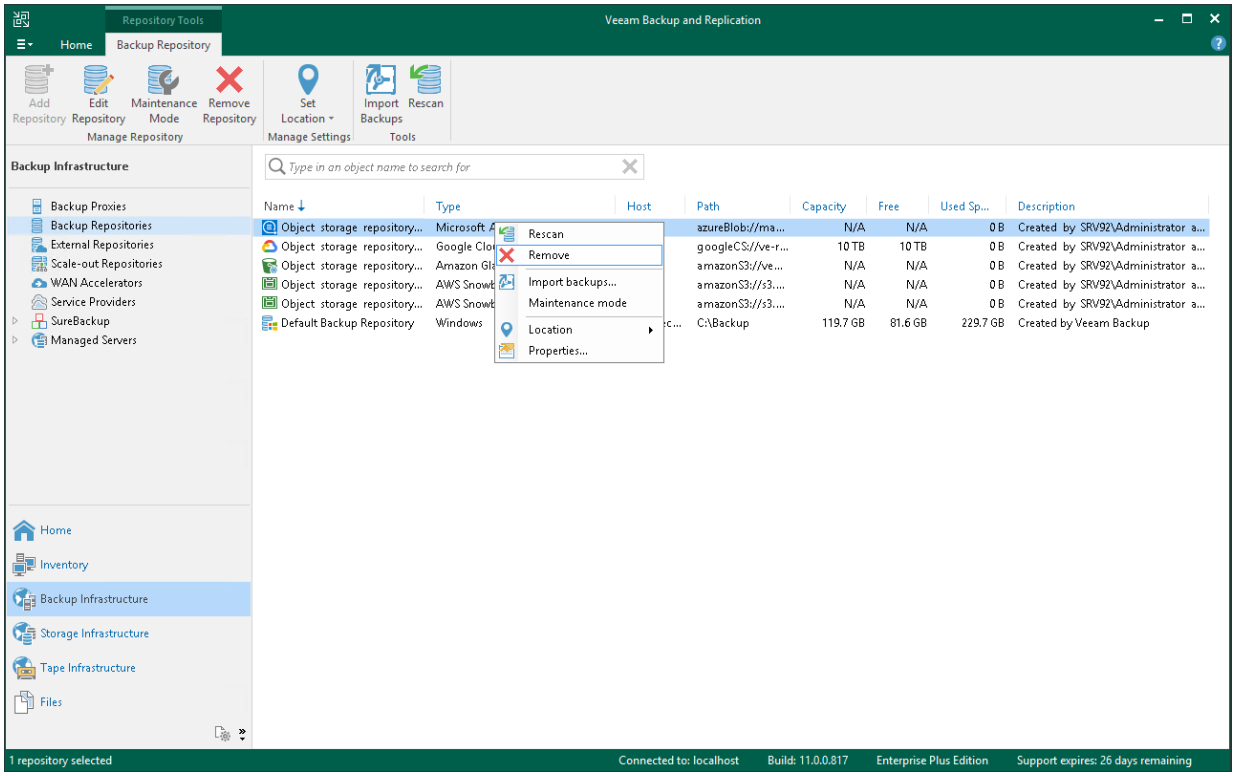
- When an object storage repository is being removed from the environment, the actual offloaded data remains completely unaffected.

  To learn how to remove data, see Removing Backups from Object Storage Repository.

To remove an object storage repository, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Backup Repositories**.

3. In the working area, select an object storage repository and click **Remove Repository** on the ribbon or right-click an object storage repository and select **Remove**.



Managing Object Storage Backups

You can manage your object storage backups in various ways: import, detach or remove object storage backups or seed them to Azure Blob storage.

Importing Object Storage Backups

Veeam Backup & Replication allows you to import backups located in object storage repositories into the Veeam Backup & Replication console so that you can perform disaster recovery operations in case the configuration of the scale-out backup repository was permanently lost and is unavailable.

Consider the following:

- The Import Backups option is available only if the object storage is not a part of the scale-out backup repository.
- Before you start importing backups, make sure to add the object storage repository that stores data you want to import.
- If you have imported backups from the object storage repository, you will not be able to select this repository as a capacity extent at the Add Capacity Tier and at the Add Archive Tier steps.
- The Import Backups option is applicable only to the Capacity Tier and Archive Tier backups. It does not support the NAS backups.

To import backups, do the following:

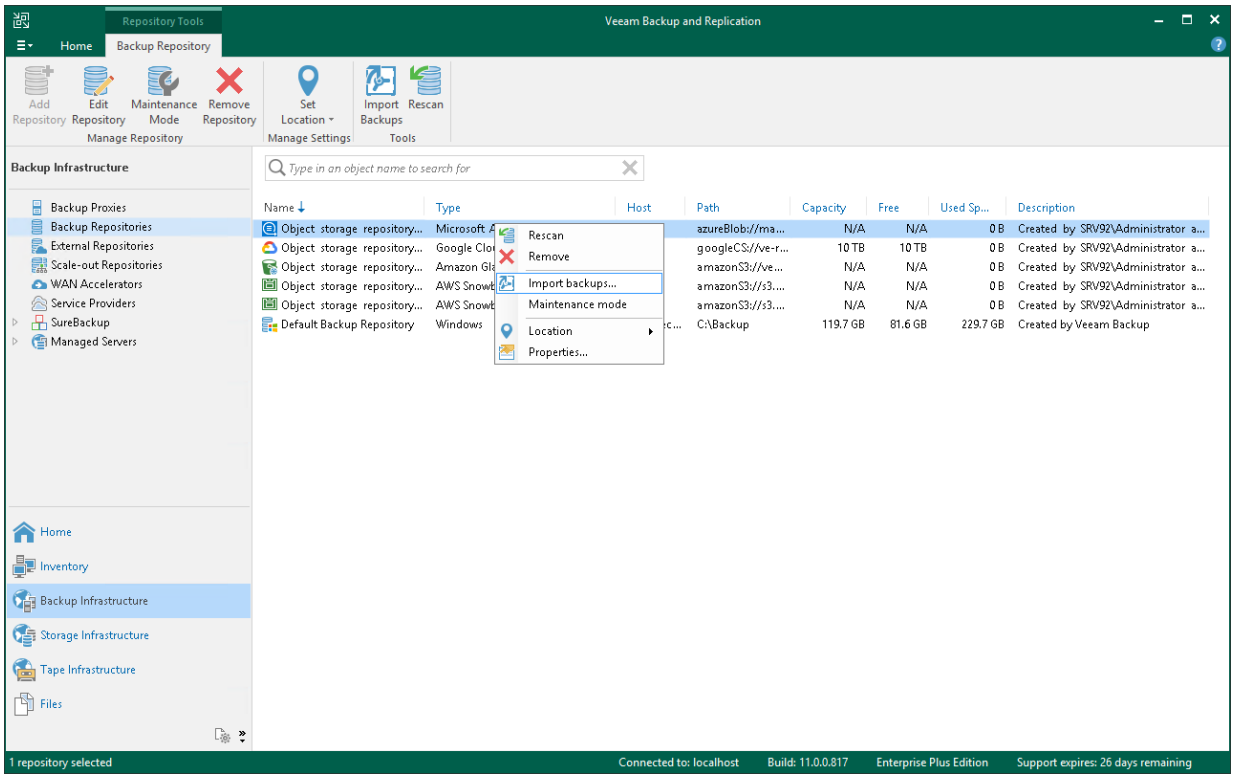
1. Launch Import Wizard
2. Specify Password

3. Wait for Import
4. Finish Working with Wizard

Step 1. Launch Import Wizard

To launch the **Import Backup** wizard, do either of the following:

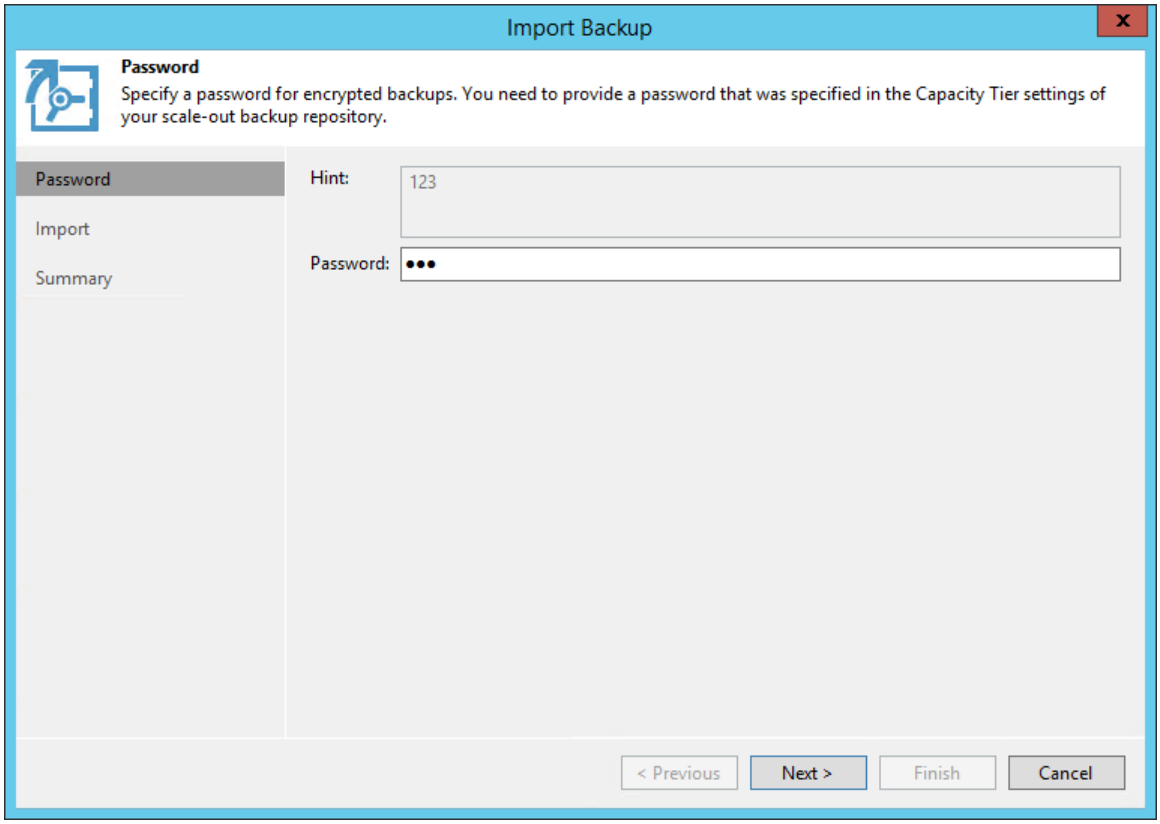
- Open the **Backup Infrastructure** view, in the inventory pane select the **Backup Repositories** node and in the preview pane, right-click object storage that contains backups that you want to import and select **Import backups**.
- Open the **Backup Infrastructure** view, in the inventory pane select the **Backup Repositories** node, in the preview pane, select object storage that contains backups that you want to import and select **Import Backups**.



Step 2. Specify Password

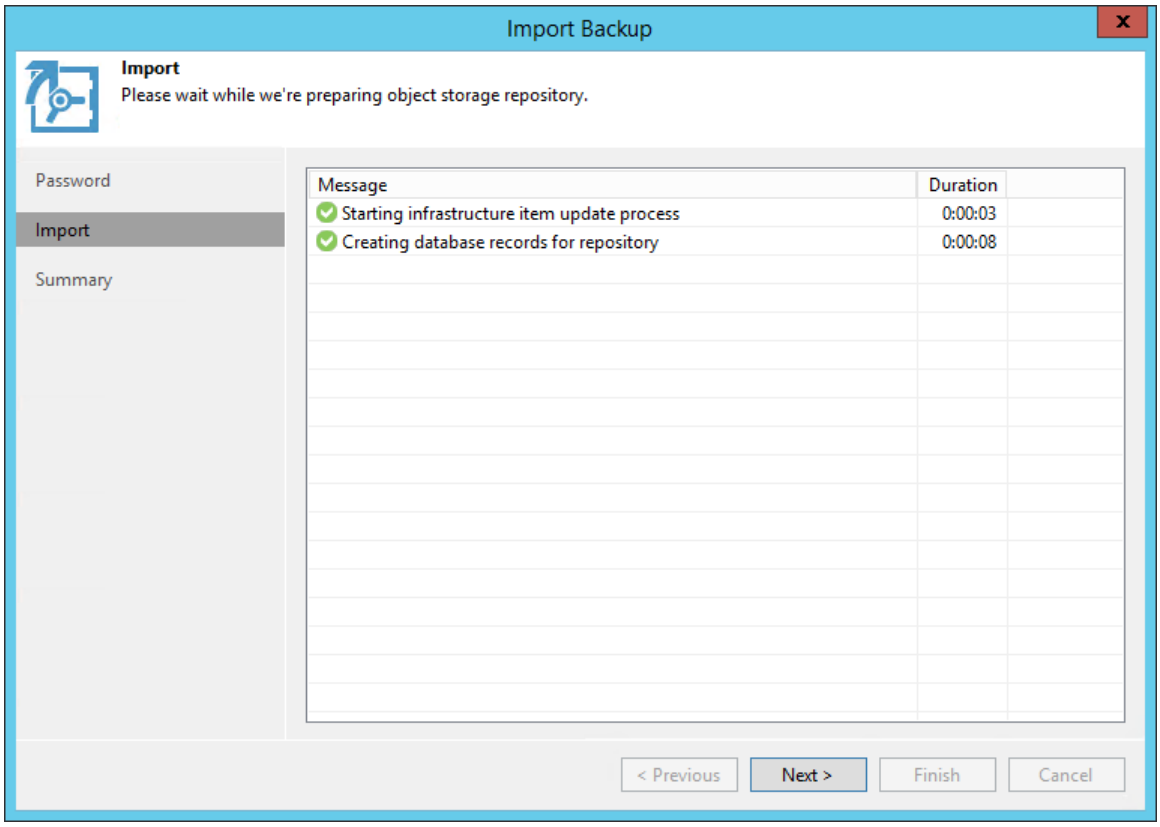
At the **Password** step of the wizard, in the **Password** field, specify the password that was used to encrypt data during offload or copy sessions.

This step is only available when importing encrypted backups.



Step 3. Wait for Import

At the **Import** step of the wizard, wait until Veeam Backup & Replication prepares a temporary database to which information about backups will be added upon import.

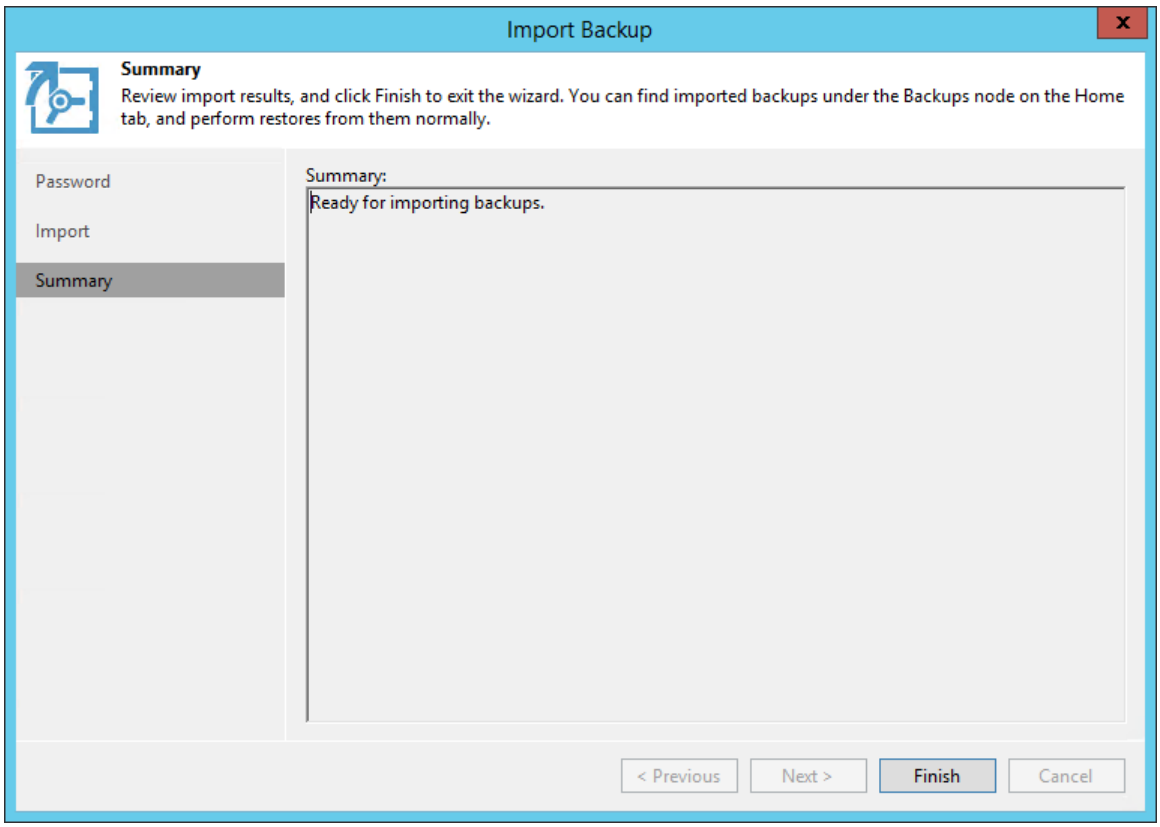


Step 4. Finish Working with Wizard

At the **Summary** step of the wizard:

1. Review details of the import operation.
2. Click **Finish** to run a **Configuration Database Resynchronize** session.

   During this session, all existing backups will be imported into the Veeam Backup & Replication console.

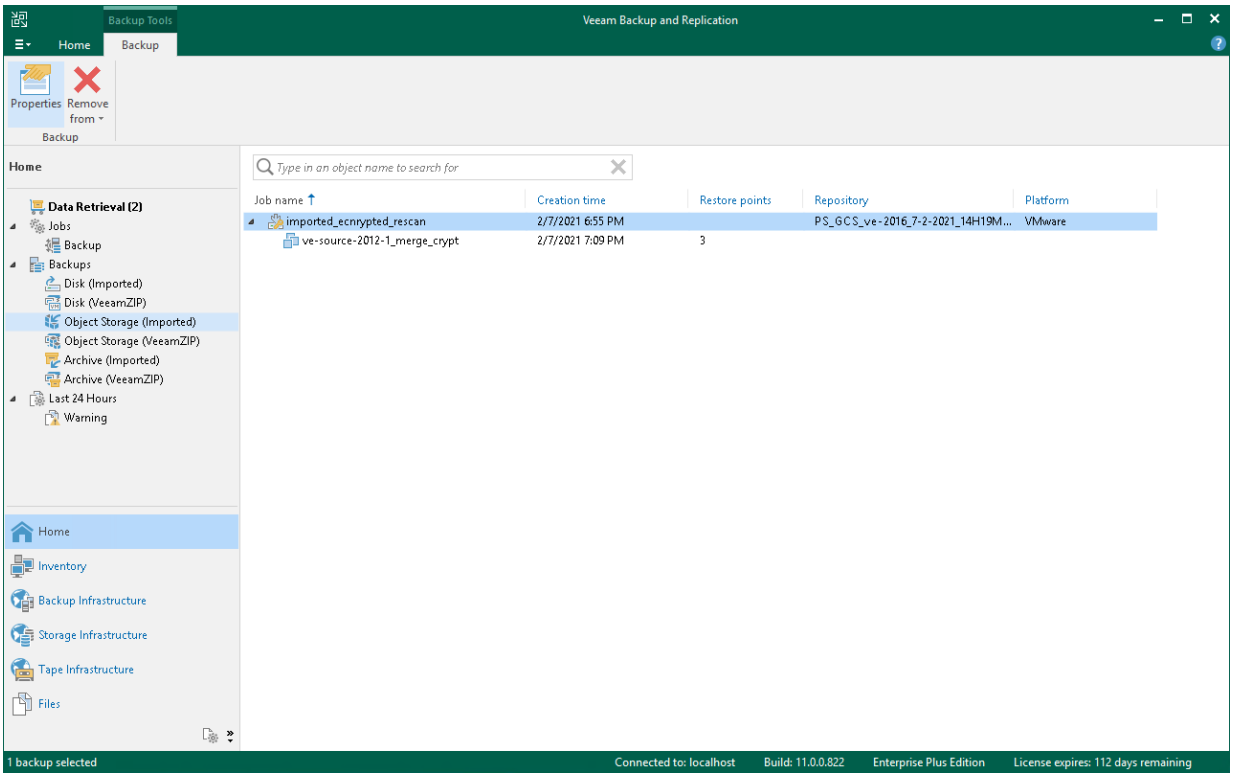


Viewing Imported Backups

To view imported backups, do the following:

1. Open the **Home** view.
2. In the navigation pane, select the **Backups** > **Object Storage (Imported)** node.

3. In the preview pane, review imported backups.



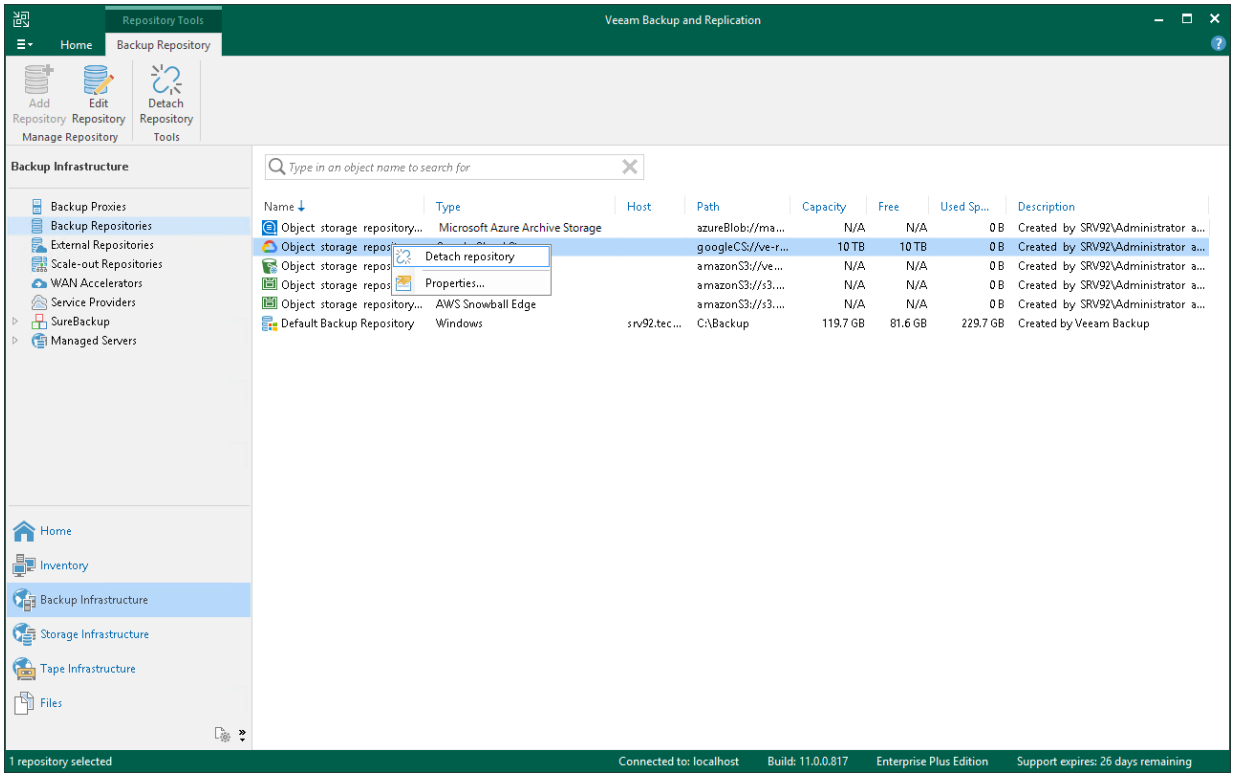
Detaching Object Storage Backups

You can detach imported object storage after you have finished working with it. After object storage is detached, the imported backups become unavailable.

To detach imported object storage, do the following:

1. Open the **Backup Infrastructure** view.
2. In the navigation pane. select the **Backup Repositories** node.

3. In the preview pane, select object storage the backups of which have been imported and click **Detach Repository** on the toolbar or right-click object storage and select **Detach Repository**.



Seeding Backups to Azure Blob Storage

After you have completed moving backup data to the Azure Data Box device, you need to ship this device back to Microsoft so that this data can be synchronized with your Azure storage account.

To prepare Azure Data Box for shipping, do the following:

1. Put an object storage repository that is associated with the Azure Data Box device into the *maintenance* mode, as described in Switching to Maintenance Mode.
2. Perform the steps described in this Microsoft article.

After mailing the device, wait until it is received by Microsoft. You will get a notification of successful data upload to your Azure storage account. After that, do the following:

1. Add an Azure Blob storage, as described in Adding Azure Blob Storage.
2. In the Capacity Tier step, change your Azure Data Box object storage to Azure Blob storage that you have created in the previous step.
3. Confirm importing backups from Azure Blob storage and wait for the synchronization job to complete.
4. Remove Azure Data Box object storage from the Veeam Backup & Replication backup infrastructure, as described in Removing Object Storage Repository.

NOTE

Keep in mind that if you delete individual VMs from your backup in the period between mailing Azure Data Box and adding Azure Blob Storage, you may need to run Active Full instead of Incremental job after final Azure synchronization. To avoid this, do not delete any VMs manually and disable the **Retention Policy for Deleted Items** option in the job settings for that period. After successful synchronization, you can re-enable this option and work as usual.

Seeding Backups to AWS Snowball Edge Storage

After you have completed moving backup data to the AWS Snowball Edge device, you need to ship this device back to Amazon so that this data can be synchronized with your storage account. For more information about returning AWS Snowball Edge, see this Amazon article.

To prepare AWS Snowball Edge for shipping, put the object storage repository that is associated with the AWS Snowball Edge device into the maintenance mode, as described in Switching to Maintenance Mode. You can learn in more detail about preparing AWS Snowball Edge for shipping in this Amazon article.

After mailing the device, wait until it is accepted by Amazon and notification of successful data upload to your storage account is received (if you subscribed for such notification). After that, do the following:

1. Add an Amazon S3 storage, as described in Adding Amazon S3 Storage. At the Specify Object Storage Settings step, select the same bucket and folder that you have selected for your AWS Snowball Edge device.
2. Open the **Edit Scale-out Backup Repository** wizard, as described in Editing Settings of Scale-Out Backup Repositories. In the Capacity Tier step, change your AWS Snowball Edge object storage to Amazon S3 storage that you have created in the previous step. Save the new setting.
3. Remove the AWS Snowball Edge object storage from the Veeam Backup & Replication backup infrastructure, as described in Removing Object Storage Repository.

IMPORTANT

Make sure that Versioning and Object Lock are not enabled on the used S3 bucket.

Adding Backup Repositories

This section describes how to add direct attached storage, network attached storage, and deduplicating storage appliances as backup repositories.

For information on how to add object storage repositories, see Adding Object Storage Repositories.

Before you add a backup repository, check prerequisites. Then use the **New Backup Repository** wizard to add the backup repository.

Before You Begin

Before you configure a backup repository, check the following prerequisites.

Dell EMC Data Domain

- Dell EMC Data Domain must meet software and/or hardware requirements. For more information, see System Requirements.

- The DD Boost license must be installed on the Dell EMC Data Domain system, DD Boost must be enabled and configured.
- The gateway server that you plan to use for work with Dell EMC Data Domain must be added to the backup infrastructure.

If the Dell EMC Data Domain storage system does not meet these requirements, you can add it as a SMB (CIFS) folder. In this case, Veeam Backup & Replication will not use the DD Boost technology to work with Dell EMC Data Domain. For more information, see Dell EMC Data Domain.

ExaGrid

- ExaGrid must meet software and/or hardware requirements. For more information, see System Requirements.
- To use ExaGrid as a backup repository, you must configure an ExaGrid share in a proper way in ExaGrid Manager. For more information, see the ExaGrid documentation.
- When you add ExaGrid servers to the Veeam backup infrastructure, and you use the UPN format for an Active Directory account user name (for example, john.doe@domain.local), make sure you enter the user name in lowercase letters only.

HPE StoreOnce

- HPE StoreOnce must meet software and/or hardware requirements. For more information, see System Requirements.
- The HPE StoreOnce Catalyst license must be installed on the HPE StoreOnce system.
- You must use a Catalyst store as a backup target.
- The gateway server that you plan to use for work with HPE StoreOnce system must be added to the backup infrastructure.
- The client account that you plan to use to connect to HPE StoreOnce must have access permissions on the Catalyst store where backup data will be kept.

If the HPE StoreOnce storage system does not meet these requirements, you can add it as a shared folder. In this case, Veeam Backup & Replication will perform target-side deduplication. For more information, see HPE StoreOnce.

Quantum DXi

- Quantum DXi must meet software and/or hardware requirements. For more information, see System Requirements.
- To use Quantum DXi as a backup repository, you must configure a Quantum DXi share in a proper way. For more information, see Quantum DXi documentation.

Storage Appliances

- Storage appliances that are used to store backup data in filer (CIFS/NFS) or block device mode (iSCSI/FC/SAS) are not supported if the backup data is offloaded to tapes and is no longer stored directly on the filer/block device (Hierarchical Storage Management with Tape tier).

  To offload data to tapes, make sure that:

  - All of the backup data is stored on the appliance altogether (that is, all of the backup chains are stored on the appliance as a whole and not scattered across multiple devices) and only copies are stored on tapes.
  - These appliances emulate a tape system (VTL) as an access protocol for Veeam Backup & Replication.

> **NOTE**
>
> Consider reading the Considerations and Limitations section that describes prerequisites for object storage.

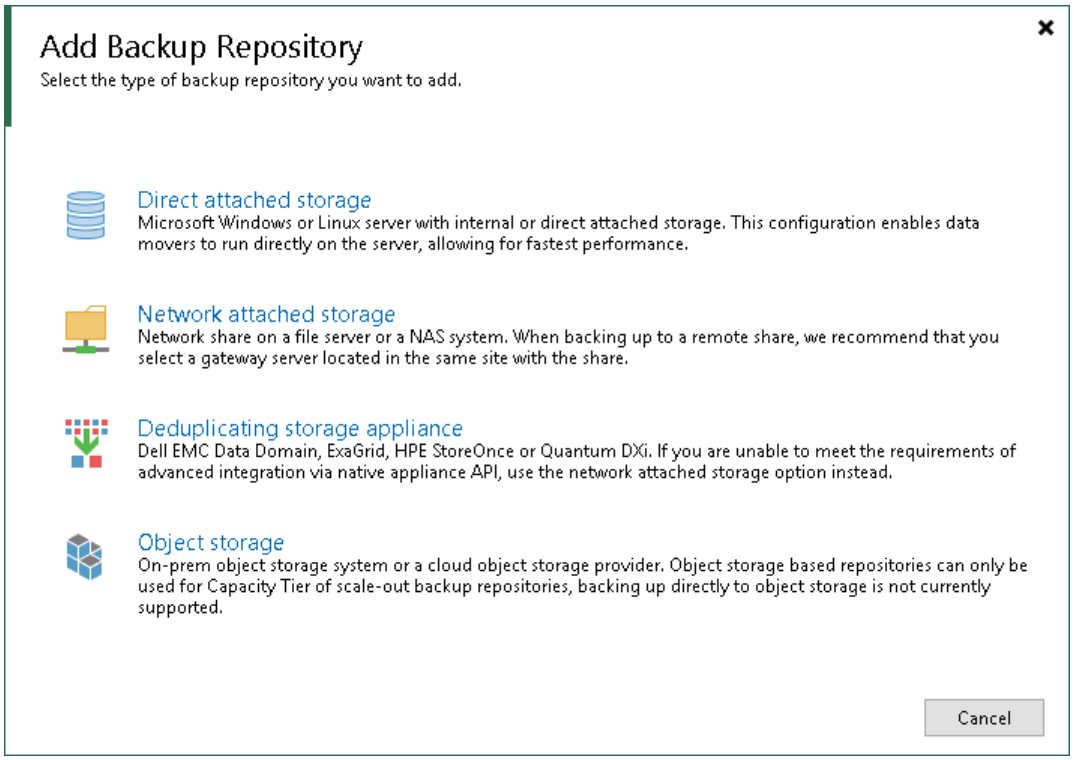
Step 1. Launch New Backup Repository Wizard

To launch the **New Backup Repository** wizard, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, right-click the **Backup Repositories** node and select **Add Backup Repository**. Alternatively, you can click **Add Repository** on the ribbon.
3. In the **Add Backup Repository** window, select the type of the backup repository you want to add.

The **New Backup Repository** wizard will guide you through steps for adding direct attached storage, network attached storage, and deduplicating storage appliances as backup repositories.

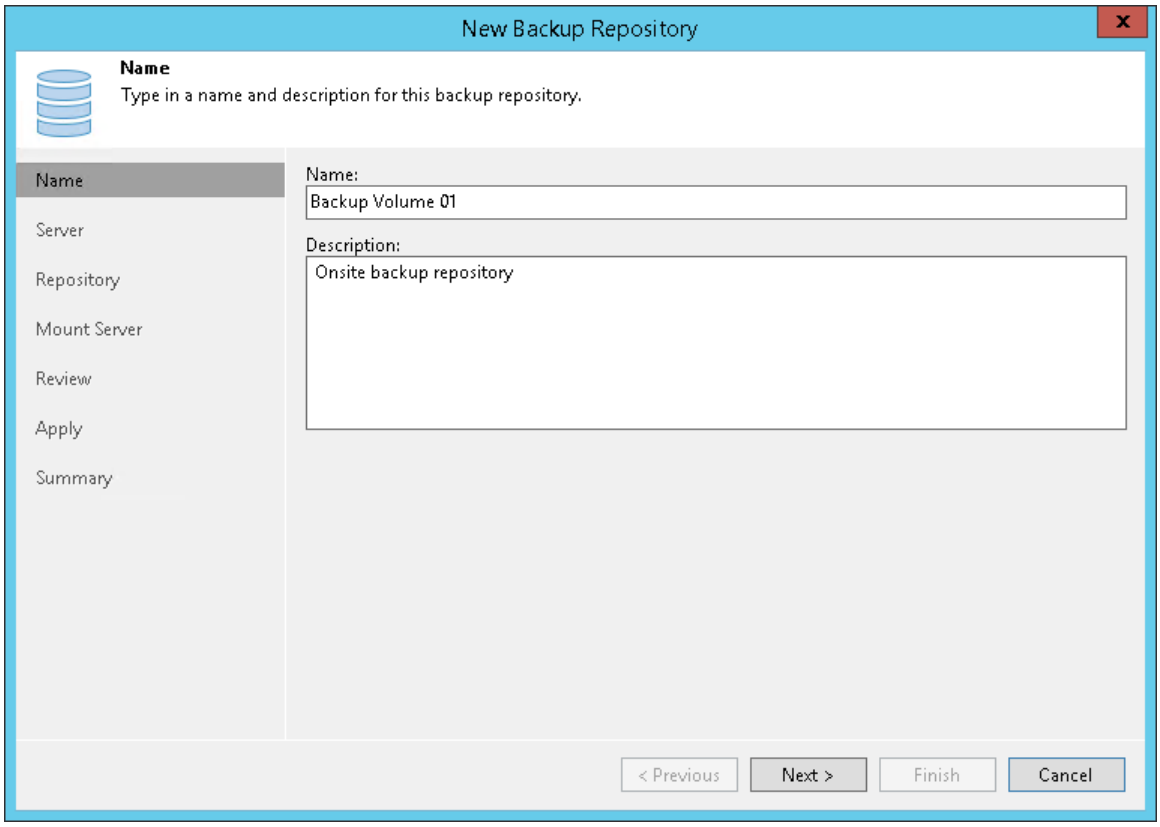
For information on how to add object storage repositories, see Adding Object Storage Repositories.



Step 2. Specify Backup Repository Name and Description

At the **Name** step of the wizard, specify a name and description for the backup repository.

1. In the **Name** field, specify a name for the backup repository.
2. In the **Description** field, provide a description for future reference. The default description contains information about the user who added the backup repository, date and time when the backup repository was added.



Step 3. Specify Server or Shared Folder Settings

Options that you can specify at the **Server** step of the wizard depend on the type of backup repository you are adding.

In this section:

- Microsoft Windows or Linux Server
- SMB Share
- NFS Share
- Dell EMC Data Domain
- ExaGrid or Quantum DXi
- HPE StoreOnce

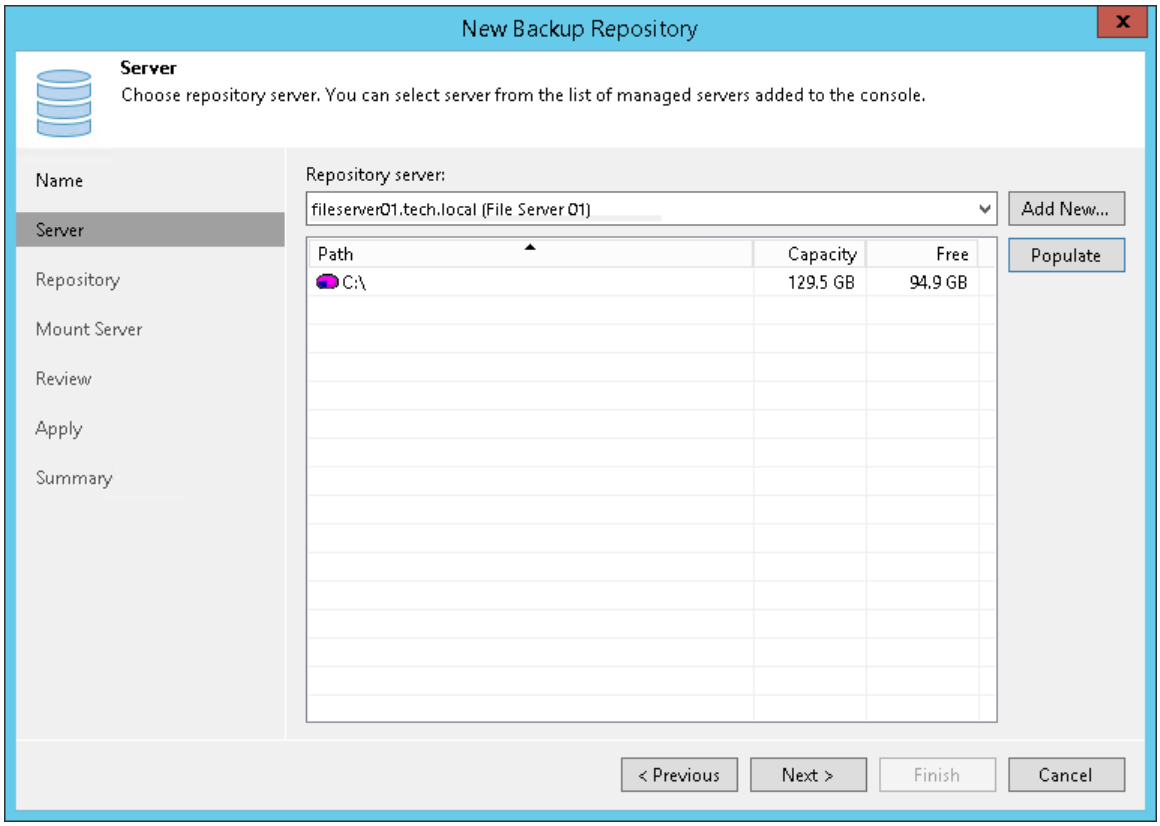
Microsoft Windows or Linux Server

To configure settings for a Microsoft Windows or Linux server:

1. From the **Repository server** list, select a Microsoft Windows or Linux server that you want to use as a backup repository. The **Repository server** list contains only those servers that are added to the backup infrastructure. If the server is not added to the backup infrastructure yet, you can click **Add New** on the right to open the **New Windows Server** or **New Linux Server** wizard.

   Note that you cannot add ExaGrid or Quantum DXi servers as Linux backup repositories. ExaGrid and Quantum DXi are integrated with Veeam Backup & Replication, and thus must be added as deduplicating storage appliances.

2. Click **Populate** to see a list of disks connected to the server, their capacity and free space.



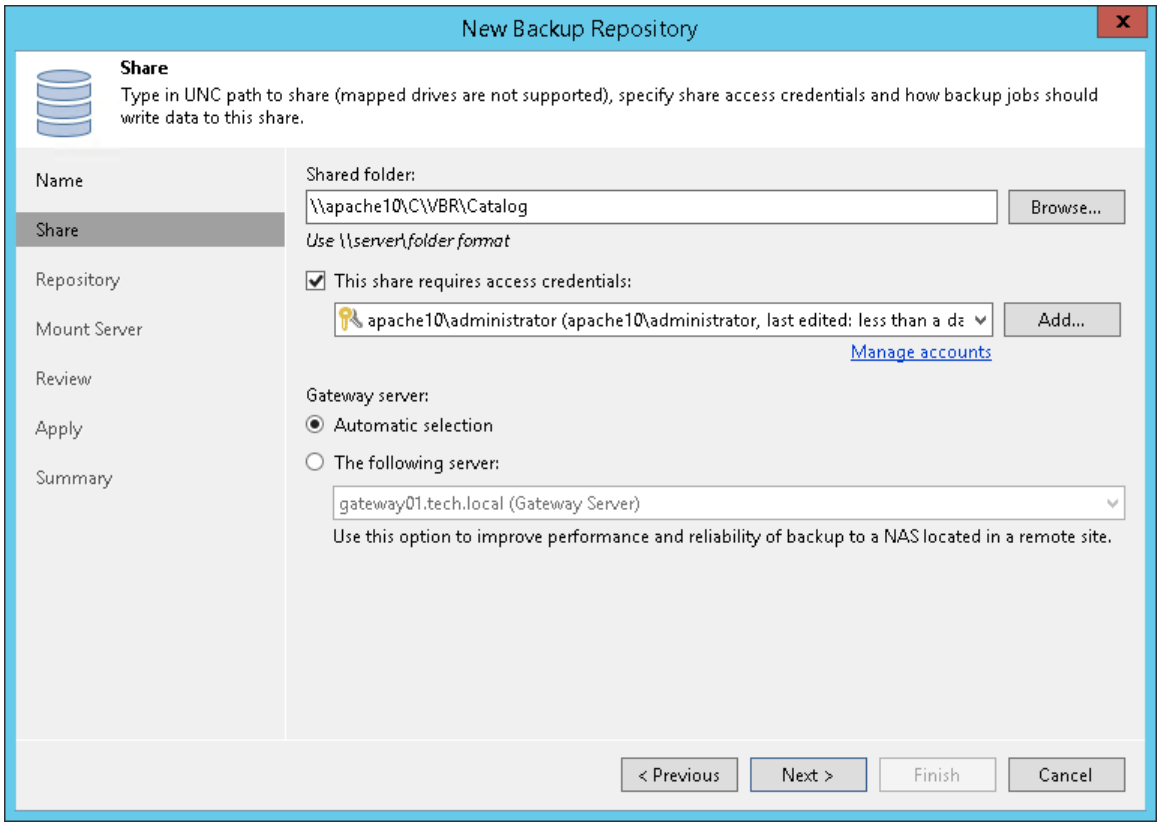
SMB Share

To configure settings for an SMB share:

1. In the **Shared folder** field, specify a UNC path to the SMB shared folder that you want to use as a backup repository.
2. If you must specify user credentials to access the shared folder, select the **This share requires access credentials** check box. From the **Credentials** list, select a credentials record for a user account that has Full Control permissions on the shared folder. Note that the username must be in the down-level logon name format. For example, *DOMAIN\username* or *HOSTNAME\username*.

   If you have not set up credentials beforehand, click the **Manage accounts** link at the bottom of the list or click **Add** on the right to add the credentials. For more information, see Managing Credentials.
3. In the **Gateway server** section, specify settings for the gateway server:
   - If a network connection between the source datastore and backup repository is fast, choose **Automatic selection**. In this case, Veeam Backup & Replication will automatically select a gateway server.
   - If you perform backup and replication over WAN or slow connections, choose **The following server**. From the list below, select a Microsoft Windows server on the target site that you want to use as a gateway server. The selected server must have a direct access to the SMB share and must be located as close to the SMB share as possible.

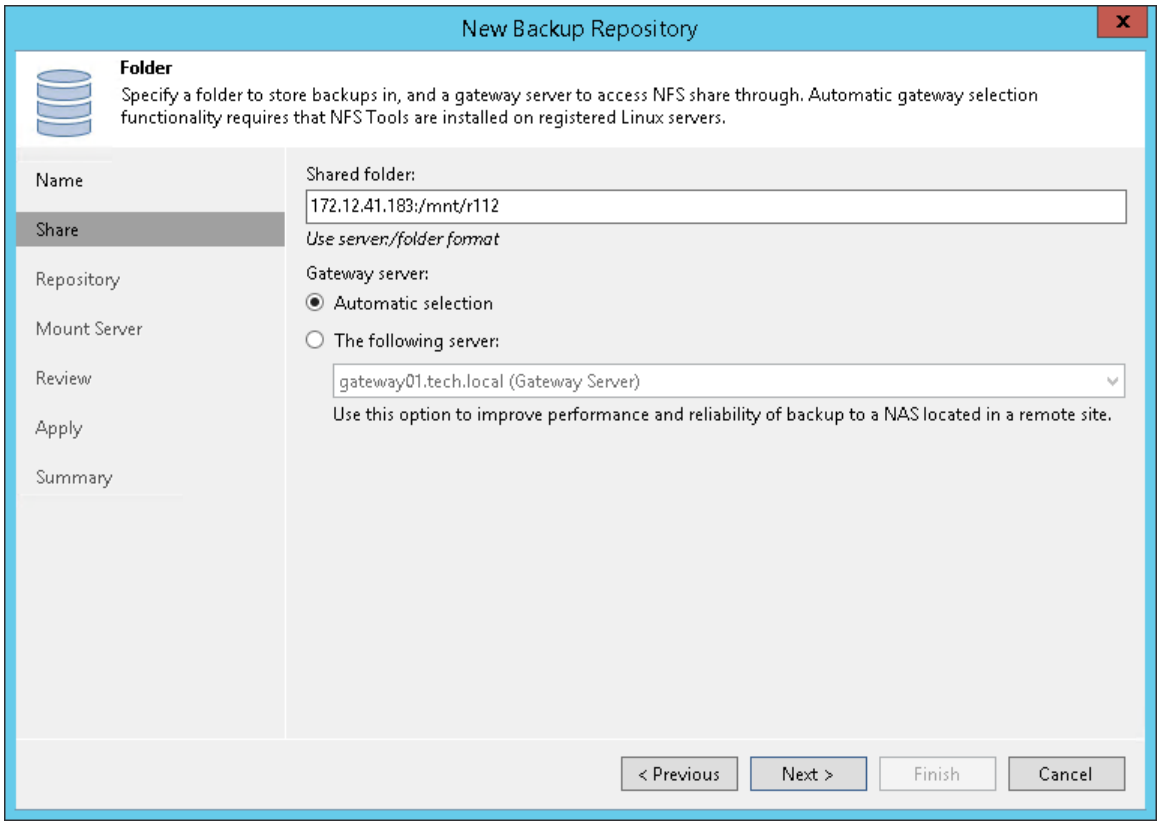
In some cases, the automatic selection mechanism may cause problems as Veeam Backup & Replication may use different gateway servers for different job sessions. For example, during one job session Veeam Backup & Replication may use a 64-bit gateway server to create a backup file. If during the next job session Veeam Backup & Replication uses a 32-bit gateway server, Veeam Backup & Replication will fail to open the created backup file in the backup repository. To overcome this situation, you must explicitly define the gateway server.



NFS Share

To configure settings for an NFS share:

1. In the **Shared folder** field, specify a path to the NFS shared folder that you want to use as a backup repository.
2. In the **Gateway server** section, specify settings for the gateway server:
   - If a network connection between the source datastore and backup repository is fast, choose **Automatic selection**. In this case, Veeam Backup & Replication will automatically select a gateway server.
   - If you perform backup and replication over WAN or slow connections, choose **The following server**. From the list below, select a Microsoft Windows or a Linux server on the target site that you want to use as a gateway server. The selected server must have a direct access to the NFS share and must be located as close to the NFS share as possible.



Dell EMC Data Domain

To configure settings for Dell EMC Data Domain:

1. Specify connection settings for Dell EMC Data Domain:
   - If Dell EMC Data Domain works over TCP, in the **Type in Data Domain server name** field enter a full DNS name or IP address of the Dell EMC Data Domain server.
   - If Dell EMC Data Domain works over Fibre Channel, select the **Use Fibre Channel (FC) connectivity** check box. In the **Type in Data Domain server name** field, enter a name of the Data Domain Fibre Channel server. To get the Data Domain Fibre Channel server name, in Data Domain System Manager open the **Data Management** > **DD Boost** > **Fibre Channel** tab.

2. In the **Credentials** field, specify credentials of the user account to connect to the Dell EMC Data Domain server or Dell EMC Data Domain Fibre Channel server. If you have not set up credentials beforehand, click the **Manage accounts** link at the bottom of the list or click **Add** on the right to add the credentials. For more information, see Managing Credentials.

   To connect to the Dell EMC Data Domain server, you must use credentials for the DD Boost User. To specify the DD Boost User account settings, in Data Domain System Manager, open the **Data Management** > **DD Boost Settings** tab.

3. To use in-flight encryption between the backup proxy and Dell EMC Data Domain, select the **Enable DDBoost encryption** check box and choose the encryption level — *Medium* or *High*. The encryption option works for Dell EMC Data Domain version 5.5 to 7.0.

4. In the **Gateway server** section, specify settings for the gateway server:

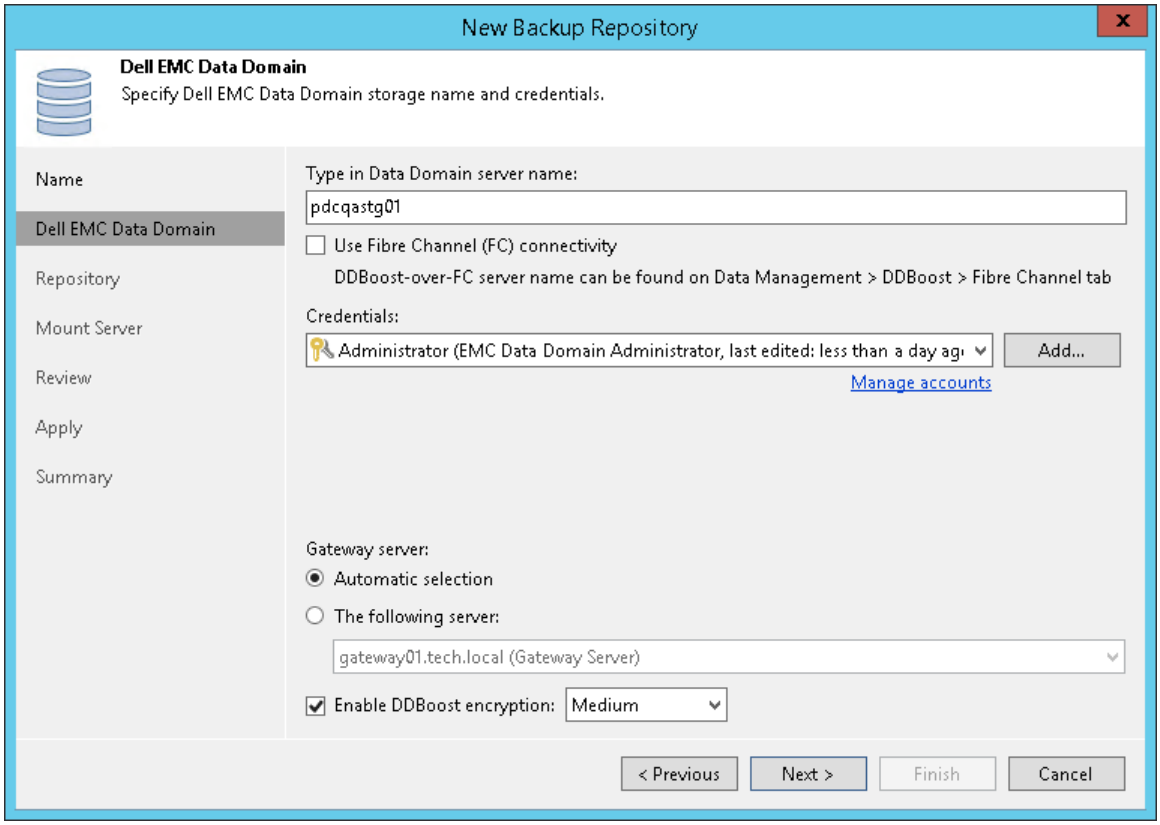
   - If a network connection between the source datastore and Dell EMC Data Domain appliance is fast, choose **Automatic selection**. In this case, Veeam Backup & Replication will automatically select a gateway server.

   - If you perform backup over WAN or slow connections, choose **The following server**. From the list below, select a Microsoft Windows server on the target site that you want to use as a gateway server. The server must have a direct access to the Dell EMC Data Domain appliance and must be located as close to the appliance as possible.

In some cases, the automatic selection mechanism may cause problems as Veeam Backup & Replication may use different gateway servers for different job sessions. For example, during one job session Veeam Backup & Replication may use a 64-bit gateway server to create a backup file. If during the next job session Veeam Backup & Replication uses a 32-bit gateway server, Veeam Backup & Replication will fail to open the created backup file in the backup repository. To overcome this situation, you can explicitly define the gateway server.

IMPORTANT

If you connect to Dell EMC Data Domain over Fibre Channel, you must explicitly define the gateway server to communicate with Dell EMC Data Domain. The server you select must be added to the backup infrastructure and must have access to the Dell EMC Data Domain appliance over Fibre Channel.

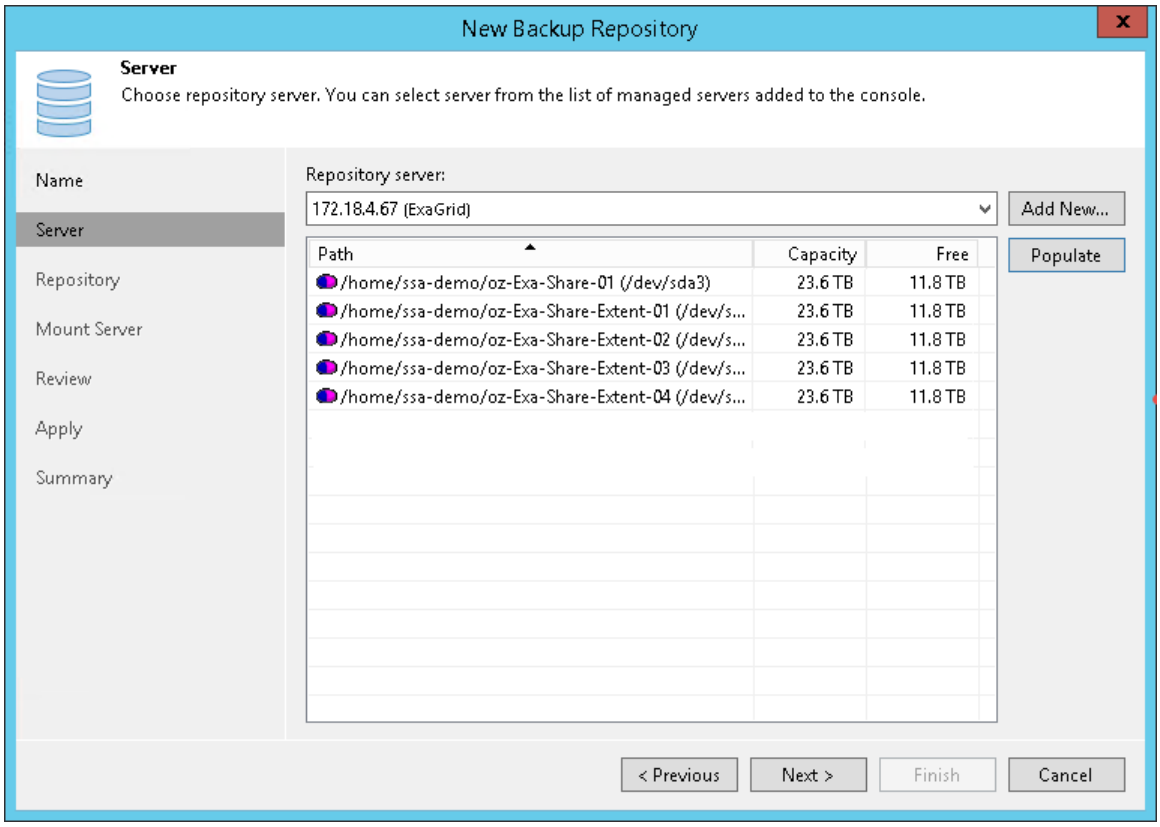


ExaGrid or Quantum DXi Deduplicating Appliance

To configure settings for ExaGrid or Quantum DXi deduplicating appliance:

1. From the **Repository server** list, select an appliance that you want to use as a backup repository. The **Repository server** list contains only those servers that are added to the backup infrastructure. If the server is not added to the backup infrastructure yet, you can click **Add New** to open the **New Linux Server** wizard. For more information, see Managing Servers.

2. Click **Populate** to see the appliance capacity and available free space.



HPE StoreOnce Deduplicating Appliance

To configure settings for HPE StoreOnce:

1. In the **Type in HPE StoreOnce server name** field, enter a full DNS name or IP address of the HPE StoreOnce appliance.
2. If HPE StoreOnce works over Fibre Channel, select the **Use Fibre Channel (FC) connectivity** check box.
3. In the **Credentials** field, specify credentials of the client account to connect to the HPE StoreOnce appliance. If you have not set up credentials beforehand, click the **Manage accounts** link at the bottom of the list or click **Add** on the right to add the credentials. For more information, see Managing Credentials.

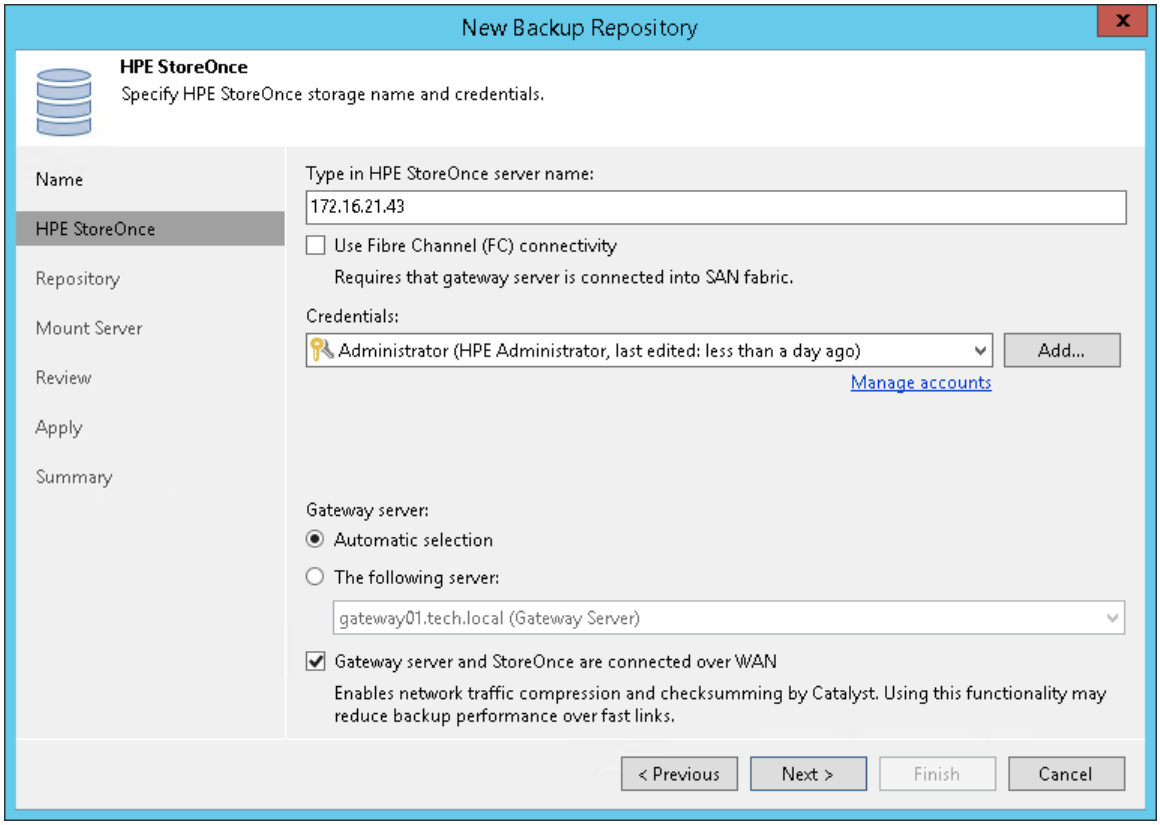
   The client account that you plan to use to connect to HPE StoreOnce must have access permissions on a Catalyst store where backup data will be kept. To check the client account permissions, in the HPE StoreOnce management console, select the Catalyst store and open the **Permissions** tab for it.
4. In the **Gateway server** section, specify settings for the gateway server:
   - If you want Veeam Backup & Replication to pick the gateway server automatically, choose **Automatic selection**. In this case, Veeam Backup & Replication will automatically select a gateway server.
   - If you want to define the gateway server explicitly, choose **The following server**. From the list below, select a Microsoft Windows server that you want to use as a gateway server.

   In some cases, the automatic selection mechanism may cause problems as Veeam Backup & Replication may use different gateway servers for different job sessions. For example, during one job session Veeam Backup & Replication may use a 64-bit gateway server to create a backup file. If during the next job session Veeam Backup & Replication uses a 32-bit gateway server, Veeam Backup & Replication will fail to open the created backup file in the backup repository. To overcome this situation, you can explicitly define the gateway server.

IMPORTANT

If you connect to HPE StoreOnce over Fibre Channel, you must explicitly define the gateway server to communicate with HPE StoreOnce appliance. The server you select must be added to the backup infrastructure and must have access to the HPE StoreOnce appliance over Fibre Channel.

5. If a WAN connection between the gateway server and the HPE StoreOnce appliance is weak, select the **Gateway server and StoreOnce are connected over WAN** check box. Veeam Backup & Replication will compress VM data transported from the gateway server to the HPE StoreOnce appliance, and calculate checksums for data blocks going from the gateway server to the HPE StoreOnce appliance.



Step 4. Configure Backup Repository Settings

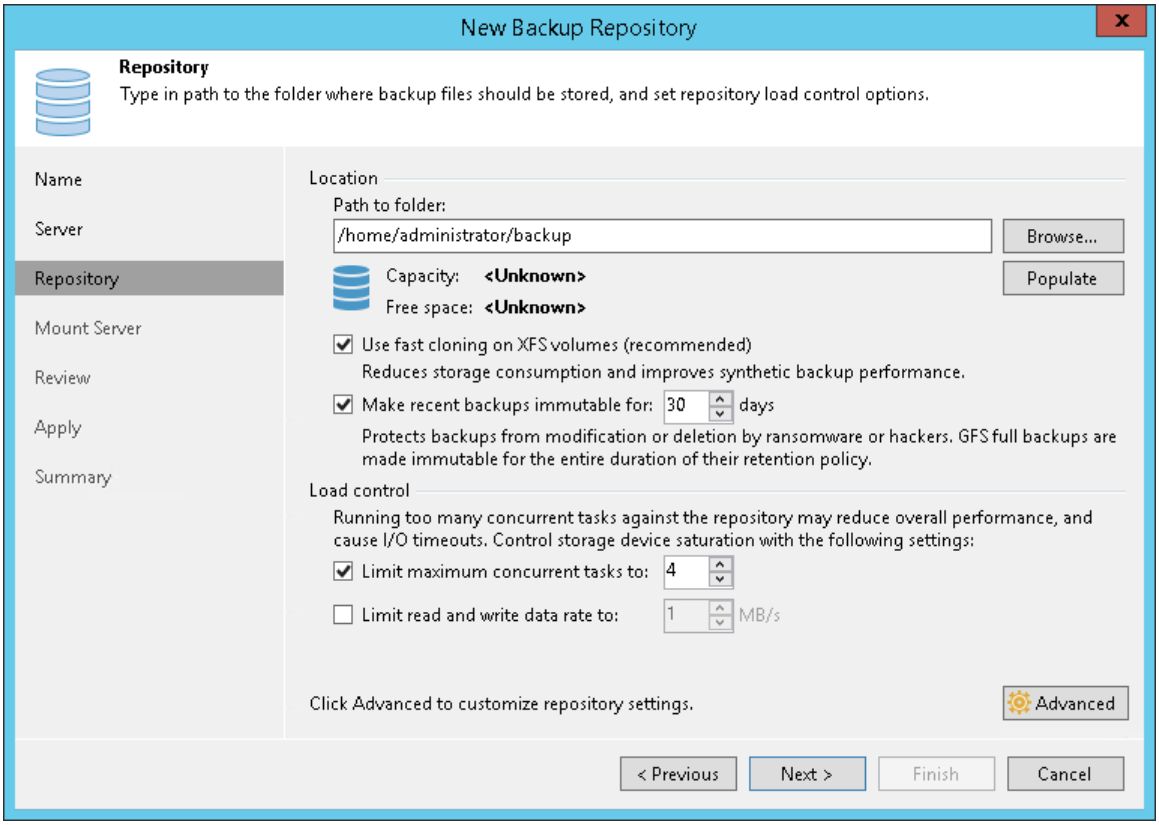
At the **Repository** step of the wizard, specify path and load control repository settings.

1. In the **Location** section, specify a path to the folder where backup files must be stored. Click **Populate** to check capacity and available free space in the selected location.
   - For Dell EMC Data Domain, click **Browse** and select a location from the list of available paths.
   - For HPE StoreOnce, select a Catalyst store from the list.
2. [For Linux repository] Select the **Use fast cloning on XFS volumes** check box to enable copy-on-write functionality. In terms of Veeam Backup & Replication, this functionality is known as Fast Clone. For more information, see Fast Clone.
3. [For hardened repository] To prohibit deletion of blocks of data from the backup repository, select the **Make recent backups immutable for** check box and specify the immutability period. For more information, see Hardened Repository.
4. Use the **Load control** section to limit the number of concurrent tasks and data ingestion rate for the backup repository. These settings will help you control the load on the backup repository and prevent possible timeouts of storage I/O operations.
   - Select the **Limit maximum concurrent tasks** check box and specify the maximum allowed number of concurrent tasks for the backup repository. If this value is exceeded, Veeam Backup & Replication will not start a new task until one of current tasks finishes. For more information, see Limiting the Number of Concurrent Tasks.

     > **IMPORTANT**
     >
     > Limitation of concurrent tasks is not supported for vCD backup jobs.

   - Select the **Limit read and write data rates to** check box and specify the maximum rate to restrict the total speed of reading and writing data to the backup repository disk. For more information, see Limiting Combined Data Rate for Backup Repositories.

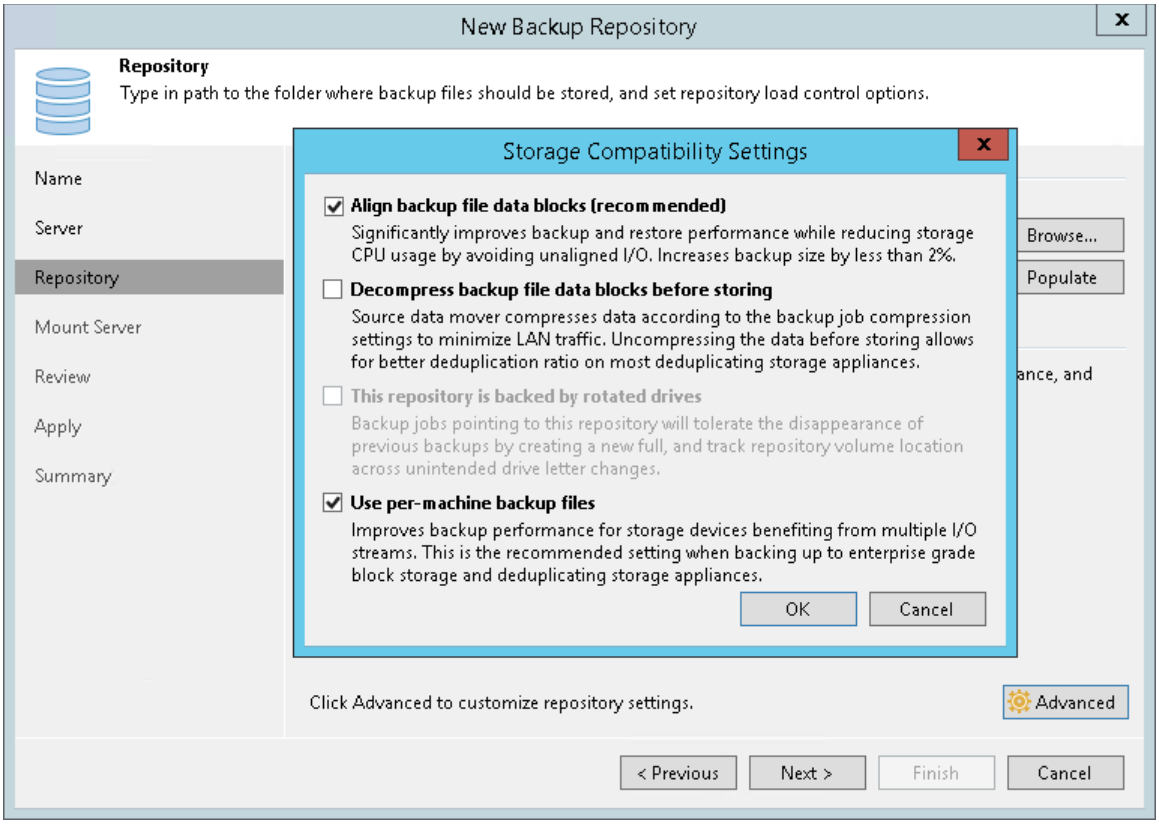
> **NOTE**
>
> The **Limit read and write data rates to** settings does not apply to health checks performed as part of backup and backup copy jobs. Even if you limit read/write rate for a backup repository, the health check will consume resources of the backup repository regardless of this setting. Bear this limitation in mind when configuring basic and health check schedules for backup and backup copy jobs.



5. Click **Advanced** to configure additional settings for the backup repository:
   - For storage systems using a fixed block size, select the **Align backup file data blocks** check box. Veeam Backup & Replication will align VM data saved to a backup file at a 4 KB block boundary.
   - When you enable compression for a backup job, Veeam Backup & Replication compresses VM data at the source side and then transports it to the target side. Writing compressed data to a deduplicating storage appliance results in poor deduplication ratios as the number of matching blocks decreases. To overcome this situation, select the **Decompress backup data blocks before storing** check box. If data compression is enabled for a job, Veeam Backup & Replication will compress VM data on the source side, transport it to the target side, decompress VM data on the target side and write raw VM data to the storage device to achieve a higher deduplication ratio.

     > **NOTE**
     >
     > Veeam Backup & Replication does not compress VM data if encryption is enabled for a job, and the **Decompress backup data blocks before storing** check box is selected in the settings of the target backup repository. Therefore, in the job statistics, you may observe a higher amount of transferred data (the **Transferred** counter) as compared to a job for which encryption is disabled. For details on job statistics, see Viewing Real-Time Statistics.

- If you plan to use rotated drives for the backup repository, select the **This repository is backed by rotated hard drives** check box. For more information, see Configuring Backup Repositories with Rotated Drives.
- To create a separate backup file for every machine in the job, select the **Use per-machine backup files** check box. This setting is recommended if you use a deduplicating storage appliance as a backup repository. Veeam Backup & Replication will write machine data to the backup repository in several streams, which will improve the backup job performance. However, in this case Veeam Backup & Replication will not deduplicate data between VMs added to the job. For more information, see Per-machine Backup Files.



Settings for Deduplicating Storage Appliances

If you use a deduplicating storage appliance as a backup repository, Veeam Backup & Replication automatically sets advanced settings to the following ones:

Dell EMC Data Domain

- The **Align backup file data blocks** option is disabled and cannot be changed.
- The **Decompress backup data blocks before storing** option is enabled.
- The **This repository is backed by rotated hard drives** option is disabled and cannot be changed.
- The **Use per-machine backup files** option is enabled.

ExaGrid

- The **Align backup file data blocks** option is disabled and must not be changed.
- The **Decompress backup data blocks before storing** option is disabled.

- The **This repository is backed by rotated hard drives** option is disabled and cannot be changed.
- The **Use per-machine backup files** option is enabled.
- **Limit max concurrent tasks** is equal to 10 (recommended).

Quantum DXi

- The **Align backup file data blocks** option is disabled and must not be changed.
- The **Decompress backup data blocks before storing** option is enabled.
- The **This repository is backed by rotated hard drives** option is disabled and cannot be changed.
- The **Use per-machine backup files** option is enabled.

HPE StoreOnce

- The **Align backup file data blocks** option is disabled and must not be changed.
- The **Decompress backup data blocks before storing** option is enabled.
- The **This repository is backed by rotated hard drives** option is disabled and cannot be changed.
- The **Use per-machine backup files** option is enabled and cannot be changed.

Step 5. Specify Mount Server Settings

At the **Mount Server** step of the wizard, specify settings for the mount server that you plan to use for file-level and application items restore.

1. From the **Mount Server** list, select a server that you want to use as a mount server. The mount server is required for file-level and application items restore. During the restore process, Veeam Backup & Replication will mount the VM disks from the backup file residing in the backup repository to the mount server. As a result, VM data will not have travel over the network, which will reduce the load on the network and speed up the restore process. For more information, see Mount Server.

   The **Mount Server** list contains only Microsoft Windows servers that are added to the backup infrastructure. If the server is not added to the backup infrastructure yet, click **Add New** on the right to open the **New Windows Server wizard**. For more information, see Adding Microsoft Windows Servers.

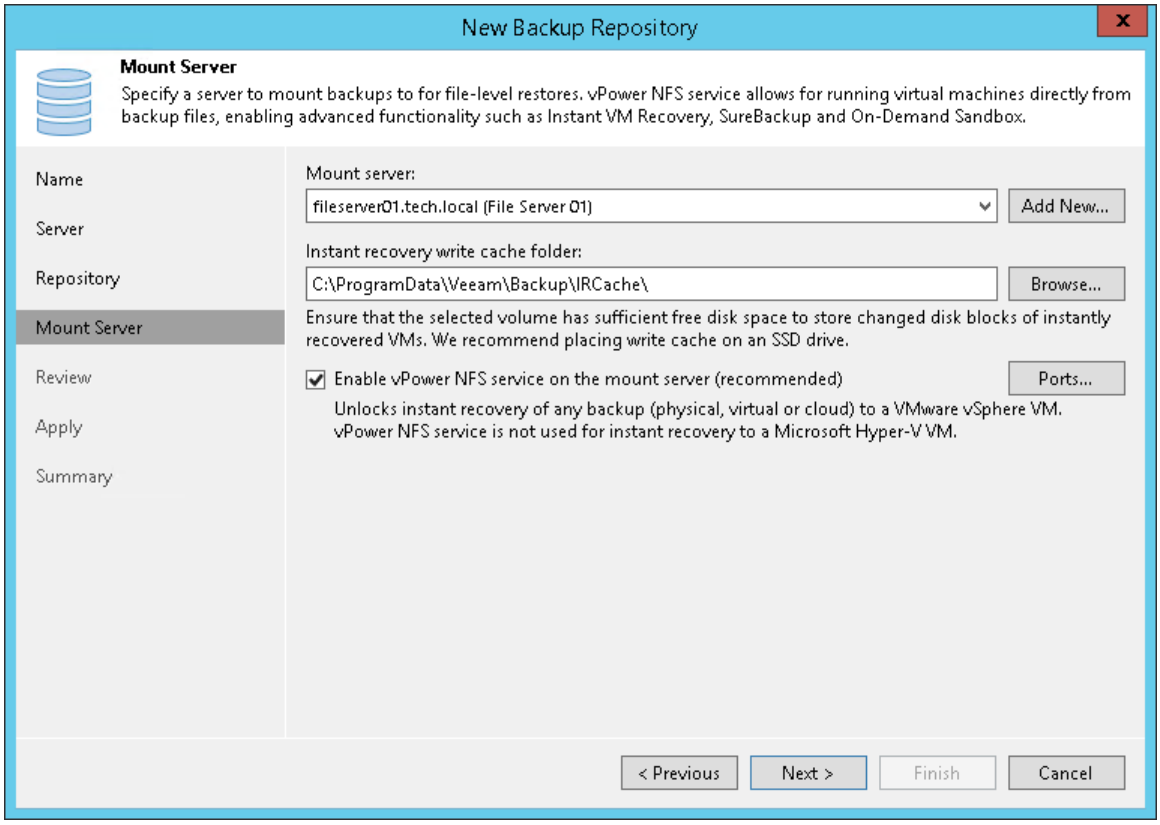
2. In the **Instant recovery write cache folder** field, specify a folder that will be used for writing cache during mount operations.

3. To make the backup repository accessible by the Veeam vPower NFS Service, select the **Enable vPower NFS service on the mount server** check box. Veeam Backup & Replication will enable the vPower NFS Service on the mount server you have selected.

4. To customize network ports used by the vPower NFS Service, click **Ports**. For information on ports used by default, see Used Ports.

IMPORTANT

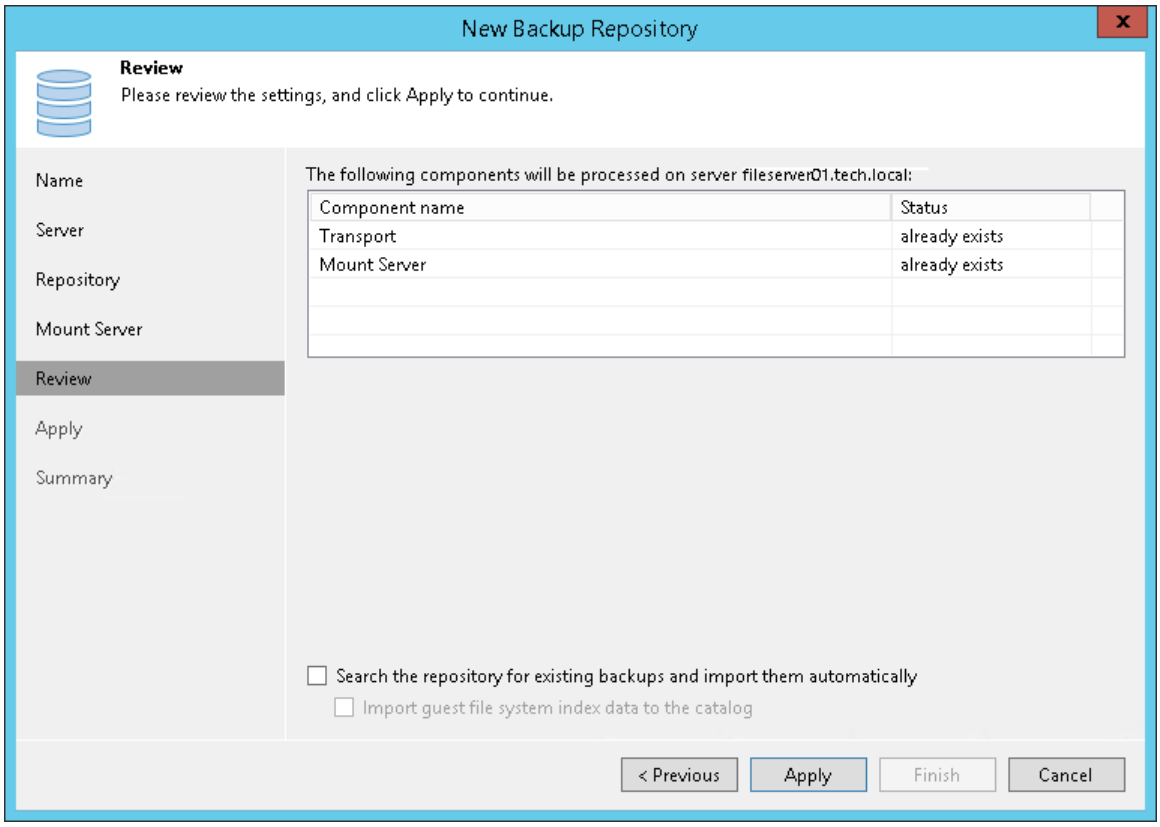
Do not enable Microsoft Windows NFS services on the machine where you install the Veeam vPower NFS Service. If Microsoft NFS services and Veeam vPower NFS Service are enabled on the same machine, both services may fail to work correctly.



Step 6. Review Properties and Components

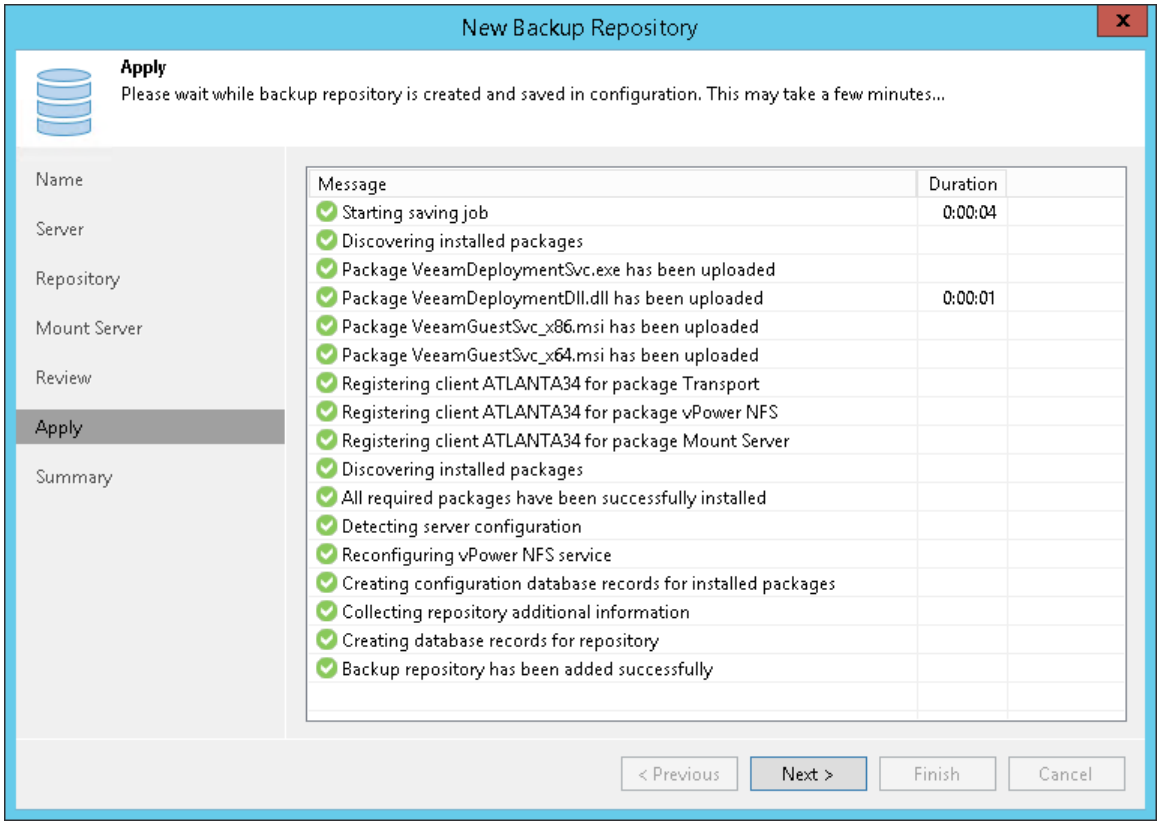
At the **Review** step of the wizard, review details of the backup repository and specify importing settings.

1. Review the backup repository settings and list of components that will be installed on the backup repository server.
2. If the backup repository contains backups that were previously created with Veeam Backup & Replication, select the **Search the repository for existing backups and import them automatically** check box. Veeam Backup & Replication will scan the backup repository to detect existing backup files and display them in the Veeam Backup & Replication console under the **Imported > Backups** node.
3. If the backup repository contains guest file system index files that were previously created by Veeam Backup & Replication, select the **Import guest file system index** check box. Index files will be imported with backup files, and you will be able to search for guest OS files inside imported backups.



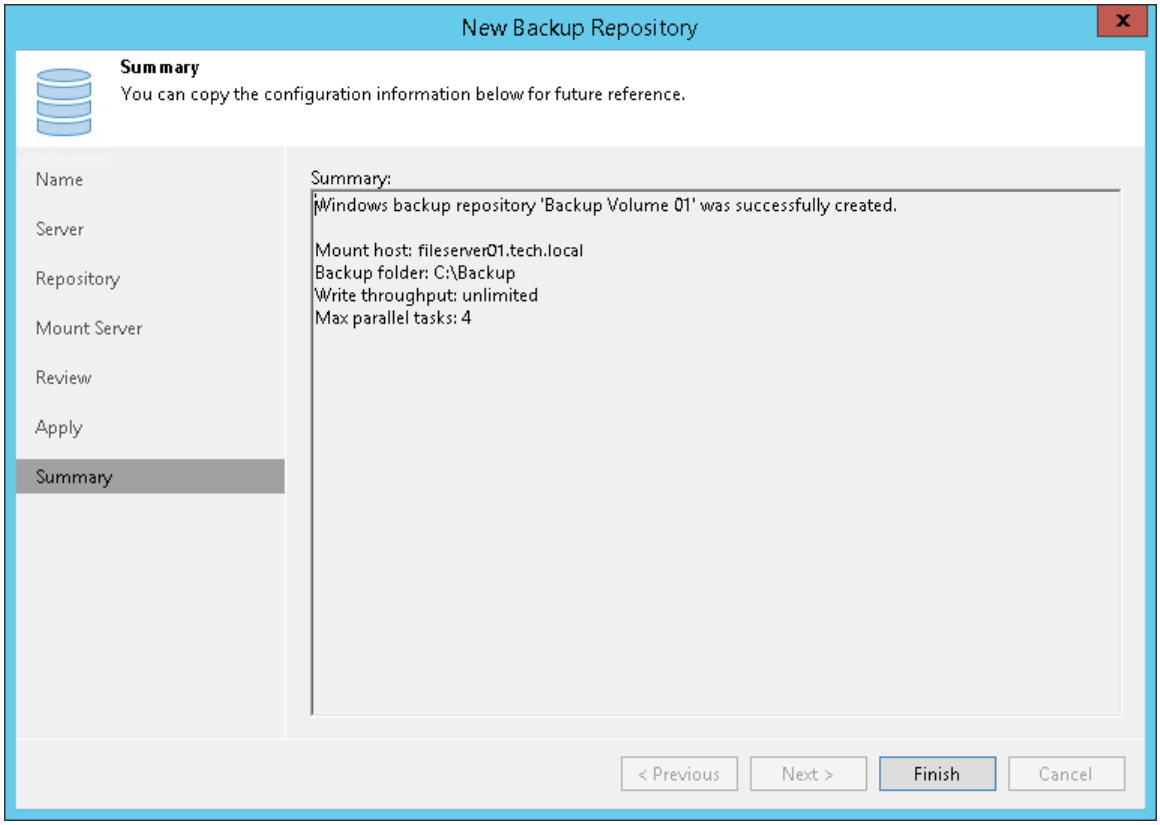
Step 7. Apply Backup Repository Settings

At the **Apply** step of the wizard, wait for Veeam Backup & Replication to install and configure all required components. Then click **Next** to complete the procedure of adding the backup repository to the backup infrastructure.



Step 8. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the added backup repository. Then click **Finish** to exit the wizard.



Managing Backup Repositories

You can manage your backup repositories and data stored in them in various ways: edit settings of the backup repositories, set up access permissions, rescan backup repositories or remove backup repositories.

In This Section

- Editing Settings of Backup Repositories
- Access Permissions
- Rescanning Backup Repositories
- Removing Backup Repositories
- Fast Clone
- Proxy Affinity

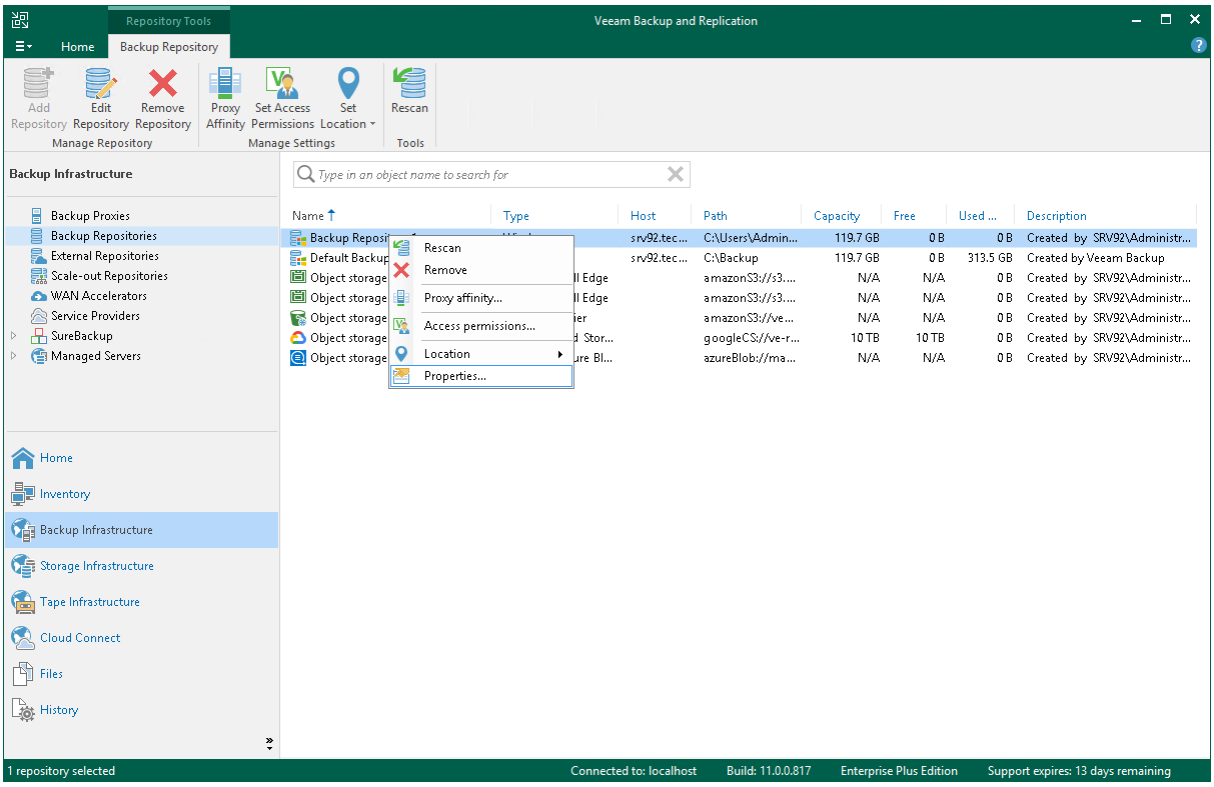
Editing Settings of Backup Repositories

You can edit settings of backup repositories that you have added to the backup infrastructure.

> **NOTE**
>
> When editing backup repository settings, you cannot change the selected backup repository server. Therefore, the **Repository server** field at the **Server** step is grayed out.

To edit settings of a backup repository:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select the **Backup Repositories** node.
3. In the working area, select the backup repository and click **Edit Repository** on the ribbon or right-click the backup repository and select **Properties**.
4. Edit the backup repository settings as required. Mind that you cannot change the selected repository server and path to the folder used for storing backups.



Access Permissions

If you want to store in the backup repository backups of virtual and physical machines created with additional components of Veeam Backup & Replication, for example, Veeam Agent for Microsoft Windows, Veeam Agent for Linux and so on, you need to set up access permissions on backup repositories.

Access permissions are granted to security principals such as users and AD groups by the backup administrator working with Veeam Backup & Replication. Users with granted access permissions can target Veeam Agent backup jobs at this backup repository and perform restore from backups located in this backup repository.

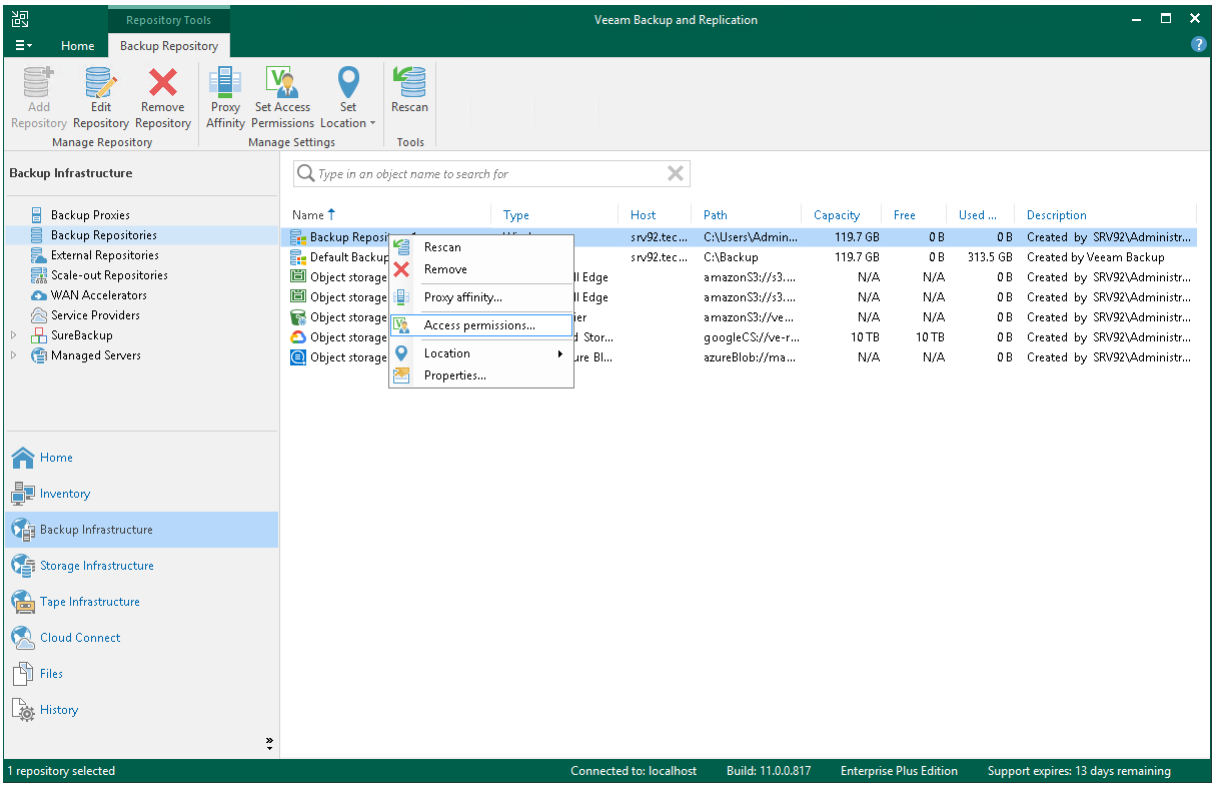
> **NOTE**
>
> If you plan to create backups in a Veeam backup repository with Veeam Agent backup jobs configured in Veeam Backup & Replication, you do not need to grant access permissions on the backup repository to users. In the Veeam Agent management scenario, to establish a connection between the backup server and protected computers, Veeam Backup & Replication uses a TLS certificate. To learn more, see the Configuring Security Settings section in the Veeam Agent Management Guide.

Right after installation, access permissions on the default backup repository are set to *Allow to everyone* for testing and evaluation purposes. If necessary, you can change these settings.

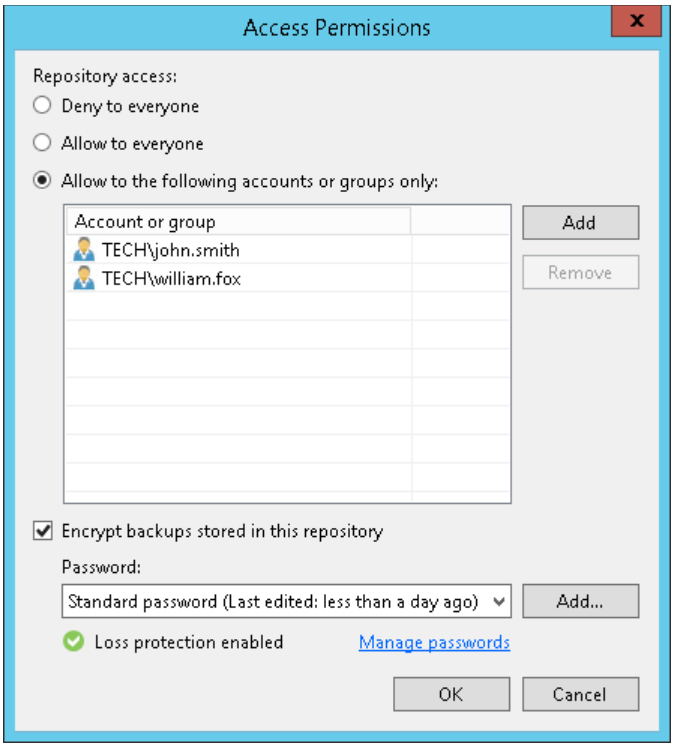
After you create a new backup repository, access permissions on this repository are set to *Deny to everyone*. To allow users to store backups in the backup repository, you must grant users with access permissions to this repository.

To grant access permissions to a security principal:

1. In Veeam Backup & Replication, open the **Backup Infrastructure** view.
2. In the inventory pane, click one of the following nodes:
   - The **Backup Repositories** node — if you want to grant access permissions on a regular backup repository.
   - The **Scale-out Repositories** node — if you want to grant access permissions on a scale-out backup repository.
3. In the working area, select the necessary backup repository and click **Set Access Permissions** on the ribbon or right-click the backup repository and select **Access permissions**. If you do not see the **Set Access Permissions** button on the ribbon or the **Access permissions** command is not available in the shortcut menu, press and hold the [CTRL] key, right-click the backup repository and select **Access permissions**.



4. In the **Access Permissions** window, specify to whom you want to grant access permissions on this backup repository:
   - **Allow to everyone** – select this option if you want all users to be able to store backups in this backup repository. Setting access permissions to *Everyone* is equal to granting access rights to the *Everyone* Microsoft Windows group (*Anonymous* users are excluded). Note, however, this scenario is recommended for demo environments only.
   - **Allow to the following accounts or groups only** – select this option if you want only specific users to be able to store backups in this backup repository. Click **Add** to add the necessary users and groups to the list.
5. If you want to encrypt Veeam Agent backup files stored in the backup repository, select the **Encrypt backups stored in this repository** check box and choose the necessary password from the field below. If you have not specified a password beforehand, click **Add** on the right or the **Manage passwords** link to add a new password. Veeam Backup & Replication will encrypt files at the backup repository side using its built-in encryption mechanism.



Rescanning Backup Repositories

You can rescan a backup repository configured in the backup infrastructure. Backup repository rescan may be required, for example, if you have archived backups from a backup repository to tape and deleted backup files in the backup repository. Or you have copied backups to the backup repository manually and want to work with them in Veeam Backup & Replication.

During the rescan operation, Veeam Backup & Replication gathers information about backups that are currently available in the backup repository and updates the list of backups in the configuration database. After the rescan operation, backups that were not in this configuration database will be shown on the **Home** view in the **Backups** > **Disk (Imported)** node. If backups are encrypted, they will be shown in the **Backups** > **Disk (Encrypted)** node.

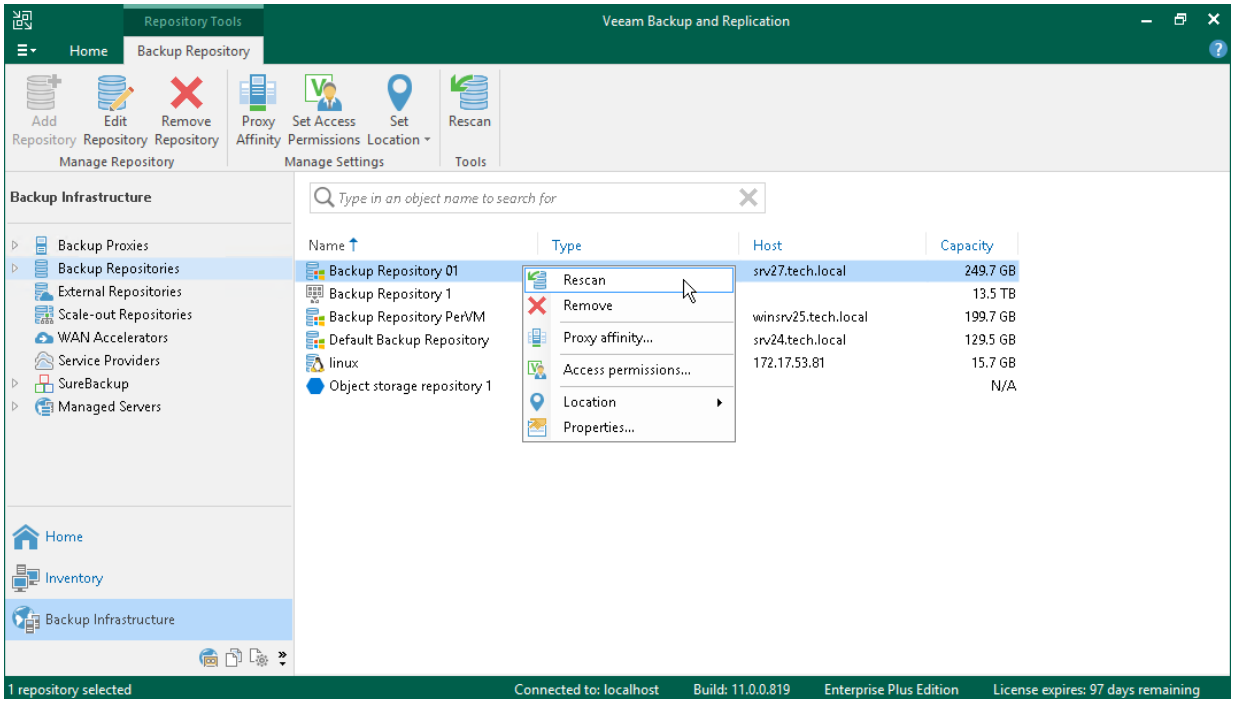
IMPORTANT

Mind the following:

- It is recommended that you stop or disable all jobs before performing the rescan.Veeam Backup & Replication skips from scanning backups created by active jobs.
- Veeam Backup & Replication will not be able to import backups automatically while performing a rescan if VBM files are not available. In this case you will have to import backups manually using the VBK files. For more information, see Importing Backups Manually.

To rescan a backup repository:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select the **Backup Repositories** node.
3. In the working area, select the backup repository and click **Rescan Repository** on the ribbon or right-click the backup repository and select **Rescan repository**.



Removing Backup Repositories

You can permanently remove a backup repository from the backup infrastructure. When you remove a backup repository, Veeam Backup & Replication unassigns the backup repository role from the server and this server is no longer used as a backup repository. The actual server remains in the backup infrastructure.

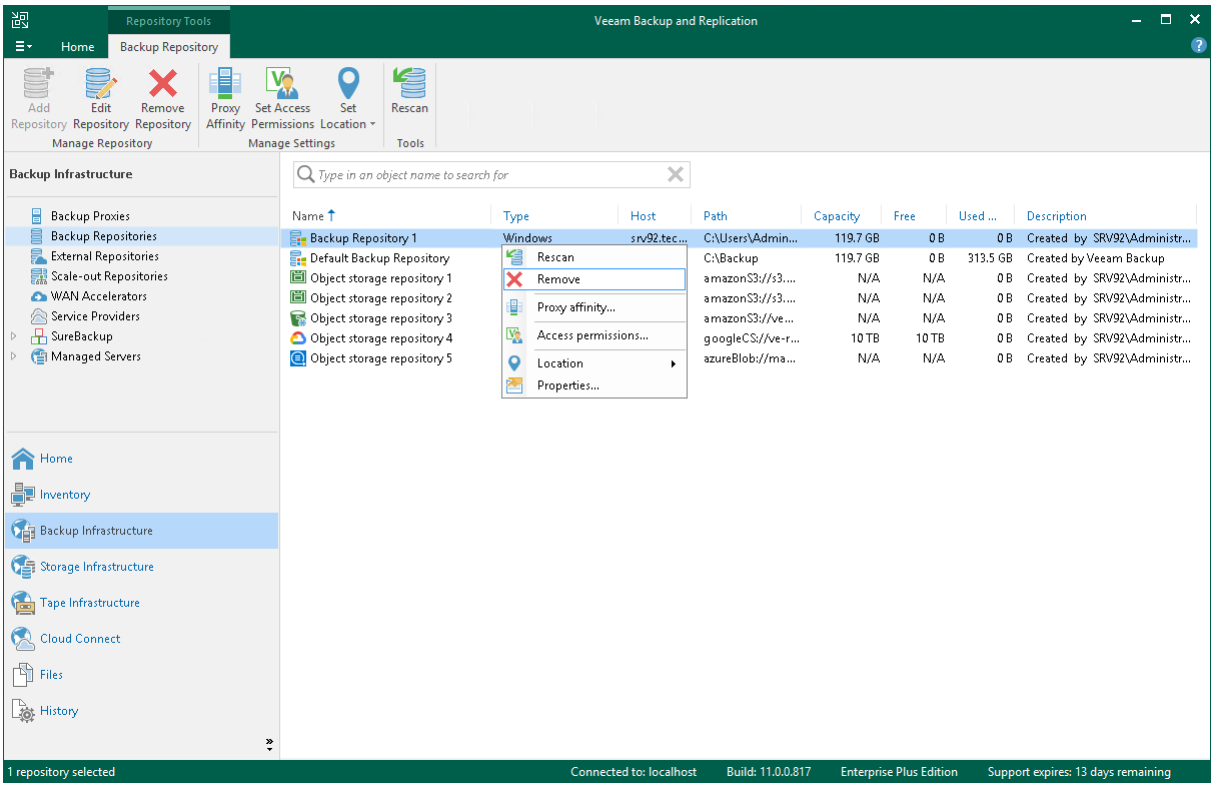
Veeam Backup & Replication does not remove backup files and other data stored in the backup repository. You can re-connect the backup repository at any time and import backups from this backup repository to Veeam Backup & Replication.

You cannot remove a backup repository that is used by any job. To remove such backup repository, you first need to delete a reference to this backup repository in the job settings.

To remove a backup repository:

1. Open the **Backup Infrastructure** view.

2. In the inventory pane, select the **Backup Repositories** node.
3. In the working area, select the backup repository and click **Remove Repository** on the ribbon or right-click the backup repository and select **Remove**.



Fast Clone

Fast Clone is the Veeam Backup & Replication technology that helps create quick file copies. Fast Clone increases the speed of synthetic backup creation and transformation, reduces disk space requirements and decreases the load on storage devices.

With this technology, Veeam Backup & Replication references existing data blocks on volumes instead of copying data blocks between files. Data blocks are copied only when files are modified.

Veeam Backup & Replication supports Fast Clone for the following types of backup repositories:

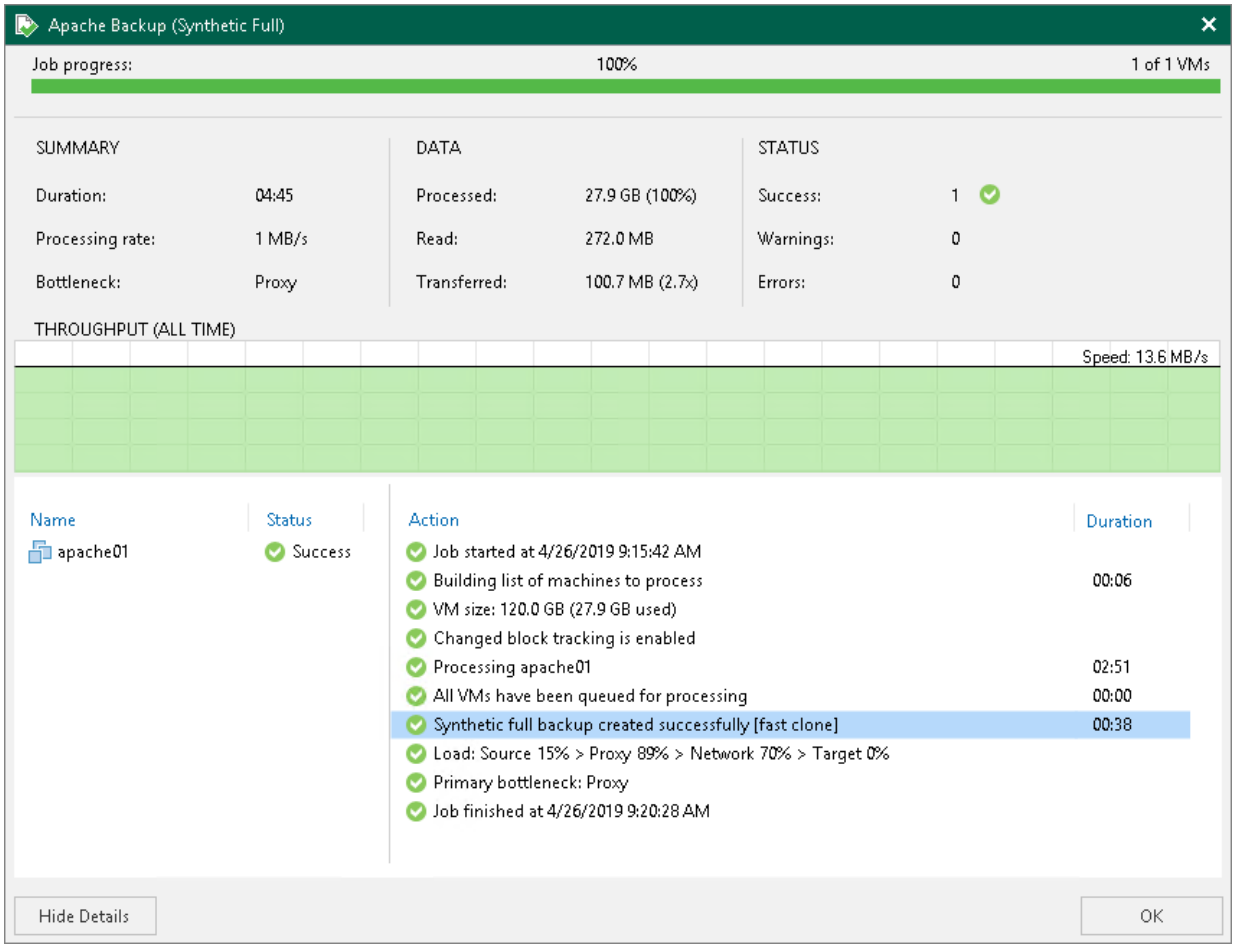
- Linux server
- Microsoft Windows server
- SMB share

Depending on the repository type, Fast Clone uses different technologies and has different requirements and limitations. For more information, see Fast Clone for Linux Repositories and Fast Clone for Microsoft Windows and SMB Repositories.

Veeam Backup & Replication uses Fast Clone for the following operations:

- **In backup jobs**: merge of backup files, creation of synthetic full backups (including GFS backups), transformation of reverse incremental backups and compact of full backup files.
- **In backup copy jobs**: merge of backup files, creation of GFS backups (synthetic method) and compact of full backup files.

When Veeam Backup & Replication performs an operation with Fast Clone, it reports this information to the session details of this operation.



Fast Clone for Linux Repositories

For Linux backup repositories, Fast Clone is based on the reflink technology.

Requirements for Linux Repositories

To use Fast Clone, Veeam Backup & Replication requires that Linux backup repositories meet the following conditions:

- Linux distribution has the reflink kernel module.

  Supported distributions: CentOS 8.2 and 8.3, Debian 10.x, RHEL 8.2 or later, SLES 15 SP2, Ubuntu 18.04 LTS and 20.04 LTS. For other distributions, Fast Clone support is experimental.

- File system is XFS.
- Cyclic redundancy check (CRC) is enabled.
- The minimum supported data block size is 1 KB. The maximum supported block size is 4KB.

Configuring a Linux Repository

To configure a Linux backup repository for work with Fast Clone, use the following XFS volume format string:

```
mkfs.xfs -b size=4096 -m reflink=1,crc=1 /dev/sda1
```

where:

- **size=4096** sets file system block size to 4096 bytes,
- **reflink=1** enables reflinking for the XFS instance (disabled by default),
- **crc=1** enables checksums, required for reflink=1 (enabled by default).

Limitations

After you have moved backup chains to a Linux backup repository with Fast Clone support, you must create active full backups for these chains to activate Fast Clone. You can also schedule the backup file compact operation instead of active full backup.

Fast Clone for Microsoft Windows and SMB Repositories

For Microsoft Windows and SMB backup repositories, Fast Clone is based on block cloning technology of Microsoft. For more information on block cloning, see Microsoft Docs.

By default, Veeam Backup & Replication uses Fast Clone for all Microsoft Windows and SMB backup repositories that meet the requirements. You can disable this option with a registry key. For more information, contact Veeam Customer Support.

Requirements for Microsoft Windows and SMB Repositories

Microsoft Windows Backup Repository

To use Fast Clone, Veeam Backup & Replication requires that Microsoft Windows backup repositories meet the following conditions:

- OS is Microsoft Windows Server 2016 (or later) or Microsoft Windows 10 Pro for Workstations (or later).
- File system is ReFS 3.1 (or later).

> **NOTE**
>
> All ReFS supported configurations must use Windows Server Catalog certified hardware. For other requirements, limitations and known issues, see this Veeam KB article.

Shared Folder Backup Repository

To use Fast Clone, Veeam Backup & Replication requires that SMB backup repositories support FSCTL_DUPLICATE_EXTENTS_TO_FILE and FSCTL_SET_INTEGRITY_INFORMATION. SMB shares configured on Microsoft Windows machines must also support the SMB 3.1.1 protocol and the ReFS 3.1 (or later) file system.

Depending on the type of the performed job, Veeam Backup & Replication also imposes the following requirements on backup infrastructure components.

Type of Job	Requirements to backup infrastructure components
Backup job	Protocol: SMB 3.1.1 OS: Microsoft Windows Server 2016 (or later) or Microsoft Windows 10 Pro for Workstations (or later) on the following backup infrastructure components: • If a gateway is selected manually: Gateway server. • If a gateway is selected automatically: Mount server associated with the backup repository, or backup server. For reverse incremental backup chains, Microsoft Windows Server 2016 (or later) or Windows 10 Pro for Workstations (or later) must additionally be installed on backup proxies assigned for the job.
Backup copy job	Protocol: SMB 3.1.1 OS: Microsoft Windows Server 2016 (or later) or Microsoft Windows 10 Pro for Workstations (or later) on the following backup infrastructure components: • If a gateway is selected manually: Gateway server. • If a gateway is selected automatically: [For direct data transport path] Mount server associated with the backup repository, or backup server. [For data transport path over WAN accelerators] Microsoft Windows Server 2016 (or later) or Microsoft Windows 10 Pro for Workstations (or later) on the target WAN accelerator.

Limitations

The following limitations apply when Veeam Backup & Replication uses Fast Clone for Microsoft Windows or SMB backup repositories:

- Veeam Backup & Replication does not use Fast Clone for backup repositories configured with Veeam Backup & Replication 9.5 or an earlier version. After upgrade, such backup repositories will work as backup repositories without Fast Clone support. To leverage Fast Clone, edit settings of such backup repositories and complete the **Edit Backup Repository** wizard without changing settings.
- After you have enabled Fast Clone for existing repositories as described in the previous paragraph or have moved backup chains to backup repositories with Fast Clone support, you must create active full backups for backup chains stored in / moved to the repositories to activate Fast Clone. You can also schedule the backup file compact operation instead of performing active full backup.
- Due to Microsoft limitations, all backup files in the backup chain must be stored on the same volume. For more information, see Restrictions and Remarks at Microsoft Docs.
- Fast Clone requires that source and destination files are stored on the same ReFS volume. If you add a backup repository with Fast Clone support as an extent to a scale-out backup repository, make sure that you enable the Data Locality placement policy for this scale-out backup repository. If backup files are stored on different extents, Fast Clone will not be used.

- Veeam Backup & Replication automatically aligns data blocks at a 4KB or 64 KB block boundary depending on the volume configuration or SMB share used storage.

  We recommend that you use ReFS volume formatted with 64 KB cluster size to provide better performance with large data volumes.

- When you copy data from a ReFS volume to another location, the file system downloads cloned data blocks. For this reason, copied data occupy more space in the target location than it used to occupy in the source location. This can happen, for example, if you evacuate an extent that supports block cloning from a scale-out backup repository and migrate VM backup data to another extent: copied data will require more space than it originally took.

- If you plan to assign the role of a backup repository to Microsoft Windows Server 2016 version 1709 (or later) or Microsoft Windows 10 Pro for Workstations (or later), mind the following limitations:

  - Fast Clone and Windows data deduplication cannot be used simultaneously. Thus, if you target a backup job to a repository supporting Fast Clone and enable Windows data deduplication, the Fast Clone technology will not be used for this job.

  - If you target a backup job to a CIFS ReFS repository and enable Windows data deduplication, the job will fail. Veeam Backup & Replication does not support such scenario.

Proxy Affinity

By default, Veeam Backup & Replication assigns backup proxies and repositories for jobs or tasks independently of each other. If you need to bind backup proxies to specific backup repositories and use them together, you can define proxy affinity settings. Proxy affinity determines what backup proxies are eligible to access a specific backup repository and read/write data from/to this backup repository.

Proxy affinity lets you control assignment of resources in the backup infrastructure and reduce administration overhead. For example, in case of a geographically distributed infrastructure, you can restrict a backup repository in the local site from communicating with backup proxies in a remote site. Or you can configure proxy affinity rules based on a connection speed between backup proxies and backup repositories.

Proxy affinity settings are specified at the level of a backup repository. By default, Veeam Backup & Replication lets all backup proxies in the backup infrastructure access the backup repository. Using proxy affinity settings, you can define a list of backup proxies that can access this backup repository.

Proxy affinity can be set up for the following types of backup repositories:

- Backup repositories
- Scale-out backup repositories
- Cloud repositories (proxy affinity settings are configured on the tenant side)

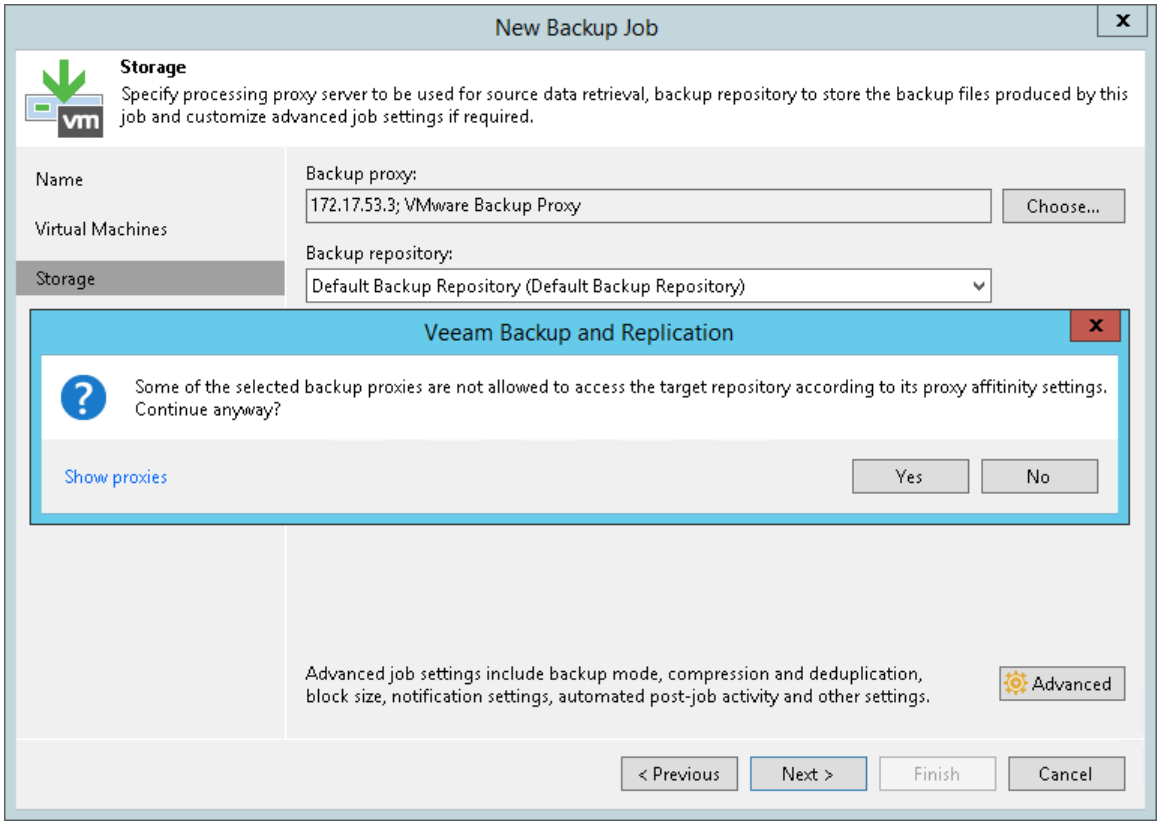
Proxy affinity rules are applied to the following types of jobs and tasks that engage backup proxies and repositories:

- Backup jobs, including VMware vCloud backup and backup jobs from storage snapshots on primary and target storage arrays
- VeeamZIP
- VM copy
- Entire VM restore
- Hard disk restore

Proxy affinity rules are not applied to replication jobs.

Proxy affinity rules are not restrictive. You can think of affinity rules as a priority list. If backup proxies from the proxy affinity list cannot be used for some reason, for example, these backup proxies are inaccessible, Veeam Backup & Replication automatically fails over to the regular processing mode. It displays a warning in the job or task session and picks the most appropriate backup proxy from the list of proxies selected for the job or task.

When you target a job at a backup repository for which proxy affinity settings are configured, you must make sure that you assign a backup proxy from the proxy affinity list for job or task processing. If you assign a backup proxy that is not bound to this backup repository, Veeam Backup & Replication will display a warning. For job processing, Veeam Backup & Replication will use the backup proxy that you define in the job settings, which may result in degraded job performance.



Proxy Affinity for Scale-Out Backup Repositories

In case of a scale-out backup repository, you can configure proxy affinity settings at the extent level. Proxy affinity settings cannot be configured at the scale-out backup repository level.

Extent selection rules have a higher priority than proxy affinity rules. Veeam Backup & Replication first selects an extent and then picks a backup proxy according to the proxy affinity rules specified for this extent.

For example, you have 2 backup proxies: *Backup Proxy 1* and *Backup Proxy 2*. You create a backup job and target it at a scale-out backup repository configured in the following way:

- Scale-out backup repository policy is set to Data Locality.
- Scale-out backup repository has 2 extents: *Extent 1* has 100 GB of free space and is bound to *Backup Proxy 1*; *Extent 2* has 1 TB of free space and is bound to *Backup Proxy 2*.

In the backup job settings, you define that *Backup Proxy 1* must be used for job processing.

When you run the backup job, Veeam Backup & Replication will store backup files to *Extent 2* since it has more free space. For job processing, it will pick *Backup Proxy 1* and will display a message in the job statistics that requirements of proxy affinity rules cannot be met.

In case of restore from a scale-out backup repository, backup files may be located on different extents. In this case, Veeam Backup & Replication picks a backup proxy according to the following priority rules (starting from the most preferable one):

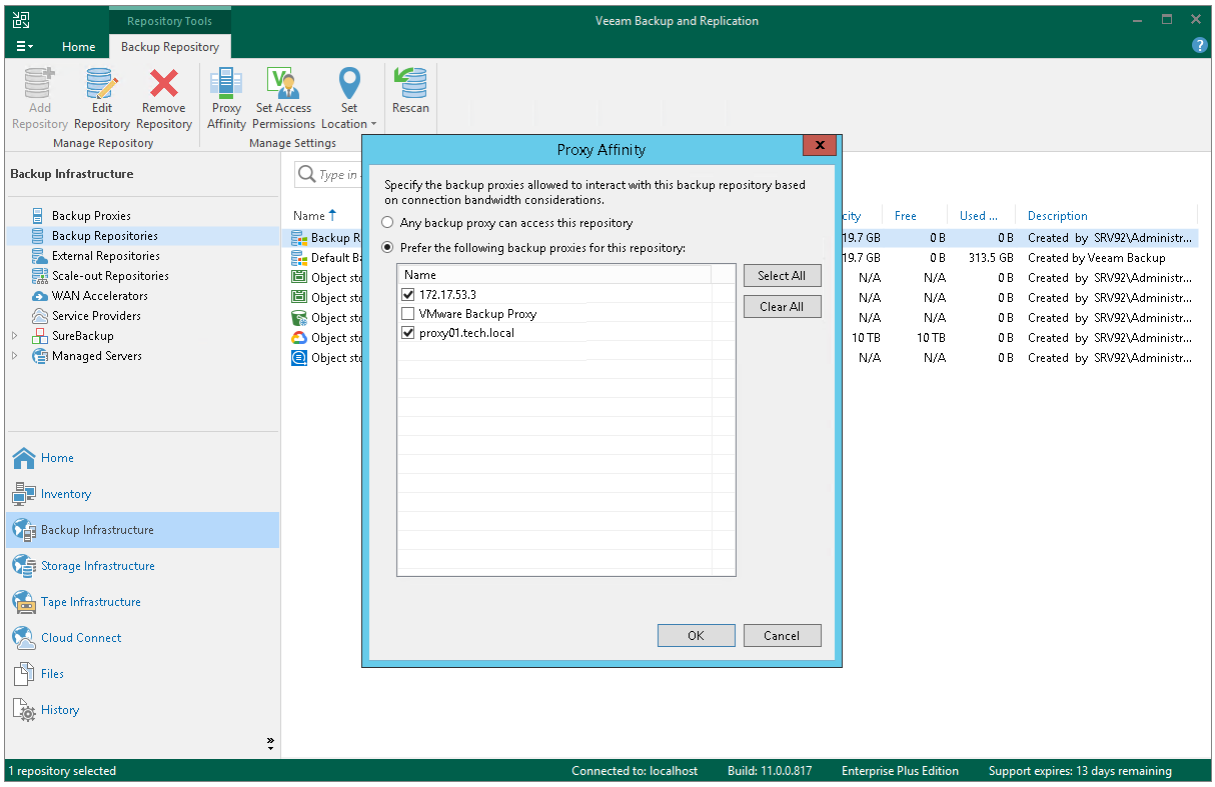
1. Backup proxy is added to the affinity list for all extents.
2. Backup proxy is added to the affinity list for the extent where the full backup file is stored.
3. Backup proxy is added to the affinity list for at least one extent.

Specifying Proxy Affinity Settings

For every backup repository, you can configure proxy affinity settings — define a list of backup proxies that can work with this backup repository. Proxy affinity binds backup proxies to specific backup repositories. When transporting data to/from the backup repository, Veeam Backup & Replication picks a backup proxy from the proxy affinity list rather than backup proxies specified in job or task settings.

To configure a proxy affinity list for a backup repository:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Backup Repositories**.
3. In the working area, select the backup repository and click **Proxy Affinity** on the ribbon or right-click the backup repository and select **Proxy affinity**.
4. In the **Proxy Affinity** window, select **Prefer the following backup proxies for this repository** and select check boxes next to the backup proxies that you want to bind to the backup repository.



External Repository

An external repository is a read-only repository. You can use Veeam Backup & Replication to copy, import and/or restore backups created by Veeam Backup for AWS, Veeam Backup for Microsoft Azure and Veeam Backup for Google Cloud Platform from external to on-premises repositories. This way, you can perform data migration between cloud, on-premises and virtual infrastructures.

Veeam Backup & Replication supports the following types of external repositories:

- Amazon S3 (with Standard storage class assigned)
- Azure Blob (Hot and Cool access tier)
- Google Cloud Platform (with Standard storage class assigned)

To start working with backups created by Veeam Backup for AWS, Veeam Backup for Microsoft Azure and Veeam Backup for Google Cloud Platform, you must add a repository that contains backups of Amazon EC2 instances, Azure VMs or Google Cloud Platform VMs to the Veeam Backup & Replication infrastructure as an external repository. After that, you can perform the following operations:

- Copy backups to on-premises repositories
- Restore EC2 instances, Azure VMs and Google Cloud Platform VMs to AWS
- Restore Azure VMs, EC2 instances and Google Cloud Platform VMs to Microsoft Azure
- Restore Google Cloud Platform VMs, EC2 instances and Azure VMs to Google Cloud
- Restore guest OS files and folders
- Export disks of EC2 instances, Azure VMs and Google Cloud Platform VMs
- Govern retention policies
- Perform Instant Recovery to VMware vSphere
- Perform Instant Recovery to Microsoft Hyper-V

> **NOTE**
>
> Mind the following:
>
> - During the process of copying backups, or restore to Amazon EC2 or Microsoft Azure, data of EC2 instances and Azure VMs may migrate from one geographic location to another. In this case, Veeam Backup & Replication displays a warning and stores a record about data migration to job or task session details. For more information, see Locations.
> - You cannot use an external repository as a target for backup or backup copy jobs.

How External Repository Works

Continue with this section to learn more on how external repositories work.

Ownership

Ownership defines what entity can own an Amazon S3, Azure Blob or Google Cloud storage repository at a time.

How Does Taking Ownership Occur

After Veeam Backup for AWS, Veeam Backup for Microsoft Azure or Veeam Backup for Google Cloud Platform has finished its initial backup job session, it becomes the rightful owner of both a storage repository and backup files in that repository.

Taking ownership of such a repository along with its backup files by the Veeam Backup & Replication client consists of the following consecutive steps:

- Step 1. Taking ownership of a repository.

  Reclaiming ownership of a repository occurs every time a client adds object storage as an external repository to the Veeam Backup & Replication console.

- Step 2. Taking ownership of backup files in the repository.

  Becoming an owner of backup files in object storage is only possible after Veeam Backup for AWS, Veeam Backup for Microsoft Azure or Veeam Backup for Google Cloud Platform launches the backup job session which is referring to backups you are trying to take ownership of (for example, backup files that are located in the repository you have added at the step one).

  During its session, Veeam Backup for AWS, Veeam Backup for Microsoft Azure or Veeam Backup for Google Cloud Platform verifies the owner of a repository and if it finds out that the owner has been changed, it changes the owner of each backup file in that repository by creating a new checkpoint that refers to a new rightful owner. Such a checkpoint will be used during subsequent sessions of a backup job to repeat owner verification.

It is possible, however, that after you add an external repository, you never launch the associated backup job again. In such a scenario, Veeam Backup & Replication will not be able to manage retention policies, but you will still be able to restore external repository data, remove backups from external repositories and perform backup copy.

Taking Ownership by Another Veeam Backup & Replication Client

Ownership of a repository along with its backup files can only be granted to one Veeam Backup & Replication client at a time.

Therefore, if a client *A* adds an external repository that has previously been added by the client *B*, the client *B* completely loses its ownership privileges.

Losing privileges means that the client *B* will no longer be able to manage retention policies. All the previously created backup copy jobs and restore sessions will be failing.

Ownership, however, can easily be reclaimed by re-adding the same object storage anew.

Cache

Veeam Backup & Replication caches data that is being retrieved from external repositories every time a backup copy job or restore session is performed.

Such an approach helps not only to reduce the number of cost-expensive operations incurred by AWS, Microsoft Azure or Google Cloud Platform, but also decrease the amount of traffic being sent over the network.

Consider the following:

- Cache is created on a gateway server while the following activities are being processed:
  - Backup copy jobs.
  - Restore sessions.
- Cache is not created upon the addition of an external repository to the Veeam Backup & Replication console.
- Cache consists of metadata of blocks being retrieved from external repositories.
- Cache is written to:
  - On a Windows-based gateway server: `C:\ProgramData\Veeam\ExternalCache`
  - On a Linux-based gateway server: `/var/veeam/ExternalCache`
- Cache is reused and updated during each subsequent execution of a backup copy job or restore session.
- Cache size limit is 10GB. Once reached, Veeam will purge obsolete cache by 20% preserving most frequently used parts. Purging is done by the maintenance job.
- Cache is removed in the following cases:
  - An external repository has been removed from the Veeam Backup & Replication console.
  - The gateway server has been changed in the settings of the external repository configuration.
  - The backup file has been removed from the external repository.
  - During the maintenance job session.
- Cache can be removed manually without affecting the backup infrastructure in any negative way.

Encryption

Backups that reside in Amazon S3 buckets, Azure Blob storage and Google Cloud storage can be encrypted by Veeam Backup for AWS, Veeam Backup for Microsoft Azure and Veeam Backup for Google Cloud Platform. Moreover, password for such encrypted backups may change on a daily basis. For example, there is a backup chain in Amazon S3 bucket that consists of 10 restore points, each of which was encrypted with different password. Therefore, there are 10 different passwords in total that have been used.

To be able to decrypt each restore point in such a backup chain without having to provide each previously used password separately, Veeam Backup & Replication implements the ability of backward hierarchical decryption.

Backward hierarchical decryption requires you to provide only the latest password so that all the previously created restore points can be decrypted as well. For example, there are three restore points: A, B, and C. The point A was encrypted with password 1, B with password 2, and C with password 3. Therefore, you will only need to know the password of the C point to decrypt points C, B, and A.

If you plan to perform data recovery operations with encrypted backups, you must provide a password for these backups in the **New External Repository** wizard:

- [For Veeam Backup for AWS] At the Bucket step of the wizard.
- [For Veeam Backup for Microsoft Azure] At the Container step of the wizard.
- [For Google Cloud Platform] At the Bucket step of the wizard.

Managing Retention Policy

A retention policy is set in the backup policy settings of Veeam Backup for AWS, Veeam Backup for Microsoft Azure and Veeam Backup for Google Cloud Platform and defines the number of restore points to keep in repositories.

Retention policies are initially managed by Veeam Backup for AWS, Veeam Backup for Microsoft Azure or Veeam Backup for Google Cloud Platform until a Veeam Backup & Replication client reclaims ownership of a repository and all the backup files in such a repository.

Once ownership is reclaimed, Veeam Backup for AWS, Veeam Backup for Microsoft Azure or Veeam Backup for Google Cloud Platform ceases to govern retention policies and the Veeam Backup & Replication client becomes responsible for removing obsolete restore points from repositories.

The restore points that fall under the retention policy will be removed upon the next successful session of the maintenance job.

When a Veeam Backup & Replication client removes an external repository from its scope, it relinquishes its ownership which then goes directly to Veeam Backup for AWS, Veeam Backup for Microsoft Azure or Veeam Backup for Google Cloud Platform until another Veeam Backup & Replication client reclaims it anew and so forth.

> **IMPORTANT**
>
> A retention policy can only be applied by the Veeam Backup & Replication client that is the rightful owner of an Amazon S3 object storage repository and its backup files.

Maintenance Job

The maintenance job is a system job that is executed automatically every 24 hours.

The maintenance job does the following:

- Purges obsolete restore points that fall under the retention policy.

  To be able to purge obsolete restore points from external repositories due to the retention policy threshold, a Veeam Backup & Replication client must be the owner of a repository and its backup files.

- Purges cache by 20% of the size limit. By default, the size limit is 10GB.
- Saves its session results to the configuration database.

  The session results can be found in the **History** view under the **System** node.

Adding External Repositories

You can add the following types of external repositories:

- External Amazon S3 Storage
- External Azure Blob Storage
- External Google Cloud Storage

Adding External Amazon S3 Storage

To add an Amazon S3 storage as an external repository, do the following:

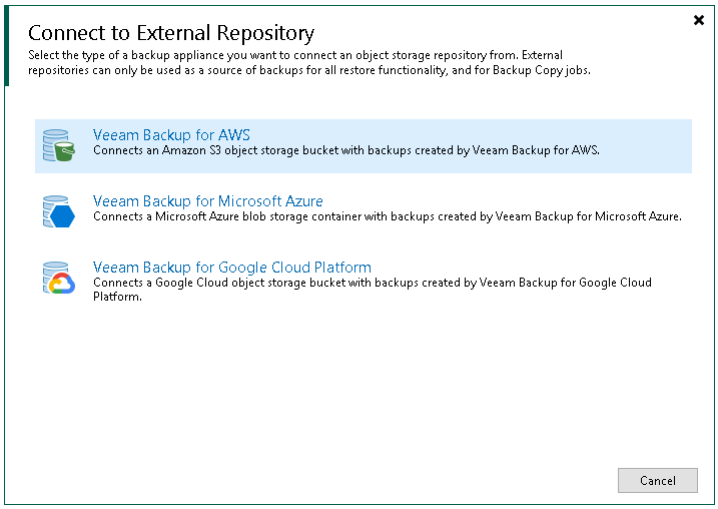
1. Launch the New External Repository wizard
2. Specify the repository name and description
3. Specify the cloud account
4. Specify cloud storage details
5. Wait for the repository to be added to the infrastructure
6. Finish working with the wizard

> **NOTE**
>
> If you added the Veeam Backup for AWS appliance to the Veeam Backup & Replication infrastructure and want to deploy a new S3 repository, see instructions in the Adding S3 Repositories section in the Integration with Veeam Backup for AWS Guide.

Step 1. Launch New External Repository Wizard

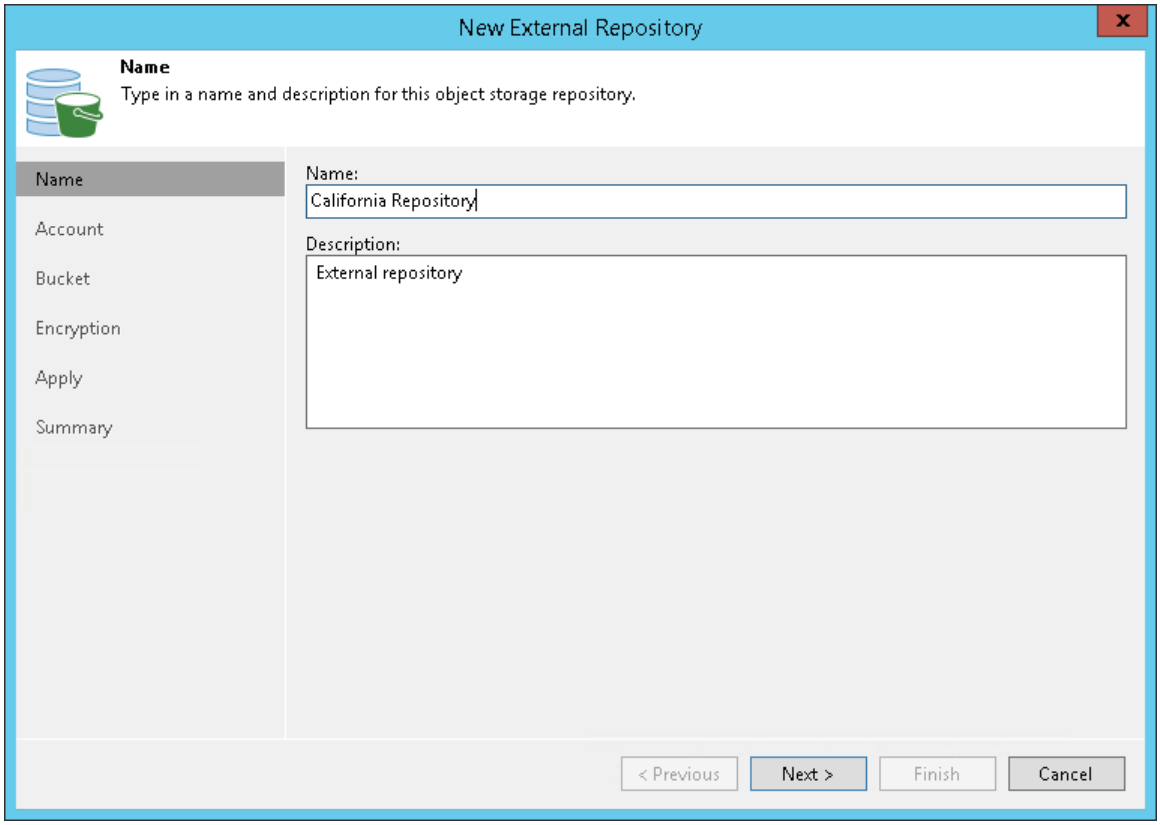
To launch the **New Backup Repository** wizard, do one of the following:

- Open the **Backup Infrastructure** view. In the inventory pane select the **External Repositories** node and click **Connect to Repository** on the ribbon. At the **Connect to External Repository** window, select **Veeam Backup for AWS**.
- Open the **Backup Infrastructure** view. In the inventory pane right-click the **External Repositories** node and select **Connect to...** At the **Connect to External Repository** window, select **Veeam Backup for AWS**.



Step 2. Specify External Repository Name

At the **Name** step of the wizard, specify a name and description for the external repository.



Step 3. Specify Cloud Repository Account

At the **Account** step of the wizard, specify Amazon S3 connection settings:

1. From the **Credentials** drop-down list, select user credentials to access an Amazon S3 bucket with Amazon EC2 instance backups.

   If you have not set up credentials beforehand in the Cloud Credentials Manager, click the **Manage cloud accounts** link or click **Add** on the right to add the necessary credentials.

2. From the **Data center region** drop-down list, select the AWS region where the Amazon S3 bucket is located: *Global*, *GovCloud (US)*, or *China*.

3. From the **Gateway server** drop-down list, select a gateway server that will be used to access the Amazon S3 bucket. We recommend that you use a gateway server, for example, if your organization has NAT or different types of firewalls and your access to the Internet is limited.

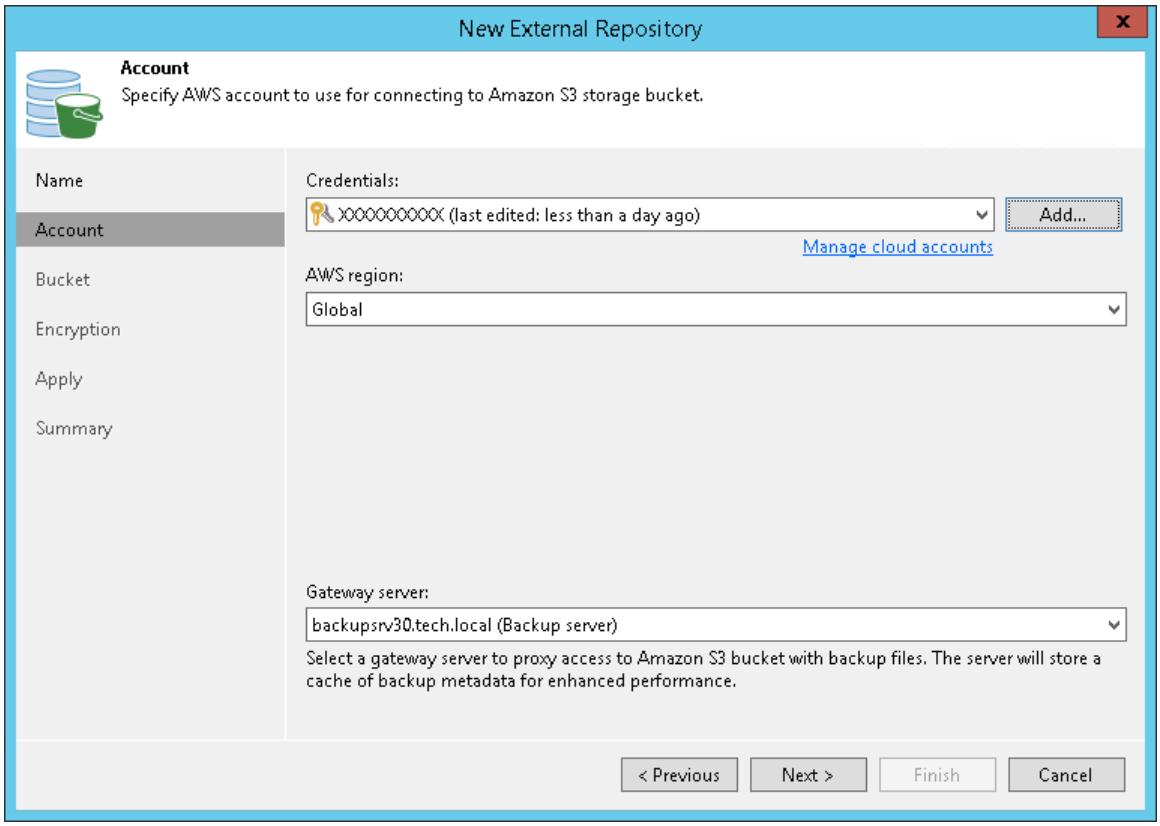
   The gateway server caches data when you copy backups or perform restore operations. The gateway server helps you decrease the amount of traffic being sent over the network and reduce data transfer costs. For more information on caching data, see Cache.

   By default, the role of a gateway server is assigned to the Veeam Backup & Replication server. If the Veeam Backup & Replication server resides in a region that differs from the Amazon region where your Amazon S3 bucket resides, choose a server that is located close to the bucket. You can choose any Microsoft Windows or Linux server that is added to your Veeam Backup & Replication infrastructure and has internet connection. Note that the server must be added to the infrastructure beforehand. For more information on how to add a server, see the Adding Microsoft Windows Servers and Adding Linux Servers sections.

NOTE

Mind the following:

- On the gateway server, Veeam Backup & Replication deploys Veeam Data Mover that handles traffic sent when you work with external repository data. If Veeam Data Mover becomes outdated, you must upgrade it as described in the Upgrading External Repositories section.
- If you choose not to use a gateway server, make sure that all scale-out repository extents have direct internet access.



Step 4. Specify Cloud Storage Details

At the **Bucket** step of the wizard, specify an Amazon S3 bucket and folder where Amazon EC2 instance backups reside:

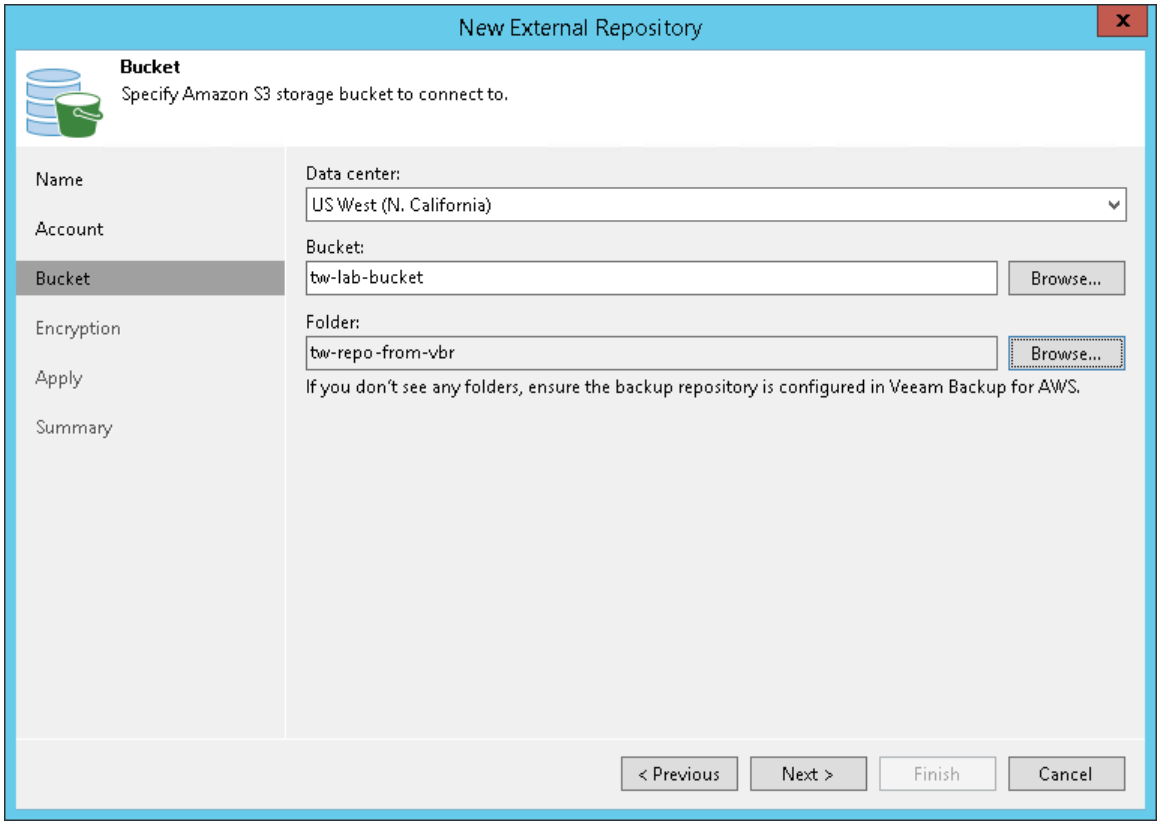
1. From the **Data center region** drop-down list, select an AWS region where the Amazon S3 bucket is located.
2. From the **Bucket** drop-down list, select the necessary Amazon S3 bucket where EC2 instance backups reside.
3. Click **Browse** to select a folder in the Amazon S3 bucket where EC2 instance backups reside.

   If you do not see the required folder, make sure that the repository you are trying to add is created on the Veeam Backup for AWS server.

NOTE

Mind the following:

- Only Standard storage class is supported.
- If another Veeam Backup & Replication client has already added the same folder, you will be prompted whether to reclaim ownership of such a folder. For more information about ownership, see Ownership.



Step 5. Configure Encryption

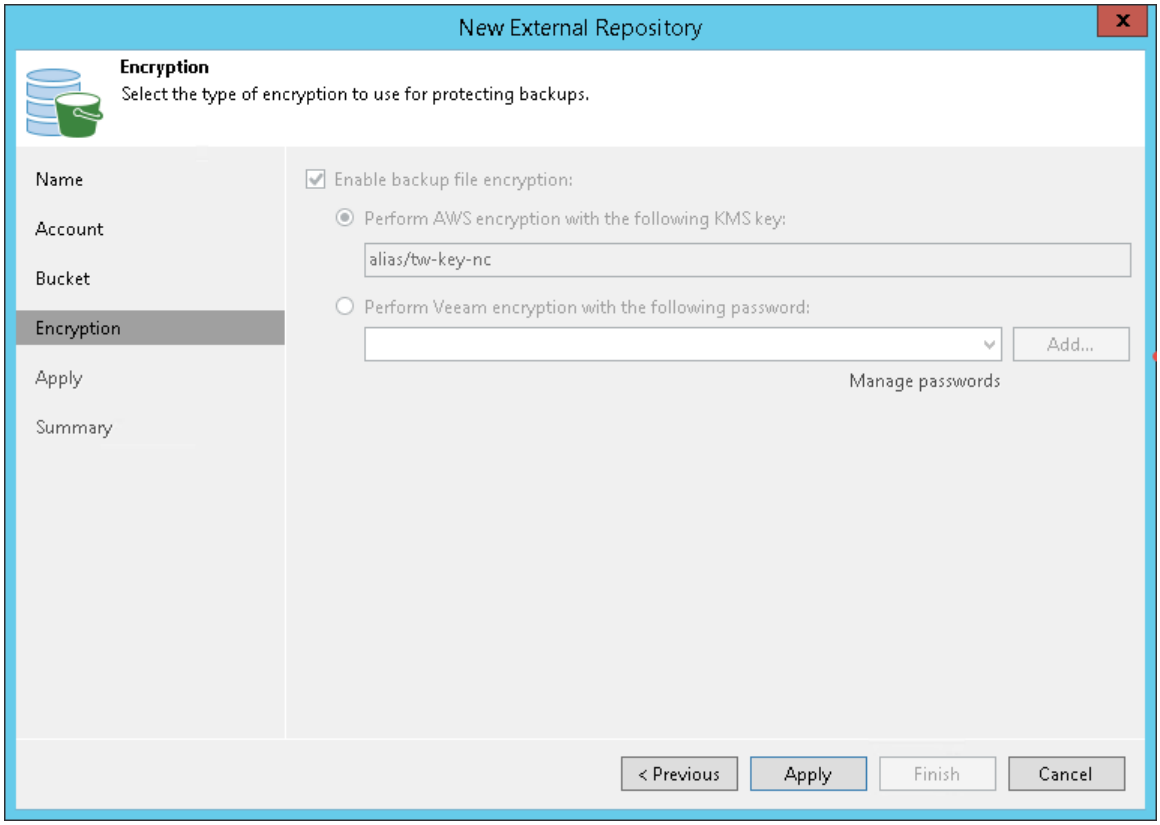
At the **Encryption** step of the wizard, provide a key for decryption:

- If data in the external repository is encrypted with a key management service (KMS) customer master key (CMK), Veeam Backup & Replication shows the used key. In this case, Veeam Backup & Replication will automatically decrypt the backups.
- If data in the external repository is encrypted with a password, select **Enable backup file encryption** and then click **Perform Veeam encryption with the following password**. From the drop-down list, select the password that must be used to decrypt the data. If the password is correct, Veeam Backup & Replication will automatically decrypt the backups.

  If you have not added the password beforehand, click the **Manage passwords** link or the **Add** button to add the necessary password. For more information on adding passwords, see Creating Passwords.

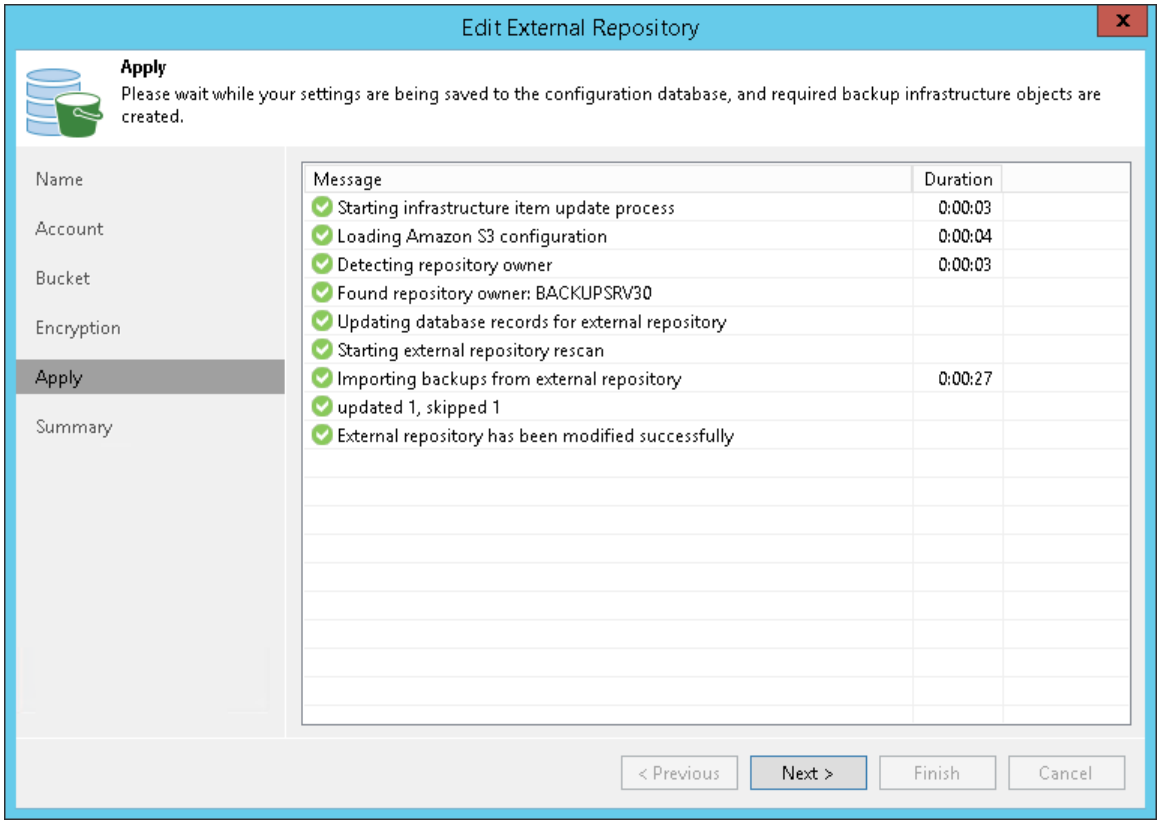
  If you do not specify the decryption password, you can do it later. For more information, see Viewing External Repository Data.

For more information on how backups are encrypted, see the S3 Repository Encryption section in the Veeam Backup for AWS User Guide.



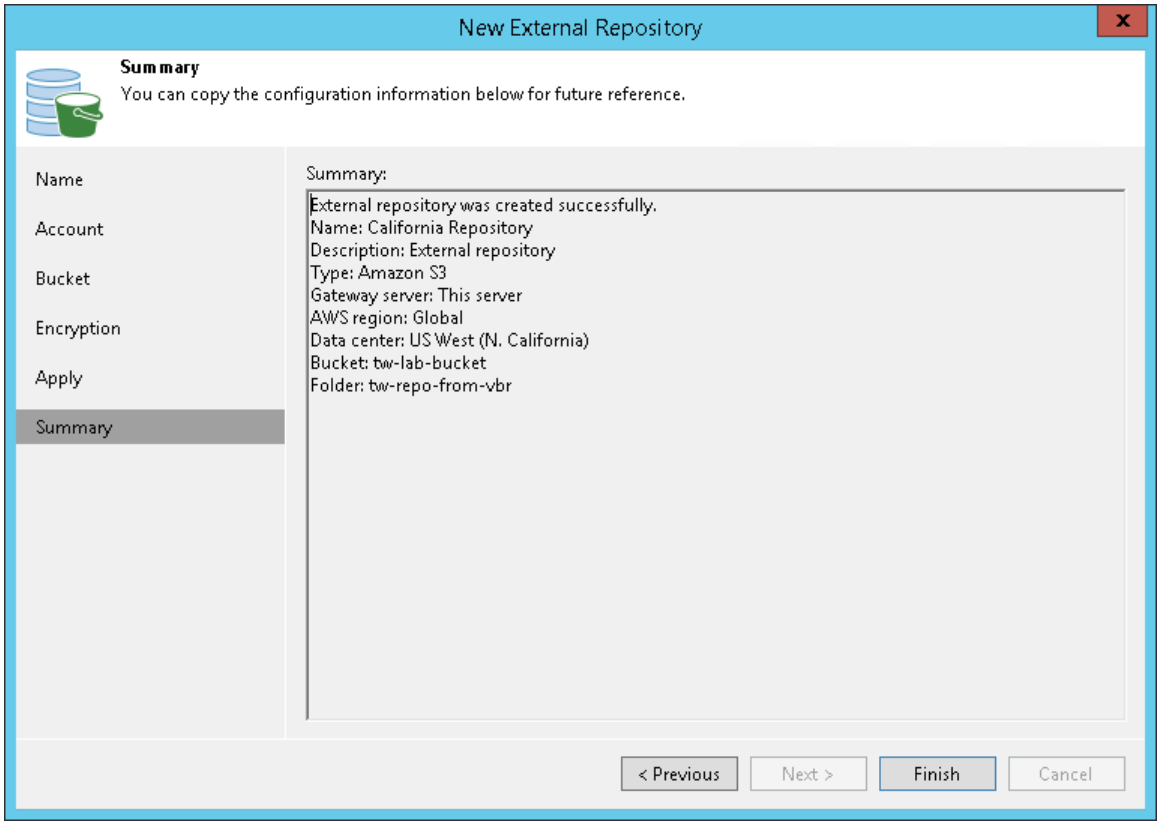
Step 6. Apply Settings

At the **Apply** step of the wizard, wait for Veeam Backup & Replication to install and configure all required components. Then click **Next** to complete the procedure of adding the external repository to the backup infrastructure.



Step 7. Finish Working with Wizard

At the **Summary** step of the wizard, review the external repository settings and click **Finish**.



Adding External Azure Blob Storage

To add an Microsoft Azure Blob storage as an external repository, do the following:

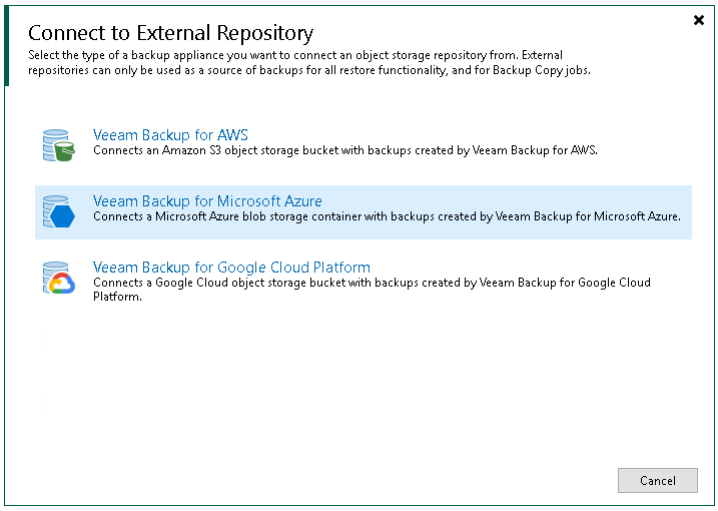
1. Launch the New External Repository wizard
2. Specify the repository name
3. Specify the cloud account
4. Select the Azure Blob container
5. Wait for the repository to be added to the infrastructure
6. Finish working with the wizard

> **NOTE**
>
> If you added the Veeam Backup for Microsoft Azure appliance to the Veeam Backup & Replication infrastructure and want to deploy a new Microsoft Azure Blob storage, see instructions in the Adding Microsoft Azure Blob Storage section in the Integration with Veeam Backup for Microsoft Azure Guide.

Step 1. Launch New External Repository Wizard

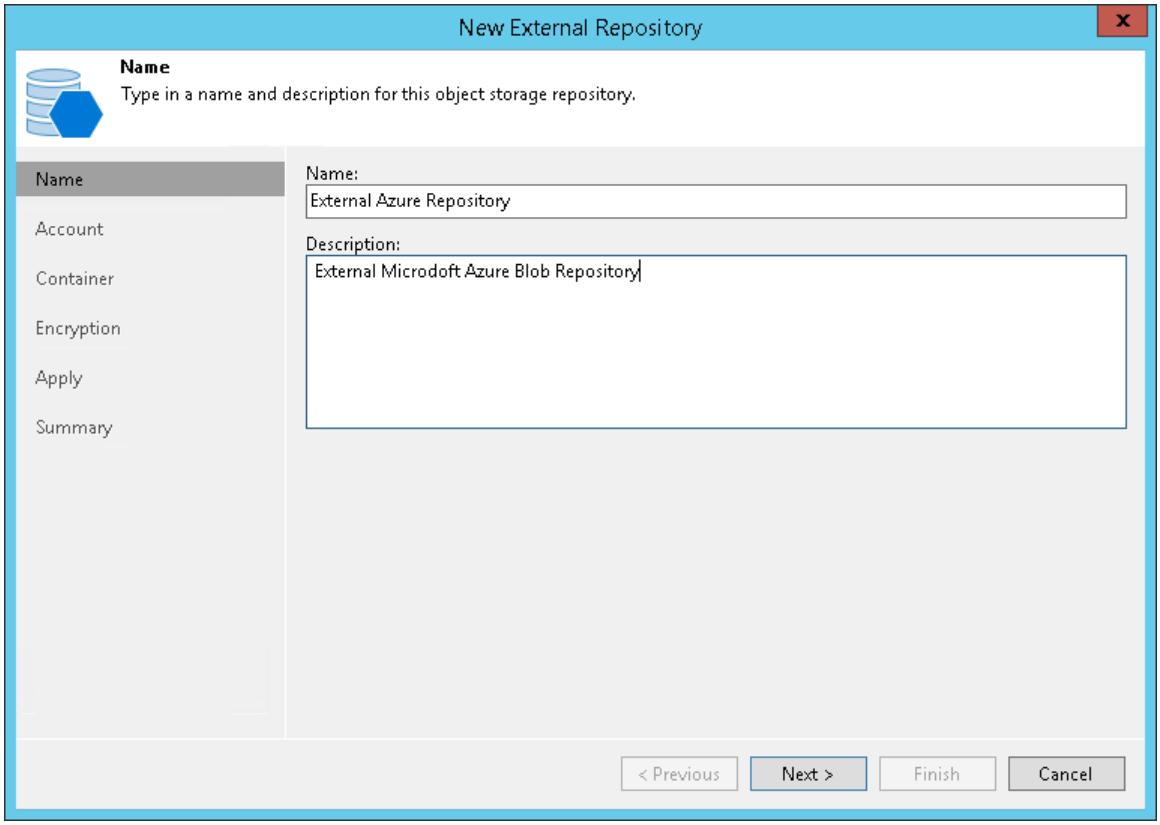
To launch the **New External Repository** wizard, do one of the following:

- Open the **Backup Infrastructure** view. In the inventory pane select the **External Repositories** node and click **Connect to Repository** on the ribbon. At the **Connect to External Repository** window, select **Veeam Backup for Microsoft Azure**.
- Open the **Backup Infrastructure** view. In the inventory pane right-click the **External Repositories** node and select **Connect to...** At the **Connect to External Repository** window, select **Veeam Backup for Microsoft Azure**.



Step 2. Specify External Repository Name

At the **Name** step of the wizard, specify a name and description for the external repository.



Step 3. Specify Cloud Repository Account

At the **Account** step of the wizard, specify settings for an account which will be used to connect to Azure Blob storage:

1. From the **Credentials** drop-down list, select user credentials to access your Azure Blob storage.

   If you have not set up credentials beforehand in the Cloud Credentials Manager, click the **Manage cloud accounts** link or click **Add** on the right to add the necessary credentials.

2. From the **Region** drop-down list, select the region type: *Azure Global, Azure Germany, Azure China, Azure Government*.

3. From the **Gateway server** drop-down list, select a gateway server that will be used to access Azure Blob storage. We recommend that you use a gateway server, for example, if your organization has NAT or different types of firewalls and your access to the Internet is limited.

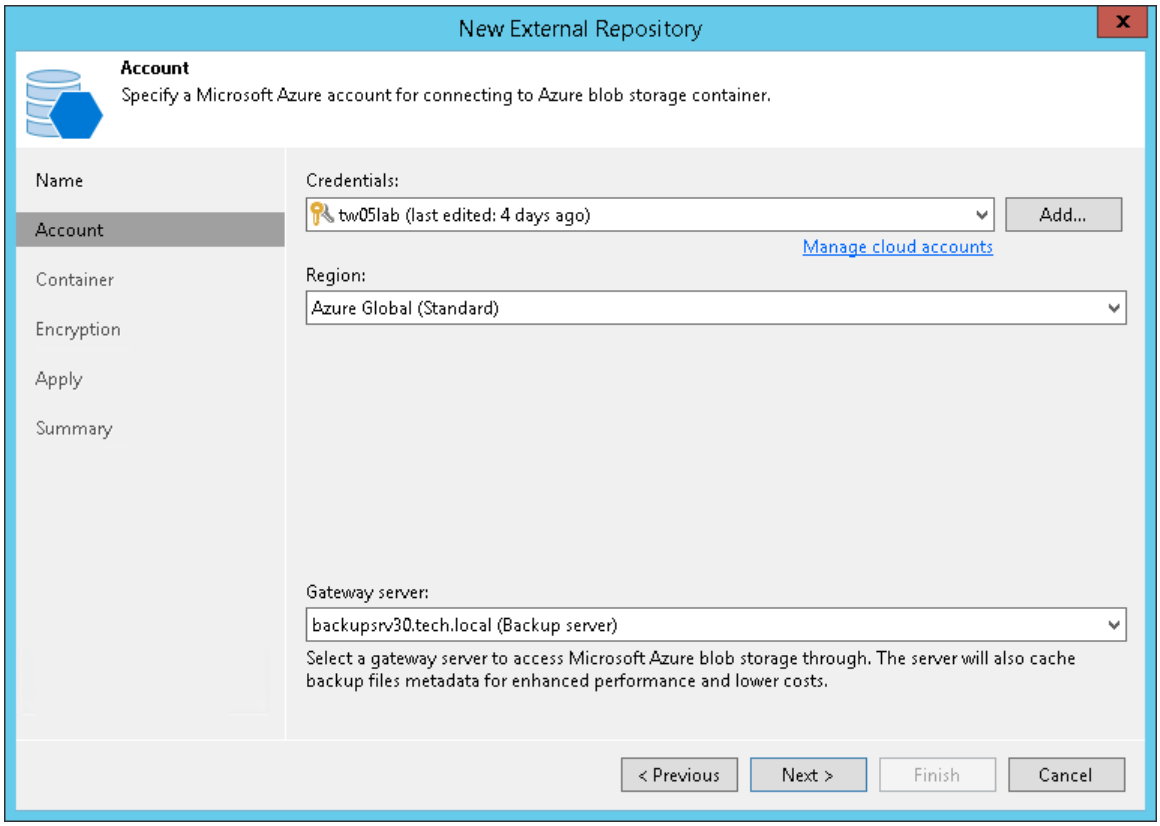
   The gateway server caches data when you copy backups or perform restore operations. The gateway server helps you decrease the amount of traffic being sent over the network and reduce data transfer costs. For more information on caching data, see Cache.

   By default, the role of a gateway server is assigned to the Veeam Backup & Replication server. If the Veeam Backup & Replication server resides in a region that differs from the Azure region where your Blob storage resides, choose a server that is located close to the storage. You can choose any Microsoft Windows or Linux server that is added to your Veeam Backup & Replication infrastructure and has internet connection. Note that the server must be added to the infrastructure beforehand. For more information on how to add a server, see the Adding Microsoft Windows Servers and Adding Linux Servers sections.

NOTE

Mind the following:

- On the gateway server, Veeam Backup & Replication deploys Veeam Data Mover that handles traffic sent when you work with external repository data. If Veeam Data Mover becomes outdated, you must upgrade it as described in the Upgrading External Repositories section.
- If you choose not to use a gateway server, make sure that all scale-out repository extents have direct internet access.



Step 4. Select Azure Blob Container

At the **Container** step of the wizard, specify Azure Blob storage container settings:

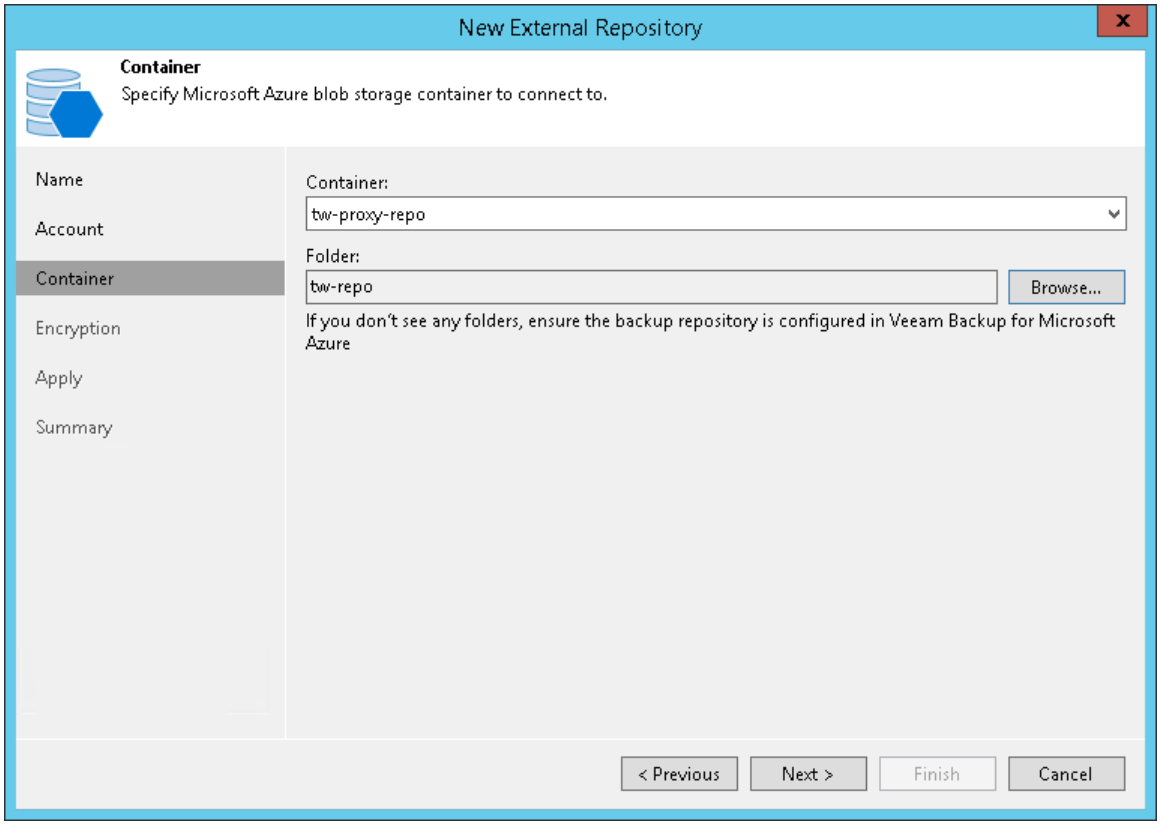
1. From the **Container** drop-down list, select a container that contains backups created by Veeam Backup for Microsoft Azure.
2. Click **Browse** to select a folder that contains backups created by Veeam Backup for Microsoft Azure.

   If you do not see the required folder, make sure that the repository you are trying to add is created on the Veeam Backup for Microsoft Azure server.

NOTE

Mind the following:

- Only Hot and Cool access tiers are supported.
- If another Veeam Backup & Replication client has already added the same folder, you will be prompted whether to reclaim ownership of such a folder. For more information about ownership, see Ownership.



Step 5. Configure Encryption

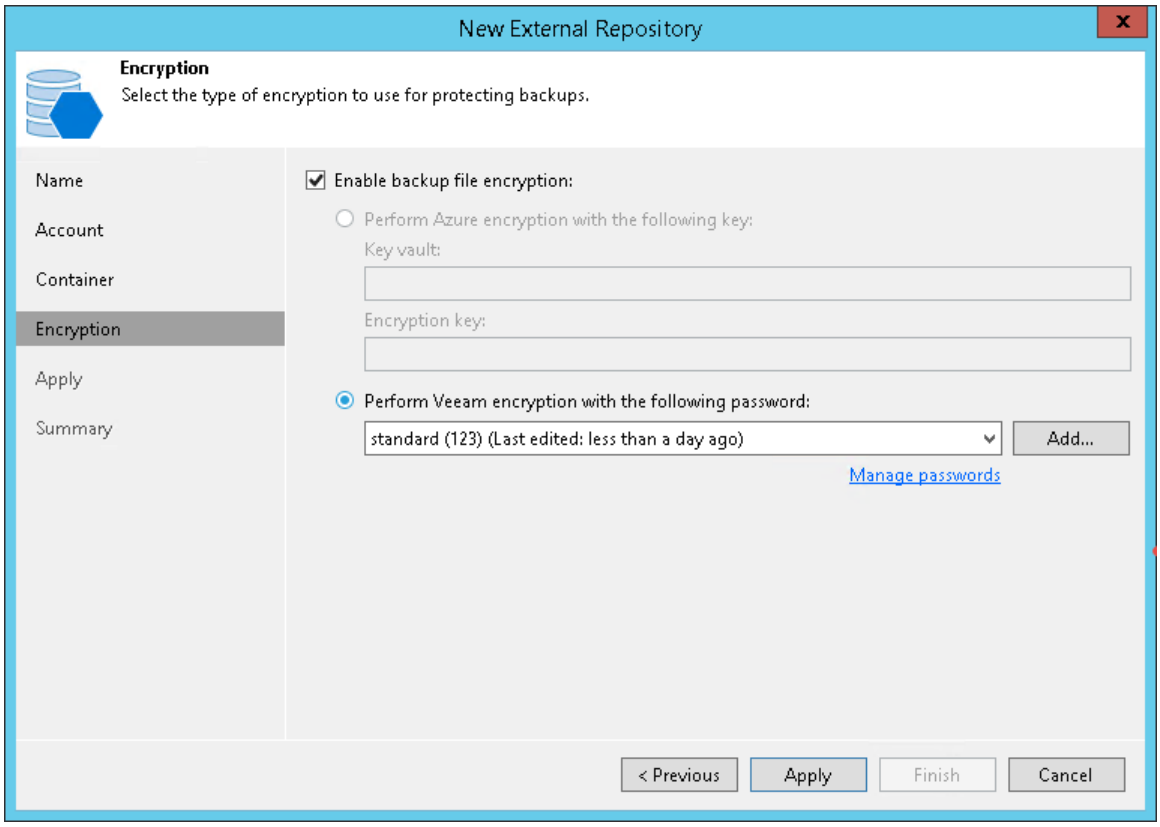
This step is available starting from Veeam Backup & Replication 11a (build 11.0.1.1261).

At the **Encryption** step of the wizard, provide a key for decryption:

- If data in the external repository is encrypted with a Key Vault key, Veeam Backup & Replication shows the used Key Vault and encryption key. In this case, Veeam Backup & Replication will automatically decrypt the backups.
- If data in the external repository is encrypted with a password, select **Enable backup file encryption** and then click **Perform Veeam encryption with the following password**. From the drop-down list, select the password that must be used to decrypt the data. If the password is correct, Veeam Backup & Replication will automatically decrypt the backups.

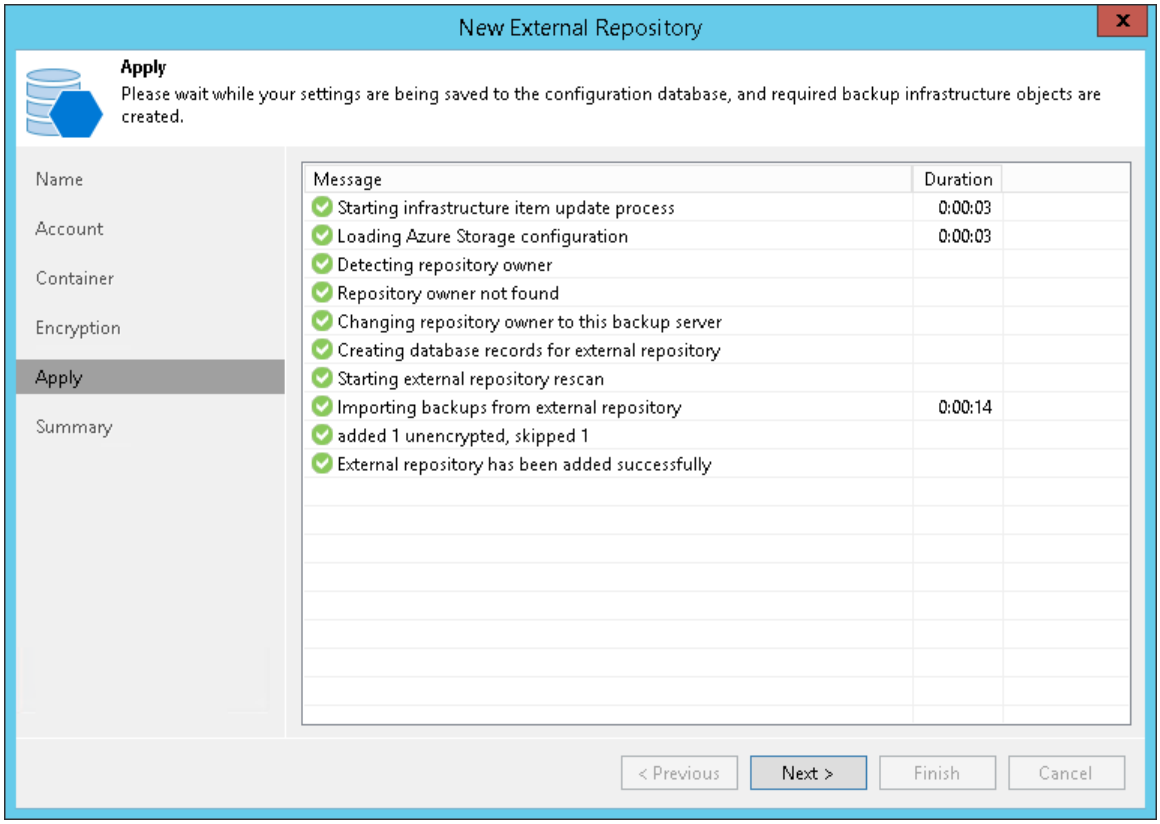
  If you have not added the password beforehand, click the **Manage passwords** link or the **Add** button to add the necessary password. For more information on adding passwords, see Creating Passwords.

  If you do not specify the decryption password, you can do it later. For more information, see Viewing External Repository Data.



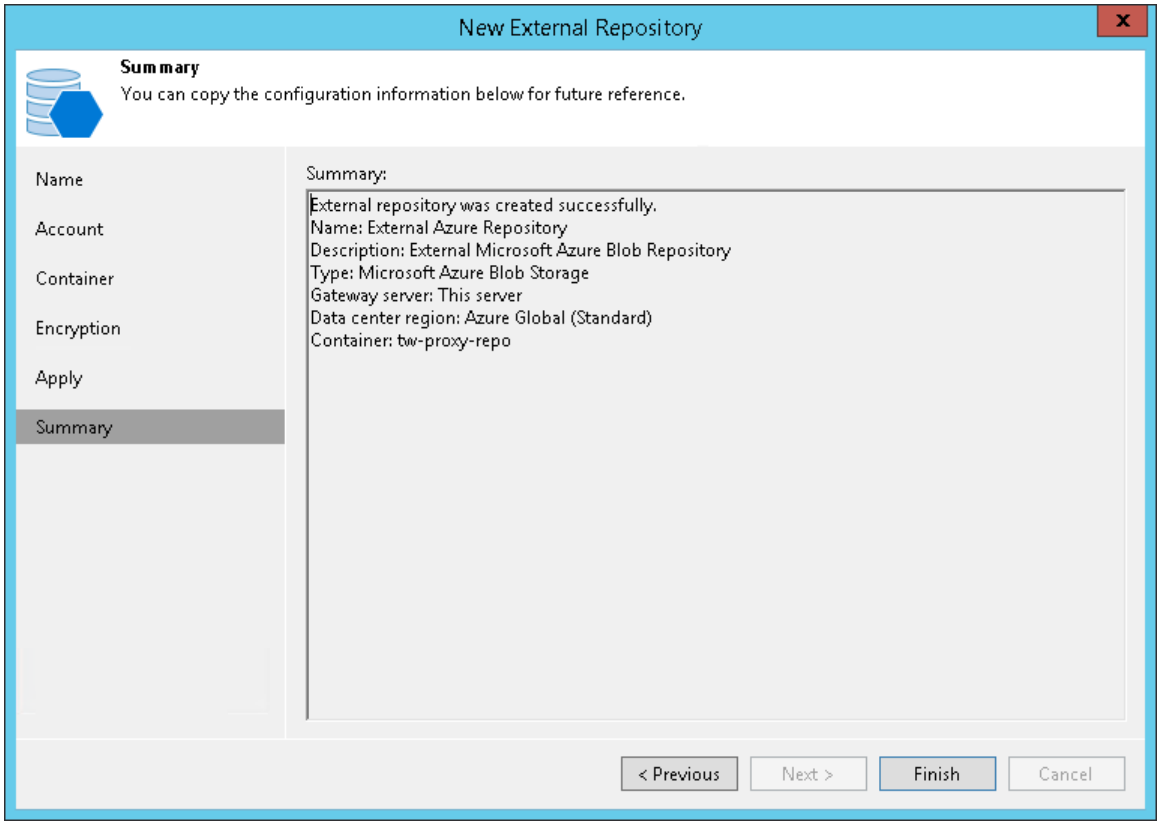
Step 6. Apply Settings

At the **Apply** step of the wizard, wait for Veeam Backup & Replication to install and configure all required components. Then click **Next** to complete the procedure of adding the external repository to the backup infrastructure.



Step 7. Finish Working with Wizard

At the **Summary** step of the wizard, review the repository configuration settings and click **Finish**.



Adding External Google Cloud Storage

To add an external Google Cloud storage repository to the backup infrastructure, do the following:

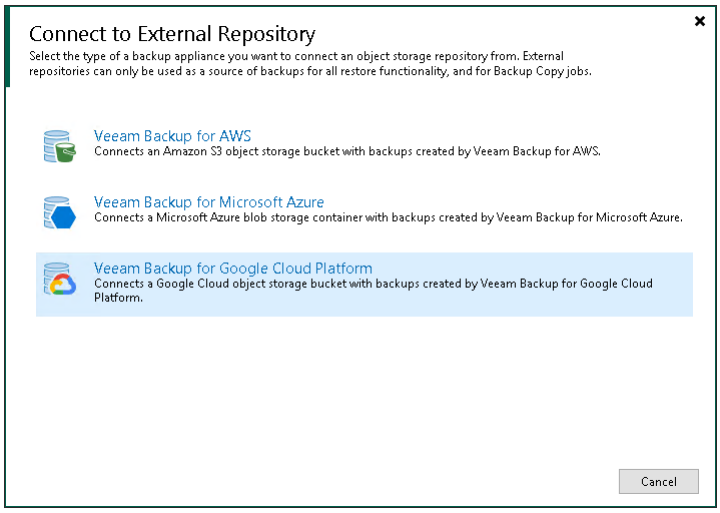
1. Launch the New External Repository Wizard
2. Specify the repository name
3. Specify the cloud account
4. Specify cloud storage details
5. Apply settings
6. Finish working with the wizard

> **NOTE**
>
> If you added the Google Cloud Platform appliance to the Veeam Backup & Replication infrastructure and want to deploy a new Cloud Storage Repository, see instructions in the Adding Standard Backup Repository section in the Integration with Google Cloud Platform Guide.

Step 1. Launch New External Repository Wizard

To launch the **New External Repository** wizard, do one of the following:

- Open the **Backup Infrastructure** view. In the inventory pane select the **External Repositories** node and click **Connect to Repository** on the ribbon. At the **Connect to External Repository** window, select **Veeam Backup for Google Cloud**.
- Open the **Backup Infrastructure** view. In the inventory pane right-click the **External Repositories** node and select **Connect to...** At the **Connect to External Repository** window, select **Veeam Backup for Google Cloud**.

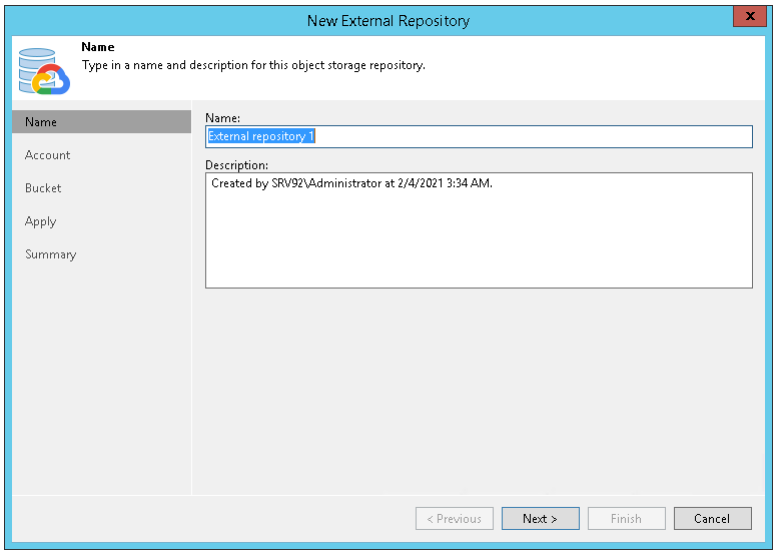


Step 2. Specify External Repository Name

At the **Name** step of the wizard, specify a name and description for the external repository.

1. In the **Name** field, enter a name for the external repository.
2. In the **Description** field, enter an optional description.

   The default description contains information about the user who added the external repository, date and time when the external repository was added.



Step 3. Specify Cloud Repository Account

At the **Account** step of the wizard, specify Google Cloud connection settings:

1. From the **Credentials** drop-down list, select user credentials to access a storage bucket with the backups.

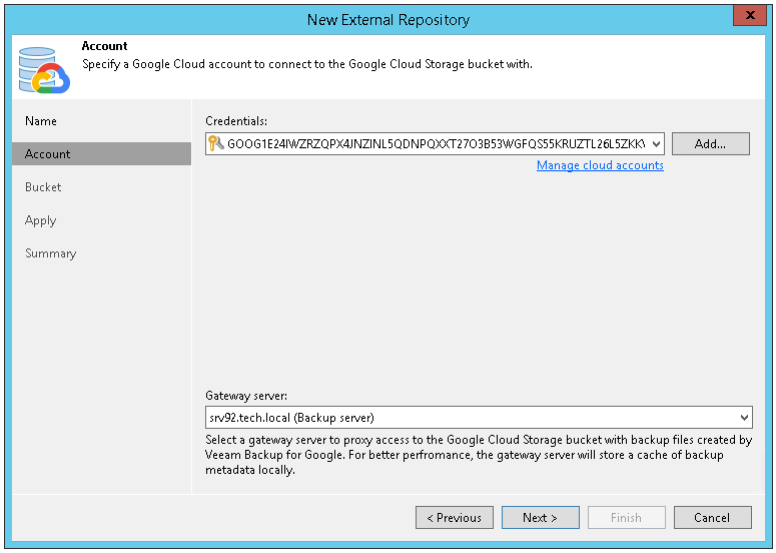
   If you have not set up credentials beforehand, click the **Manage cloud accounts** link or click **Add** on the right and add the necessary credentials, as described in Google Cloud Accounts .

2. From the **Gateway server** drop-down list, select a server that will be used to access the Google Cloud storage.

   You can select any Microsoft Windows or Linux server that is added to your backup infrastructure and has internet connection. You may want to use a gateway server, for example, if your organization has NAT or different types of firewalls and your access to the internet is limited. For more information on how to add such a server to your environment, see Adding Microsoft Windows Servers and Adding Linux Servers respectively. By default, the role of a gateway server is assigned to the machine where Veeam Backup & Replication is installed.

   On the gateway server, Veeam Backup & Replication deploys Veeam Data Mover. Veeam Data Mover is responsible for handling ingress/egress requests that are sent to/from the gateway server when you work with the external repository data in Veeam Backup & Replication. If Veeam Data Mover becomes outdated, you must upgrade it as described in the Upgrading External Repositories section.

   Gateway servers store cached data. For more information, see Cache.



Step 4. Specify Cloud Storage Details

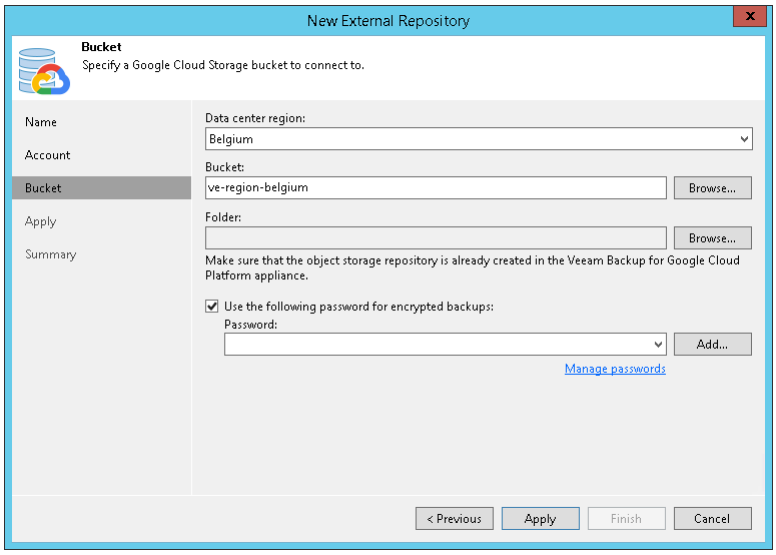
This step is available starting from Veeam Backup & Replication 11a (build 11.0.1.1261).

At the **Bucket** step of the wizard, specify an Google Cloud bucket and folder where Google Cloud instance backups reside:

1. From the **Data center region** drop-down list, select a region.
2. From the **Bucket** drop-down list, select a bucket.
3. In the **Select Folder** field, select a cloud folder where the data will be stored. To do it, click Browse and select an existing folder.
4. If the folder contains encrypted backups, select the **Use this password for encrypted backups** check box and provide a password. If you skip this step for encrypted backups, Veeam Backup & Replication will add such backups to the **External Repository (Encrypted)** node. For more information, see Viewing External Repository Data.

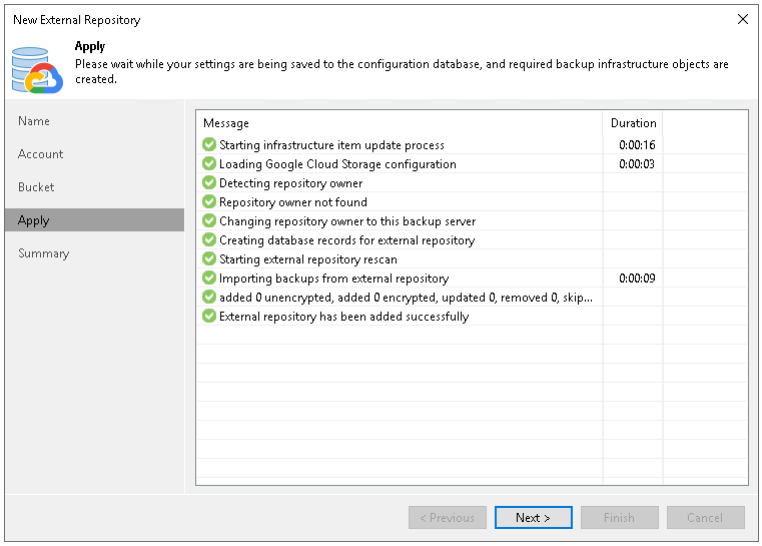
   For more information about encryption, see Encryption.

> **NOTE**
>
> Mind the following:
>
> - Only Standard storage class is supported.
> - If another Veeam Backup & Replication client has already added the same folder, you will be prompted whether to reclaim ownership of such a folder. For more information about ownership, see Ownership.



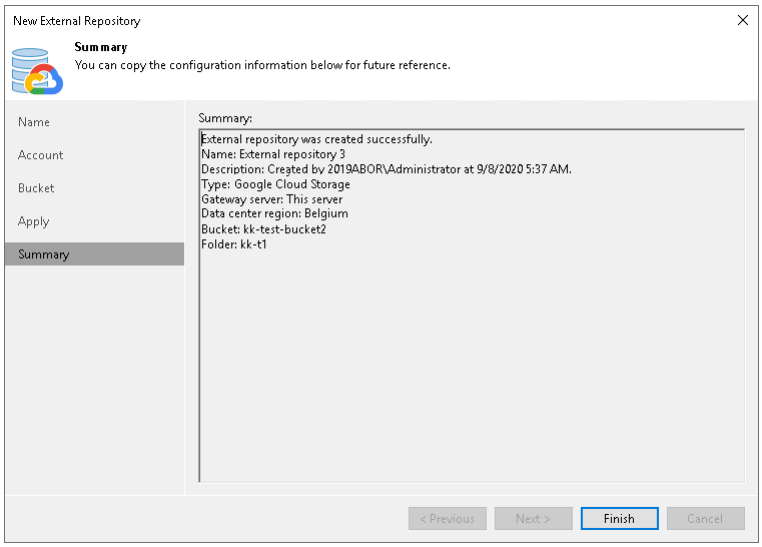
Step 5. Apply Settings

At the **Apply** step of the wizard, wait until Veeam Backup & Replication applies settings and completes adding the external repository. Then click **Next**.



Step 6. Finish Working with Wizard

At the **Summary** step of the wizard, review the external repository settings and click **Finish**.



Managing External Repositories

You can manage your external repositories in various ways: rescan external repositories to synchronize their state with the state of an object storage, upgrade external repositories by uploading a new version of Veeam Data Mover to the gateway server, edit settings of the external repositories, or remove the external repositories. For more information about these options, see the following sections.

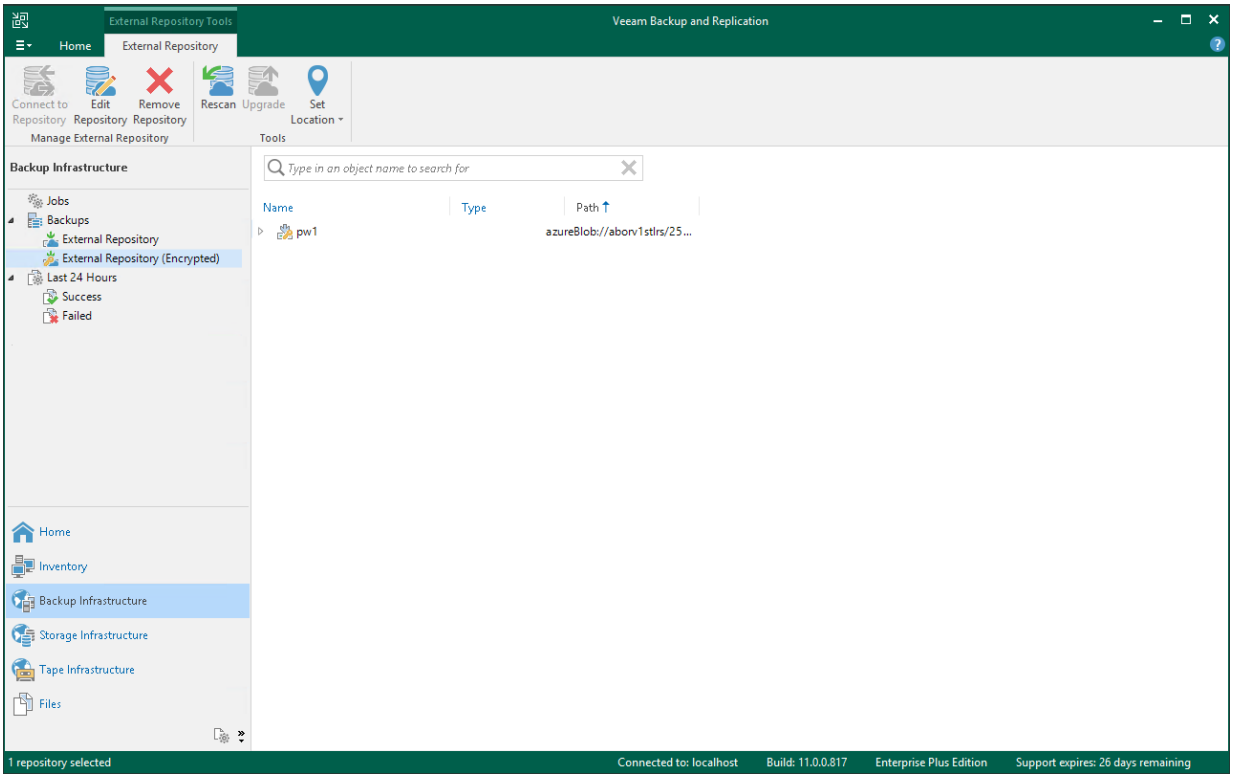
Viewing External Repository Data

After you add an external repository to the backup infrastructure, you can view backups in **External Repository** and **External Repository (Encrypted)** nodes of the **Backups** node in the **Home** view.

- In the **External Repository** node, Veeam Backup & Replication displays:
  - Amazon EC2 instance backups that were decrypted at the Bucket step of the **New External Repository** wizard.
  - Microsoft Azure VMs that were decrypted at the Container step of the **New External Repository** wizard.
  - Google Cloud Platform VM instances that were decrypted at the Bucket step of the **New External Repository** wizard.
- In the **External Repository (Encrypted)** node, Veeam Backup & Replication displays:
  - EC2 instance backups that were encrypted by Veeam Backup for AWS.
  - Microsoft Azure VMs that were encrypted by Veeam Backup for Microsoft Azure.

- Google Cloud Platform VM instances that were encrypted by Veeam Backup for Google Cloud Platform.

To decrypt backups, select a backup policy that created the backups you want to decrypt, click **Specify Password** on the ribbon, provide a password and click **OK**.



Rescanning External Repositories

To synchronize the state of an external repository with the state of an object storage (Amazon S3, Azure Blob or Google Cloud Storage), you can use the rescan feature. During rescan, Veeam Backup & Replication fetches newly created restore points and other required metadata.

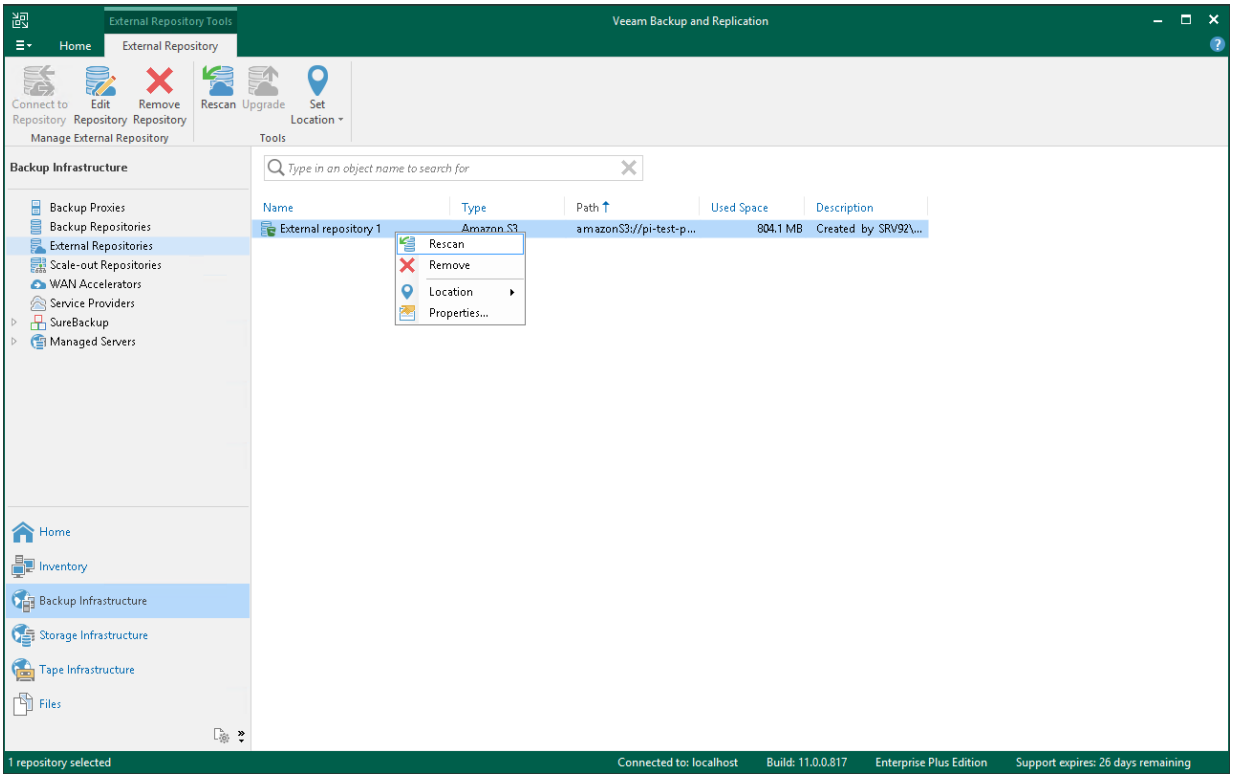
Consider the following:

- Rescan is done automatically in the following cases:
  - After you add an external repository to the backup infrastructure.
  - Every 24 hours.
  - After a backup chain is modified in the object storage. For example, if a restore point is added or deleted.
- Rescan session results are saved to the configuration database and can be found in the **History** view under the **System** node.

To rescan external repositories manually:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, click **External Repositories**.
3. Select a repository you want to rescan and click **Rescan** on the ribbon menu or right-click a repository and select **Rescan**.

If you have more than one external repository added to the scope, you may want to rescan all the repositories altogether. For that, right-click the root **External Repositories** node in the navigation pane and select **Rescan**.



Upgrading External Repositories

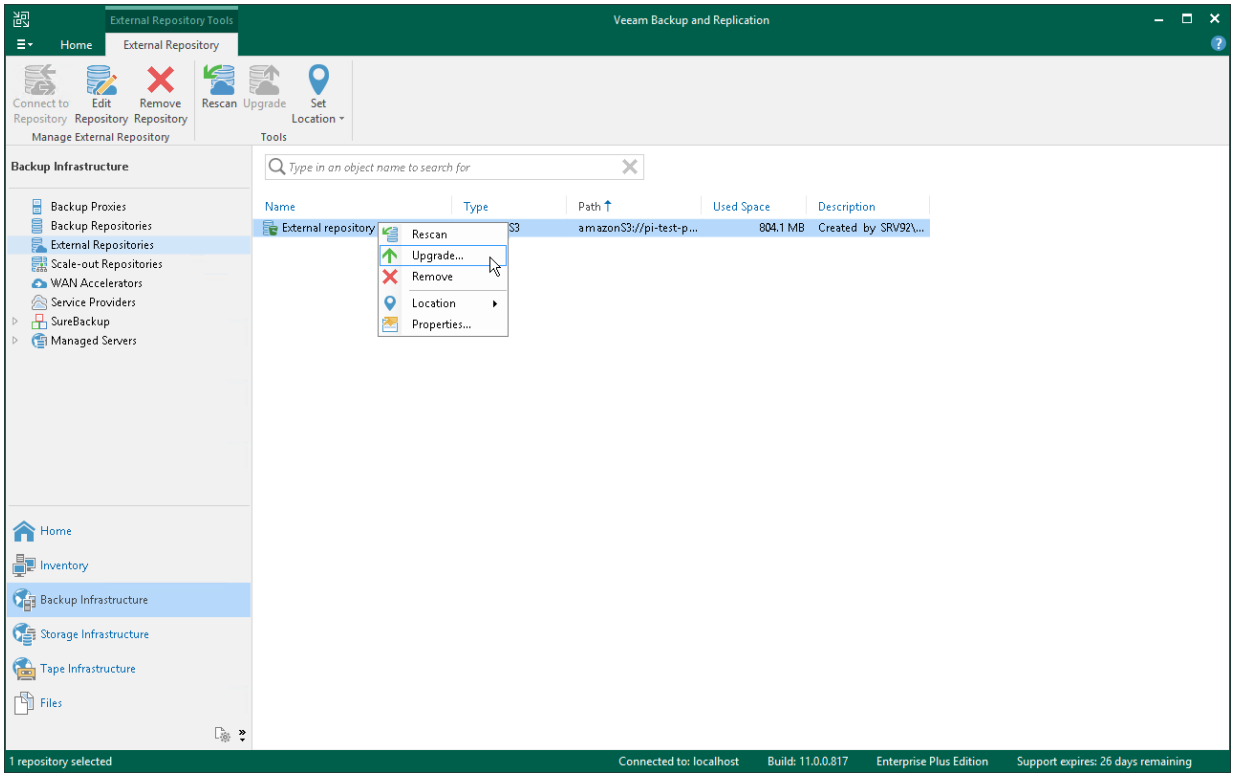
When you upgrade an external repository, a new version of Veeam Data Mover is uploaded to the gateway server. Veeam Data Mover is responsible for handling ingress/egress requests that are sent to/from the gateway server during working with the external repository data in Veeam Backup & Replication.

Upload of Veeam Data Mover is done directly to a gateway server which you specify at the Account step of the **New External Repository** wizard.

To upgrade an external repository:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, click **External Repositories**.

3. Select a repository you want to upgrade and click **Upgrade** on the ribbon menu or right-click a repository and select **Upgrade**.



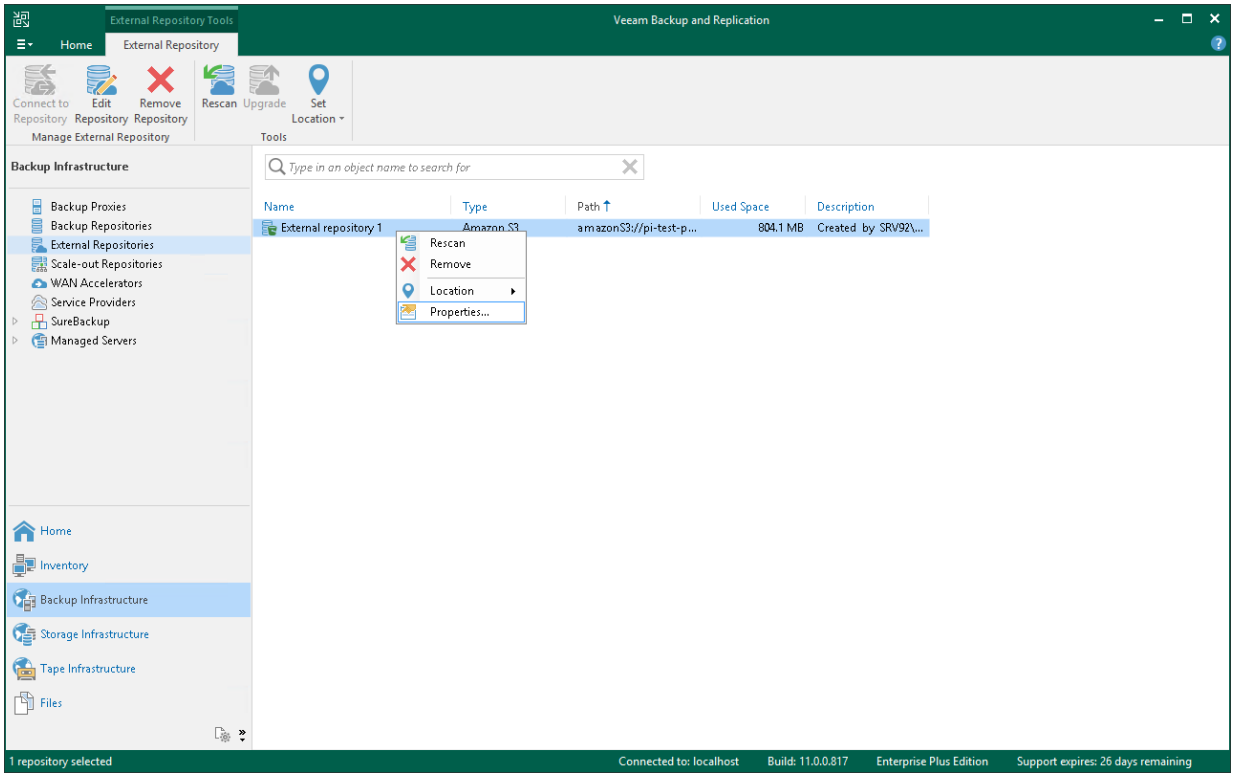
Editing Settings of External Repositories

To edit settings of an external repository:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, click **External Repositories**.
3. In the working area, select an external repository and click **Edit Repository** on the ribbon or right-click the external repository and select **Properties**.

4. Follow the steps of the **Edit External Repository** wizard and edit settings as required.

   Mind that some settings cannot be modified and will remain disabled during editing.



Removing External Repositories

You can remove an external repository from the backup infrastructure.

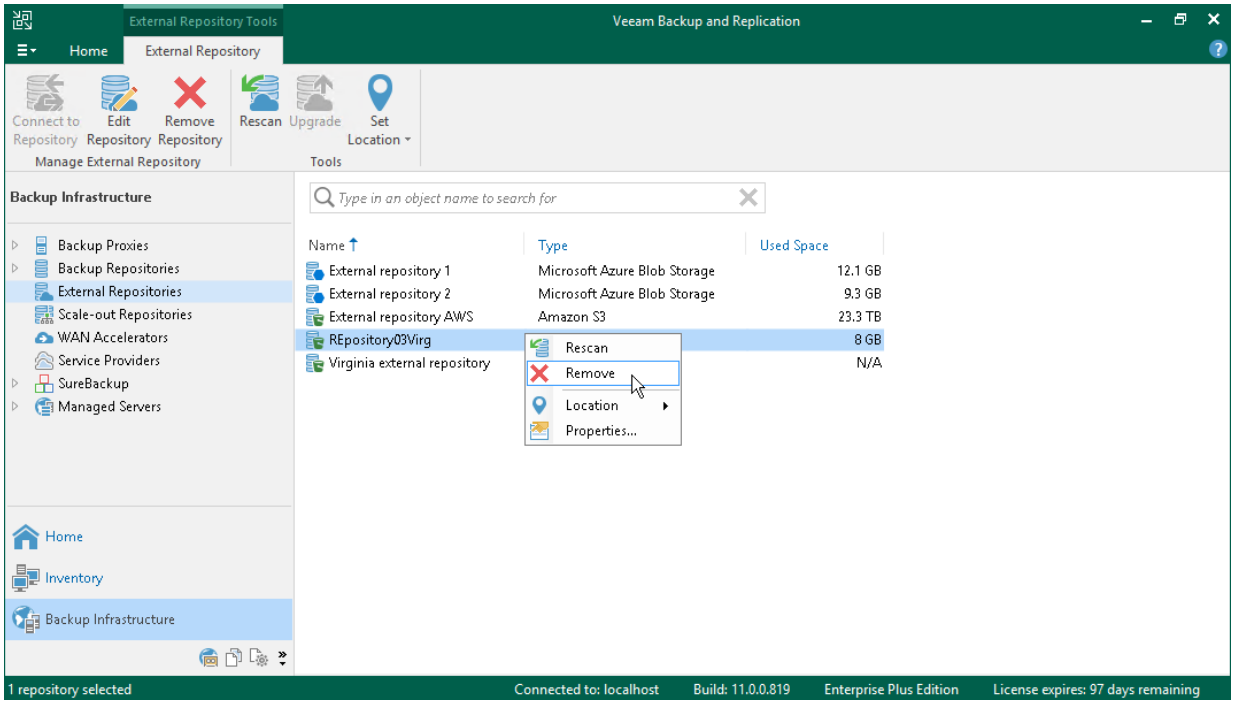
When you remove an external repository, Veeam Backup & Replication does the following:

- Relinquishes ownership.
- Removes associated external repository records from the configuration database.
- Removes associated cache from the gateway server.

To remove an external repository:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, click **External Repositories**.

3. In the working area, select an external repository and click **Remove Repository** on the ribbon or right-click the external repository and select **Remove**.



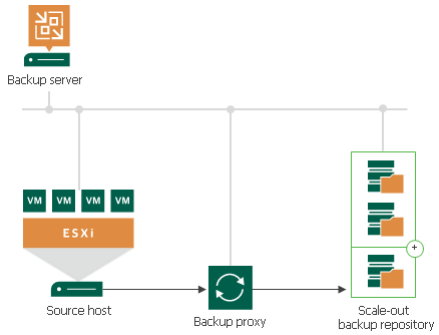
Scale-Out Backup Repository

A scale-out backup repository is a repository system with horizontal scaling support for multi-tier storage of data. The scale-out backup repository consists of one or more backup repositories called performance tier, and can be expanded with object storage repositories for long-term and archive storage: capacity tier and archive tier. All the storage devices and systems inside the scale-out backup repository are joined into a system, with their capacities summarized.

> **NOTE**
>
> Scale-out backup repository is included in the Veeam Universal License. When using a legacy socket-based license, Enterprise or higher edition is required.
>
> If you configure a scale-out backup repository and then downgrade to the Standard license, you will not be able to run jobs targeted at the scale-out backup repository. However, you will be able to perform restore from the scale-out backup repository.

The main benefits of this feature are as follows:

- It provides a convenient way of managing the backup storage.
- A scale-out backup repository can be expanded at any moment: if the extents of your scale-out backup repository run out of space, you can add a new extent to the existing scale-out backup repository. For example, if backup data grows and the backup repository reaches the storage limit, you can add a new storage system to the scale-out backup repository. The free space on this storage system will be added to the capacity of the scale-out backup repository. As a result, you will not have to move backups to a backup repository of a larger size.
- It supports any backup target supported by Veeam: Windows or Linux servers with local or DAS storage, network shares, deduplicating storage appliances. All the features of any storage device or system are preserved.
- It allows you to set up granular performance policy. For more information, see Backup File Placement.
- It provides practically unlimited cloud-based storage capacity: you can instruct Veeam Backup & Replication to offload data from extents to the cloud for long-term storage. For details, see Capacity Tier.



A scale-out backup repository can comprise different tiers, or logical levels of storage.

- Performance tier is the level used for the fast access to the data. It consists of one or more backup repositories called performance extents.

  For more information, see Performance Tier.

- Capacity tier is an additional level for storing data that needs to be accessed less frequently. However, you still can restore your data directly from it. The capacity tier consists of one cloud-based or on-premises object storage repository called capacity extent.

  For more information, see Capacity Tier.

- Archive tier is an additional level for archive storage of infrequently accessed data. Applicable data from the capacity tier can be transported to the archive tier. For restore from the archive tier, data must undergo preparation process.

  For more information, see Archive Tier.

A scale-out backup repository can be used for the following types of jobs and tasks:

- Backup jobs
- Backup copy jobs

  You can copy backups that reside in scale-out backup repositories and store backup copies in scale-out backup repositories

- Veeam backups for Amazon and Microsoft Azure (via backup copy jobs)
- VeeamZIP tasks
- Backup jobs created by Veeam Agent for Linux 2.0 or later
- Backup jobs created by Veeam Agent for Unix 2.0 or later
- Backup jobs created by Veeam Agent for Microsoft Windows 2.0 or later
- Backup jobs created by Veeam Agent for Mac
- NAS backup jobs
- Backup jobs created by Veeam Backup for Nutanix AHV

Backup files stored in the scale-out repository can be used for all types of restores, replication from backup and backup copy jobs. You can verify such backups with SureBackup jobs. The scale-out backup repository can be used as a staging backup repository for restore from tape media. Files restored from the tape media are placed to the extents according to data placement policy configured for the scale-out backup repository. For more information, see Backup File Placement.

To deploy a scale-out backup repository, you must configure one or more backup repositories and add them to a scale-out backup repository as extents. You can mix backup repositories of different types in one scale-out backup repository:

- Microsoft Windows backup repositories
- Linux backup repositories
- Shared folders
- Deduplicating storage appliances

For example, you can add a Microsoft Windows backup repository and deduplicating storage appliance to the same scale-out backup repository.

Limitations for Scale-Out Backup Repositories

Scale-out backup repositories have the following limitations:

- You cannot use a scale-out backup repository as a target for the following types of jobs:
  - Configuration backup job
  - Replication jobs (including replica seeding)
  - VM copy jobs
  - Veeam Agent backup jobs created by Veeam Agent for Microsoft Windows 1.5 or earlier
  - Veeam Agent backup jobs created by Veeam Agent for Linux 1.0 Update 1 or earlier
- You cannot add a backup repository as an extent to the scale-out backup repository if any job of unsupported type is targeted at this backup repository or if the backup repository contains data produced by jobs of unsupported types (for example, replica metadata). To add such backup repository as an extent, you must first target unsupported jobs to another backup repository and remove the job data from the backup repository in question.
- Scale-out backup repositories do not support rotated drives. If you enable the **This repository is backed by rotated hard drives** setting on an extent, Veeam Backup & Replication will ignore this setting and will work with such repository as with a standard extent.
- If a backup repository is added as an extent to the scale-out backup repository, you cannot use it as a regular backup repository. You cannot target jobs to this backup repository. Instead, you have to target jobs to the configured scale-out backup repository.
- You cannot add a scale-out backup repository as an extent to another scale-out backup repository.
- You cannot add a backup repository as an extent if this backup repository is already added as an extent to another scale-out backup repository.
- You cannot add a backup repository as an extent if this backup repository is already used as a backup destination by vCloud Director organizations.
- You cannot add a backup repository in which some activity is being performed (for example, a backup job or restore task) as an extent to the scale-out backup repository.
- If you apply the **Forget** or **Remove from disk** options to a missing restore point in a scale-out backup repository, the backup file associated with the missing restore point will be deleted from capacity tier and archive tier on the next offload and archiving job run.
- If you use Enterprise edition of Veeam Backup & Replication, you can create two scale-out backup repositories.

  For each scale-out backup repository, you can add one object storage repository as a capacity extent and four standard backup repositories as performance extents: three active, and one inactive (that is put to the maintenance mode). You can add inactive extents, for example, if any of active extents has no free space, and you want to evacuate backup data from it.

  If you add four performance extents and do not put any of them to the maintenance mode, the jobs targeted at the scale-out backup repository will fail.

  Veeam Universal License and Enterprise Plus editions have no limitations on the number of scale-out backup repositories or performance extents.
- If you want to use the Extract utility to work with backup files located on any of the extents of your scale-out backup repository, make sure that incremental and full backup files are located on the same extent.

- You can use the Veeam Backup Validator utility only for backups stored in the performance tier which consists of backup repositories. Make sure that incremental and full backup files are located on the same extent.
- To let Veeam Backup & Replication automatically import backups during rescan of a scale-out backup repository, names of VBM files and paths to VBM files (starting from the backup repository root to VBM files, not including the root itself) must contain only allowed characters:
  - Alphanumeric characters: a-zA-Z0-9
  - Special characters: _-.+=@^

  Names of VBM file and paths to VBM files must not contain spaces.

  If a name of a VBM file or a path to a VBM file contains prohibited characters, Veeam Backup & Replication will fail to import such backup during rescan of the scale-out backup repository. To import such backup, you can replace prohibited characters with the underscore character, for example: `C:\My Repository\Backup_Job\Backup_Job.vbm`. You do not need to rename the actual backup files.
- Veeam Backup & Replication does not split one backup file across multiple extents.
- An object storage repository added as a capacity tier in a scale-out backup repository can not be used for storing NAS backups. To archive NAS backup files to an object storage repository, assign the object storage repository as an archive repository when you create a file share backup job.
- You cannot assign the role of a cache repository for file share backup to a scale-out backup repository and its extents. To learn more about cache repository, see NAS Backup Support.
- If a repository is being used as a cloud repository, you cannot add it as an extent of a scale-out backup repository.
- [For Nutanix AHV VM backups] Due to specifics of backup jobs for AHV VMs, Veeam Backup for Nutanix AHV always creates a separate backup chain for each VM added to a backup job. Thus, even if you clear the **Use per-machine backup files** check box in the advanced settings of a scale-out backup repository, backups of multiple AHV VMs are not stored in a single backup file.

Immutability

Veeam Backup & Replication allows you to prohibit deletion of data from the extents of the scale-out backup repository by making that data temporarily immutable. It is done for increased security: immutability protects your data from loss as a result of attacks, malware activity or any other injurious actions.

You can enable the immutability feature for any tier of scale-out backup repository.

To learn how immutability works with performance tier of the scale-out backup repository, see Hardened Repository.

To learn how immutability works with capacity tier of the scale-out backup repository, see Immutability for Capacity Tier.

To learn how immutability works with archive tier of the scale-out backup repository, see Immutability for Archive Tier.

Preparing to Use Immutability

In order to use immutability, you need to enable the *Object Lock* and *Versioning* features on your S3 bucket at the time you create the bucket. Keep in mind that most vendors allow enabling *Object Lock* only at the moment of creating the bucket.

For more information on enabling the *Object Lock* and *Versioning* features, see these Amazon articles: How do I create an S3 Bucket? and How do I enable or suspend versioning for an S3 bucket?.

For considerations and limitations on immutability, see Considerations and Limitations.

Performance Tier

Performance tier of a scale-out backup repository is the level used for fast access to the data.

The performance tier of a scale-out backup repository can comprise one or more performance extents. A performance extent is a backup repository (except an object storage repository) added to the scale-out backup repository. The list of the performance extents is displayed at the Add Performance Extents step of the **New Scale-out Backup Repository** wizard.

On every performance extent, Veeam Backup & Replication creates the `definition.erm` file. This file contains a description of the scale-out backup repository and information about its performance extents.

The backup repository added to the scale-out backup repository as a performance extent ceases to exist as an individual backup repository.

When a backup repository is added as a performance extent, some of its original settings are kept, and some are not. The following settings are kept, or inherited:

- Number of tasks that can be performed simultaneously
- Read and write data rate limit
- Data decompression settings
- Block alignment settings

The following settings are not inherited:

- Rotated drive settings. Rotated drive settings are ignored and cannot be configured at the level of the scale-out backup repository.
- Per-machine backup file settings. Per-machine settings can be configured at the level of the scale-out backup repository.

Limitations

The same limitations that are specific for certain types of backup repositories apply to the performance extents. For example, if you add Dell EMC Data Domain as a performance extent to the scale-out backup repository, you will not be able to create a backup chain longer than 60 points in this scale-out backup repository.

The performance extents of the scale-out backup repository should be located in the same site. Technically, you can add to the scale-out backup repository the performance extents that reside in different sites. However, in this case Veeam Backup & Replication will have to access VM backup files on storage devices in different locations, and the backup performance will degrade.

Backup File Placement

Veeam Backup & Replication stores backup files on all performance extents of the scale-out backup repository.

When you configure a scale-out backup repository, you must set the backup file placement policy for it. The backup file placement policy describes how backup files are distributed between extents. You can choose one of the two policies:

- Data locality
- Performance

You can also select an extent for backup file placement, which has its nuances if you set the Performance policy for the scale-out backup repository:

- Extent Selection
- Extent Selection for Backup Repositories with Performance Policy

Keep in mind that at the beginning of a job, Veeam Backup & Replication retrieves the actual free space on a scale-out backup repository and estimates the size of a restore point to be created. Then the estimated restore point size is subtracted from the actual free space to determine the estimated free space. This estimated free space is then used by all further concurrently running jobs targeted at the same scale-out backup repository. For details, see Backup Size Estimation.

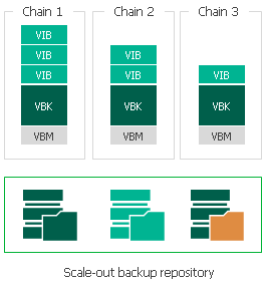
The backup file placement policy is not strict. If the necessary extent is not accessible, Veeam Backup & Replication will disregard the policy limitations and attempt to place the backup file to the extent that has enough free space for the backup file.

For example, you have set the *Performance* policy for the scale-out backup repository and specified that full backup files must be stored on *Extent 1* and incremental backup files must be stored on *Extent 2*. If before an incremental backup job session *Extent 2* goes offline, the new incremental backup file will be placed to *Extent 1*.

> **NOTE**
>
> Veeam Backup & Replication places backups of Microsoft SQL transaction logs and Oracle archived logs to the extent configured for storing incremental backup files. If such extent is not accessible, Veeam Backup & Replication will attempt to place log backups to any other extent that has enough free space.

Data Locality

If you set the *Data locality* policy for a scale-out backup repository, all backup files that belong to the same backup chain are stored on the same extent of the scale-out backup repository.



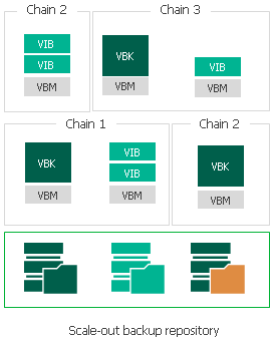
The *Data locality* policy does not put any limitations to backup chains. A new backup chain may be stored on the same extent or another extent. For example, if you create an active full backup, Veeam Backup & Replication may store the full backup file to another extent, and all dependent incremental backup files will be stored together with this full backup file.

However, if you use a deduplicating storage appliance as an extent to the scale-out backup repository, Veeam Backup & Replication will attempt to place a new full backup (active or synthetic) to the extent where the full backup from the previous backup chain resides. Such behavior will help increase the data deduplication ratio.

> **TIP**
>
> If you plan to use Fast Clone on performance extents with volumes formatted with ReFS or XFS, select Data locality.

Performance

If you set the *Performance* policy for a scale-out backup repository, full backup files and incremental backup files that belong to the same backup chain are stored on different extents of the scale-out backup repository. If necessary, you can explicitly specify on which extents full backup files and incremental backup files must be stored.



The *Performance* policy can improve performance of transformation processes if you use raw data devices as extents. When Veeam Backup & Replication performs transformation, it needs to access a number of backup files in the backup repository. If these files are located on different storage devices, the I/O load on the devices hosting backup files will be lower.

If you set the *Performance* policy, you must make sure that the network connection between extents is fast and reliable. You must also make sure all extents are online when the backup job, backup copy job or a restore task starts. If any extent hosting backup files in the current backup chain is not available, the backup chain will be broken, and Veeam Backup & Replication will not be able to complete the task. To avoid data loss in this situation, you can enable the **Perform full backup when required extent is offline** option for the scale-out backup repository. With this option enabled, Veeam Backup & Replication will create a full backup instead of incremental backup if some files are missing from the backup chain.

Extent Selection

To select an extent for backup file placement, Veeam Backup & Replication checks the following conditions:

1. Availability of extents on which backup files reside. If some extent with backup files from the current backup chain is not accessible, Veeam Backup & Replication will trigger a full backup instead of incremental (if this option is enabled). For more information, see Adding Backup Repository Extents.
2. Backup placement policy set for the scale-out backup repository.
3. Load control settings — maximum number of tasks that the extent can process simultaneously.
4. Amount of free space available on the extent — the backup file is placed to the extent with the most amount of free space.

5. Availability of files from the current backup chain — extents that host incremental backup files from the current backup chain (or current VM) have a higher priority than extents that do not host such files.

Extent Selection for Backup Repositories with Performance Policy

If you set the *Performance* policy for the scale-out backup repository, Veeam Backup & Replication always stores full backup files and incremental backup files that belong to the same backup chain on different extents. To choose the extent to which a backup file can be stored, Veeam Backup & Replication applies this policy and policies mentioned above.

For example, a scale-out backup repository has 2 extents that have 100 GB and 200 GB of free space. You set the *Performance* policy for the scale-out backup repository and define that all types of backup files (full and incremental) can be placed on both extents.

When a backup job runs, Veeam Backup & Replication picks the target extent in the following manner:

1. During the first job session, Veeam Backup & Replication checks to which extent a full backup file can be stored. As both extents can host the full backup file, Veeam Backup & Replication checks which extent has more free space, and picks the extent that has 200 GB of free space.
2. During incremental job session, Veeam Backup & Replication checks to which extent an incremental backup file can be stored. As both extents can host the incremental backup file, Veeam Backup & Replication picks the extent that does not store the full backup file — the extent that has 100 GB of free space.

Backup Size Estimation

At the beginning of a job, Veeam Backup & Replication retrieves the actual free space on a scale-out backup repository and estimates the size of a restore point to be created. Then the estimated restore point size is subtracted from the actual free space to determine the estimated free space. This estimated free space is then used by all further concurrently running jobs targeted at the same scale-out backup repository.

Veeam Backup & Replication assumes that the following amount of space is required for backup files:

- The size of the first full backup file is equal to 50% of source VM data.
- The size of further full backup files is equal to 100% of the previous full backup file size.
- The size of the first incremental backup file is equal to 10% of the previous full backup file size.
- The size of further incremental backup files in the backup chain is equal to 100% of the previous incremental backup file size.

  In case of reverse incremental backup chains, during incremental job sessions Veeam Backup & Replication allocates 10% of the previous full backup file size on the extent where a rollback file is placed and additional 10% on the extent where the full backup file resides.

This mechanism is also applied to backup files created with backup copy jobs.

Mind the following:

- On every extent of a scale-out backup repository, Veeam Backup & Replication reserves 1% of storage space to guarantee correct update of backup metadata files (VBM) and success of merge operations.
- Make sure that you have enough free space on the extent where the full backup file resides. Veeam Backup & Replication requires 10% of the size of the full backup file to perform merge operations in the backup chain. If the disk space is low, merge operations may fail.

- The actual free space value is only captured at the start of a job targeted at the scale-out backup repository while no other jobs actively use the same scale-out backup repository. For more information, see the Veeam KB2282 article.

Removing Performance Extents from Scale-Out Repositories

You can remove a performance extent from the scale-out backup repository, for example, if you do not want to store backup files on the underlying storage anymore. When you remove a performance extent, Veeam Backup & Replication puts the relevant backup repository to the maintenance mode and unassigns the performance extent role from it. The backup repository ceases to exist as a part of the scale-out backup repository and becomes an individual backup repository.

If the performance extent contains backup files, it is strongly recommended that you perform the following actions before you remove the extent:

1. Put the performance extent to the maintenance mode.
2. Evacuate backups from the extent. For more information, see Evacuating Backups from Performance Extents.

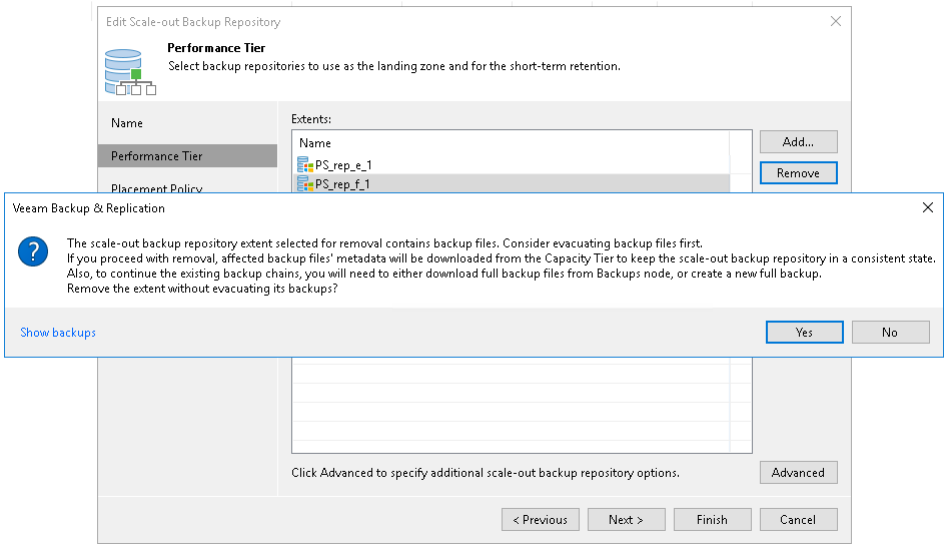
In this case, backup files will be moved to other performance extents of the scale-out backup repository, and the backup chain will remain consistent. If you do not evacuate backups from the performance extent, the backup chain may get broken as some restore points will be missing from it.

To remove a performance extent from the scale-out backup repository:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, click **Scale-out Repositories**.
3. In the working area, select the scale-out backup repository and click **Edit Scale-out Repository** on the ribbon or right-click the scale-out backup repository and select **Properties**.
4. Move to the **Performance tier** step of the wizard.
5. In the **Extents** list, select the performance extent and click **Remove**.

If the performance extent contains backup files, Veeam Backup & Replication will suggest evacuating them. To evacuate files, click **No**, close the wizard and evacuate backup files. For more information, see Evacuating Backups from Performance Extents.

If you do not want to evacuate the backup files, click **Yes** and proceed with the wizard.



Capacity Tier

Capacity tier is an additional tier of storage that can be attached to a scale-out backup repository. Data from the scale-out backup repository performance extents can be transported to the capacity tier for long-term storage.

This feature is most useful if:

- You are running out of storage space.
- Your organization policies allow you to store only a certain amount of data on your extents, while the outdated data should be stored elsewhere.
- You seek to store data on several sites to ensure its safety in case of a disaster.

With capacity tier you can:

- Move inactive backup chains to the capacity extent, as described in Moving Backups to Capacity Tier and Manually Moving Backups to Capacity Tier.
- Copy new backup files as soon as these files are created, as described in Copying Backups to Capacity Tier.
- Download data that was moved from the capacity extent back to the performance extents, as described in Downloading Data from Capacity Tier.
- Restore your data. For more information, see Restore from Capacity Tier. In particular, you can promptly restore data from the capacity tier in case of disaster without creating a scale-out backup repository anew. For more information about this feature, see Importing Object Storage Backups.

The capacity tier consists of a single capacity extent. The capacity extent can be either cloud-based object storage repository or on-premises object storage repository, such as:

- S3-compatible object storage repository
- Amazon S3
- Microsoft Azure Blob storage
- Microsoft Azure Data Box
- IBM Cloud Object Storage
- Google Cloud Object Storage

Before an object storage repository can be configured as the capacity extent, it must be added to Veeam Backup & Replication. For more information, see Adding Object Storage Repositories.

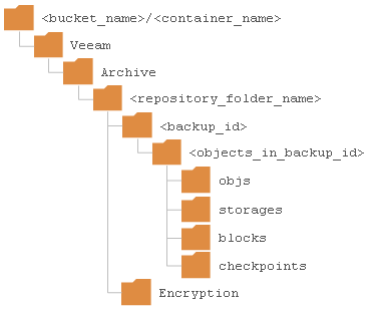
The capacity extent is displayed in the scale-out backup repository wizard, on the Add Capacity Tier step.

For information on configuring capacity tier and synchronizing capacity tier data, see Add Capacity Tier.

> **NOTE**
>
> Before you start using the capacity tier, make sure to check the pricing plans of your cloud storage provider to avoid additional costs for offloading and downloading backup data.

Capacity Extent Structure

When data is being transferred to the capacity extent, Veeam Backup & Replication creates and maintains the following structure of directories:



Directory	Description	Misc
<bucket_name> or <container_name>	A bucket or container name. Buckets and containers must be created in advance.	N/A
Veeam/Archive/	Standard folders created by Veeam Backup & Replication.	
<repository_folder_name>	A repository folder that you create when adding a new capacity extent.	
<backup_id>	Contains objects in a backup.	These folders are automatically removed during data removal.
<objects_in_backup_id>	An identifier of an object in a backup. • If a backup was created using the per-machine method, then each VM will be placed to its own directory. • If a backup was created as a single storage, then all the VMs will be placed to a unique directory.	
objs	Contains meta information and other auxiliary data.	
storages	Contains a replicated version of offloaded backup files with metadata that also remain on the source extents.	
blocks	Contains offloaded data blocks created by the offload session, as described in Data Transfer. The maximum data block size is configured in the storage optimization settings.	

Directory	Description	Misc
checkpoints	Contains meta information about the state of offloaded backup chains. Such meta information is updated upon each successful offload session.	
Encryption	Contains required information to work with encrypted backups.	N/A

Immutability for Capacity Tier

To make the data immutable, Veeam Backup & Replication uses the Object Lock technology provided by Amazon and some S3-compatible providers. Once imposed, the object lock prohibits deletion of data from the capacity extent until the immutability expiration date comes.

You can enable the immutability feature when adding or editing an Amazon S3 or S3-compatible object storage repository as a capacity extent, at the Bucket step of the Adding Amazon S3 Object Storage or Adding S3 Compatible Object Storage wizards. The immutability expiration date is specified at the same point.

Backups are immutable only during the immutability period set in the bucket even if their retention policy allows for longer storage. This also applies to the following types of backups:

- Backups with GFS flags
- Backups created with VeeamZIP jobs
- Exported backup files

The immutable data within the capacity extent cannot be subject to the following operations:

- Manual removal of data, as described in Removing Backups from Capacity Tier.
- Removal of data by the retention policy, as described in Retention Policy.
- Removal of data using any cloud service provider tools.
- Removal of data by the cloud service provider technical support department.
- Removal of data by the **Remove deleted items data after** option, as described in Maintenance Settings.

> **TIP**
>
> There are two ways to migrate your data from a non-immutable bucket to an immutable one. For more information, see section Migrating Data Between Different Buckets.

Block Generation

To reduce I/O operations and associated costs, Veeam Backup & Replication will automatically add from 1 to 10 days to the immutability expiration date. This period is called Block Generation. You do not have to configure it, the Block Generation setting is applied automatically.

For example, if you set your immutability period to 30 days, Veeam Backup & Replication will add from 1 to 10 days to specific objects to reduce I/O operations with the storage over time. This will not change the retention and their effective immutability. It is a background optimization. Thus, if you need 30 days immutability period, set the period to 30 days.

Block Generation works in the following way. When the first data block (a full backup) arrives, its immutability period is set to 30 + 10 = 40 days. The first full backup starts its generation, that will be appended with the incremental backups. All the incremental backups within the generation (i.e. within the 10-days period) will have the same immutability expiration date as the full backup. For instance, a data block that was offloaded on day 9 will have the same immutability expiration date as a data block offloaded on day 1. Thus we ensure that the immutability period for all the data blocks within a generation is no less than 30 days.

Consider this example: within one forward incremental backup chain, a full backup file can not be removed before an incremental backup file. On the other hand, an incremental backup file makes no sense without relevant full backup file. So the immutability period is extended for all data blocks in the backup chain.

The Block Generation period was introduced in order to reduce the number of requests to the capacity tier, thereby saving traffic and reducing costs that might be incurred by your storage provider. With 10 days of immutability automatically added, it means there is no need to extend the immutability period for the incremental backups in forward chain and for the unchanged blocks of current full backups in reverse chain offloaded within those 10 days. On the 11th day, though, the immutability period will have to be extended (in order to ensure that the immutability period for all the data blocks within a generation is no less than 30 days).

The immutable blocks of data may be reused during the offload. Veeam Backup & Replication continues to keep reused or dependent blocks of data locked by continuously assigning them to new generations, thereby extending their immutability expiration period. This is valid both for forward and reverse incremental backup chains.

The extension of immutability works differently in different cases, as follows:

Forward incremental backup chain

- New full backup file in the new generation:
  - Immutability is extended for the data blocks that are being reused from the old backup chain,
  - Immutability is set anew for the new blocks of the new full backup file.
- New incremental backup file in the new generation:
  - Immutability is extended for all the data blocks from the current backup chain,
  - Immutability is set anew for the new blocks of the latest incremental backup file.

Reverse incremental backup chain

- New full backup file in the new generation:
  - Immutability is extended for the data blocks that are being reused from the previous full backup file,
  - Immutability is set anew for the new blocks of the new full backup file.
- Current full backup file in the new generation:
  - Immutability is extended for all the data blocks of this full backup file that are already stored in the capacity tier,
  - Immutability is set anew for the new blocks of this full backup file.

Encryption for Capacity Tier

Veeam Backup & Replication allows you to encrypt offloaded data. This helps you protect the data from an unauthorized access.

You can enable data encryption in the following ways:

- When you create a backup or backup copy job

- When you add a capacity tier extent to your scale-out backup repository

To get benefits of both encryption levels, you can use job-level and capacity tier encryption within the same object storage. Both encryption levels allow you to keep your data from an unauthorized access, but capacity tier encryption allows you to encrypt backup chain metadata and restore points.

Job-level Encryption

Before data is offloaded to capacity tier, Veeam Backup & Replication checks if encryption is enabled in the job settings. If encryption is enabled, data encrypted by the job is not decrypted or decompressed. It is offloaded to capacity tier as is.

Capacity Tier Encryption

With the **Encrypt data uploaded to object storage** setting selected, the entire collection of blocks along with the metadata will be encrypted while being offloaded regardless of the jobs' encryption settings. If you have both job-level and capacity tier encryption enabled, already encrypted backup data will be encrypted again before being uploaded to capacity tier.

If capacity tier encryption has been disabled, backup data encrypted by the job's settings will be uploaded unmodified to capacity tier.

> **TIP**
>
> Make sure you enable encryption when you add a capacity tier extent. Otherwise, unencrypted backups, offloaded to capacity tier, may be reused for synthetic backup creation.
>
> If you enable encryption after you have already offloaded data to capacity tier, Veeam Backup & Replication will not encrypt previously offloaded backup chains.

Managing Capacity Tier

You can manage your capacity tier and the offloaded data in the following ways:

- Download previously offloaded data from the capacity extent back to the performance extents.
- Migrate data to another capacity tier.
- Reduce the number of cost-based operations incurred by your cloud storage provider and decrease the amount of traffic being sent over the network when moving or copying data to the capacity tier. For more information, see Indexes.
- Configure the retention policy.
- Exclude the capacity extent from the scale-out backup repository scope.

Data Transfer

Veeam Backup & Replication offers the following opportunities of data transfer to and from the capacity extent:

- Copy policy: the backups are copied to the capacity extent automatically.
- Move policy: the inactive backup chains can be transferred to the capacity extent.
- Previously offloaded data can also be downloaded from the capacity tier back to the performance extents.

To manage data transfer to and from the capacity extent, Veeam Backup & Replication uses system sessions.

The following types of backup files can be moved or copied to capacity tier:

- Regular backups (except transaction logs)
- Veeam backups for Amazon, Google and Microsoft Azure (via backup copy jobs)
- Backups created with Veeam Agent for Microsoft Windows, Veeam Agent for Linux, Veeam Agent for Unix or Veeam Agent for Mac
- Backups created with Veeam Plug-ins for Enterprise Applications (Oracle RMAN, SAP HANA, SAP on Oracle)
- Backups created with Veeam Backup for Nutanix AHV
- Backups created with VeeamZIP jobs
- Exported backup files
- Orphaned backups

Imported backups can be copied or moved to capacity tier only manually.

In This Section

- Copying Backups to Capacity Tier
- Moving Backups to Capacity Tier
- Manually Moving Backups to Capacity Tier
- Downloading Data from Capacity Tier
- Viewing Capacity Tier Sessions Statistics

Copying Backups to Capacity Tier

Veeam Backup & Replication allows you to copy backups from the performance extents of your scale-out backup repository to the capacity extent as soon as these backups are created.

To enable data copy, make sure to select the **Copy backups to object storage as soon as they are created** option, as described in Add Capacity Tier.

To copy data to the capacity extent, Veeam Backup & Replication uses a copy session that is executed automatically upon completion of a backup (or backup copy) job that is targeted to a scale-out backup repository that was extended with the capacity extent.

A complete name of each copy session is built up of the backup (or backup copy) job name plus *Offload* postfix. That is, if your backup (or backup copy) job name is *Amazon*, the copy session name will be *Amazon Offload*.

Before your data can be copied to the capacity tier, Veeam Backup & Replication performs the following mandatory verifications and required actions:

- Verifies whether performance extents are available and have not been put into the maintenance mode.
- Builds and maintains indexes to verify whether data that is being moved is unique and has not been offloaded earlier.

Backup files with metadata are created as described in Moving Backups to Capacity Tier.

Once the backup (or backup copy) job is complete, Veeam Backup & Replication initiates a new copy session which simply extracts data blocks and metadata from each new backup file (.VBK, .VIB, .VRB) created on any of the performance extents of your scale-out backup repository and copies these blocks and metadata to the capacity extent, thereby making an identical replica of your backup data.

Having such replica gives you the ability to quickly restore data as of the latest state in case of trouble with any backup files, any unexpected failure of any of your performance extents, or even of the entire scale-out backup repository, as described in Restore Scenarios.

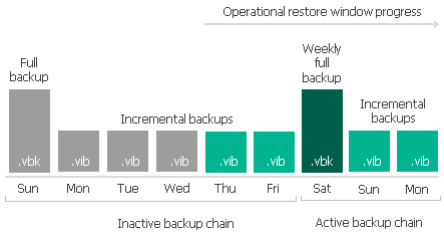
Combining Copy and Move Operations

You can combine both the **Copy backups to object storage as soon as they are created** option and the **Move backups to object storage as they age out of the operational restores window** option. In such a scenario, a copy session will be copying newly created backups right upon creation.

Once the backup chain becomes inactive (for example, sealed) and exceeds the operational restore window, data blocks will be removed from each associated backup file in such an inactive backup chain and only metadata will be preserved. Such a behavior mimics data movement, but instead of moving data that was already copied, Veeam Backup & Replication simply purges associated data blocks from the performance extents, thereby saving traffic and reducing costs that might be incurred by your storage provider for performing read/write operations.

The following figure shows an example in which both options are enabled, suggesting that each backup file has been copied to object storage upon its creation.

The backup chain on the left becomes inactive after a new full backup file is created and consists of one .VBK file and five .VIB files. Since only the first four backup files (represented as grey blocks) in this inactive backup chain exceed the operational restore window, Veeam Backup & Replication removes blocks of data from these four files only and leaves the other two .VIB files (created on *Thursday* and *Friday*) as they are until the operational restore window moves forward, and so on.



After copy is complete, the new **Object Storage** node appears in the **Home** view, under the **Backups** node and shows backups that have been copied to the capacity extent.

> **NOTE**
>
> The copy is not performed during prohibited hours specified in the scale-out backup repository backup window configuration. You can configure the backup window at the Add Capacity Tier step of the **New Scale-out Backup Repository** wizard.

Moving Backups to Capacity Tier

To collect backup files that belong to inactive backup chains from the performance extents and move them to the capacity tier, Veeam Backup & Replication uses an offload session which is executed automatically every 4 hours.

To enable data movement, make sure to select the **Move backups to object storage as they age out of the operational restores window** option, as described in Add Capacity Tier.

A complete name of each offload session is built up of the scale-out backup repository name plus *Offload* postfix. That is, if your scale-out backup repository name is *Amazon*, the offload session name will be *Amazon Offload*.

The offload session manages the following:

- Validation Process
- Data Transfer

Validation Process

Before your data can safely be moved to the capacity tier, Veeam Backup & Replication performs the following mandatory verifications and required actions:

- Verifies whether data that is about to be moved belongs to an inactive backup chain.

  For more information, see Backup Chain Detection.

- Verifies whether performance extents are available and have not been put into the maintenance mode.

  Consider that data will not be offloaded from Linux-based performance extents that have internet access through HTTP(S) proxy. All Linux-based performance extents configured in your scale-out backup repository must have direct access to the internet.

- Verifies whether the capacity extent has not been put into the maintenance or the seal mode.

  For more information, see Switching to Maintenance Mode and Switching to Seal Mode.

- Verifies whether configuration parameters that define how and when inactive backup chains must be moved to the capacity extent are met.

  Such parameters are configured as described in Add Capacity Tier.

- Builds and maintains indexes to verify whether data that is being moved is unique and has not been offloaded earlier.

  For more information, see Indexes.

- Synchronizes the backup chain state between the performance and capacity extents to maintain retention policies.

  For more information, see Retention Policy.

Data Transfer

After the validation process is complete, the *SOBR Offload* session does the following:

- Collects backup files that have passed verification.

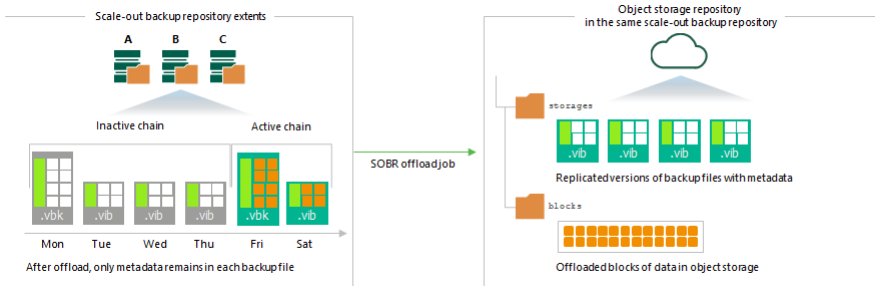
  Such verified backup files are collected from all the performance extents added to a scale-out backup repository.

- Extracts data blocks from these files and moves these blocks to the capacity extent leaving the backup files only with metadata (i.e. free of data blocks).

  Such backup files (without data blocks) will remain on the performance extents and will also be replicated to the capacity extent.

Having a copy of the dummy files on your performance extents allows you to:

- Download back to the performance extents the data that has been moved to the capacity extent, as described in Downloading Data from Capacity Tier.
- Restore data back to production servers, to Azure or Amazon EC2 cloud platforms or to your virtual machine, as described in Restore from Capacity Tier.



In this example, there is a scale-out backup repository and an object storage repository in the same scale-out backup repository.

The performance extent A has an inactive backup chain consisting of one VBK file and three VIB files, that is, four restore points in total. Each of these files comprises metadata and the actual blocks of data. During the offload session, Veeam Backup & Replication will collect the actual blocks of data from all the backup files (VBK and VIB) and offload these blocks to the object storage repository.

Each offloaded block might be of different size, which is defined during configuring storage optimization. The offloaded blocks are placed to the blocks directory in your capacity extent.

Backup files with metadata (or dummy files) will be replicated to the capacity extent and will be placed to the storage directory.

Such approach will be applied to all inactive backup chains that satisfy validation criteria.

After offload is complete, the new **Object Storage** node appears in the **Home** view, under the **Backups** node and shows backups that have been moved to the capacity extent.

> **NOTE**
>
> The offload is not performed during prohibited hours specified in the scale-out backup repository backup window configuration. You can configure the backup window at the Add Capacity Tier step of the New Scale-out Backup Repository wizard.

Offload Session Statistics

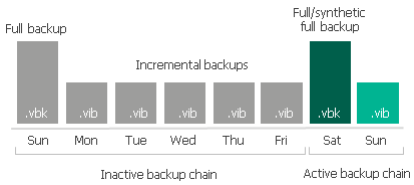
The offload session results are saved to the configuration database and available for viewing, as described in Viewing Offload Job Session Results.

Backup Chain Detection

Veeam Backup & Replication will transfer to the capacity extent only those restore points that belong to inactive backup chains. To ensure a backup chain is inactive, Veeam Backup & Replication verifies its state. This does not apply to the copy policy: all newly created restore points are copied immediately.

When a backup job runs for the first time, Veeam Backup & Replication creates an initial full backup file. It contains complete information about the VMs that are being backed up. With each subsequent backup job session, new incremental backup files are created. They contain only the changes that have occurred since the last backup session.

For forward incremental backup method, the active backup chain is the one that has not yet been sealed with a new full backup file.



To transform an active backup chain into inactive, a new active full (or synthetic full) backup file must be created for this chain. This can be done either manually, as described in Performing Active Full Backup. Else, you can configure a schedule according to which new active or synthetic full backups will be created automatically, as described in Active Full Backup and Synthetic Full Backup.

Once a new full backup file is created and the offload session is being executed, Veeam Backup & Replication collects all the restore points (full and incremental) that were created prior to the latest active full, verifies that they belong to an inactive chain, and prepares them to be moved to the capacity extent. This process is called detect. For more information, see Moving Backups to Capacity Tier.

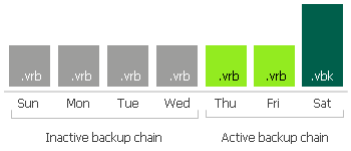
The same applies to backup chains created by backup copy jobs. Veeam Backup & Replication will transfer to the capacity extent only those restore points that belong to an inactive backup copy chain. Note that if you enable backup copy GFS, Veeam Backup & Replication implements the forward incremental retention policy. If you disable backup copy GFS, Veeam Backup & Replication implements the forever forward incremental retention policy.

Note that Veeam Backup & Replication will not transfer to the capacity tier the corrupted restore points and the files dependent on those. For more information on the corrupted restore points, see Health Check for Backup Files.

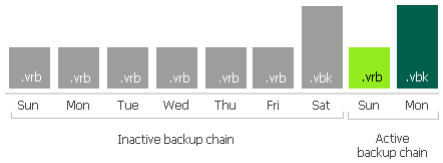
The same applies to the backup chains created with the reverse-incremental method. In this case, all the .VRB files starting from the third restore point will be considered inactive automatically. Thus, you do not need to create an active full (nor synthetic full) backup manually unless you want to offload all the restore points including the most recent .VBK file and the first two .VRB files.

Consider the following examples:

- Four VRB files are inactive and can be offloaded:



- Six VRB files and a VBK file belong to an inactive chain and can be offloaded:



NOTE

Mind that a full backup file and the first two incremental backup files (that is, two .VRB files that immediately follow the most recent .VBK file) will not be offloaded until another full backup file is created successfully.

The structure of the backup chains can be different. That depends on whether your backups were created using the per-machine method (for more information, see Per-Machine Backup Files) or as a single storage, with all VMs placed into a single file. The type of the backup chain structure does not affect the offload process.

For more information on how Veeam Backup & Replication creates and manages backup chains, see Backup Chain.

Manually Moving Backups to Capacity Tier

The **Move to capacity tier** option allows you to manually offload selected backup files to the capacity extent.

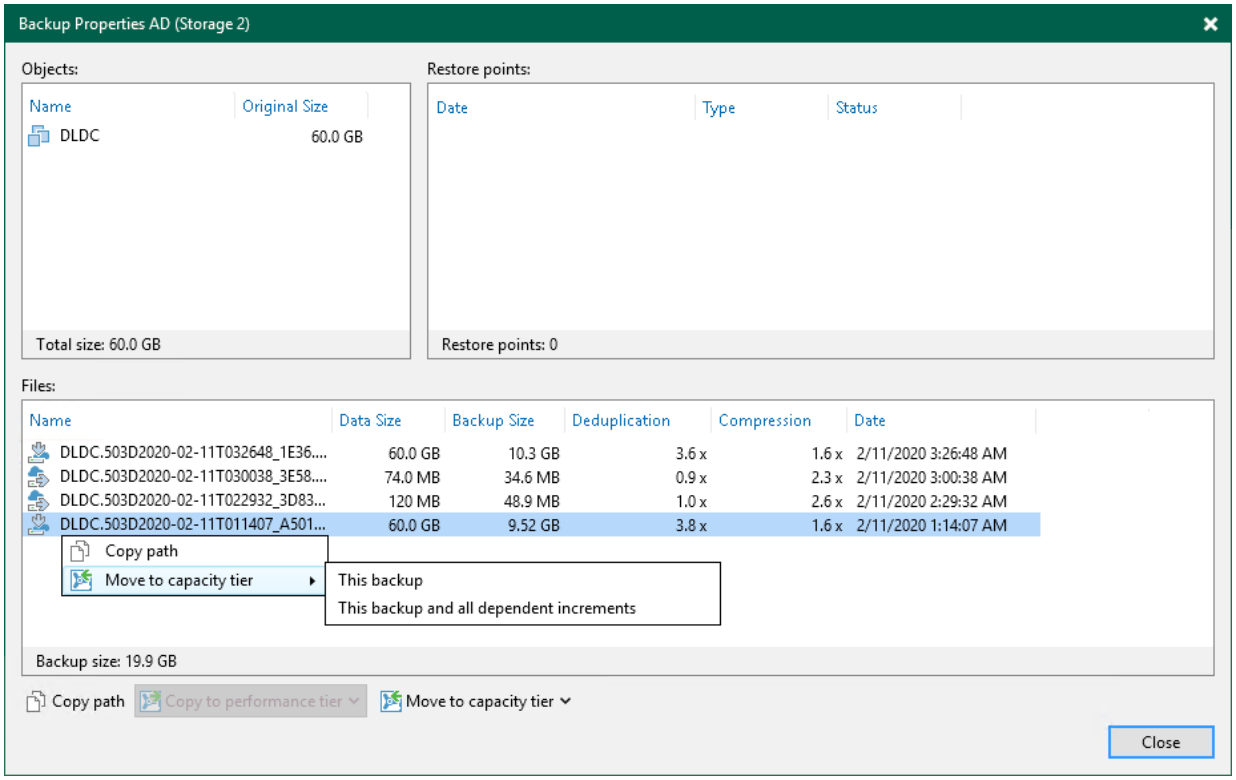
Consider that backup files you want to offload must belong to an inactive backup chain. For more information, see Backup Chain Detection.

To move backup data to capacity tier, do the following:

1. Open the **Home** view.
2. In the inventory pane, select the **Backups > Disk** node.
3. In the working area, right-click a backup job and select **Properties**.
4. In the **Properties** window, right-click a backup file that you want to offload and select **Move to capacity tier**.

   Alternatively, you can use the **Move to capacity tier** control at the bottom.
5. Select either of the following options:
   - For VIB*/VBK* backup files:
     - **This backup and all dependent increments** — to offload the selected backup along with its increments or to remove the copied blocks of data from the selected backup and its increments.
   - For VBK backup files:
     - **This backup** — to offload the selected full backup only or to remove the copied blocks of data from such a full backup.

- **This backup and all dependent increments** — to offload the selected full backup along with its increments or to remove the copied blocks of data from the selected backup and its increments.



Downloading Data from Capacity Tier

In Veeam Backup & Replication, you can manually download offloaded backups back to the performance extents.

You can download one backup at a time using the **Copy to Performance Tier** option, or get all offloaded backups in bulk using the **Download** feature. For more information, see below Downloading Single Backup Chain and Downloading All Backups, respectively.

For more information on how data is downloaded, see How Downloading from Capacity Tier Works.

Downloading Single Backup Chain

To download previously offloaded backup data back to the performance extents, one backup at a time, do the following:

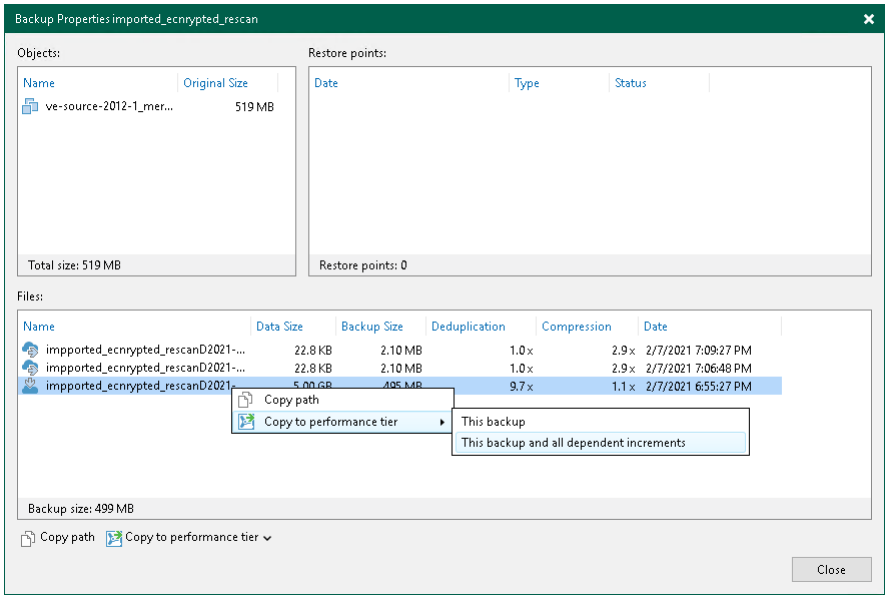
1. Open the **Home** view.
2. In the inventory pane, select the **Backups > Object Storage** node.
3. In the working area, right-click a backup job and select **Properties**.
4. In the **Properties** window, right-click an offloaded backup file and select **Copy to performance tier**.

   Alternatively, you can use the **Copy to performance tier** control at the bottom.
5. Select either of the following options:
   - For VIB/*VBK* backup files:

- **This backup and all dependent increments** — to copy the selected backup along with its increments.

- For VBK backup files:
  - **This backup** — to copy the full backup only.
  - **This backup and all dependent increments** — to copy the selected backup along with its increments.

> **NOTE**
>
> To remove copied blocks from the performance extents, use the **Move to capacity tier** option, as described in Moving Inactive Backup Chains to Capacity Tier.

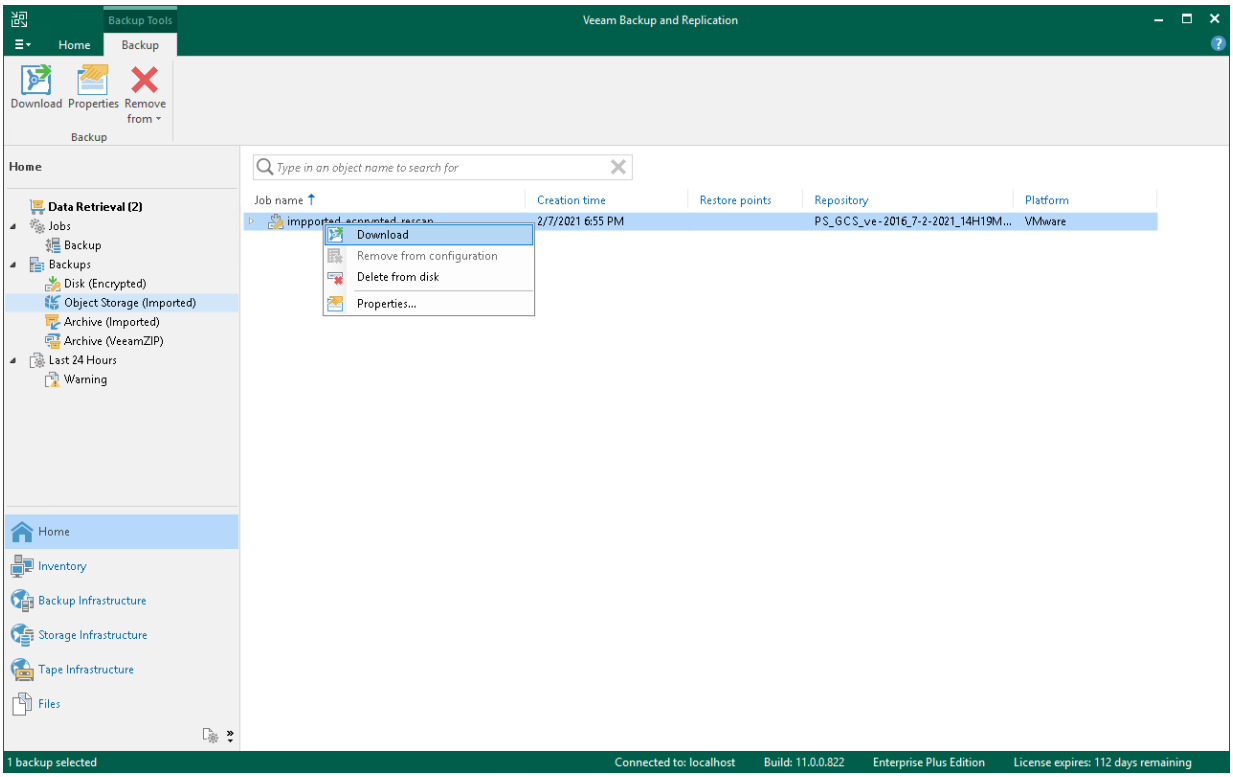


Downloading All Backups

To download all offloaded backups back to the performance extents, do the following:

1. Open the **Home** view.
2. In the inventory pane, select the **Backups > Object Storage** node.
3. On the toolbar, click **Download** or in the working area, right-click a backup job and select **Download**.
   - If the latest backup chain is already present on your performance extent, you will be asked if you wish to download all other backup files. Click **Yes** if you wish to do so or **No** to cancel the download.
   - If the latest backup chain is not on your performance extent yet, you will be asked which backup files you would like to download. Click **All Backups** to download all backup files or **Latest Only** to download just the latest backup chain.

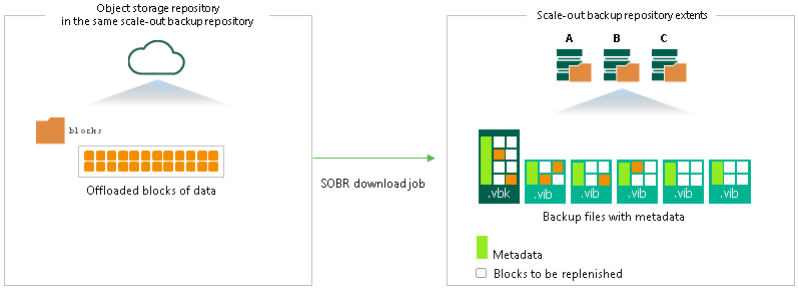
> **NOTE**
>
> To remove downloaded data from the performance extents, use the **Move to capacity tier** option, as described in Moving Inactive Backup Chains to Capacity Tier.



How Downloading from Capacity Tier Works

To download data from the capacity extent back to the performance extents, Veeam Backup & Replication uses the SOBR Download job.

The SOBR Download job is triggered right after you select the **Copy to performance tier** option; it collects offloaded blocks of data from the capacity extent and copies them back to the performance extents. For more information, see Downloading from Capacity Tier.



Consider the following:

- Before copying data blocks, Veeam Backup & Replication verifies whether any of such blocks exist on any of the performance extents of your scale-out backup repository. If found, Veeam Backup & Replication will reuse the existing blocks instead of downloading the exact same data from the capacity extent.

> **NOTE**
>
> If the password in the backup job has been changed, in some cases Veeam Backup & Replication may not be able to reuse the performance tier data in the course of the download.

- If a performance extent is unable to accommodate data being copied due to lack of free storage space, Veeam Backup & Replication will find another extent in the associated scale-out backup repository that has sufficient storage capacity to receive the data. If your scale-out backup repository has no performance extents other than the one running out of space, the copy will not be possible.
- If you have removed any of the performance extents from a scale-out backup repository without evacuating backup files with metadata, the copy will not be possible until the files with metadata are downloaded back to the performance extents in the course of the rescan process. For more information on the rescan process, see Rescanning Scale-Out Repositories.

  Backup files with metadata are created as described in Moving Inactive Backup Chains to Capacity Tier.
- The SOBR download session results are saved to the configuration database and available for viewing, as described in Viewing Download Job Session Results.

Viewing Capacity Tier Sessions Statistics

Continue with this section to learn more about:

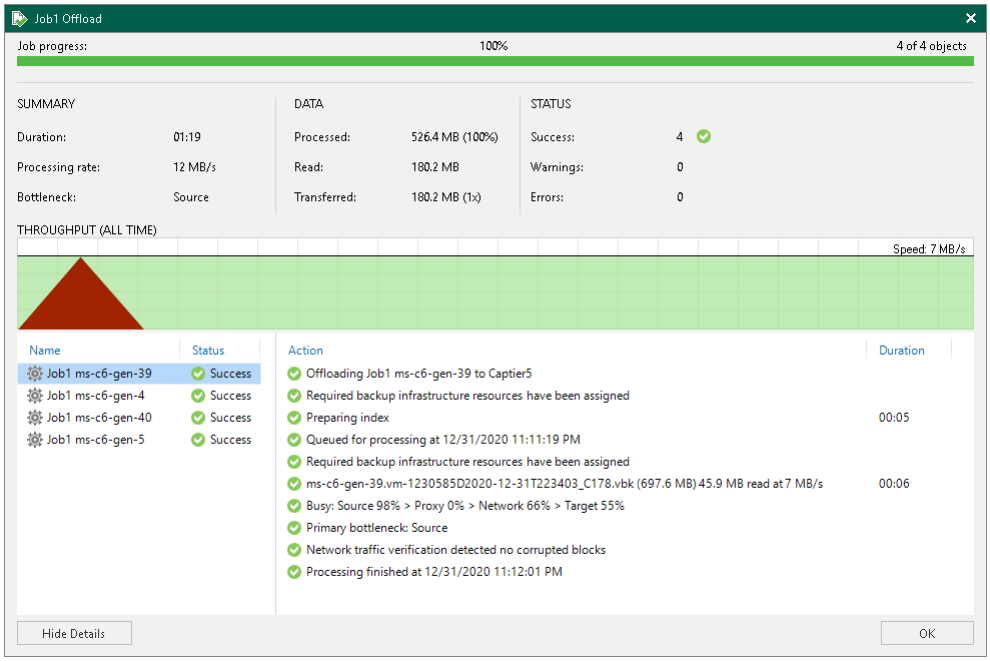
- Viewing Offload Session Results
- Viewing Copy Session Results
- Viewing Download Session Results

Viewing Offload Session Results

To review offload session results, do the following:

1. Open the **History** view.
2. In the inventory pane, select the **Storage management** node.
3. In the working area, right-click an offload session and select **Statistics**.

For more information about moving inactive backup chains to the capacity tier, see Moving Inactive Backup Chains to Capacity Tier.



Veeam Backup & Replication displays offload job session statistics for the following counters:

- The **Job progress** bar shows percentage of the offload session completion.
- The **Summary** box shows general information about the offload session:
  - **Duration** – duration of the offload session.
  - **Processing rate** – average speed of VM data processing. This counter is a ratio between the amount of data that has actually been read and the offload session duration.
  - **Bottleneck** – bottleneck in the data transmission process. To learn more about bottlenecks, see Performance Bottlenecks.
- The **Data** box shows information about processed data:
  - **Processed** – total size of all VM disks processed by the offload session.
  - **Read** – the amount of data read from the extents.
  - **Transferred** – the amount of data transferred from the extents to object storage.
- The **Status** box shows information about the job results. This box informs how many tasks have completed with the *Success*, *Warning* and *Error* statuses (1 task per 1 VM).
- The pane in the lower-left corner shows a list of objects processed by the offload session.
- The pane in the lower-right corner shows a list of operations performed during the session. To see a list of operations for a specific object, click the object in the pane on the left. To see a list of operations for the whole offload session, click anywhere on the blank area in the left pane.

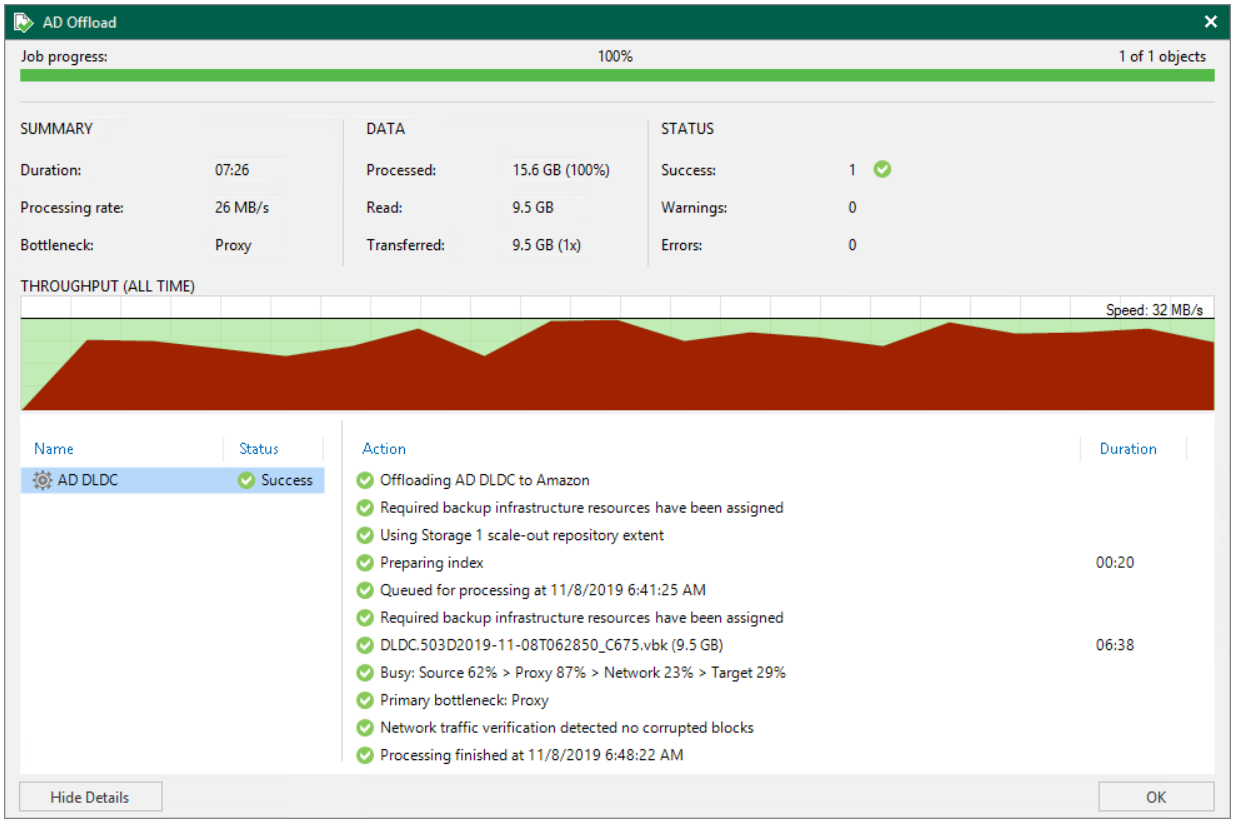
Viewing Copy Session Results

To review copy session results, do the following:

1. Open the **History** view.

2. In the inventory pane, select the **Storage management** node.
3. In the working area, right-click a copy session and select **Statistics**.

For more information about moving inactive backup chains to the capacity tier, see Moving Inactive Backup Chains to Capacity Tier.



Veeam Backup & Replication displays offload job session statistics for the following counters:

- The **Job progress** bar shows percentage of the copy session completion.
- The **Summary** box shows general information about the copy session:
  - **Duration** – duration of the copy session.
  - **Processing rate** – average speed of VM data processing. This counter is a ratio between the amount of data that has actually been read and the offload session duration.
  - **Bottleneck** – bottleneck in the data transmission process. To learn more about bottlenecks, see Performance Bottlenecks.
- The **Data** box shows information about processed data:
  - **Processed** – total size of all VM disks processed by the copy session.
  - **Read** – the amount of data read from the extents.
  - **Transferred** – the amount of data transferred from the extents to object storage.
- The **Status** box shows information about the copy results. This box informs how many tasks have been completed with the *Success*, *Warning* and *Error* statuses (1 task per 1 VM).
- The pane in the lower-left corner shows a list of objects processed by the copy session.

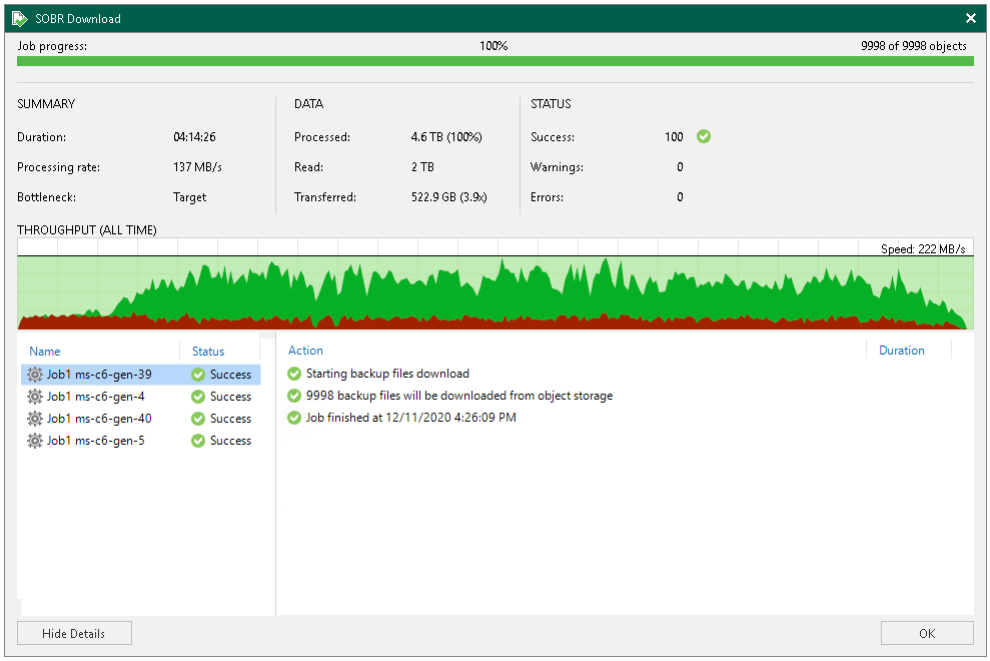
- The pane in the lower-right corner shows a list of operations performed during the session. To see a list of operations for a specific object, click the object in the pane on the left. To see a list of operations for the whole copy session, click anywhere on the blank area in the left pane.

Viewing Download Session Results

To review **SOBR Download** job session results, do the following:

1. Open the **History** view.
2. In the inventory pane, select the **Storage management** node.
3. In the working area, right-click a **SOBR Download** session and select **Statistics**.

For more information about the SOBR Download session, see How Downloading Data from Capacity Tier Works.



Veeam Backup & Replication displays **SOBR Download** job session statistics for the following counters:

- The **Job progress** bar shows percentage of the job completion.
- The **Summary** box shows general information about the job session:
  - **Duration** – duration of the job session.
  - **Processing rate** – average speed of data processing. This counter is a ratio between the amount of data that has actually been read and the job session duration.
  - **Bottleneck** – bottleneck in the data transmission process. To learn more about bottlenecks, see Performance Bottlenecks.
- The **Data** box shows information about processed data:
  - **Processed** – total size of data blocks being downloaded from object storage repository plus blocks (if any) being taken from the extents of your scale-out backup repository.
  - **Read** – the amount of data read from both the object storage repository and extents of your scale-out backup repository.
  - **Transferred** – the amount of data downloaded from object storage.

- The **Status** box shows information about the job results. This box informs how many tasks have completed with the *Success*, *Warning* and *Error* statuses.
- The pane in the lower-left corner shows a list of objects processed by the job.
- The pane in the lower-right corner shows a list of operations performed during the session. To see a list of operations for a specific object, click the object in the pane on the left. To see a list of operations for the whole job session, click anywhere on the blank area in the left pane.

Migrating Data To Another Capacity Tier

You can migrate your data between object storage repositories. After that, you can add a new object storage repository as a capacity extent instead of the previous one.

The migration option may be useful if you want, for instance, switch your provider.

Such a migration can be done under the following scenarios:

- Migrating data between different cloud providers

  Is used to migrate data from one object storage provider to another. For example, you may have been using an Amazon S3 object storage repository for quite some time and decided to change it to Azure Blob storage or any other storage provider.

- Migrating data between different buckets

  Is used to migrate data between different buckets of the same cloud provider.

> **IMPORTANT**
>
> Migrating data between object storage repositories is available only at the capacity tier. Archive tier does not support such scenarios.

Migrating Data Between Different Cloud Providers

To migrate data located in an object storage repository between different cloud providers, do the following:

1. Download data from an object storage repository back to the performance extents, as described in Downloading Data from Capacity Tier.
2. Add a new object storage repository to the Veeam Backup & Replication environment, as described in Adding Object Storage Repositories.
3. Change an object storage repository to a new one that you have created at the previous step, as described in Add Capacity Tier.
4. Copy or move your data to a new object storage repository, as described in Moving to Capacity Tier.

Migrating Data Between Different Buckets

To migrate data located in an object storage repository between different buckets of the same cloud provider, do the following:

1. Use any available 3rd-party tool to copy cloud data to a different bucket of your cloud storage.
2. Add a new object storage repository to the Veeam Backup & Replication environment, as described in Adding Object Storage Repositories.

3. Change the capacity extent to a new one that you have created at the previous step, as described in Add Capacity Tier.
4. Synchronize your data, as described in Synchronizing Capacity Tier Data.

Keep in mind that you can copy cloud data only to a bucket without the *Versioning* feature enabled. If you need to migrate your data from a non-immutable bucket to an immutable one, use one of the following ways:

- Perform the scenario described in Migrating Data Between Different Cloud Providers.
- If your source bucket had the *Versioning* feature enabled, copy your data to a bucket without the *Versioning* feature, then manually enable the *Versioning* and the *Object Lock* features. However, instead of migrating data, you can enable the *Versioning* and the *Object Lock* features on the existing bucket, and proceed to use this bucket as an immutable one. This only applies to object storage that support the *Versioning* and the *Immutability* features. Before you enable the *Object Lock* feature on the existing bucket, make sure your vendor allows enabling the *Object Lock* feature after the bucket is already created.

Use command line interface (CLI) to copy from an immutable bucket to a non-versioned one.

Indexes

To reduce the number of cost-based operations incurred by your cloud storage provider and to decrease the amount of traffic being sent over the network when moving or copying data to object storage, Veeam Backup & Replication uses indexes.

Index behavior is as follows:

- Indexes are created (or updated) during each offload or copy session and consist of hash values of blocks that are being transferred to the capacity tier. These hashes are retrieved from meta information of your backup files (VBK, VIB, or VRB).
- Indexes are stored in the *ArchiveIndex* directory that is located on the Performance Tier.

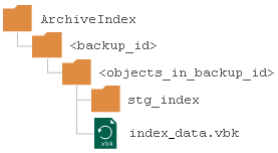
  On each subsequent offload/copy session, Veeam Backup & Replication reuses these indexes to verify whether new blocks that are about to be transferred to the capacity tier have not been offloaded earlier. Verification is done by comparing existing indexes hashes with that of a block being transferred.
- If backups are created using the per-machine method, indexes are built per backup chain and cannot have any cross references to any other backup chains. If all VM data is backed up to a single storage, indexes are built for the whole backup.
- Indexes are updated every time a backup chain is modified.

  For example, some data might have been removed due to the retention policy threshold, or you may have removed it manually. Both scenarios will modify your indexes upon the next successful offload or copy session to maintain consistency.
- Corrupted indexes can be rebuilt by using the **Rescan** feature, as described in Rescanning Scale-Out Repositories.

  Once an index is rebuilt, Veeam Backup & Replication will have to wait for 24 hours before it can offload any data again. This is necessary to comply with the eventual consistency model of S3 compatible object storage repositories.

ArchiveIndex Directory Structure

When Veeam Backup & Replication creates indexes, it also creates and maintains the following directory structure on each extent.



Directory	Description
ArchiveIndex	The root directory for keeping indexes.
<backup_id>	Contains objects in a backup file.
<objects_in_backup_id>	An identifier of an object in a backup file. • If a backup was created using the per-machine method, then each VM will be placed to its own directory. • If a backup was created as a single storage, then all the VMs will be placed to a unique directory.
stg_index	Contains indexes of offloaded backup files (VBK, VIB, or VRB).
index_data.vbk	Contains meta information on hash values stored in index files.

Retention Policy

Retention policy defines the number of restore points to keep on your performance extents and is configured at the Specify Backup Storage Settings step of the **New Backup Job** wizard.

You can manage retention policies to remove obsolete restore points both from the performance extents and the capacity tier.

The restore points that fall under the retention policy will be removed both from the performance and capacity extents in the following manner:

- An earliest restore point will be removed from the backup chain on the associated extent.
- Data blocks that correspond to the restore point that is being removed will be purged from the capacity extent upon the next offload or copy session.

  Make sure that the capacity extent has not been put into the maintenance mode, as this mode prevents synchronization of the state of the backup chain in the performance tier with that of the capacity tier.

  For more information about the move and copy offload sessions, see Moving Inactive Backup Chains to Capacity Tier and Copying Backups to Capacity Tier respectively.

- Immutable blocks of data are removed after the immutability period is over.

  When a retention policy encounters immutable copied/moved blocks of data, it removes such blocks from the associated backup files on the extents only and then updates indexes, informing Veeam Backup & Replication that these blocks no longer exist and must be removed from the capacity tier once mutable.

  For more information about immutability, see Immutability.

- Associated indexes will be updated for consistency purposes.

  For more information, see Indexes.

Excluding Capacity Extent from Scale-Out Repositories

You can exclude a capacity extent from the scale-out backup repository scope, for example, if you no longer want to use any third party services to store your data.

Consider that after you exclude an object storage repository that is being used as a capacity extent and is storing offloaded backup data, Veeam Backup & Replication automatically puts the excluded object storage repository into the maintenance mode. Once a repository is in the maintenance mode, you will not be able to restore your data from it. To switch back to normal, you will have to re-add that repository as a capacity extent and synchronize existing backup chains with your performance extents. After the synchronization is complete, the existing backups will become available as Imported. For more information, see Synchronizing Capacity Tier Data.

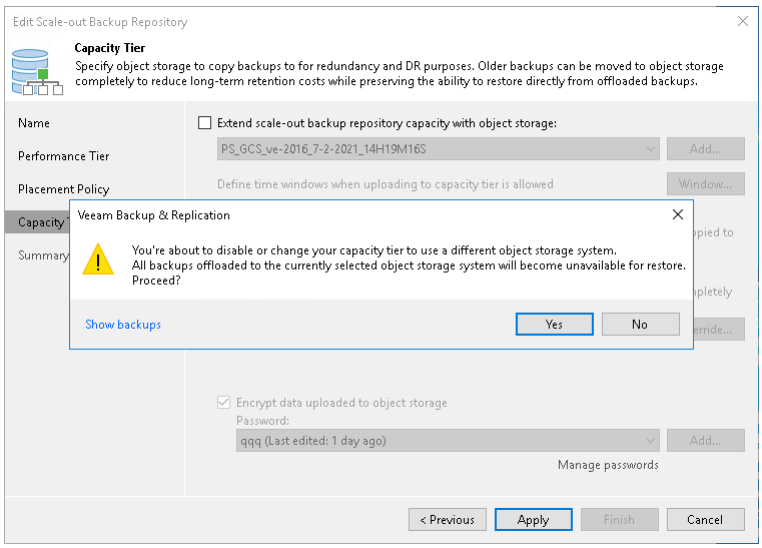
> **NOTE**
>
> If you have an archive extent in the same scale-out backup repository, it will also switch to the maintenance mode when you exclude a capacity extent.

To exclude a capacity extent from the scale-out backup repository scope, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, click **Scale-out Repositories**.
3. In the working area, select a scale-out backup repository and click **Edit Scale-out Repository** on the ribbon, or right-click a scale-out backup repository and select **Properties**.
4. Move to the **Capacity tier** step of the wizard.

5. Clear the **Extend scale-out backup repository capacity with object storage** check box.

   You will be asked to confirm the action in the pop-up dialog. After that, the object storage repository will be immediately put into the maintenance mode.



Restore from Capacity Tier

You can restore your data directly from the capacity tier back to production servers or to Azure or Amazon EC2 cloud platforms. Capacity tier data recovery does not differ from that of a standard backup data recovery and can be performed by using any of the following methods:

- Instant Recovery to VMware vSphere
- Instant Recovery to Microsoft Hyper-V
- Entire VM Recovery
- VM Files Restore
- Virtual Disks Restore
- Guest OS File Recovery
- Instant Disk Recovery
- Application Items Restore
- Exporting Backups

Data recovery can also be done directly to Amazon EC2 or Microsoft Azure, as described in the Restore to Amazon EC2 and Restore to Microsoft Azure sections, respectively.

Restore Scenarios

This section explains possible restore scenarios from the capacity tier.

Unavailability of Backup Files

If one of the offloaded backup files becomes unavailable on any of the extents in a scale-out backup repository, you can restore it by doing the following:

1. Rescan a scale-out backup repository.

   For more information on how to rescan a scale-out backup repository, see Rescanning Scale-Out Repositories.

2. Copy data from the capacity tier to the performance tier.

   For more information on how to copy data, see Copying to Performance Tier.

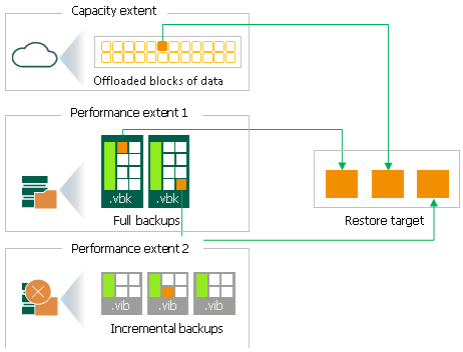
NOTE

The performance or capacity extents must not be in the maintenance mode.

Unavailability of Extents

A performance extent in a scale-out backup repository may become unavailable or be in the maintenance mode. To restore data in such case, you can use any method described in Data Recovery.

For example, you are restoring a virtual machine consisting of three data blocks, of which two blocks reside on the *Extent 1* and another required block is stored on the *Extent 2* which is unavailable. In such scenario, Veeam Backup & Replication gets two blocks from the *Extent 1* and another required block from the capacity tier.

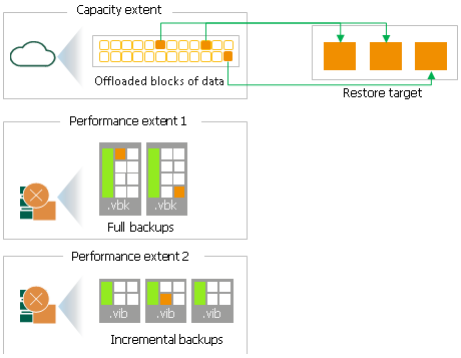


Unavailability of Scale-Out Backup Repository

To restore data if the entire scale-out backup repository is unavailable, you can use any method described in Data Recovery.

If the entire scale-out backup repository becomes unavailable, Veeam Backup & Replication restores data from the capacity tier only.

For example, both performance extents that store required backup files to restore a virtual machine are not available. In such a scenario, Veeam Backup & Replication restores data from the capacity tier only.



Unavailability of Backup Server

To get access to backups in the capacity tier in case the entire configuration of the backup server becomes corrupted and your scale-out backup repositories are no longer available, you can:

- Restore the configuration of the backup server from the configuration backup, as described in Configuration Backup and Restore.
- Import backups from the capacity tier, as described in Importing Backups.

Archive Tier

Archive tier is an additional tier of storage that can be attached to a scale-out backup repository. Data from the capacity extent of the scale-out backup repository can be transported to the archive tier for archive storage.

Storing archived data in the archive tier is cheaper than in the capacity tier. However, restoring data from the archive tier is longer and more expensive compared to the capacity tier. Data must be prepared for restore from the archive tier.

This feature is most useful in the following cases:

- You have a lot of rarely (no more than once a quarter) accessed data that has to be stored in an archive.
- You want to save costs and/or space on storing archived data.

The following types of backup files are suitable for archive storage:

- Backup files with GFS flags assigned
- VeeamZIP backup files
- Exported backup files
- Orphaned backups

The archive tier consists of a single archive extent. One of the following cloud-based object storage repositories providing "cold" data storage can be added as the archive extent:

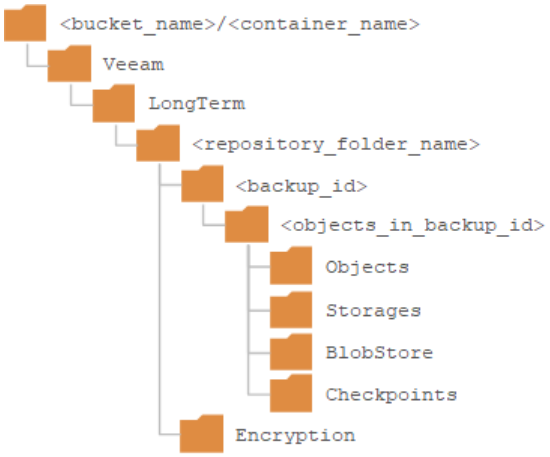
- Amazon S3 Glacier
- Microsoft Azure Archive Storage

Before an object storage repository can be configured as the archive extent, it must be added to Veeam Backup & Replication backup infrastructure. For more information, see Adding Amazon S3 Glacier Storage or Adding Azure Archive Storage.

You can add the archive extent to your scale-out backup repository and configure its settings on the Add Archive Tier step of the **New Scale-out Backup Repository** wizard.

Archive Extent Structure

When data is being transferred to the archive extent, Veeam Backup & Replication creates and maintains the following structure of directories:



Directory	Description	Misc
<bucket_name> or <container_name>	A bucket or container name. Buckets and containers must be created in advance.	N/A
Veeam/LongTerm/	Standard folders created by Veeam Backup & Replication.	
<repository_folder_name>	A repository folder that you create when adding a new capacity extent.	
<backup_id>	Contains objects in a backup.	These folders are automatically removed during data removal.
<objects_in_backup_id>	An identifier of an object in a backup.	
Objects	Contains meta information and other auxiliary data.	
Storages	Contains a replicated version of offloaded backup files with metadata.	
BlobStore	Contains data blobs created by the archiving session, as described in Archive Tier Data Transfer.	

Directory	Description	Misc
Checkpoints	Contains meta information about the state of archived backup chains. Such meta information is updated upon each successful archiving session.	
Encryption	Contains required information to work with encrypted backups.	N/A

Limitations for Archive Tier

The archive tier has the following limitations:

- Each type of repository that can be added as an archive extent is compatible only with corresponding type of repository configured as a capacity extent:

  a. To add Amazon S3 Glacier as an archive extent, select Amazon S3 repository in the capacity tier settings. If you do not have an Amazon S3 repository as a capacity extent of your scale-out backup repository, you must first configure it.

  b. To add Microsoft Azure Archive storage as an archive extent, select Microsoft Azure Blob repository in the capacity tier settings. If you do not have a Microsoft Azure Blob repository as a capacity extent of your scale-out backup repository, you must first configure it.

  For more information on configuring capacity extents, see Add Capacity Tier.

- Migrating data to another archive tier is not supported.
- You cannot transport backup files from performance tier directly to the archive tier. Transportation to the archive tier is available only from the capacity tier.
- Microsoft Azure Archive storage does not support immutability.
- Imported backups cannot be offloaded to archive tier.
- Incremental backup files cannot be stored in the archive tier.

Immutability for Archive Tier

> **NOTE**
>
> Currently, the immutability feature is available only for Amazon S3 Glacier object storage repository. Microsoft Azure Archive storage does not support immutability.

Veeam Backup & Replication allows you to prohibit deletion of data from the archive extent by making that data temporarily immutable. It is done for increased security: immutability protects your data from loss as a result of attacks, malware activity or any other injurious actions.

To make the data immutable, Veeam Backup & Replication uses the Object Lock technology provided by Amazon and some S3-compatible providers. Once imposed, the object lock prohibits deletion of data from the archive extent until the immutability expiration date comes. Keep in mind that in order to use immutability, you need to enable the *Object Lock* and *Versioning* features on your S3 bucket at the time you create the bucket. For more information, see Preparing to Use Immutability.

You can enable the immutability feature when adding or editing an Amazon S3 Glacier object storage repository as an archive extent, at the Bucket step of the Adding Amazon S3 Glacier Storage wizard.

When you enable immutability for the archive tier, keep in mind that only the settings of the Amazon S3 Glacier repository will be taken into account. The settings of the capacity tier repository and of the original data blocks will be ignored.

For Amazon S3 Glacier, all the types of files that are suitable for archive storage can be made immutable:

- Backup files with GFS flags assigned: in case GFS retention is extended in the backup job or backup copy job settings, the immutability period for existing backup files will be prolonged at the end of the archiving session. For more information about GFS retention policy, see Long-Term Retention Policy (GFS).
- VeeamZIP backup files with specified retention (deletion date). For more information, see Creating VeeamZIP Backups.
- Exported backup files with specified retention (deletion date). For more information, see Exporting Backups.

The immutability period of a backup file will be equal to its retention period at the moment of archiving. If the retention period is not specified for VeeamZIP backup files or exported backup files, such files will not be made immutable.

Managing Archive Tier

You can manage your archive tier and the archived data in the following ways:

- Move outdated backups from the capacity extent to the archive extent.
- Delete outdated backup files.
- Exclude the archive extent from the scale-out backup repository scope.

Archive Tier Data Transfer

Outdated backup files from the capacity tier can be moved to the archive tier. After that, the transported files are deleted from the capacity tier. They can stay in the performance tier or be deleted from there depending on the data transfer policy of the scale-out backup repository:

- In case the copy policy is selected (the outdated backup files are copied from the performance tier to the capacity tier), the original file stays in the performance tier.
- In case the move policy is selected (the outdated backup files are transferred from the performance tier to the capacity tier), the original file disappears from the performance tier.

Data transfer policy	Performance tier	Capacity tier	Archive tier
Copy	Original file	(nothing)	Copy of the file
Move	(nothing)	(nothing)	The transferred file

Proxy Appliances

Data transfer from the capacity extent to the archive extent is done through proxy appliances — temporary virtual machines. The template for all the proxy appliances is set up at the **Proxy Appliance** step of the Adding Amazon S3 Glacier or Adding Azure Archive Storage wizard.

You can set up the size of the virtual machine and cloud resources where the proxy appliance will be created. You can specify the proxy appliance settings manually, or use the default settings.

After the archiving job is finished, all the proxy appliances are automatically deleted. If the job ends prematurely, the proxy appliances will be deleted as well. Also, any proxy appliance can be deleted if there are no more tasks for it.

Archiving Job

The process of moving backup files to the archive tier is called an archiving job. The archiving job takes place in a separate job session. The archiving job launches periodically, under the same default schedule as the offload job from the performance tier to the capacity tier.

New backup files are offloaded to the archive tier and outdated ones are cleaned up during the same archiving session. You can archive backup files if they meet all the following conditions:

- The type of backup file is appropriate. For details, see Archive Tier.
- Backup files are offloaded to the capacity tier.
- Backup files have not been archived yet.
- Backup files belong to inactive backup chains.
- Backup files have been created N days ago according to **Archive backup files older than N days** setting of the Add Archive Tier step of the wizard.

  > **NOTE**
  >
  > If you set the value in the **Archive backup files older than N days** field to 0, the backup files will be moved to the archive tier during the nearest archiving session, and will no longer participate in the capacity tier data deduplication.

- If you have selected the **Archive backups only if the remaining retention time is above minimal storage period** check box of the Add Archive Tier step of the wizard:

  a. for Amazon S3 Glacier storage: only backup files with retention period no less than 90 days.

  b. for Microsoft Azure Archive storage: only backup files with retention period no less than 180 days.

Backup files that meet all the following conditions will be cleaned up:

- Backup files that do not have corresponding original file on the performance tier
- Backup files with expired or unspecified immutability period. For more information, see Immutability for Archive Tier.

After the archiving job is completed, the transported files are deleted from the capacity tier. It happens during the next offload to capacity tier session.

Cleanup Tasks

Cleanup tasks exist within archiving jobs. They are launched if there are only outdated backup files to delete during a certain job session, and no new backup files are to be offloaded to the archive tier.

Cleanup tasks do not require proxy appliances.

Excluding Archive Extent from Scale-Out Backup Repository

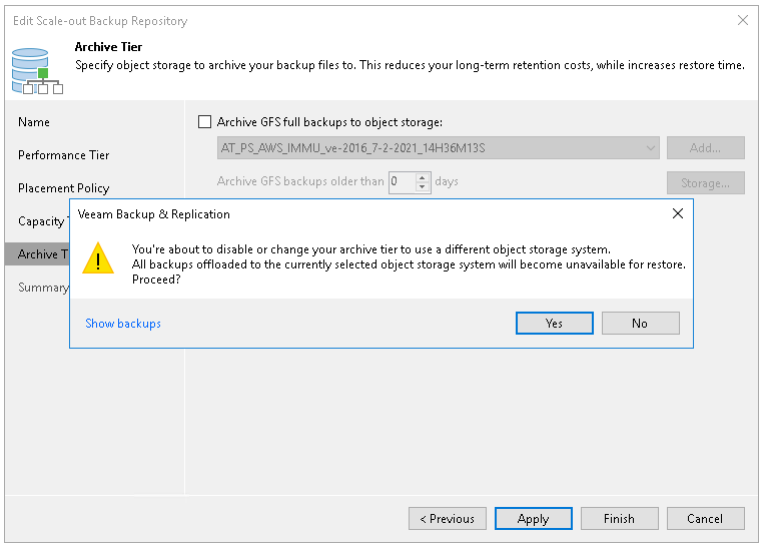
You can exclude an archive extent from the scale-out backup repository scope, for example, if you no longer want to use any third party services to store your data.

Consider that after you exclude an object storage repository that is being used as an archive extent and is storing archived backup data, Veeam Backup & Replication automatically puts the excluded object storage repository into the maintenance mode. Once a repository is in the maintenance mode, you will not be able to restore your data from it. To switch back to normal, you will have to re-add that repository as an archive extent and synchronize existing backup chains with data in this scale-out backup repository. After the synchronization is complete, the existing backups will become available as Imported.

To exclude an archive extent from the scale-out backup repository scope, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, click **Scale-out Repositories**.
3. In the working area, select a scale-out backup repository and click **Edit Scale-out Repository** on the ribbon, or right-click a scale-out backup repository and select **Properties**.
4. Move to the **Archive tier** step of the wizard.
5. Clear the **Archive GFS full backups to object storage** check box.

   You will be asked to confirm the action in the pop-up dialog. After that, the object storage repository will be immediately put into the maintenance mode.



Restore from Archive Tier

You can restore data from the archive storage. To do so, you have to receive a temporary access to the data first. Thus, the process of restore from the archive tier consists of two consecutive parts:

1. Data retrieval

2. Restoring retrieved data

Data Retrieval

Data retrieval is the process of receiving temporary access to archived data, so that it can be restored.

The process of retrieving data from the archive tier takes course in a separate retrieval job session. It is completed when a restore point is available for reading and restore. For information on how to launch the retrieval job, see Retrieving Backup Files.

When the retrieval job is complete, the retrieved data will be available for a certain period of time, during which you can restore it to the capacity tier. However, you can extend the availability period.

> **NOTE**
>
> If you launch restore job when retrieval job is not over yet, the restore job will be pending until the retrieval job is complete.

Retrieval cost varies depending on the desired speed of the process and on the targeted period of the data accessibility. See below the options for different vendors.

Retrieval Options for Amazon S3 Glacier

Amazon provides the following options for data retrieval. The indicated time is approximate. For more information, see this Amazon article.

- Expedited: the most expensive method. Retrieval takes 1-5 minutes.

  Keep in mind that this option will be unavailable for data stored in Glacier Deep Archive class.

  A backup file can be stored in Glacier Deep Archive if the following conditions are met:

  a. The **Use Glacier Deep Archive for backups with the retention policy over 180 days** check box was selected on the Specify Object Storage Settings step of the Adding Amazon S3 Glacier wizard.

  b. The backup file retention period is 180 days or more.

  c. If the **Store archived backups as standalone fulls check** box was not selected on the Add Archive Tier step of the **New Scale-Out Backup Repository** wizard, that is, the reuse of the data blocks was allowed, all the backup files in the chain after the one stored in Deep Archive will be automatically considered as belonging to Deep Archive as well.

- Standard: retrieval time 3-5 hours for Amazon S3 Glacier and within 12 hours for Amazon S3 Glacier Deep Archive.

- Bulk: the least expensive method. Retrieval time within 5-12 hours for Amazon S3 Glacier and within 48 hours for Amazon S3 Glacier Deep Archive.

Expiration of the Retrieved Data for Amazon S3 Glacier

Estimated time of data expiration is calculated from the moment of the data retrieval launch. Required number of days is added to that moment.

However, during the retrieval job Veeam Backup & Replication constantly requests the S3 API to move the expiration time forward, until the job is completed and all the data blocks are ready.

For more information on calculating the estimated expiration time, see this Amazon article.

For information on how to extend the expiration time, see Extending Expiration Time.

Retrieval Options for Azure Archive Storage

Azure provides the following options for data retrieval. The indicated time is approximate. For more information, see this Microsoft article.

- High Priority: most expensive method. Retrieval may take less than one hour.
- Standard Priority: least expensive method. Retrieval takes up to 15 hours.

Expiration of the Retrieved Data for Azure Archive Storage

Unlike Amazon S3 Glacier, Azure API does not temporarily change the storage class of a backup file from Archive Tier to Hot. Instead, a temporary copy of a backup file is created in Hot storage class. Since the deletion of this temporary copy is managed by Veeam Backup & Replication, the expiration time is quite accurate (within ten minutes, which is the frequency of the deletion process).

For information on how to extend the expiration time, see Extending Expiration Time.

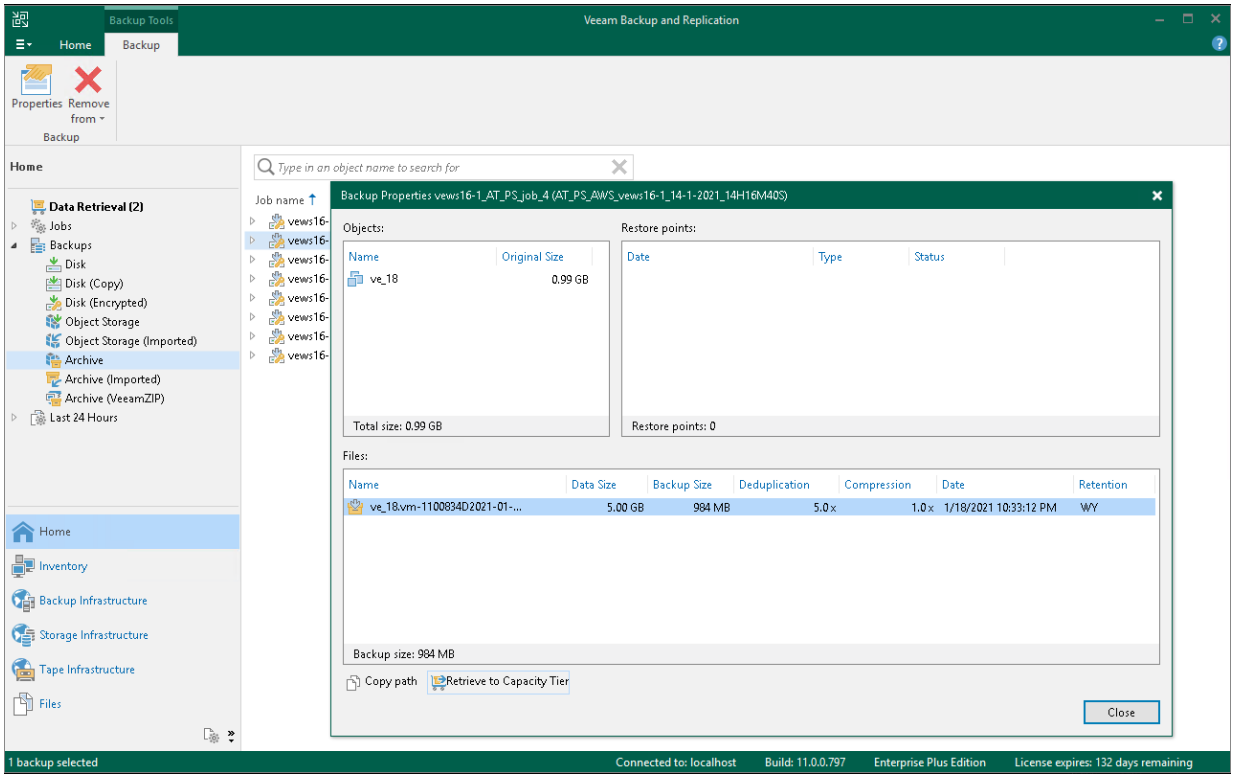
Retrieving Backup Files

To start the retrieval of archived backup files, use the **Retrieve Backup** wizard.

Step 1. Launch Retrieve Backup Wizard

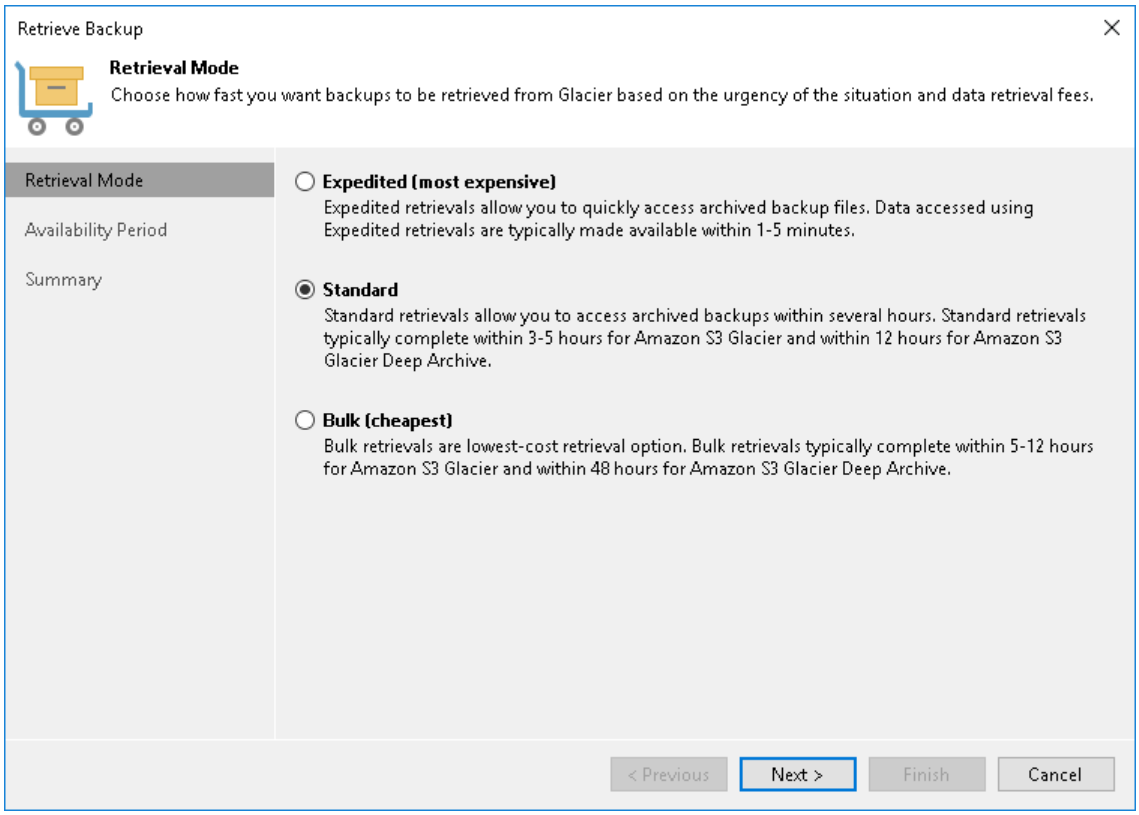
To launch the retrieval job, do one of the following:

- Open the **Home** view. In the inventory pane select **Archive**. In the working area, select the backup job whose files you want to retrieve and click **Properties** on the ribbon. In the **Backup Properties** window, click on **Retrieve to Capacity Tier**.
- Open the **Home** view. In the inventory pane select **Archive**. In the working area, select the VM whose guest OS files you want to restore and click **Guest Files** on the ribbon. In this case, you will proceed immediately to the **Select Restore Point** step of the Virtual Machine Files Restore wizard. If this restore point has not been retrieved yet, you will be prompted to launch the **Retrieve Backup** wizard.



Step 2. Select Retrieval Mode

At the **Retrieval Mode** step of the wizard, select the desired retrieval option. For information on the retrieval modes for different archive storage options, see Data Retrieval.

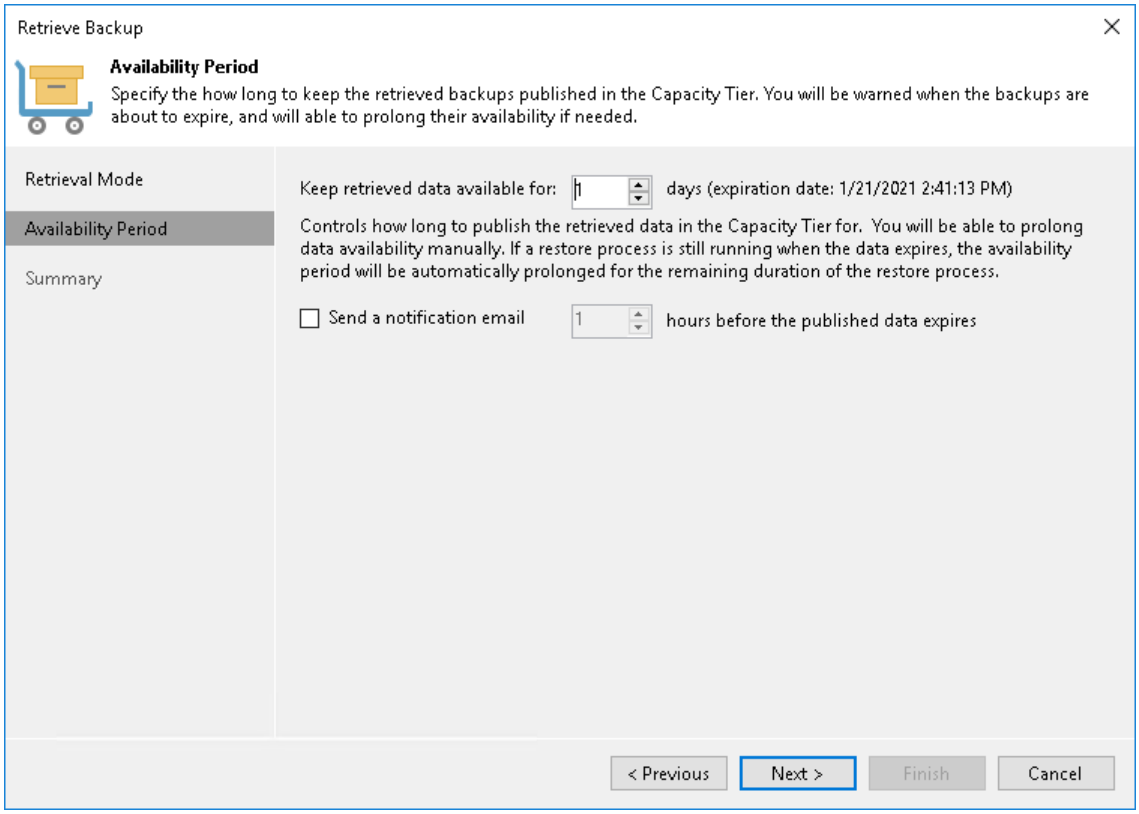


Step 3. Select Availability Period

At the **Availability Period** step of the wizard, select the desired availability period of the retrieved backup files. During that period you will be able to restore the data.

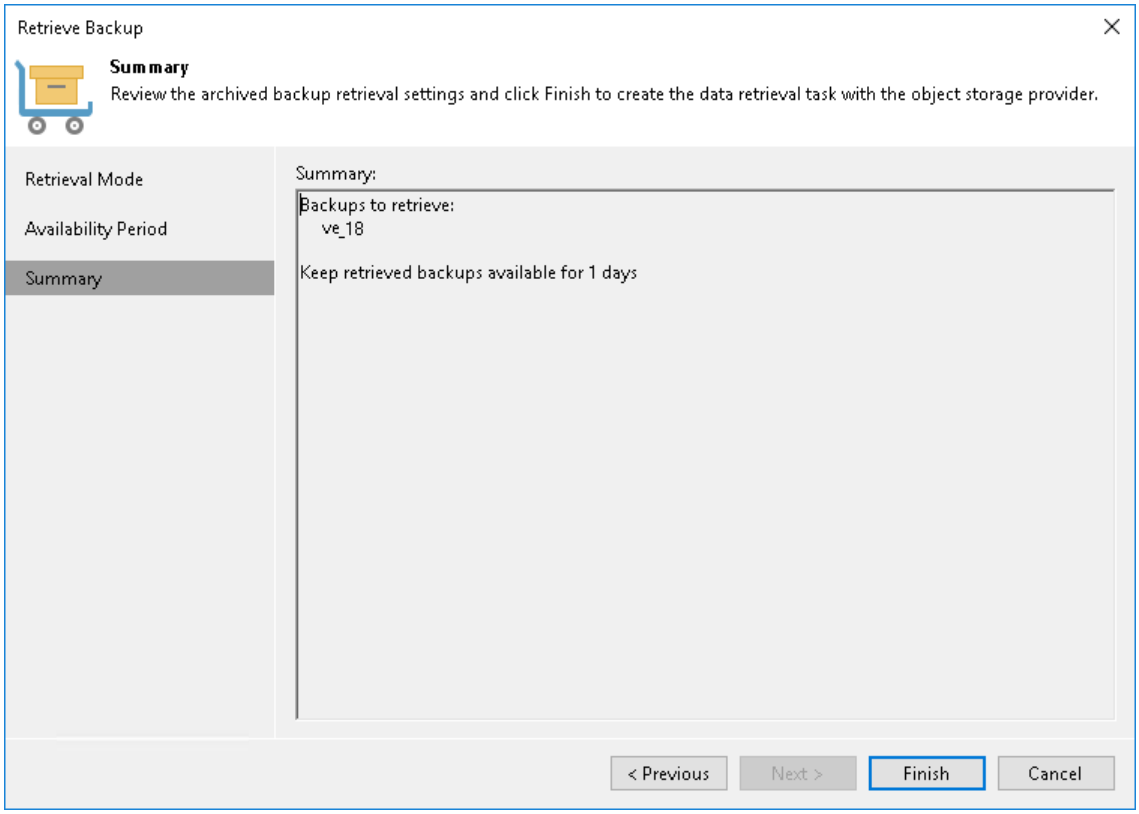
If you want to receive a notification that the availability period is about to end, select the **Send a notification email N hours before the published data expires** check box and choose the desired time for the notification.

> **TIP**
>
> You will be able to extend that period if necessary. For information on how to extend backup files availability, see Extending Data Availability.



Step 4. Finish Working with Wizard

At the **Summary** step of the wizard, review the retrieval settings. Click **Finish** to exit the wizard.



Extending Data Availability

The availability period of the retrieved data can be prolonged:

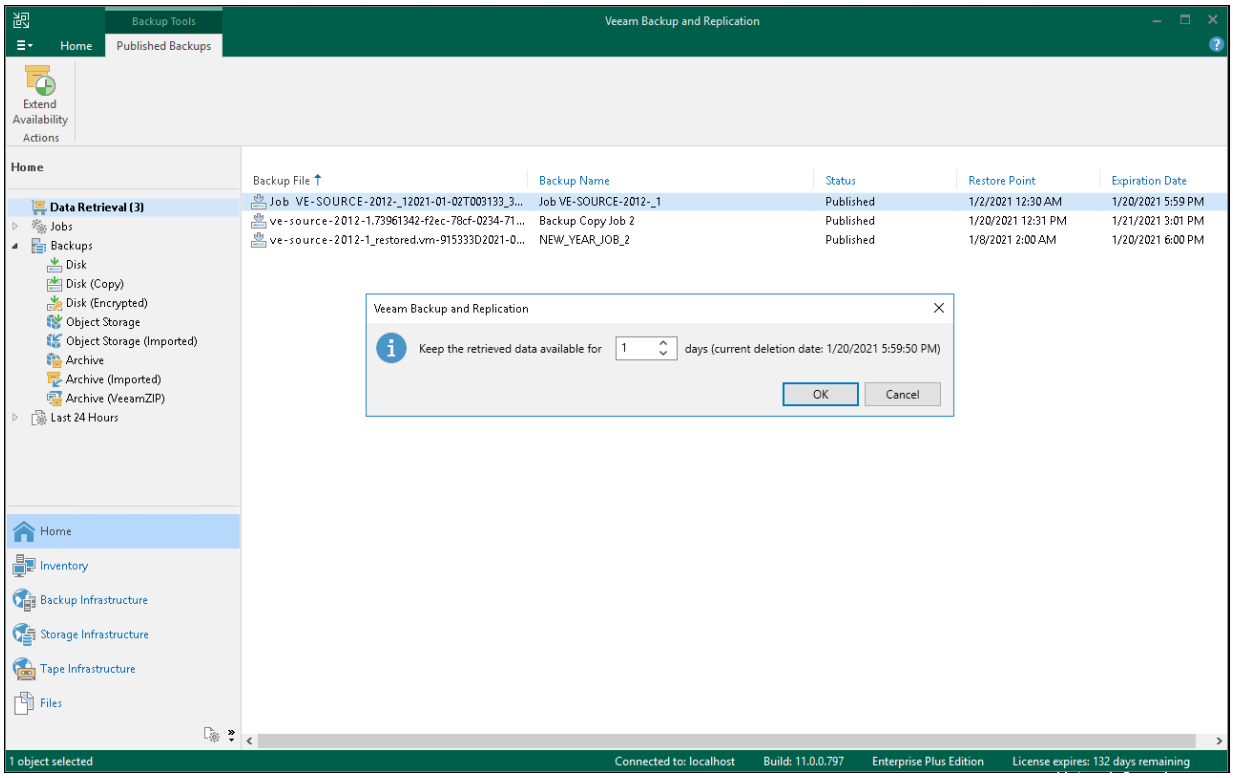
- Manually. If the retrieval job still continues at the moment of manual extension, the new extension time is automatically added to the job.
- Automatically. This happens in cases when the expiration time is close, but the restore process still continues.

At the Select Availability Period step of the **Retrieve Backup** wizard, you can request a notification that the data availability is about to expire. If you extend the expiration time, the notification will be postponed accordingly.

To extend the availability period, do the following:

- Open the **Home** view. In the inventory pane select **Data Retrieval**. In the working area, select the retrieved backup file whose availability period you want to extend, and click **Extend Availability** on the ribbon. You will be prompted to select the number of days you want to add to the availability expiration time.

- Open the **Home** view. In the inventory pane select **Archive**. In the working area, select the backup job for whose retrieved files you want to extend availability and click **Properties** on the ribbon. In the **Backup Properties** window, click on **Extend Availability Period.** You will be prompted to select the number of days you want to add to current expiration time.



Restoring Retrieved Data

You can restore your data directly from the archive tier back to production servers or to Azure or Amazon EC2 cloud platforms. The retrieved data recovery does not differ from standard backup data recovery and can be performed by using any of the following methods:

- Instant Recovery to VMware vSphere
- Instant Recovery to Microsoft Hyper-V
- Entire VM Recovery
- VM Files Restore
- Virtual Disks Restore
- Guest OS File Recovery
- Instant Disk Recovery
- Application Items Restore
- Exporting Backups

You can also recover data directly to Amazon EC2 or Microsoft Azure. For more information on these types of restore, see sections Restore to Amazon EC2 and Restore to Microsoft Azure.

Adding Scale-Out Backup Repositories

Before you add a scale-out backup repository, check prerequisites. Then use the **New Scale-out Backup Repository** wizard to configure the scale-out backup repository.

Before You Begin

Before you add a scale-out backup repository to the backup infrastructure, check the following prerequisites:

- Backup repositories that you plan to add as performance extents to the scale-out backup repository must be added to the backup infrastructure. For more information, see Adding Backup Repositories.
- You must check limitations for scale-out backup repositories. For more information, see Limitations for Scale-Out Backup Repositories.
- You can add only one capacity extent per scale-out backup repository.
- An object storage repository cannot be added as part of two or more different scale-out backup repositories at the same time.
- If the selected object storage contains offloaded backup data, you will be offered to synchronize this data with your performance extents.

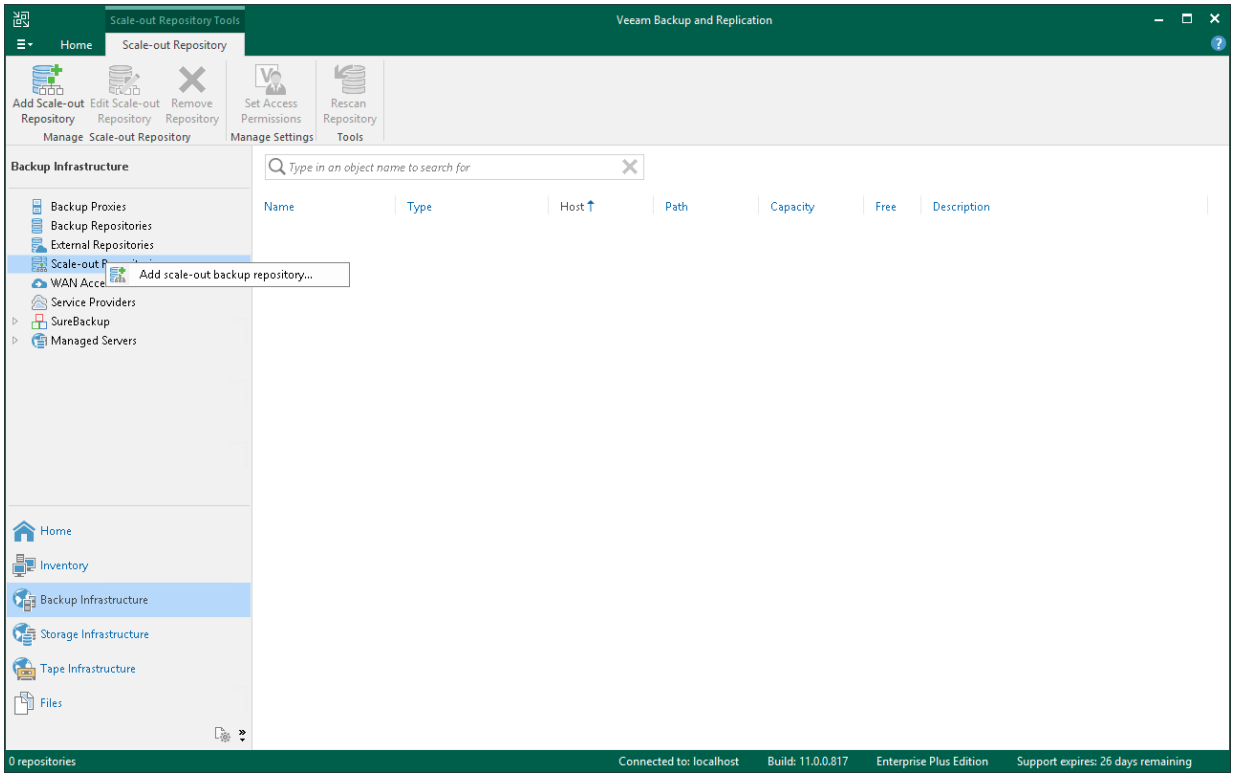
  For more information, see Synchronizing Capacity Tier Data.
- Object storage that contains imported backups cannot be added as a capacity extent.

  For more information, see Importing Object Storage Backups.

Step 1. Launch New Scale-Out Backup Repository Wizard

To launch the **New Scale-out Backup Repository** wizard, do one of the following:

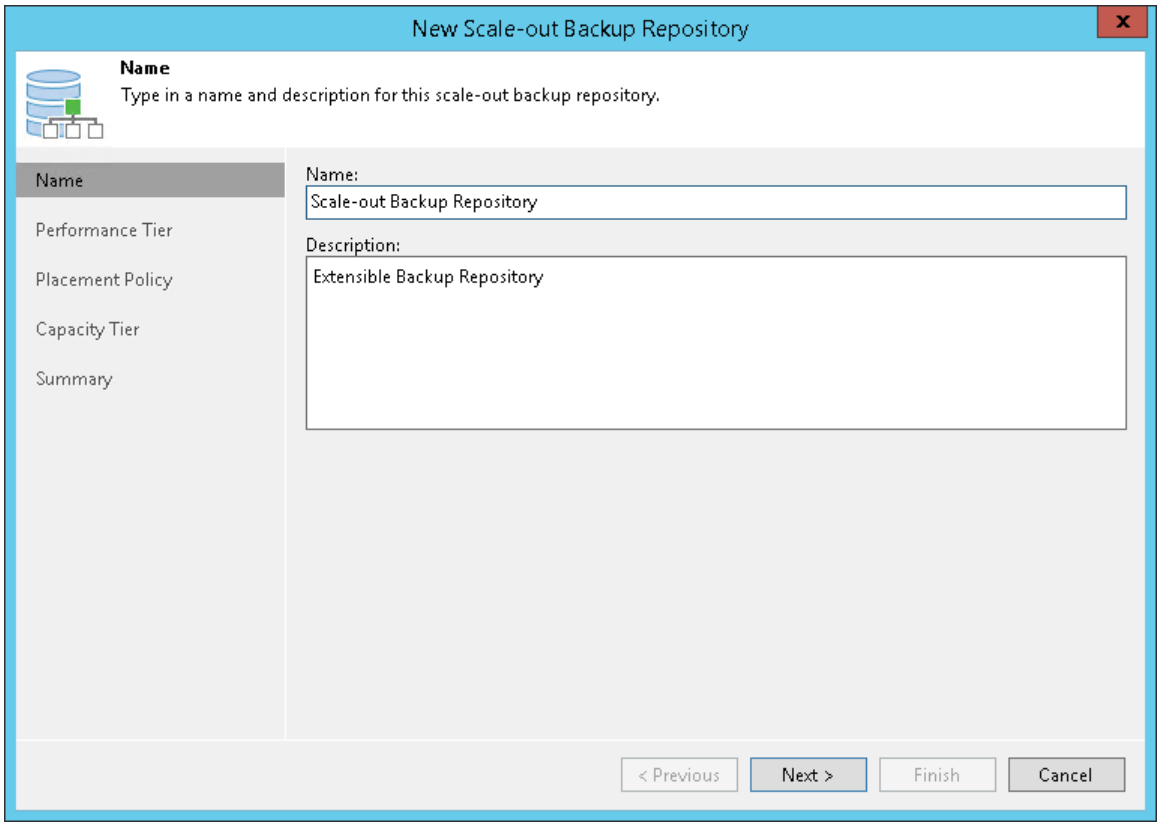
- Open the **Backup Infrastructure** view. In the inventory pane select **Scale-out Repositories** and click **Add Scale-out Repository** on the ribbon.
- Open the **Backup Infrastructure** view. In the inventory pane right-click **Scale-out Repositories** and select **Add Scale-out Backup Repository**.



Step 2. Specify Scale-Out Backup Repository Name

At the **Name** step of the wizard, specify a name and description for the scale-out backup repository.

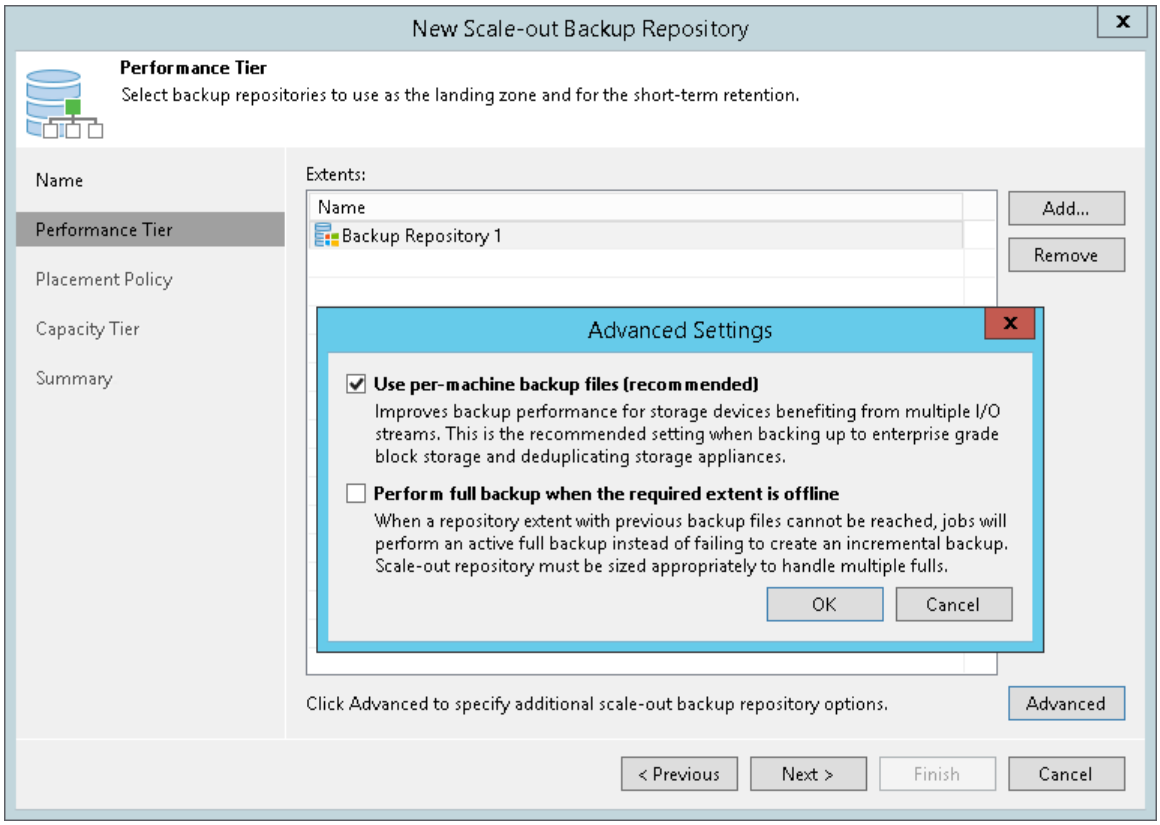
1. In the **Name** field, specify a name for the scale-out backup repository.
2. In the **Description** field, provide a description for future reference. The default description contains information about the user who added the backup repository, date and time when the backup repository was added.



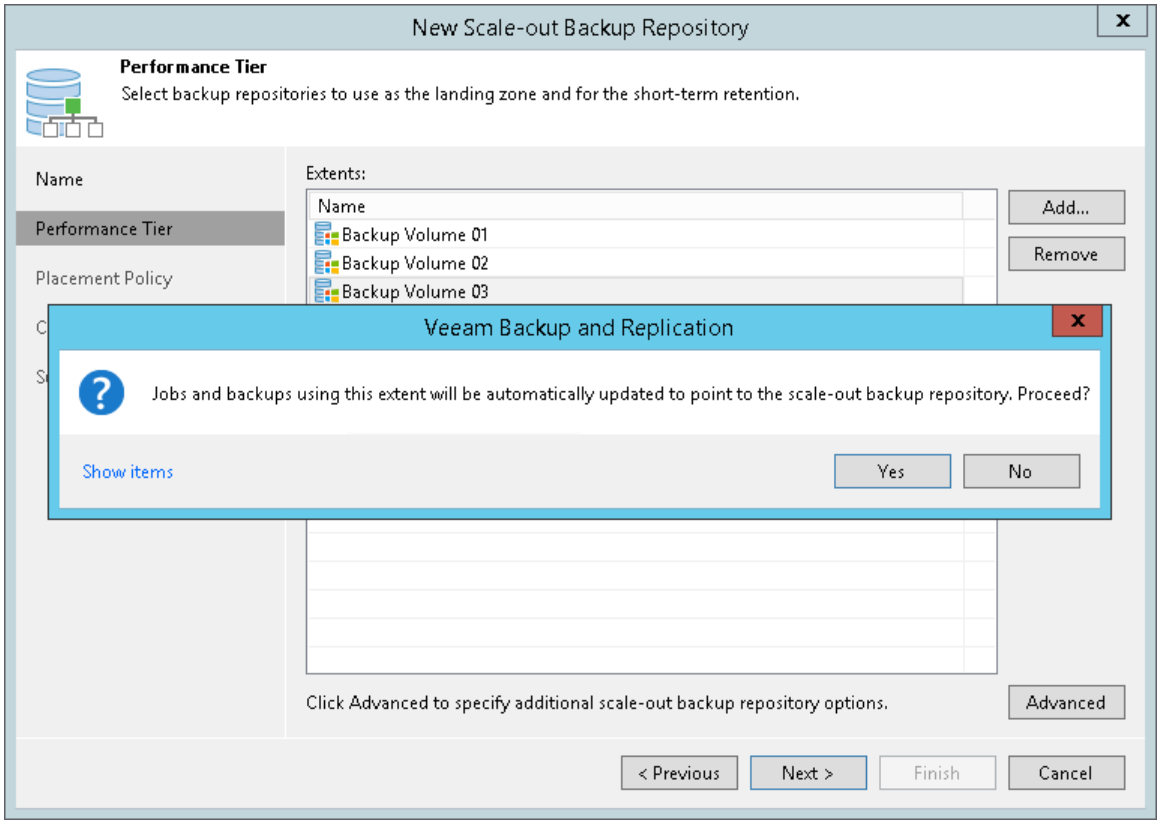
Step 3. Add Performance Extents

At the **Performance Tier** step of the wizard, specify which backup repositories you want to add as performance extents, and configure options for the scale-out backup repository.

1. On the right of the **Extents** list, click **Add**.
2. In the **Extents** window, select check boxes next to backup repositories that you want to add as performance extents.
3. Click **OK**.
4. At the lower right corner of the **Extents** list, click **Advanced**.
5. Specify advanced options for the scale-out backup repository:
   a. Select the **Use per-machine backup files** check box if you want to create a separate backup chain for every machine in the job. With this option enabled, during one backup job session Veeam Backup & Replication will produce a number of backup files — one per every machine, and will write these files to the backup repository in multiple streams simultaneously. It is recommended that you enable this option to achieve better storage and compute resource utilization, especially if you use as a backup repository a deduplicating storage appliance that supports multiple write streams.
   b. To preserve the consistency of backup chains in the scale-out backup repository, select the **Perform full backup when required extent is offline** check box. If an extent that contains previous restore points from the current backup chain gets offline, the backup chain will be broken. Veeam Backup & Replication will not be able to add a new incremental backup file. With this option enabled, Veeam Backup & Replication will create a full backup file instead of an incremental backup file. If you enable this option, you must make sure that you have enough free space in the scale-out backup repository to host a full backup file.



If a backup repository that you add as a performance extent is already used by jobs of supported type or there are backups pointing at the backup repository (for example, independent backups created with VeeamZIP), Veeam Backup & Replication will offer you to update a link to the backup repository in the job properties. Click **Yes** to update the link and target the jobs and backups at the scale-out backup repository. If you click **No**, you will not be able to pass to the next steps of the wizard.

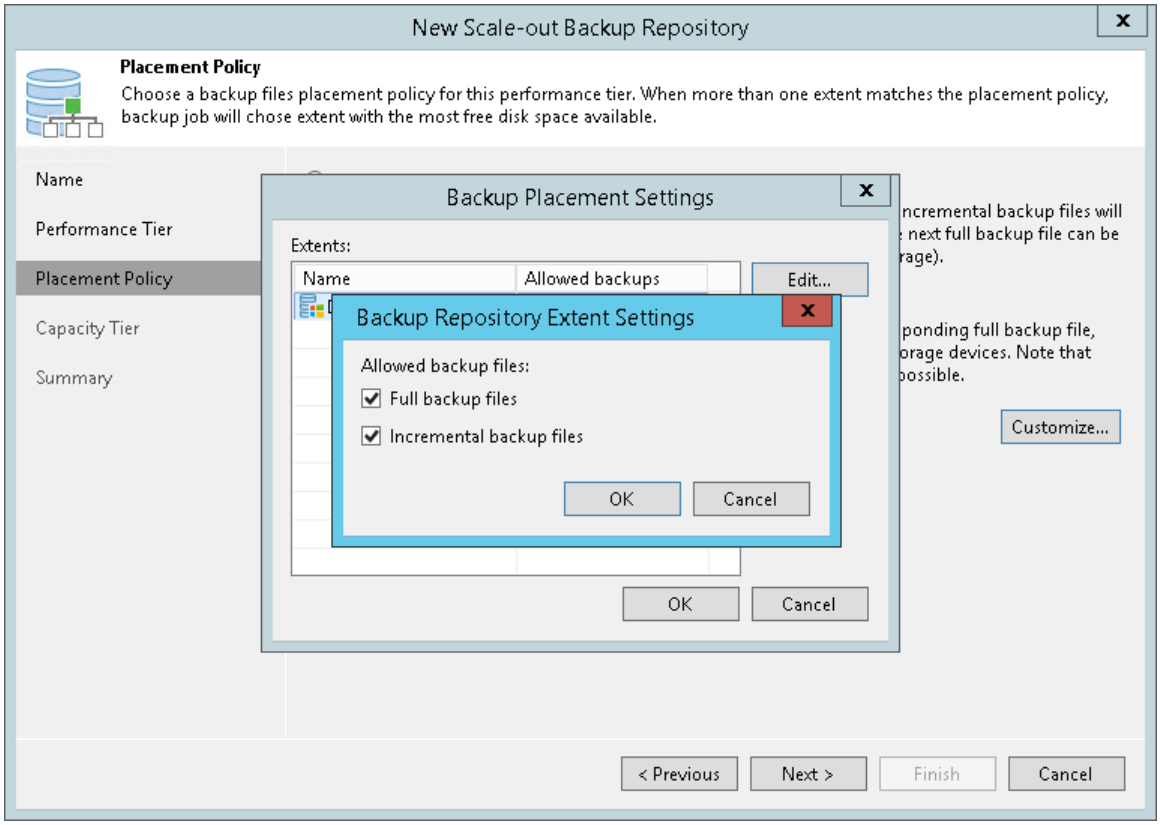


Step 4. Specify Backup Placement Policy

At the **Policy** step of the wizard, specify how you want to store backup files on performance extents of the scale-out backup repository.

1. Set the backup file placement policy for the scale-out backup repository:
   - Select **Data locality** if you want to store backup files that belong to the same backup chain together. In this case, a full backup file and subsequent incremental backup files will be stored to the same performance extent of the scale-out backup repository. The new backup chain may be stored to the same performance extent or to another performance extent (unless you use a deduplicating storage appliance as a performance extent).
   - Select **Performance** if you want to store full and incremental backup files to different performance extents of the scale-out backup repository. If you set the Performance policy, you must make sure that the network connection is fast and reliable so that Veeam Backup & Replication can access all backup files from the backup chain.

   For more information, see Backup File Placement.
2. If you select the **Performance** policy, you can restrict which types of backup files can be stored on a specific performance extent. For example, if you have added three performance extents to the scale-out backup repository, you may want to store full backup files on one extent and incremental backup files — on the other two extents.
   a. Click **Customize**.
   b. In the **Backup Placement Settings** window, select a performance extent and click **Edit**.
   c. Select a check box next to the type of backup files that you want to store on the extent: **Full backup files** or **Incremental backup files**. By default, Veeam Backup & Replication can store both full and incremental backup files on the same extent.



Step 5. Add Capacity Tier

Before you add a capacity tier, check the prerequisites.

At the **Capacity Tier** step of the wizard, select an object storage repository that you want to add as a capacity extent. Then specify when to move and/or copy data.

> **TIP**
>
> If you already have a scale-out backup repository in your backup infrastructure and you want to add capacity tier, select the scale-out backup repository, click **Edit Scale-out Repository** on the ribbon or right-click the scale-out backup repository and select **Properties**. In the **Edit Scale-out Backup Repository** wizard go to the **Capacity Tier** step and proceed with the steps below.

To configure capacity tier, do the following:

1. Select the **Extend scale-out backup repository capacity with object storage** check box.
2. Select an object storage repository to which you want to offload your data.

   Make sure that this storage has been added to your environment in advance. In case an object storage repository has not been configured, click **Add** and follow the associated steps of the wizard, as described in Adding Object Storage Repository.
3. Click **Window** and specify when it is allowed or prohibited to move or copy data to object storage.
4. Select the **Copy backups to object storage as soon as they are created** check box to copy new backups as soon as they are created, as described in Copying Backups to Capacity Tier.

   When selecting this option, you will be asked whether to copy all backup files that you might already have on any of the performance extents, or only those that have been created recently.

   If you select **Latest**, only backup files that belong to the last active backup chain will be copied from each of the performance extents. If you select **All**, Veeam Backup & Replication will copy all backup files that belong to all backup chains located on any of the specified extents.
5. Select the **Move backups to object storage as they age out of the operational restores window** check box to move inactive backup chains to the capacity extent, as described in Moving Backups to Capacity Tier.

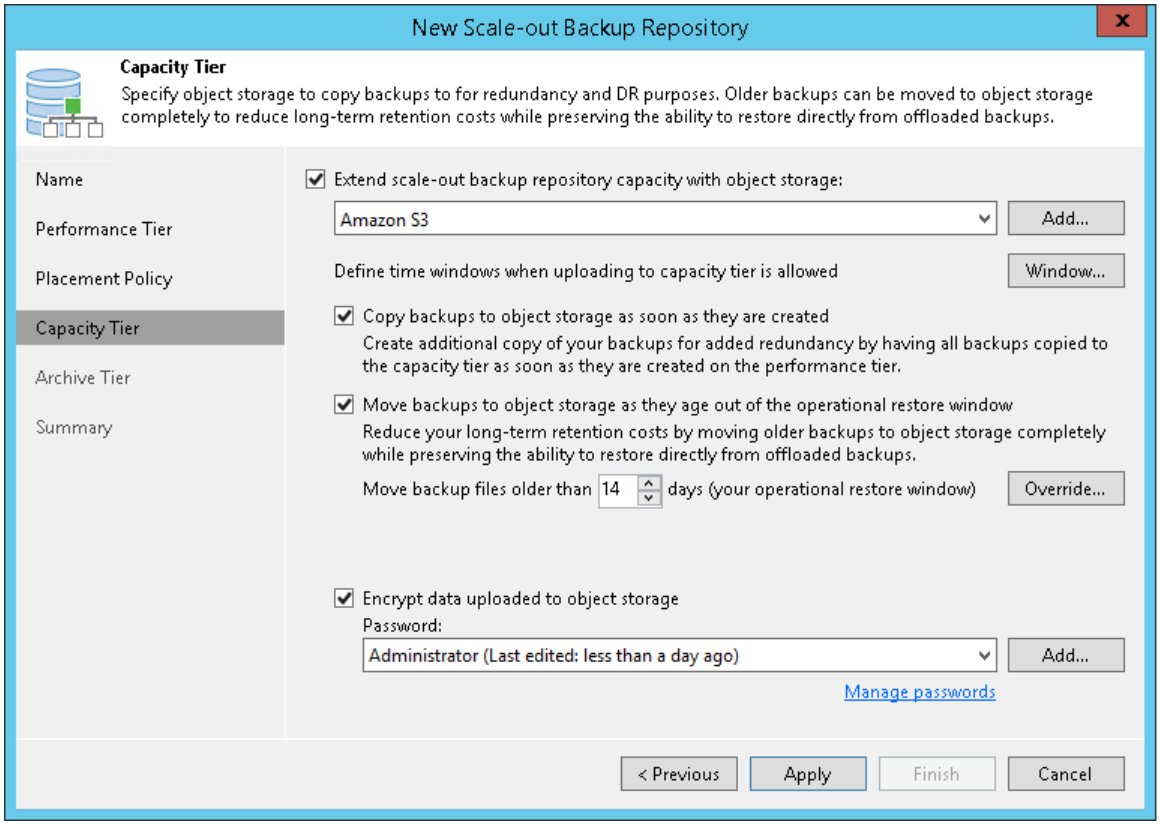
   In the **Move backup files older than X days** field, specify the operational restore window to define a period after which inactive backup chains on your performance extents would be considered outdated and, therefore, should be moved to the capacity extent. Consider that "0" is an acceptable value, which you can specify to offload inactive backup chains on the same day they are created.

   To override behavior of moving old backups, click **Override**, select the **Move oldest backup files sooner if scale-out backup repository is reaching capacity** check box and define a threshold in percent to force data transfer if a scale-out backup repository has reached the specified threshold.
6. To offload data encrypted, select **Encrypt data uploaded to object storage** and provide a strong password. With this option selected, the entire collection of blocks along with the metadata will be encrypted while being offloaded.

   If you have not created the password beforehand, click **Add** or use the **Manage passwords** link to specify a new password.

TIP

You can combine both the **Copy backups to object storage as soon as they are created** option and the **Move backups to object storage as they age out of the operational restores window** option, as described in Copying Backups to Capacity Tier.



Synchronizing Capacity Tier Data

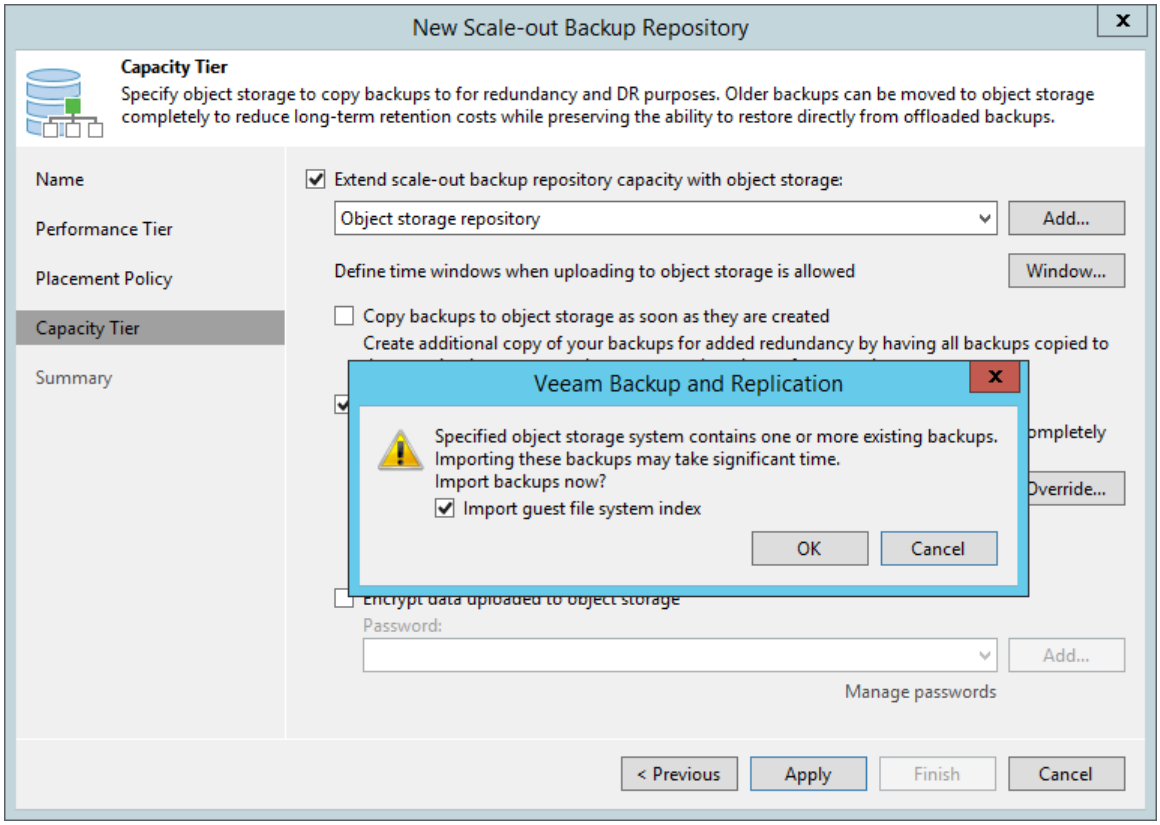
When you add as a capacity extent an object storage repository that contains offloaded backup data, you will be prompted to synchronize this data with the performance extents in the scale-out backup repository.

Consider the following:

- An object storage repository can only be added as a capacity extent after existing data (if any) is synchronized.
- During synchronization, Veeam Backup & Replication downloads backup files with metadata located in object storage to the performance extents that are part of the scale-out backup repository that is being added.

  These files are created as described in Moving Backups to Capacity Tier.
- Extents to which backup data is going to be downloaded (synchronized), will be selected automatically, depending on the available resources.
- The actual data blocks will not be downloaded and will continue to remain in object storage.
- When synchronizing encrypted storage, make sure to provide the same exact password with which the data was encrypted.

After the synchronization is complete, the associated backup files located in object storage will become available as Imported and will be displayed in the **Home** view, under the **Object Storage (Imported)** node in the inventory pane.



Step 6. Add Archive Tier

At the **Archive Tier** step of the wizard, select an object storage repository that you want to add as an archive extent and specify when to move and/or copy data.

> **TIP**
>
> If you have a compatible type of a repository configured as a capacity extent, you can add an archive extent to the existing scale-out backup repository. To do so, select the scale-out backup repository, click **Edit Scale-out Repository** on the ribbon or right-click the scale-out backup repository and select **Properties**. In the **Edit Scale-out Backup Repository** wizard go to the **Archive Tier** step and proceed with the steps below.

Consider the following:

- The Archive Tier step of the wizard will appear only if you have a compatible type of repository configured as a capacity extent. For more information, see Limitations for Archive Tier.
- You can add only one archive extent per scale-out backup repository.

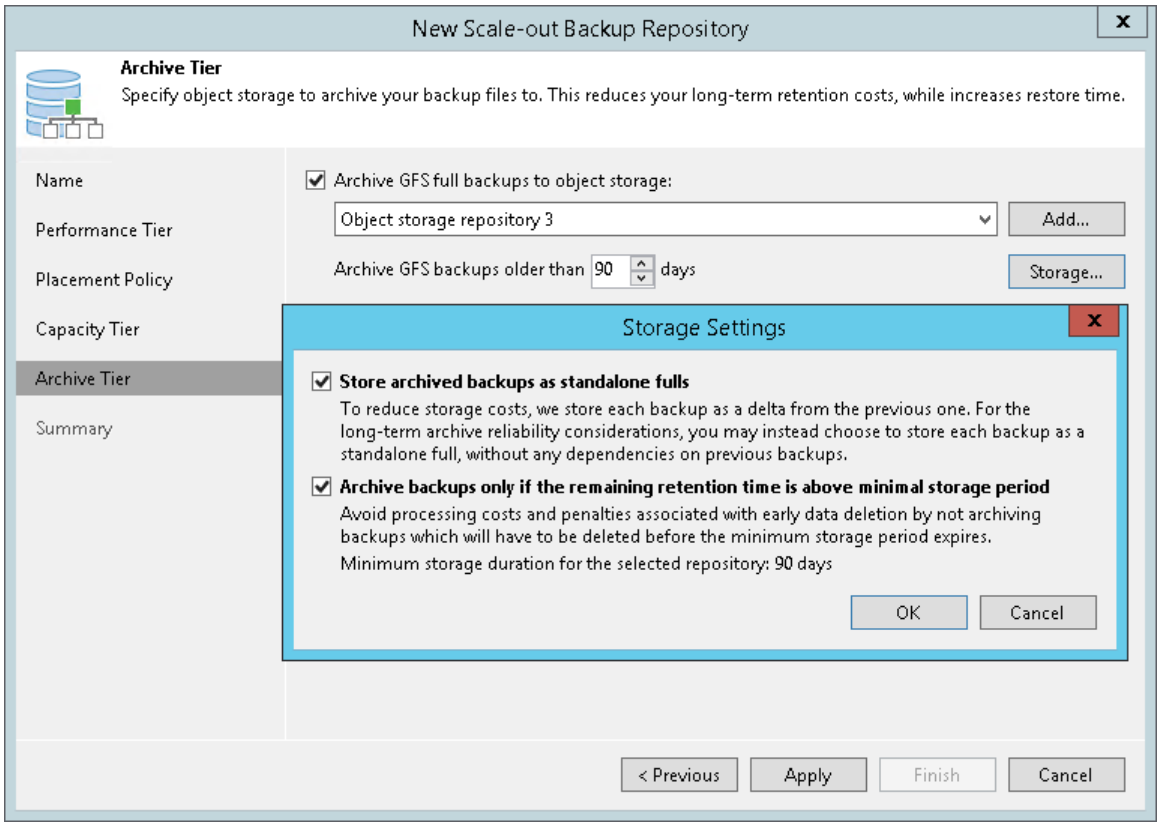
To configure the archive extent, do the following:

1. Select the **Archive GFS full backups to object storage** check box. In the drop-down list, select one of available object storage repositories or click **Add** to add a new one.
2. In the **Archive GFS backups older than N days** field, specify the operational restore window to define a period after which inactive backup chains on your capacity extent would be considered outdated and, therefore, should be moved to the archive extent. Consider that "0" is an acceptable value, which you can specify to archive inactive backup chains on the same day they are created.

You can use the default storage settings or specify them manually. For that, click **Storage**.

- Select the **Store archived backups as standalone fulls** check box to forbid reuse of the data blocks.
- Select the **Archive backups only if the remaining retention time is above minimal storage period** check box to specify which data blocks can be transported to the archive tier.

When you add as an archive extent an object storage repository that contains offloaded backup data, you will be prompted to synchronize existing backup chains with data in this scale-out backup repository. After the synchronization is complete, the existing backups will become available as Imported and will be displayed in the **Home** view, under the **Archive (Imported)** node in the inventory pane.

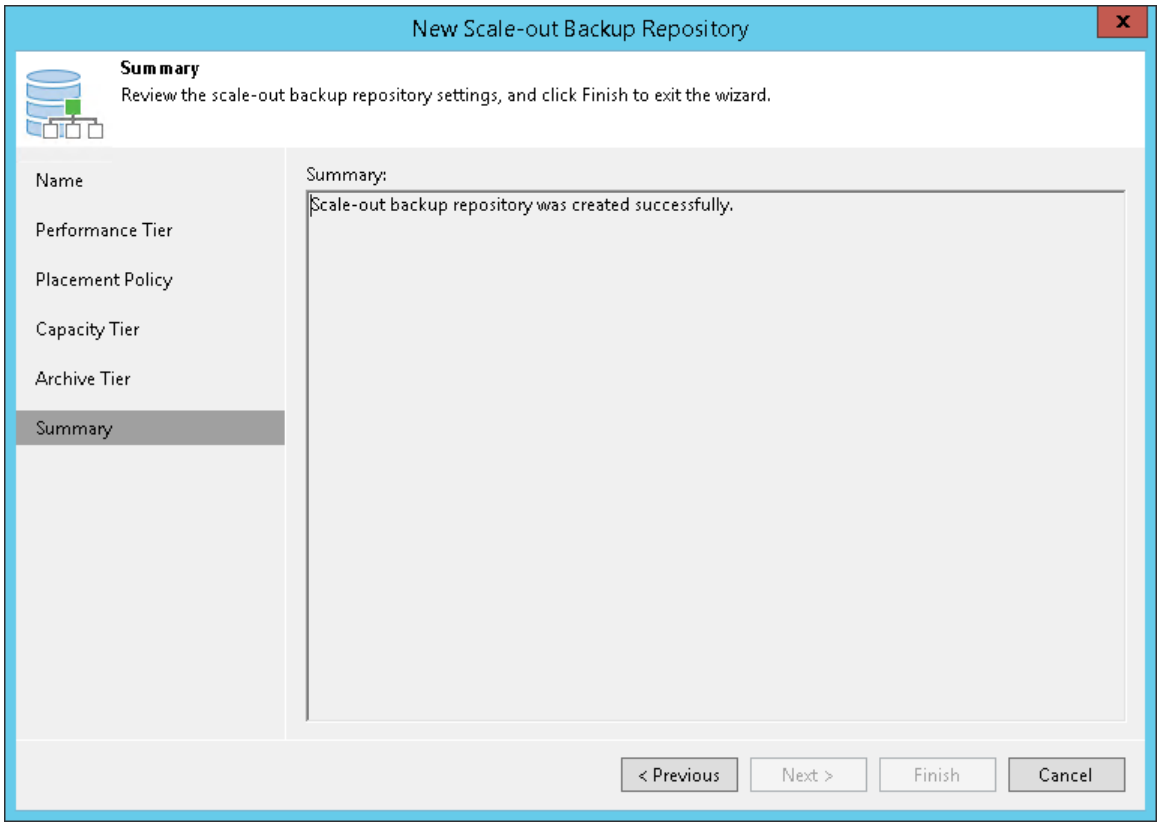


Step 7. Finish Working with Wizard

At the **Summary** step of the wizard, complete the procedure of scale-out backup repository configuration.

Wait for the scale-out backup repository to be added to the backup infrastructure. The process may take some time.

1. Review details of the scale-out backup repository.
2. Click **Finish** to exit the wizard.



Managing Scale-Out Backup Repositories

You can manage your scale-out backup repositories and the data therein in various ways: edit settings of the scale-out backup repository, rescan scale-out backup repositories automatically or manually, discover on which performance extent of the scale-out backup repository a particular backup file is stored, extend scale-out backup repository or remove certain performance extents from it, perform service actions or remove the scale-out backup repository.

Editing Settings of Scale-Out Backup Repositories

You can edit settings of the scale-out backup repository, for example, if you want to change the backup file placement policy or specify other advanced settings for the backup repository.

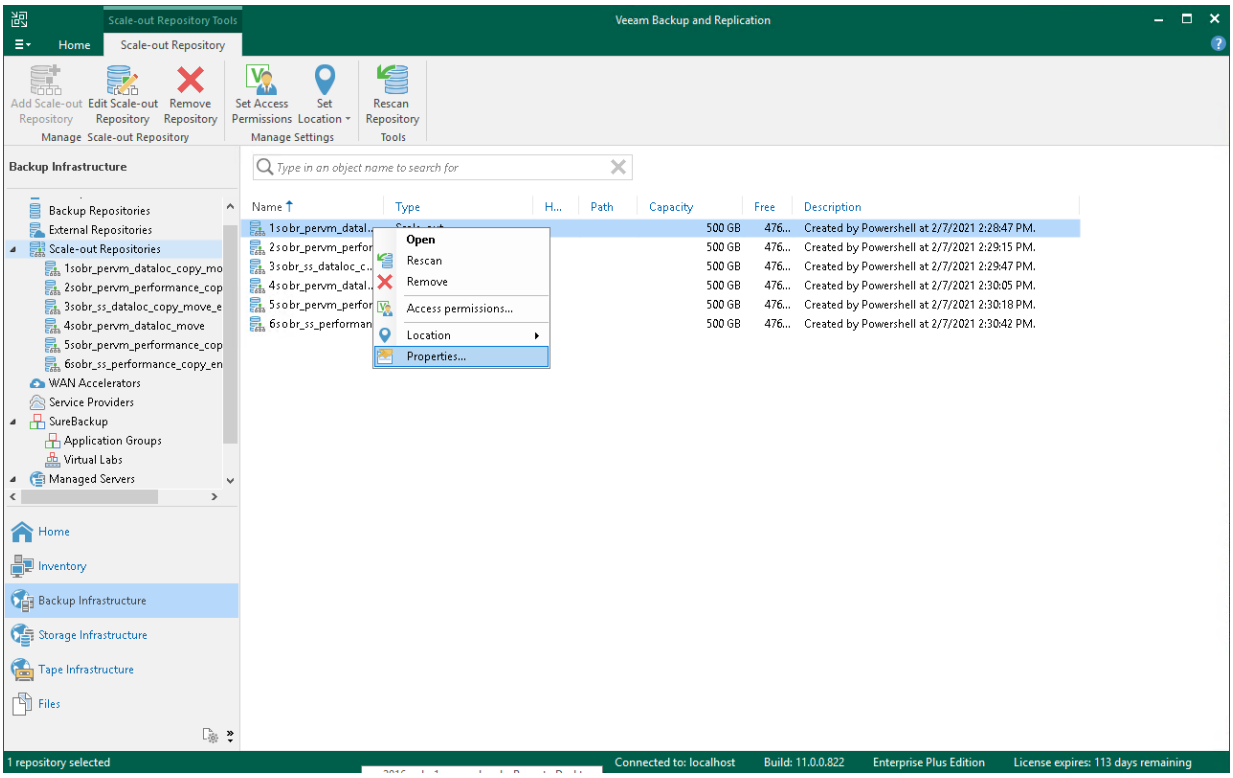
Mind the following:

- If you enable or disable the **Use per-machine backup file** option, Veeam Backup & Replication will apply new settings after a new full backup file is created.

- If you enable or disable the **Perform full backup when required extent is offline** option, Veeam Backup & Replication will apply the new settings starting from the next session of the job targeted at this scale-out backup repository.
- If you change the backup file placement policy settings, Veeam Backup & Replication will apply the new settings starting from the next session of the job targeted at this scale-out backup repository.

To change the scale-out backup repository settings:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, click **Scale-out Repositories**.
3. In the working area, select the scale-out repository and click **Edit Scale-out Repository** on the ribbon or right-click the scale-out backup repository and select **Properties**.
4. Follow the steps of the **Edit Scale-out Backup Repository** wizard and edit settings as required.



Rescanning Scale-Out Repositories

Veeam Backup & Replication periodically rescans scale-out backup repositories. During the rescan process, it gets the following information:

- State of every performance extent added to the scale-out backup repository: online or offline.
- Status of Veeam Data Movers on extents: up-to-date or outdated.
- Space available in the scale-out backup repository.

The rescan operation is performed automatically by a rescan process that works permanently in the background. The process is started every 24 hours. It can be also started when a new task session starts, and the Veeam Backup Service requires information about the infrastructure to be refreshed.

In addition to the automated rescan process, you can manually start rescan of the scale-out backup repository. Backup repository rescan may be helpful, for example, if you want to discover backup files that were manually relocated from one performance extent to another one.

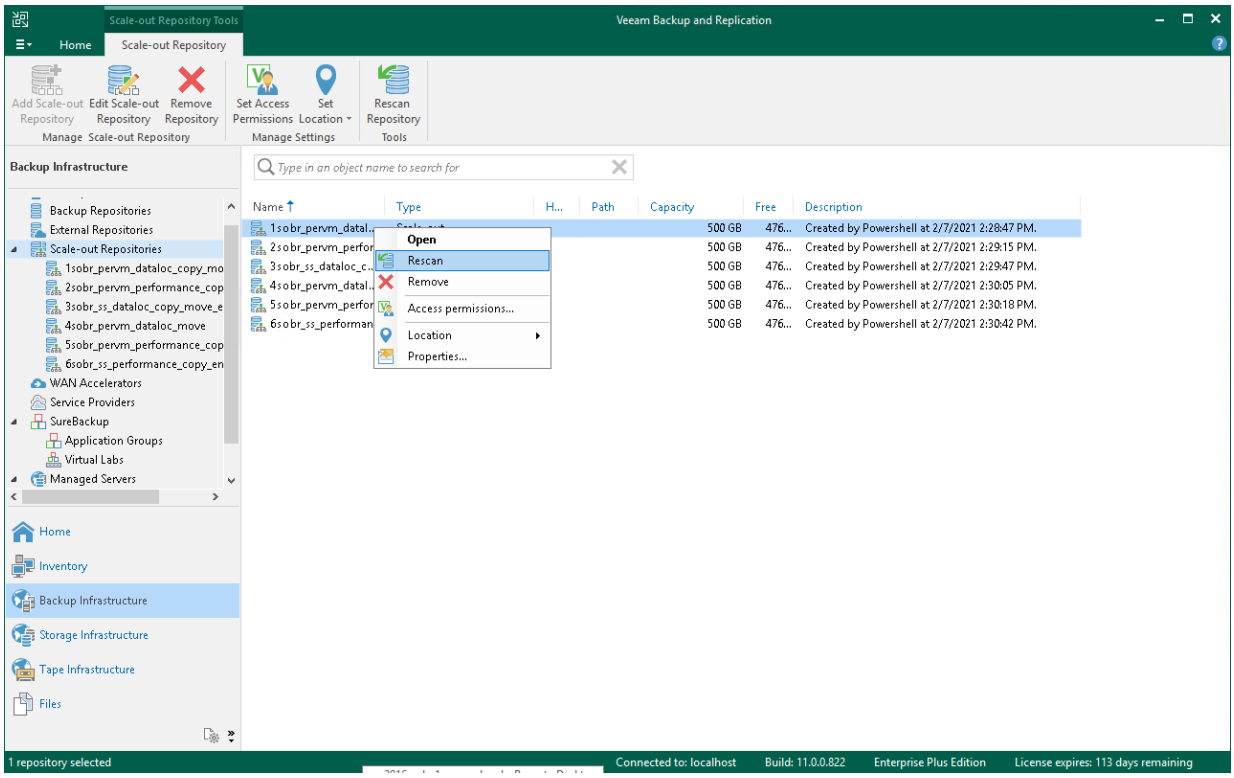
Consider the following:

- Information about backup files location is updated only if you perform manual rescan of scale-out backup repositories.
- Veeam Backup & Replication rescans scale-out backup repositories when you perform backup files import.
- To successfully rediscover relocated backups files created by backup copy jobs, make sure to disable these jobs manually prior to rescanning.

  For more information, see Disabling and Removing Jobs.

To start the rescan process:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane select **Scale-out Repositories**.
3. In the working area, select the scale-out repository and click **Rescan Repository** on the ribbon or right-click the scale-out backup repository and select **Rescan**.



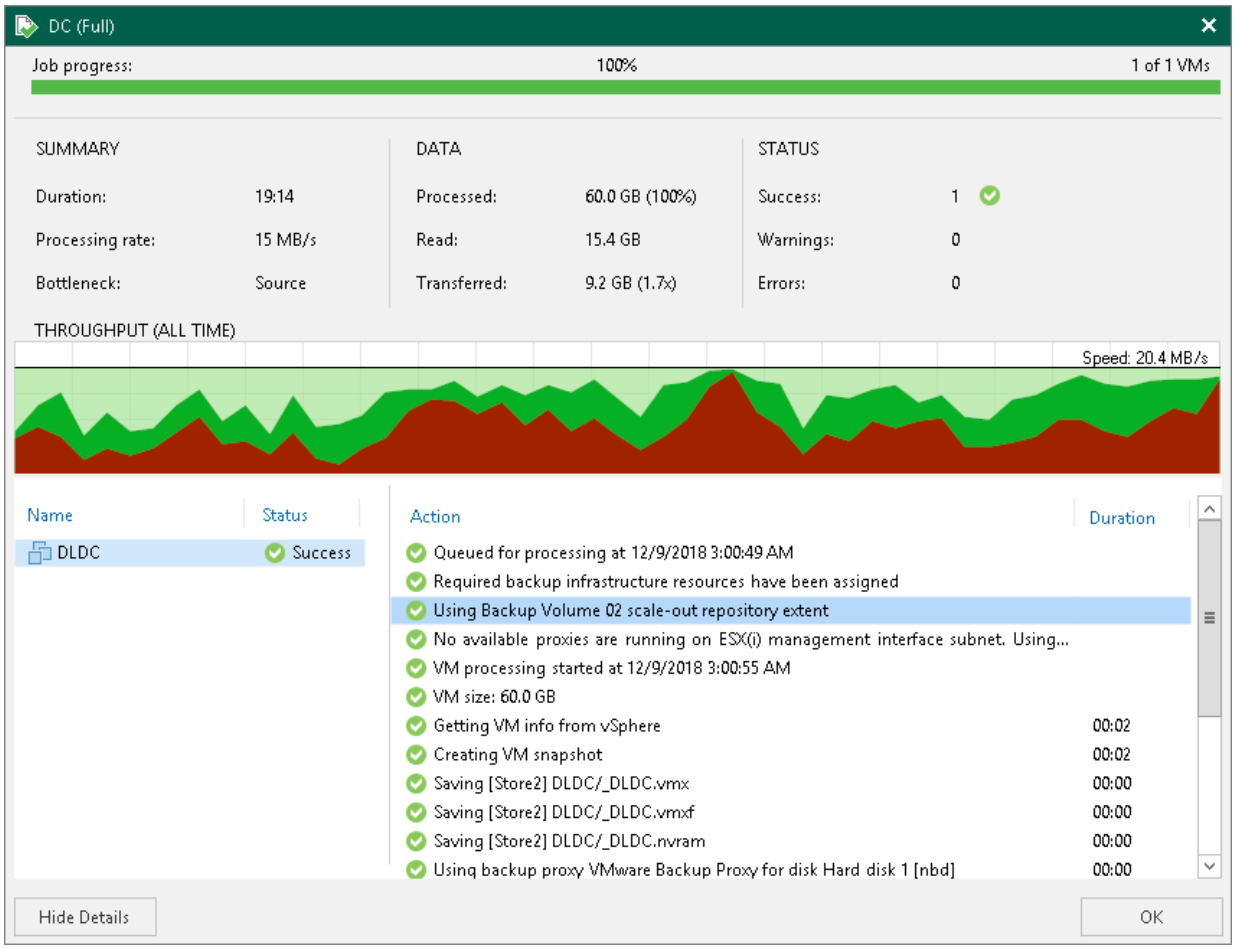
Discovering Backups in Scale-Out Backup Repositories

To discover on which performance extent of the scale-out backup repository a particular backup file is stored, you can examine the job session statistics or check the backup properties.

To view the job session statistics:

1. Open the **Home** view.
2. In the inventory pane, click **Backup** under **Jobs**.

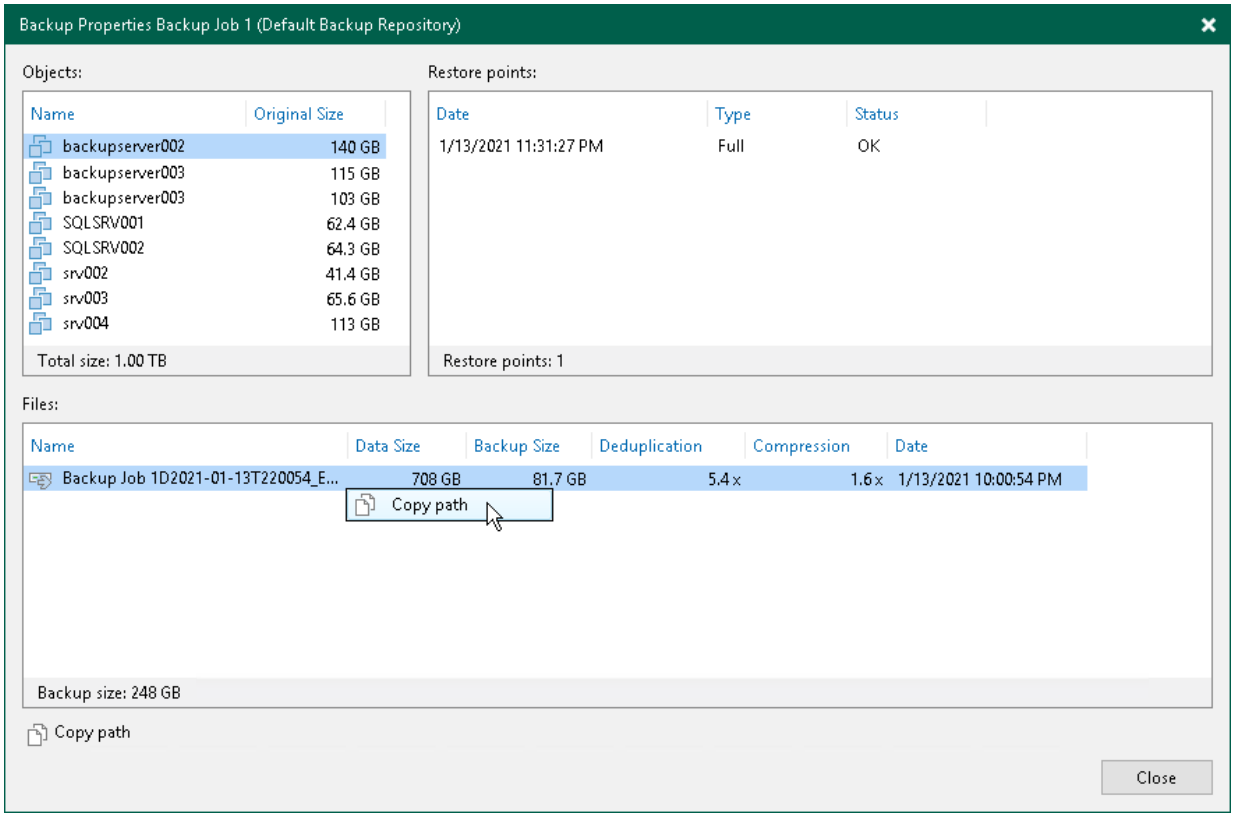
3. In the working area, right-click the job and select **Statistics**.
4. In the bottom left pane of the window, click the VM name. In the **Action** pane, locate the message: *Using N scale-out repository extent*.



To view the backup properties:

1. Open the **Home** view.
2. In the inventory pane, select **Disk** under **Backups**.
3. In the working area, right-click the backup and select **Properties**.

4. Veeam Backup & Replication will display the extent where the backup file resides in the headline of the **Backup Properties** window. To see the path to the backup file, right-click the job and select **Copy path**.



Backup State Indicators

Restore points icons help you understand current state of a restore point in a scale-out backup repository.

For more information on icons and what they indicate, see Restore Points Icons.

Extending Scale-Out Repositories

You can add a backup repository as a performance extent to the scale-out backup repository at any time. For example, the scale-out backup repository may run low on space, and you will need to add storage capacity to it.

To add a backup repository as an extent to the scale-out backup repository:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, click **Scale-out Repositories**.
3. In the working area, select the scale-out repository and click **Edit Scale-out Repository** on the ribbon or right-click the backup repository and select **Properties**.
4. Move to the **Performance tier** step of the wizard.
5. Click **Add**.

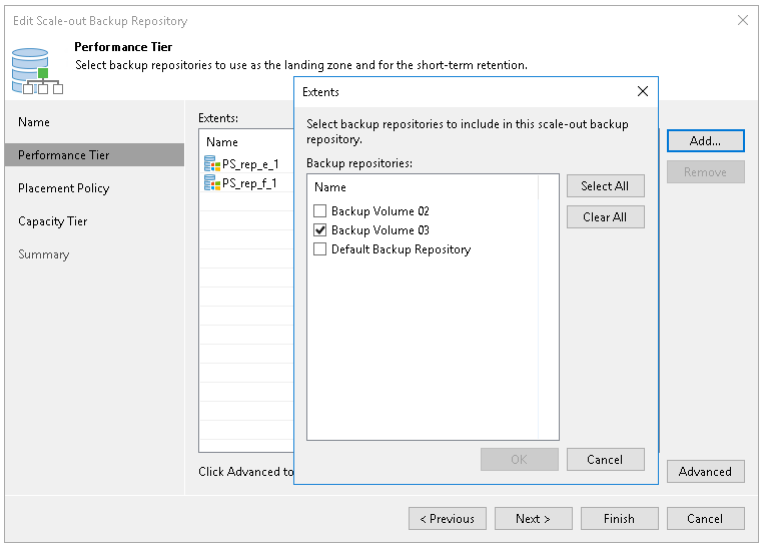
6. In the **Extents** window, select a check box next to the backup repository that you want to add as a performance extent to the scale-out backup repository.

   If a backup repository that you add as a performance extent is already used by jobs of supported type or there are backups pointing at the backup repository (for example, independent backups created with VeeamZIP), Veeam Backup & Replication will offer you to update a link to the backup repository in the job properties. Click **Yes** to update the link and target the jobs and backups at the scale-out backup repository. If you click **No**, you will not be able to pass to the next steps of the wizard.

7. Pass through the next wizard steps and finish working with the wizard. The new performance extent will be added to the scale-out backup repository.

NOTE

After you add a backup repository to the scale-out backup repository as a performance extent, you will not be able to use it as an individual backup repository.



Service Actions with Scale-Out Backup Repositories

In some cases, you may want to perform service actions with scale-out backup repository extents. For example, you need to upgrade the backup repository server and add more memory to it. Or you want to replace a storage device backing the extent and need to relocate backup files.

You can perform the following service actions with extents of a scale-out backup repositories:

- Switching to maintenance mode
- Evacuating backups from extents

Switching to Maintenance Mode

Maintenance Mode

Veeam Backup & Replication allows you to put any of the scale-out backup repository extents into the *maintenance* mode. You can use this mode if you need to perform service actions, such as upgrading an extent or installing a patch on it. Putting an extent into the maintenance mode is mandatory to evacuate backups, as described in Evacuating Backups from Extents.

The extent in the maintenance mode operates with the limited functionality:

- Veeam Backup & Replication does not start new tasks targeted at this extent.
- You cannot restore VM data from backup files residing on the extent. You also cannot restore VM data from backup files residing on other extents if a part of the backup chain resides on the extent in the maintenance mode.

When you switch to the maintenance mode, Veeam Backup & Replication launches the *Repository Maintenance* job. The *Repository Maintenance* job checks the status of jobs and tasks targeted at the extent and puts the extent to one of the following modes:

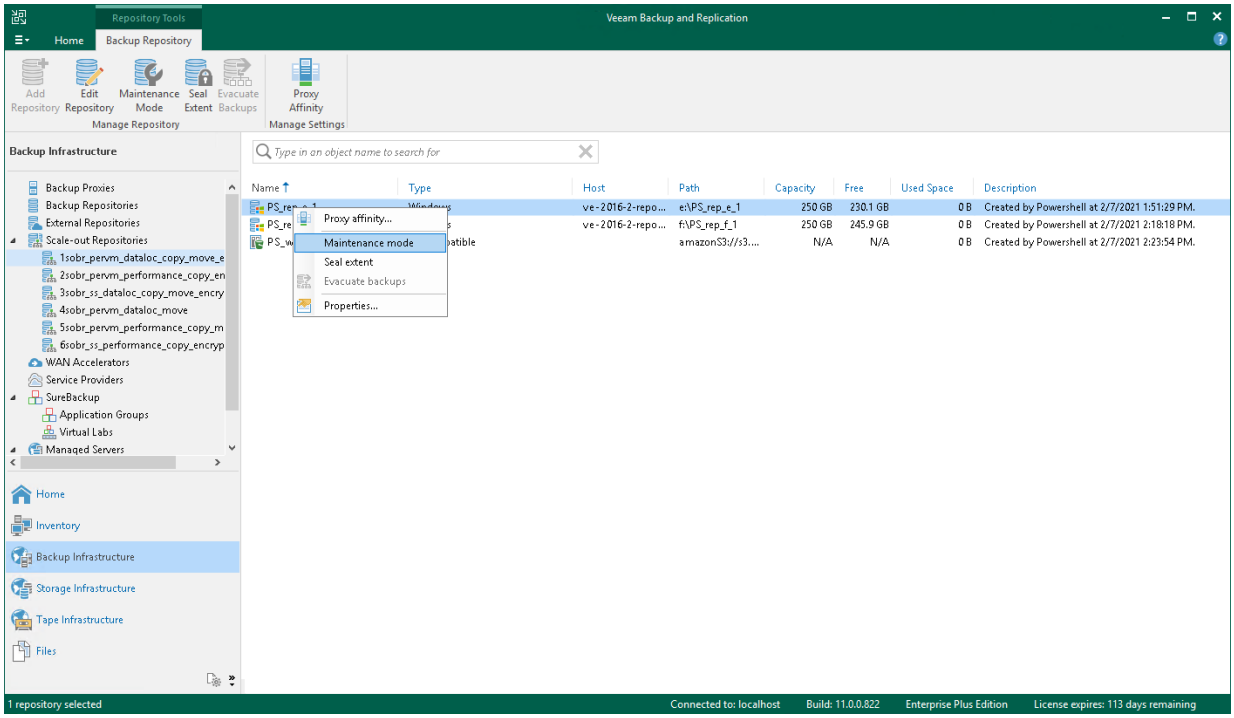
- If no tasks using the extent are currently running, the job puts the extent to the maintenance mode immediately.
- If the extent is busy with any task, for example, a backup job, the job puts the extent to the *Maintenance pending* state and waits for the task to complete. When the task is complete, the extent is put to the maintenance mode.

> **NOTE**
>
> If you put to the maintenance mode an extent that has a part of an active backup chain on it, the backup job targeted to the scale-out backup repository will fail. To avoid this, create a full backup manually.

To put an extent into the maintenance mode:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select the scale-out backup repository under **Scale-out Repositories**.
3. In the working area, select the extent and click **Maintenance Mode** on the ribbon or right-click the extent and select **Maintenance mode**.

To remove the extent from the maintenance mode, select the extent and click **Maintenance Mode** on the ribbon or right-click the extent and select **Maintenance mode** once again.



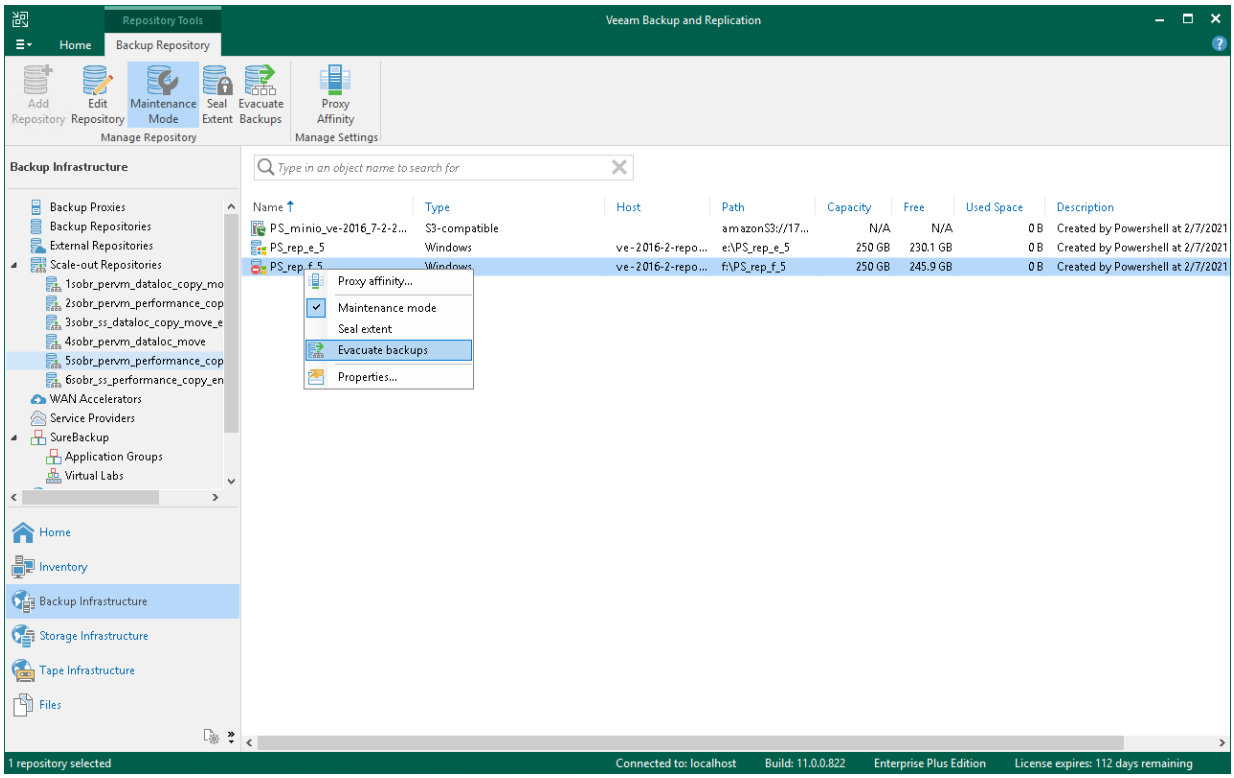
Evacuating Backups from Performance Extents

If you want to remove an extent from the scale-out backup repository, you first need to evacuate backups from this extent. When you evacuate backups, Veeam Backup & Replication moves backup files from the extent in question to other extents within the same scale-out backup repository.

You must put the extent into the maintenance mode before you evacuate backups from it. For more information, see Switching to Maintenance Mode.

To evacuate backup files from an extent:

1. [Recommended] Stop and disable jobs targeted at the extent from which you plan to evacuate backups.
2. Open the **Backup Infrastructure** view.
3. In the inventory pane, select the scale-out backup repository under **Scale-out Repositories**.
4. In the working area, select the extent and click **Evacuate Backups** on the ribbon or right-click the extent and select **Evacuate backups**.
5. If you have disabled jobs, enable them.



When selecting the target performance extent for evacuated files, Veeam Backup & Replication attempts to keep to the backup placement settings specified for remaining extents. For example, you have 3 performance extents in the scale-out backup repository with the following backup file placement settings:

- On *Extent 1*, full backup files are stored.
- On *Performance Extents 2* and *3*, incremental backup files are stored.

If you evacuate backup files from *Performance Extent 2*, Veeam Backup & Replication will relocate them to *Performance Extent 3*.

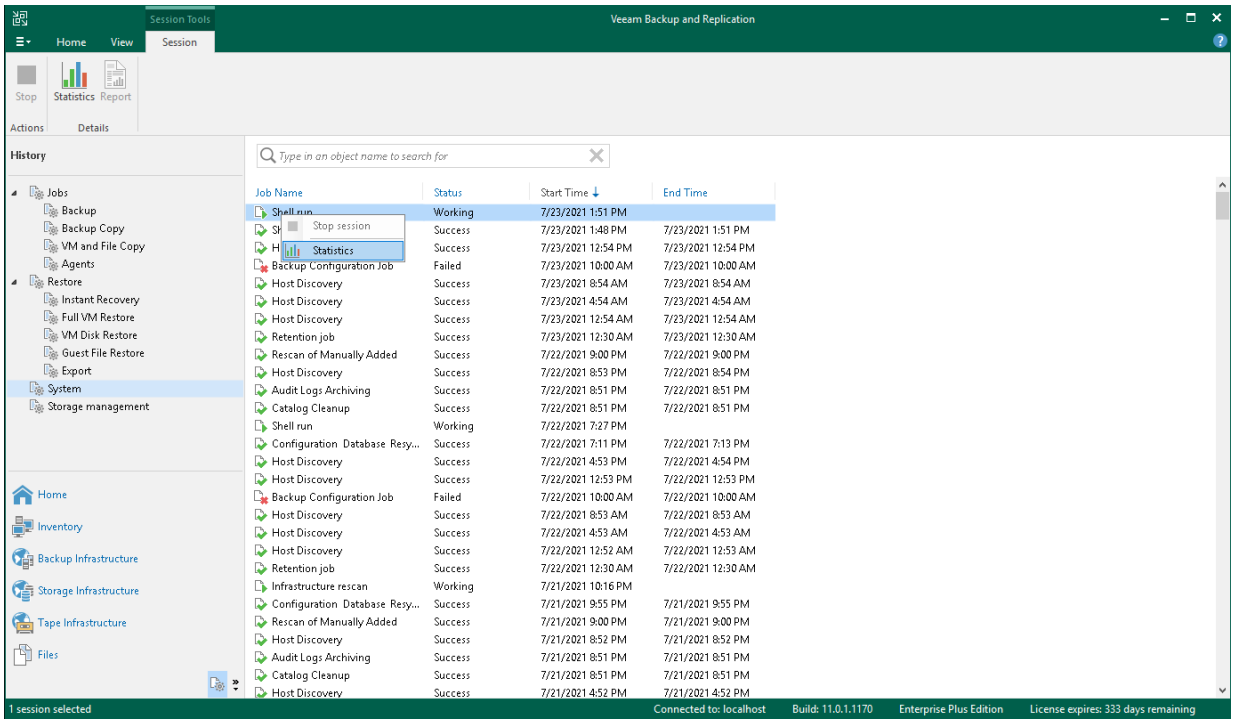
> **IMPORTANT**
>
> Evacuating backup files from an extent formatted with ReFS will hydrate synthetic full backups created by Fast Clone. Synthetic full backups will be evacuated as regular full backups and will take up their full size at a new repository.

After you evacuate backups, you can proceed to removing the performance extent from the scale-out backup repository. For more information, see Removing Performance Extents from Scale-Out Repositories.

Monitoring Evacuating Backups

To monitor backups evacuation, do the following:

1. Open the **History** view.
2. In the inventory pane, click **System**.
3. In the working area, select the evacuation session and click **Statistics** on the ribbon or right-click the evacuation session and select **Statistics**.



Switching to Seal Mode

Veeam Backup & Replication allows you to put any of the scale-out backup repository extents into the *Seal* mode.

Sealing up scale-out backup repository extents allows you to gradually remove data located on these extents by applying a retention policy. You can use this feature to gracefully stop using some of your extents and exclude them from the scale-out backup repository configuration.

After the extent is sealed, no further data is saved to the extent, and only read operations such as restore, merge and remove are allowed.

Backup jobs that are targeted to a scale-out backup repository with the sealed extents that store active backup chains will be forced to create a new active full backup on the next run. However, if backup jobs targeted at such scale-out backup repository were created outside of Veeam Backup & Replication, you must trigger active full backup job manually. The new active full will be saved to another available extent in the scale-out backup repository scope, thereby forming a new active backup chain. The extent to which the new active full is going to be saved is chosen automatically by Veeam Backup & Replication, depending on the available resources.

All restore points that exceed the specified retention period will be continuously removed from the sealed extents on each subsequent backup session.

Consider the following:

- An extent can be put into both the maintenance and the seal modes at the same time.

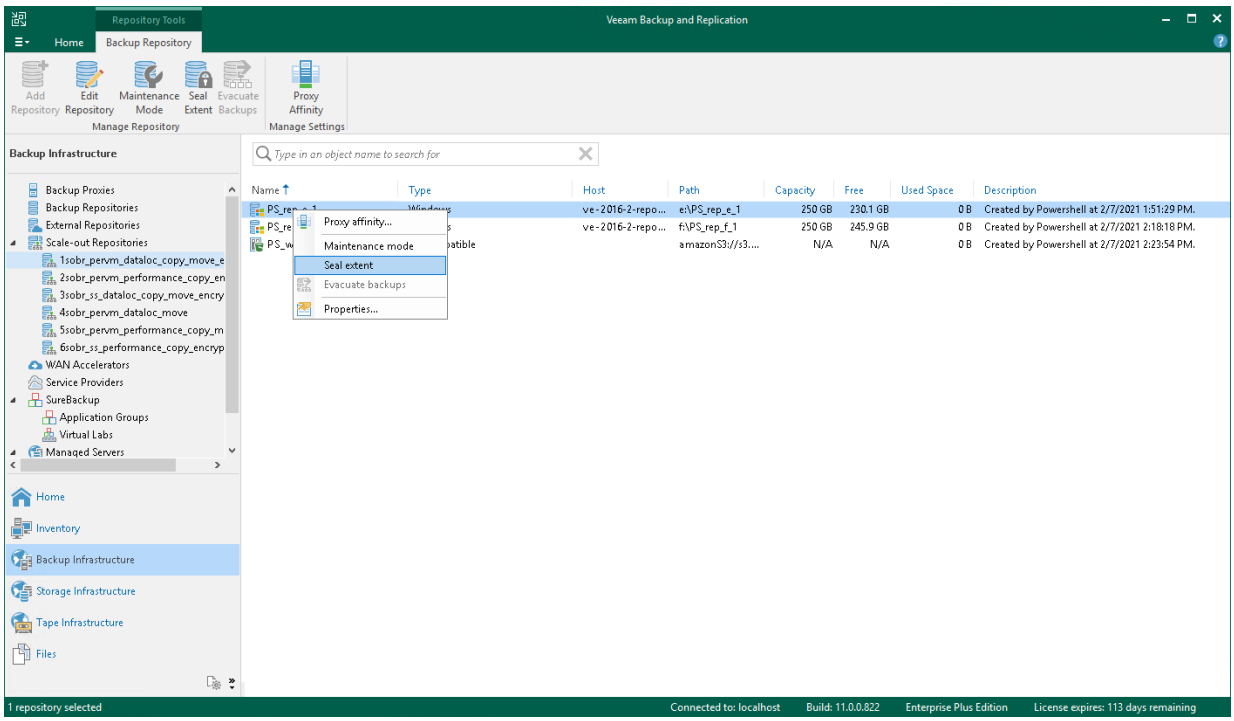
  When both modes are applied, the maintenance mode overrides seal mode.

- When you put an extent into the seal mode, you will still be able to restore, merge and remove data without exiting this mode.

To put an extent into the seal mode:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select the scale-out backup repository under the **Scale-out Repositories node.**
3. In the working area, select the extent and click **Seal Mode** on the ribbon or right-click the extent and select **Seal mode**.

To remove the extent from the seal mode, select the extent and click **Sealed** on the ribbon or right-click the extent and select **Sealed**.



Receiving Scale-Out Backup Repository Reports

Veeam Backup & Replication is capable of sending reports that contain information about processing results of your scale-out backup repositories data.

Consider the following:

- Reports are sent only after you have enabled and configured email notifications, as described in Configuring Global Email Notification Settings.
- Reports are sent daily at time specified in global notification settings.
- Reports are sent for all notification types selected in global notification settings, such as *Success*, *Warning* and *Failure*.
- The title of a report is built up of "*Scale-out Backup Repository*" + a repository name. That said, if your scale-out backup repository name is Amazon, then the report title would be *Scale-out Backup Repository Amazon*.

Each report is divided into sections and contains the following information:

- **Performance Tier** (upper-left) section:
  - **Used Space**. Shows the used disk space of your scale-out backup repository.
  - **Capacity**. Shows the total storage capacity of your scale-out backup repository.
- **Capacity Tier** (upper-right) section:
  - **Used Space**. Shows the occupied storage space in your object storage repository.
  - **Space Limit**. Shows the space limit (if any). A space limit is specified when adding a new object storage repository, as described in Adding Object Storage Repositories.
- **Performance Tier** (middle) section:
  - **Extent**. Shows extents of a scale-out backup repository.
  - **Capacity**. Shows the total storage capacities of your performance extents.
  - **Used Space**. Shows the amount of disk space used on your extents.
  - **Status**. Shows the status of each extent, as described in Description of Report Statuses.
- **Capacity Tier** (lower) section:
  - **Extent**. Shows the name of the capacity extent.
  - **Space Limit**. Shows the space limit (if any).
  - **Used Space**. Shows the occupied storage space in your capacity extent.
  - **Status**. Shows the status of the capacity extent, as described in Description of Report Statuses.

    If an automatic offload job session exits with any status other than *Success*, you will see the associated status message in this field. For more information about the offload job, see Moving Inactive Backup Chains to Object Storage.

Description of Report Statuses

The following table lists possible combinations of *Warning* and *Error* messages shown under the **Status** column of a report.

If none of the conditions listed in the **Extent state** column is true, then the report status will be shown as *Success*.

Extent type	Extent state	Status message	Report type
Performance tier	Maintenance mode	Maintenance mode	Warning
	Threshold limit exceeded. Threshold is specified in the **Backup storage** section, as described in Specifying Other Notification Settings.	Reaching capacity	Warning
	Unavailable	Offline	Error
Capacity tier	Maintenance mode	Maintenance mode	Warning
	Space limit exceeded. Space limit is specified when adding a new object storage repository, as described in Adding Object Storage Repositories.	Out of capacity	Error
	Unavailable	Offline	Error
	Threshold limit exceeded. Threshold is specified in the **Backup storage** section, as described in Specifying Other Notification Settings.	Reaching capacity	Warning

Report Examples

Success Reports

The following figure shows an example of a report consisting of two performance extents (*Backup Volume 01* and *Backup Volume 02*); both share *253.3 GB* of storage capacity, of which *52.4 GB* is occupied.

Both extents have OK status, which means that neither extent was put into the maintenance mode, nor has any of these performance extents exceeded the allowed threshold limit.

This report also includes the **Capacity Tier** section consisting of a capacity extent with no **Space Limit** applied. This capacity extent stores *29.6 GB* of data and has the *OK* status.

Scale-out backup repository Amazon — **Success**
Remote Object Storage

Wednesday, December 19, 2018 4:04:12 AM

Performance Tier		Capacity Tier	
Used Space	52.4 GB	**Used Space**	29.6 GB
Capacity	253.3 GB (70% free)	**Space Limit**	Not set

Performance Tier

Extent	Capacity	Used Space	Status
Backup Volume 01	126.7 GB (67% free)	30.3 GB	OK
Backup Volume 02	126.7 GB (73% free)	22.1 GB	OK

Capacity Tier

Extent	Space Limit	Used Space	Status
Amazon S3 Object Storage	Not set	29.6 GB	OK

Warning Reports

The figure below demonstrates a report with the *Warning* status.

As per example, the *Backup Volume 01* performance extent has been put into the maintenance mode, and the *Backup Volume 02* performance extent has exceeded the allowed threshold both of which have caused a report to be generated with the *Warning* status.

Scale-out backup repository Amazon — **Warning**
Remote Object Storage

Wednesday, December 19, 2018 4:58:35 AM

Performance Tier		Capacity Tier	
Used Space	82.0 GB	**Used Space**	29.6 GB
Capacity	253.3 GB (58% free)	**Space Limit**	Not set

Performance Tier

Extent	Capacity	Used Space	Status
Backup Volume 01	126.7 GB (43% free)	60.0 GB	Maintenance mode
Backup Volume 02	126.7 GB (73% free)	22.1 GB	Reaching capacity

Capacity Tier

Extent	Space Limit	Used Space	Status
Amazon S3 Object Storage	Not set	29.6 GB	OK

Error Reports

In the figure below a report has been generated with the *Error* status caused by the *Amazon S3 Object Storage* performance extent which has exceeded its allowed space limit.

Scale-out backup repository Amazon — **Error**
Remote Object Storage

Thursday, December 20, 2018 3:55:02 AM

Performance Tier		Capacity Tier	
Used Space	82.0 GB	**Used Space**	29.6 GB
Capacity	253.3 GB (58% free)	**Space Limit**	2.0 GB (0% free)

Performance Tier

Extent	Capacity	Used Space	Status
Backup Volume 01	126.7 GB (43% free)	60.0 GB	OK
Backup Volume 02	126.7 GB (73% free)	22.1 GB	OK

Capacity Tier

Extent	Space Limit	Used Space	Status
Amazon S3 Object Storage	2.0 GB (0% free)	29.6 GB	Out of capacity

Removing Backups from Capacity or Archive Tier

To remove moved or copied backups from capacity or archive extent, use the **Delete from disk** feature, as described in Deleting Backups from Object Storage Repositories.

Consider the following:

- When removing offloaded backup files from the backup chain that was created with the per-machine method, the associated blocks of data will be removed from the capacity or archive extent altogether.

  For more information about per-machine backups, see Per-Machine Backup Files.

- When removing offloaded backup files from the backup chain that was created as a single storage backup file, then nothing will be removed until either of the following occurs:
  - All the VMs were removed from the backup.
  - The backup itself was removed.
- Immutable backups cannot be removed.

  For more information, see Immutability.

- If the capacity or archive extent has been put into the maintenance mode, the removal of data from such a repository is not possible until the extent is removed from the maintenance mode.

  For more information, see Switching to Maintenance Mode.

- [For capacity tier] When backup is removed from the capacity extent, associated indexes will be removed as well.

  For more information, see Indexes.

- During data removal, the entire folder structure starting from the repository folder (<backup_id>) will be completely purged.

  For more information on how Veeam Backup & Replication stores data in the capacity extent, see Capacity Extent Structure.

- If backup files with metadata that are located on your extents have been removed locally in any way other than by using the Deleting from Disk feature, Veeam Backup & Replication will not be able to synchronize the backup chain state with that of the capacity or archive extent. Therefore, the offloaded blocks of data will continue to remain in cloud storage. To remove such blocks, use your cloud platform abilities.

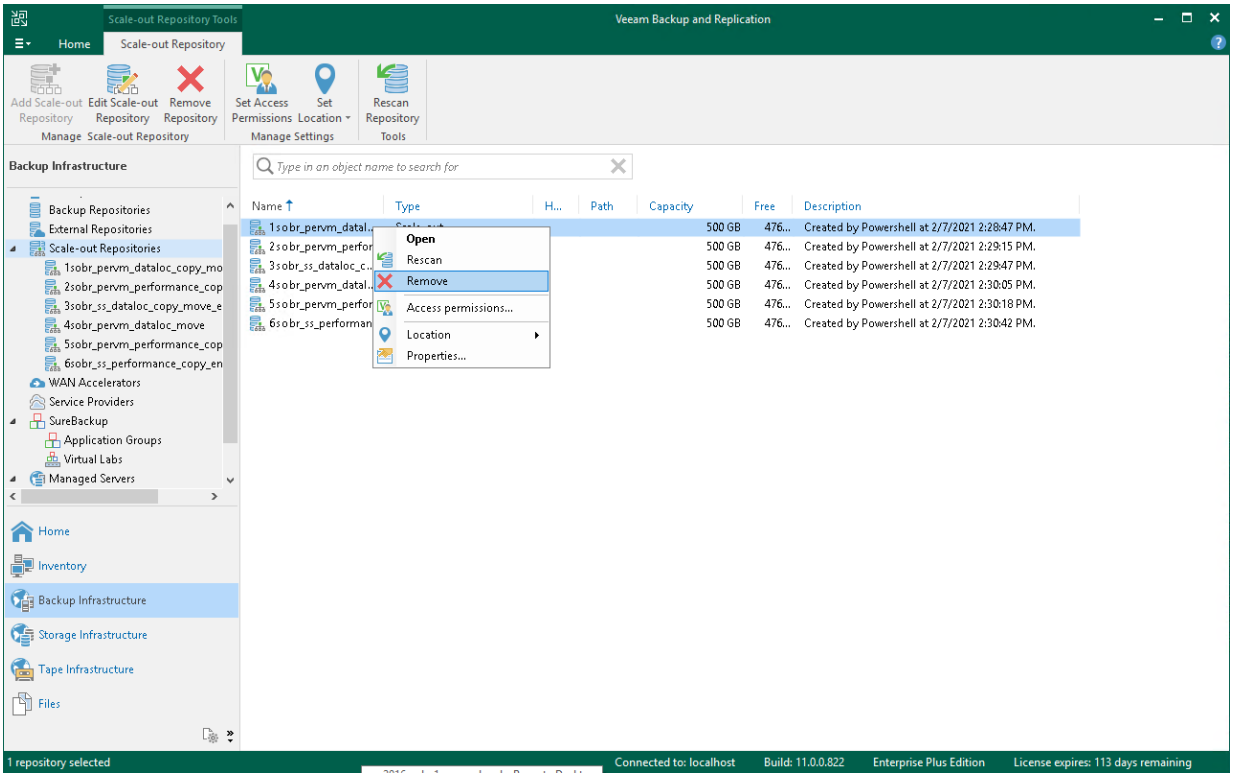
Removing Scale-Out Backup Repositories

You can remove the scale-out backup repository at any time. When you remove the scale-out backup repository, Veeam Backup & Replication unassigns the performance extent role from all the backup repositories configured into it, and they become individual backup repositories. Backup files are not removed from backup repositories — they remain on the disk.

You cannot remove a scale-out backup repository if at least one job is targeted at it. First, you need to retarget jobs to another backup repository in the backup infrastructure.

To remove a scale-out backup repository:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, click **Scale-out Repositories**.
3. In the working area, select the scale-out repository and click **Remove Repository** on the ribbon or right-click the backup repository and select **Remove**.



Guest Interaction Proxy

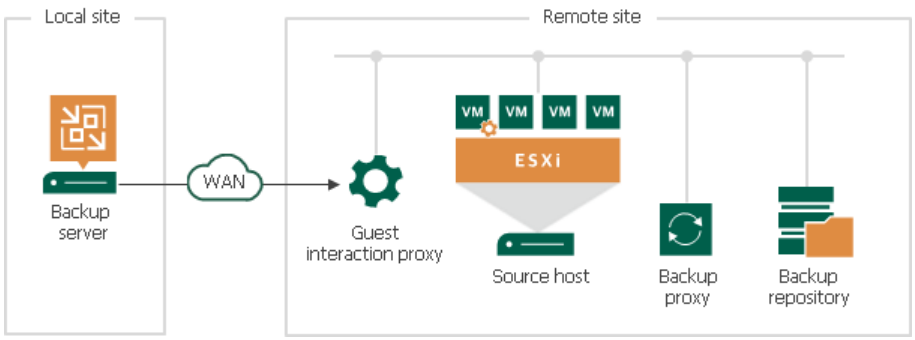
The guest interaction proxy is a backup infrastructure component that sits between the backup server and processed VM. This component is needed if the backup or replication jobs perform the following processing of VMs:

- Application-aware processing
- Guest file system indexing
- Transaction logs processing

To interact with the VM guest OS, Veeam Backup & Replication needs either to install non-persistent runtime components or use (if necessary, install) persistent agent components in each VM. The task of deploying these components in a VM is performed by the guest interaction proxy. For more information on the components, see Non-Persistent Runtime Components and Persistent Agent Components.

> **NOTE**
>
> The guest interaction proxy functionality is included in the Veeam Universal License. When using a legacy socket-based license, Enterprise or higher edition is required.

The guest interaction proxy allows you to communicate with the VM guest OS even if the backup server and processed VM run in different networks.



> **IMPORTANT**
>
> The guest interaction proxy deploys the non-persistent runtime components or persistent agent components only in Microsoft Windows VMs. In VMs with another guest OS, the non-persistent runtime components or persistent agent components are deployed by the backup server.

Guest Interaction Proxy Deployment

You can use multiple guest interaction proxies to improve performance. Multiple guest interaction proxies will deploy non-persistent runtime components or persistent agent components in VMs faster compared to the same operation performed by one guest interaction proxy.

In a backup infrastructure with multiple remote sites, you can deploy a guest interaction proxy in each site. This can reduce load on the backup server and produce less traffic between the backup server and remote site.

Requirements for Guest Interaction Proxy

A machine performing the role of a guest interaction proxy must meet the following requirements:

- The role of a guest interaction proxy can be assigned to a Microsoft Windows server (physical or virtual).
- You must add the machine to the Veeam Backup & Replication console as a managed server.
- Guest interaction proxy must have either a LAN or VIX connection to the VM that will be processed. You do not have to set up both connections — only one connection is required. For more information about setting up a connection to the VM, see this Veeam KB article.

The guest interaction proxy role can be performed by any machine that meets the requirements, including backup proxy, backup repository, WAN accelerator or backup server.

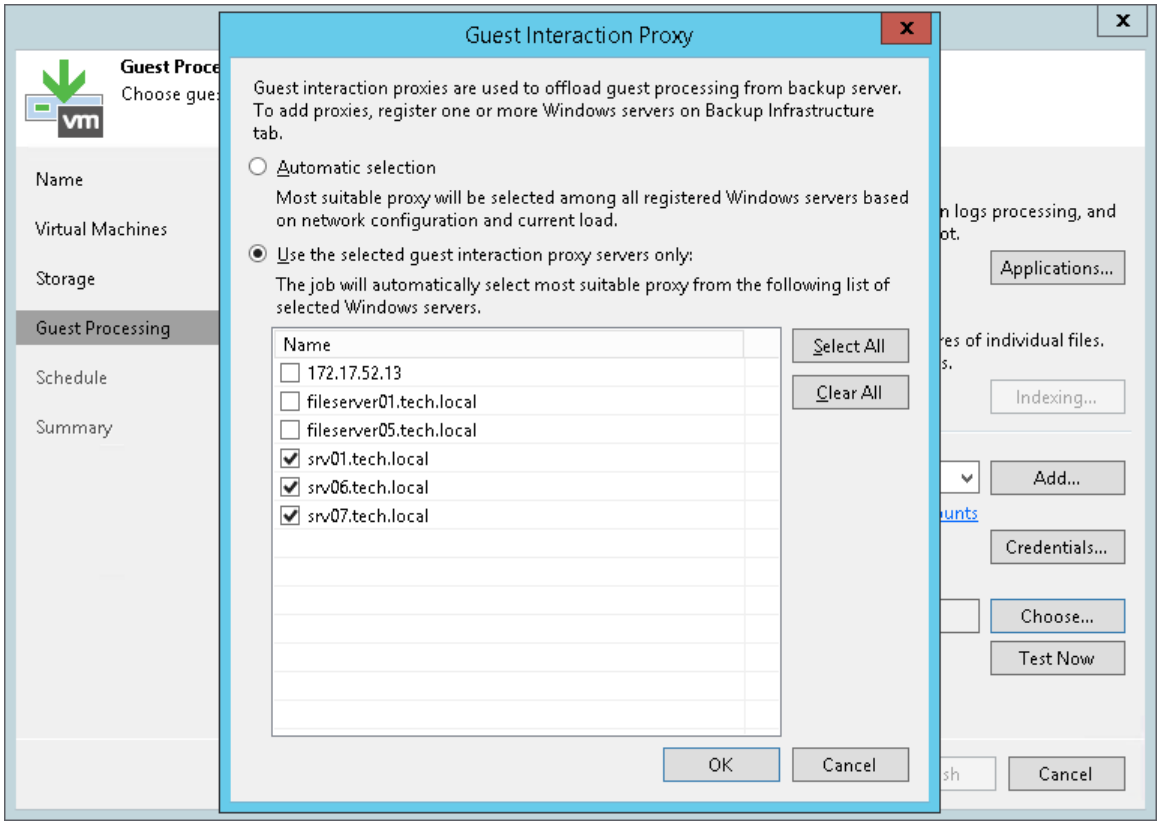
Guest Interaction Proxy Selection

When you add a Microsoft Windows machine to the backup infrastructure, Veeam Backup & Replication deploys Veeam Data Mover on it. Veeam Data Mover includes the components responsible for deployment of non-persistent runtime components or persistent agent components during guest OS interaction.

To assign a guest interaction proxy for the job, you must select a Microsoft Windows machine that will perform the role of the guest interaction proxy at the **Guest Processing** step of the backup or replication job wizard. You can assign the guest interaction proxy manually, or let Veeam Backup & Replication do it automatically. Veeam Backup & Replication uses the following priority rules to select the guest interaction proxy:

1. A machine in the same network as the protected VM that does not perform the backup server role.
2. A machine in the same network as the protected VM that performs the backup server role.
3. A machine in another network that does not perform the backup server role.
4. A machine in another network that performs the backup server role.

If Veeam Backup & Replication finds several available machines of equal priority, it selects the less loaded machine. The load is defined by the number of tasks that the machine already performs.



Failover from Guest Interaction Proxy to Backup Server

If the guest interaction proxy fails to connect to a Microsoft Windows VM, the guest interaction proxy will not be able to access the VM and deploy non-persistent runtime components or persistent agent components in it. In this case, the backup server will take over the role of guest interaction proxy and deploy the non-persistent runtime components or persistent agent components in the VM.

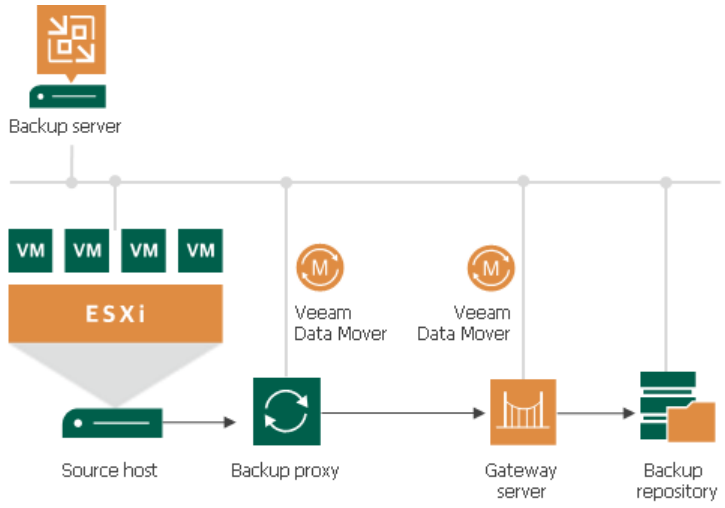
Gateway Server

A gateway server is an auxiliary backup infrastructure component that "bridges" the backup server and backup repository. The gateway server is required if you deploy the following types of backup repositories in the backup infrastructure:

- Shared folder backup repositories
- Dell EMC Data Domain deduplicating storage appliance
- HPE StoreOnce deduplicating storage appliance

Such backup repositories cannot host Veeam Data Movers— Veeam components that establish a connection between a backup proxy and backup repository (in case of backup jobs) or between backup repositories (in case of backup copy jobs). To overcome this limitation, Veeam Backup & Replication uses gateway servers.

In the backup infrastructure, a gateway server hosts the target Veeam Data Mover. Veeam Backup & Replication establishes a connection between the source Veeam Data Mover and target Veeam Data Mover, and transports data from/to backup repositories through gateway servers.



Gateway Servers Deployment

The role of a gateway server can be assigned to a Microsoft Windows machine added to the backup infrastructure. To configure a gateway server, you must first add a machine that you plan to use as a gateway server to the backup infrastructure using the **New Windows Server wizard**. After that, you must go through the **Add New Backup Repository** wizard and define gateway server settings. You can select a gateway server explicitly or instruct Veeam Backup & Replication to select it automatically.

If you select a gateway server explicitly, Veeam Backup & Replication uses the selected machine as a gateway server and performs synthetic operations on it. The machine to which you assign the role of a gateway server must be located as close to the backup repository as possible. However, if you use a deduplicating storage appliance with source-side data deduplication, it is reasonable to assign the role of a gateway server to a machine that is located closer to the backup proxy. This will help you reduce the amount of traffic traveling over the network. For more information, see Dell EMC Data Domain and HPE StoreOnce.

If you instruct Veeam Backup & Replication to select the gateway server automatically, Veeam Backup & Replication uses the following backup infrastructure components:

Type of job	Component used as gateway server	Component used as gateway server for synthetic operations
Backup job	Backup proxy that was assigned first to process VM data for a backup job.*	Synthetic operations are performed on the mount server associated with the backup repository. If the mount server is not accessible, Veeam Backup & Replication fails over to the backup server.
Backup copy job	• Direct data path: mount server associated with the backup repository. If the mount server is not accessible, Veeam Backup & Replication fails over to the backup server. • Over WAN accelerators: source and/or target WAN accelerator (depending on the shared folder backup repository location).	Synthetic operations are performed on the mount server associated with the backup repository. If the mount server is not accessible, Veeam Backup & Replication fails over to the backup server. These rules are applied to the direct data path and processing over WAN accelerators.
Tape job	If there is a direct connection between a backup repository and tape device, the role of a gateway server is assigned to the tape server. Otherwise, the role of a gateway server is assigned to the backup server.	Synthetic operations are performed on the mount server associated with the backup repository. If the mount server is not accessible, Veeam Backup & Replication fails over to the backup server.
Veeam Agent backup job	Backup server.	Synthetic operations are performed on the backup server.
Backup job created by Veeam Plug-in for Oracle RMAN/SAP HANA/SAP on Oracle	Backup server.	—
Restore operations	Backup proxy used for a restore operation*.	—
Replication from backup	Target backup proxy assigned for a replication operation*.	—

* If the backup proxy is a Linux machine, the role of the gateway server is assigned to the mount server associated with the backup repository. If the mount server is not accessible, Veeam Backup & Replication fails over to the backup server.

When you enable automatic gateway selection, Veeam Backup & Replication may use one or several gateway servers to process VMs in backup jobs. The number of gateway servers depends on the backup repository settings:

- If the **Use per-machine backup files** option is disabled, Veeam Backup & Replication selects one gateway server for the whole backup repository.
- If the **Use per-machine backup files** option is enabled, Veeam Backup & Replication selects a gateway server per every VM in the job.

For example, a backup job processes 2 VMs. The job is targeted at a backup repository for which the **Use per-machine backup files** option is enabled. In this case, Veeam Backup & Replication will detect which backup proxies were used to process VMs in the job. If VMs were processed with 2 different backup proxies, Veeam Backup & Replication will assign the role of gateway servers to these backup proxies. If VMs were processed with the same backup proxy, Veeam Backup & Replication will assign the role of a gateway server to this backup proxy, and will use it for both VMs in the job.

For scale-out backup repositories, Veeam Backup & Replication uses one gateway server per every extent. The rules of gateway server selection are described above.

Requirements for Gateway Servers

A machine that performs the role of a gateway server must meet the following requirements:

- The gateway server can run on a Microsoft Windows or Linux machine. The machine must meet the system requirements. For more information, see System Requirements.
- You must add the machine to the Veeam Backup & Replication console as a managed server.
- The machine must have access to the backup repository — shared folder, Dell EMC Data Domain or HPE StoreOnce.

Limitations for Gateway Servers

The following limitations apply to a machine that performs the role of a gateway server:

- For deduplicating storage appliances working over Fibre Channel, you must explicitly select a gateway server that will communicate with the appliance. As a gateway server, you must use a Microsoft Windows machine that is added to the backup infrastructure and has access to the appliance over Fibre Channel.
- For HPE StoreOnce deduplicating storage appliances, you must assign the role of a gateway server to a 64-bit machine.

Related Topics

- Scale-Out Backup Repository
- Specifying Server or Shared Folder Settings

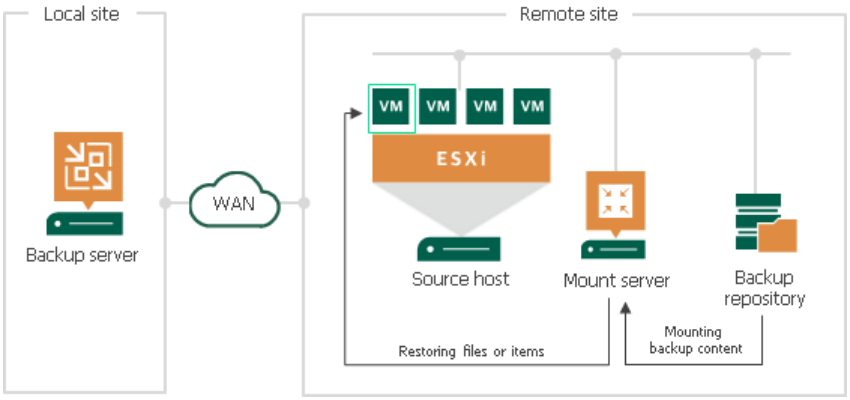
Mount Server

The mount server is a server required for restores that work with guest OS files and application items. To access files or items stored in a backup file, Veeam Backup & Replication mounts the content of the backup to the mount server. Only after the content is mounted, Veeam Backup & Replication can get files and copy them to the restore destination.

The mount server is required if you perform the following operations:

- Guest OS file recovery
- Application items restore
- Secure restore
- Instant file share recovery

To reduce the load on the network and speed up the restore process, the mount server must be located in the same site as the backup repository where backup files are stored. In this case, you will be able to keep the traffic in one site. If the mount server is located in some other site, the data will travel across the network between the sites.



NOTE

In some scenarios, Veeam Backup & Replication can mount content of backups to machines other than mount servers. For more information, see File-Level Restore Scenarios and Mount Points.

Mount Server Deployment

The mount server is created for every backup repository and is associated with it. When you configure a backup repository, you specify to which server you want to assign the role of the mount server.

You can assign the mount server role to any 64-bit Microsoft Windows machine added to the backup infrastructure. This machine and the backup repository must be located as close to each other as possible. If you have several sites, we recommend you configure at least one mount server in each site.

By default, Veeam Backup & Replication suggests assigning the mount server role to the following infrastructure components depending on the OS of the backup repository:

- For Microsoft Windows backup repositories, Veeam Backup & Replication suggests the backup repository server itself.

- For Linux, shared folder backup repositories and deduplicating storage appliances, Veeam Backup & Replication suggests the backup server.

> **NOTE**
>
> Mind the following:
>
> - For scale-out backup repositories, you must specify the mount server for every extent.
> - For cloud repositories and hosts that store replicas or backups from storage snapshots, the mount server role is assigned to the backup server. For such repositories, you cannot assign the mount server role to a different machine.

Mount Server Services and Components

Mount servers run light-weight services that take a few seconds to deploy. Deployment is fully automated. Veeam Backup & Replication installs the following services:

- **Veeam Mount Service** mounts backups and replicas for file-level access, browsing the guest file system and restoring guest OS files and application items.
- **Veeam Data Mover** handles traffic.
- **Veeam vPower NFS Service** (if you enable it when configuring the mount server).

Requirements for Mount Servers

A machine that performs the role of a mount server must meet the following requirements:

- The machine must meet the system requirements. For more information, see System Requirements.
- The mount server must have access to the backup repository with which it is associated and to the original VM (the VM to which you restore files or application items). For restore from storage snapshots, the mount server must also have access to the ESXi host on which the temporary VM is registered.

Veeam Data Movers

Veeam Data Mover performs data processing tasks on behalf of Veeam Backup & Replication, such as retrieving source machine data, performing data deduplication and compression, and storing backed-up data on the target storage.

For Microsoft Windows servers, Veeam Data Movers are persistent, that is, Veeam Data Mover is uploaded and installed on a server only once. Veeam Backup & Replication automatically installs Veeam Data Mover when you add a Microsoft Windows server to the backup infrastructure.

For Linux servers, Veeam Data Movers can be persistent or non-persistent. Non-persistent Veeam Data Mover is uploaded and removed each time Veeam Backup & Replication addresses a server.

For Veeam Data Mover to be persistent, you must specify an account with root or equivalent to root permissions when adding a Linux server. Persistent Veeam Data Movers are required for the following backup infrastructure components:

- Hardened repositories
- Backup proxies

For other backup infrastructure components based on Linux servers, Veeam Data Movers can be persistent or non-persistent. If you do not want to provide root or equivalent to root permissions, specify an account with non-root permissions. In this case, Veeam Data Movers will be non-persistent. Veeam Backup & Replication will upload and start Veeam Data Movers through the SSH connection when Veeam Backup & Replication addresses the server.

> **NOTE**
>
> In Veeam Backup & Replication version 11a, persistent Veeam Data Movers are required for Linux servers with the backup proxy role assigned. When you upgrade to Veeam Backup & Replication version 11a, Veeam Backup & Replication automatically updates all Linux backup proxies and installs Veeam Data Movers on them. Veeam Backup & Replication does not automatically update Linux servers that have other roles. If you want to install Veeam Data Movers on such servers, open the Edit Linux Server wizard for the necessary server and click **Finish**.

Requirements and Limitations for Veeam Data Movers

Before you use Veeam Data Movers, consider the following requirements and limitations for Linux and Microsoft Windows servers:

- For Linux servers:
  - Linux server version must be 64-bit. For more information, see System Requirements. Note that Perl is required only for non-persistent Veeam Data Movers. Check the full list of required Perl modules in this Veeam KB article.
  - Veeam Backup & Replication does not deploy Veeam Data Movers on deduplicating storage appliances based on Linux.
  - If you want to use persistent Veeam Data Movers, the user account specified for the server must have root or elevated to root permission. Otherwise, Veeam Data Movers will be non-persistent, that is, the Linux server will not host Veeam Data Movers permanently.
  - When you add a Linux server using single-use credentials, the user account must still have elevated to root permission in order for Veeam Data Movers to be persistent.

- For Microsoft Hyper-V and Microsoft Windows servers: file and printer sharing must be enabled in network connection settings of the added server.

Veeam vPower NFS Service

The vPower technology enables the following features:

- Recovery verification
- Instant Recovery
- Instant Disk Recovery
- Staged restore
- Universal Application-Item Recovery (U-AIR)
- Multi-OS file-level restore

The key construct of the vPower technology is the vPower NFS Service. The vPower NFS Service is a Microsoft Windows service that runs on a Microsoft Windows machine and enables this machine to act as an NFS server.

On the vPower NFS server, Veeam Backup & Replication creates a special directory — the vPower NFS datastore. When you start a VM or a VM disk from a backup, Veeam Backup & Replication "publishes" VMDK files of the VM from the backup on the vPower NFS datastore. Technically, Veeam Backup & Replication emulates the presence of VMDK files on the vPower NFS datastore — the VMDK files themselves are still located in the backup file in the backup repository.

The vPower NFS datastore is then mounted to the ESXi host. As a result, the ESXi host can "see" backed-up VM images with the help of the vPower NFS datastore and work with them as with regular VMDK files. The emulated VMDK files function as pointers to the real VMDK files in the backup repository.

> **IMPORTANT**
>
> Veeam vPower NFS datastores are service datastores that can be used for vPower operations only. You cannot use them as regular VMware vSphere datastores — for example, you cannot place files of replicated VMs on such datastores.

vPower NFS Server Location

If you store backups on a Microsoft Windows backup repository, it is strongly recommended that you enable the vPower NFS server on this backup repository. In this case, Veeam Backup & Replication will be able to set up a direct connection between the backup repository and ESXi to which the vPower NFS datastore is mounted.

The Veeam vPower NFS Service can also run on any Microsoft Windows server in the backup infrastructure, including the backup server itself. However, in this case the recovery verification performance may decrease. The connection between the ESXi host and backup repository will be split into two parts:

1. From ESXi host to the vPower NFS server
2. From the vPower NFS server to the backup repository

vPower-Specific Settings

To establish a connection between the ESXi host and vPower NFS server, you must make sure that the ESXi host has a proper network interface configuration and can access the vPower NFS server.

> **NOTE**
>
> Veeam Backup & Replication uses IP address authorization to restrict access to the vPower NFS server. By default, the vPower NFS server can be accessed only by the ESXi host that provisioned the vPower NFS datastore. However, you can disable this option with a registry key. For more information, contact Veeam Customer Support.

When connecting to the vPower NFS server, the ESXi host uses a VMkernel interface. For this reason, the ESXi host must have a VMkernel interface. Otherwise, Veeam Backup & Replication will fail to mount the vPower NFS datastore on the ESXi host.

WAN Accelerators

WAN accelerators are dedicated components that Veeam Backup & Replication uses for WAN acceleration. WAN accelerators are responsible for global data caching and data deduplication. For more information, see WAN Acceleration.

Log Shipping Servers

Log shipping servers are dedicated components that Veeam Backup & Replication uses for backup of Microsoft SQL Server transaction logs and Oracle archive logs. For more information, see Microsoft SQL Server Logs Backup and Oracle Logs Backup.

Tape Servers

Tape servers are dedicated components responsible for transferring data between data source and tape device. For more information, see the Veeam Backup & Replication User Guide.

NDMP Servers

If your NAS device supports the NDMP protocol, you can back up data from it to tape. To do this, you need to add the NAS device as an NDMP server. For more information, see the Veeam Backup & Replication User Guide.

Veeam Backup Enterprise Manager

Veeam Backup Enterprise Manager is an optional component intended for distributed enterprise environments with multiple backup servers. Veeam Backup Enterprise Manager federates backup servers and offers a consolidated view of these servers through a web browser interface. You can centrally control and manage all jobs through a single "pane of glass", edit and clone jobs, monitor job state and get reporting data across all backup servers. Veeam Backup Enterprise Manager also enables you to search for VM guest OS files in all current and archived backups across your backup infrastructure, and restore these files in one click.

Veeam Backup Enterprise Manager Deployment

Veeam Backup Enterprise Manager can be installed on a physical or virtual machine. You can deploy it on the backup server or use a dedicated machine.

Veeam Backup Enterprise Manager Services and Components

Veeam Backup Enterprise Manager uses the following services and components:

- **Veeam Backup Enterprise Manager Service** coordinates all operations of Veeam Backup Enterprise Manager, aggregates data from multiple backup servers and provides control over these servers.
- **Veeam Backup Enterprise Manager Configuration Database** is used by Veeam Backup Enterprise Manager for storing data. The database instance can be located on a SQL Server installed either locally (on the same machine as Veeam Backup Enterprise Manager Server) or remotely.
- **Veeam Guest Catalog Service** replicates and consolidates VM guest OS file system indexing data from backup servers added to Veeam Backup Enterprise Manager. Index data is stored in Veeam Backup Enterprise Manager Catalog (a folder on the Veeam Backup Enterprise Manager Server) and is used to search for VM guest OS files in backups created by Veeam Backup & Replication.

Deployment Scenarios

Veeam Backup & Replication can be used in virtual environments of any size and complexity. The architecture of the solution supports on-site and off-site data protection, operations across remote sites and geographically dispersed locations. Veeam Backup & Replication provides flexible scalability and easily adapts to the needs of your virtual environment.

Before you install Veeam Backup & Replication, familiarize yourself with common deployment scenarios and carefully plan your backup infrastructure layout.

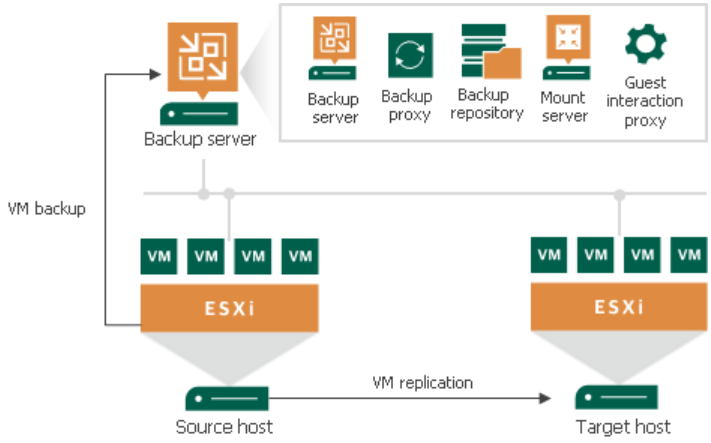
Simple Deployment

The simple deployment scenario is good for small virtual environments or the purpose of the Veeam Backup & Replication evaluation. In this scenario, Veeam Backup & Replication and all services needed for data protection tasks are installed on a single Windows-based machine.

> **NOTE**
>
> If you decide to use the simple deployment scenario, it is recommended that you install Veeam Backup & Replication on a VM. This will enable you to use the Virtual appliance transport mode and, as a result, LAN-free data transfer. For details, see Transport Modes.

The machine where Veeam Backup & Replication is installed performs the following roles:

- Backup server that coordinates all jobs, controls their scheduling and performs other administrative activities.
- Default backup proxy that handles job processing and transfers backup traffic.
- Default backup repository where backup files are stored. During installation, Veeam Backup & Replication checks volumes of the machine on which you install the product and identifies a volume with the greatest amount of free disk space. On this volume, Veeam Backup & Replication creates the `Backup` folder that is used as the default backup repository.
- Mount server that is needed for restoring of VM guest OS files.
- Guest interaction proxy that is needed for application-aware processing, guest file system indexing and transaction log processing.



Veeam Backup & Replication is ready for use right after the installation. The only thing you must do is add VMware vSphere servers that you plan to use as source and target for backup, replication and other activities. For details, see Adding VMware vSphere Servers.

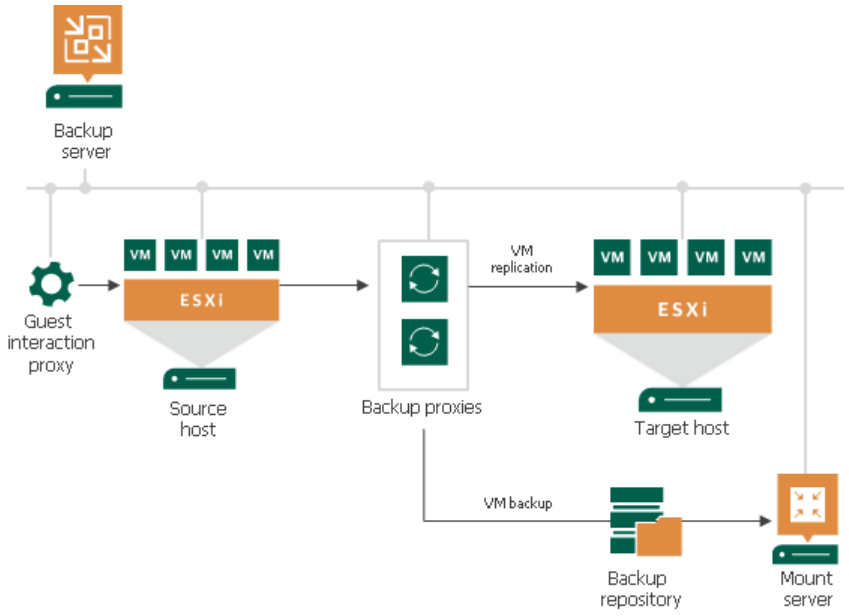
The drawback of a simple deployment scenario is that all data is handled and stored on the backup server. For medium-sized or large-scale environments, the capacity of a single backup server may not be enough. To take the load off the backup server and balance it throughout your backup infrastructure, we recommend that you use the advanced deployment scenario. For details, see Advanced Deployment.

Advanced Deployment

In large-scale virtual environments with a large number of jobs, the load on the backup server is heavy. In this case, it is recommended that you use the advanced deployment scenario that moves the backup workload to dedicated backup infrastructure components. The backup server here functions as a "manager" for deploying and maintaining backup infrastructure components.

The advanced deployment includes the following components:

- Virtual infrastructure servers — VMware vSphere hosts used as source and target for backup, replication and VM copy.
- Backup server — a configuration and control center of the backup infrastructure.
- Backup proxy — a "data mover" component used to retrieve VM data from the source datastore, process it and deliver to the target.
- Backup repository — a location used to store backup files, VM copies and auxiliary replica files.
- Dedicated mount servers — component required for VM guest OS files and application items restore to the original location.
- Dedicated guest interaction proxies — components used to deploy the non-persistent runtime components or persistent agent components in Microsoft Windows VMs.



With the advanced deployment scenario, you can easily meet your current and future data protection requirements. You can expand your backup infrastructure horizontally in a matter of minutes to match the amount of data you want to process and available network throughput. Instead of growing the number of backup servers or constantly tuning job scheduling, you can install multiple backup infrastructure components and distribute the backup workload among them. The installation process is fully automated, which simplifies deployment and maintenance of the backup infrastructure in your virtual environment.

In virtual environments with several proxies, Veeam Backup & Replication dynamically distributes backup traffic among these proxies. A job can be explicitly mapped to a specific proxy. Alternatively, you can let Veeam Backup & Replication choose the most suitable proxy. In this case, Veeam Backup & Replication will check settings of available proxies and select the most appropriate one for the job. The proxy server to be used should have access to the source and target hosts as well as to the backup repository to which files will be written.

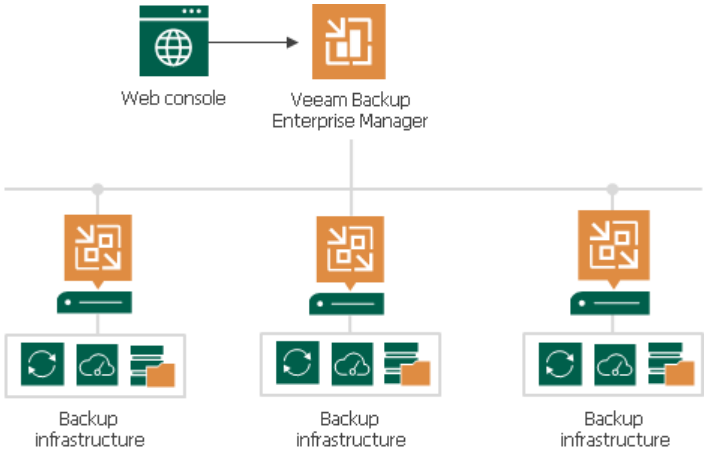
The advanced deployment scenario can be a good choice for backing up and replicating off-site. You can deploy a backup proxy in the production site and another one in the disaster recovery (DR) site, closer to the backup repository. When a job is performed, backup proxies on both sides establish a stable connection, so this architecture also allows for efficient transport of data over a slow network connection or WAN.

To regulate backup load, you can specify the maximum number of concurrent tasks per proxy and set up throttling rules to limit proxy bandwidth. The maximum number of concurrent tasks can also be specified for a backup repository in addition to the value of the combined data rate for it.

Another advantage of the advanced deployment scenario is that it contributes to high availability — jobs can migrate between proxies if one of them becomes overloaded or unavailable.

Distributed Deployment

The distributed deployment scenario is recommended for large geographically dispersed virtual environments with multiple Veeam Backup & Replication servers installed across different sites. These backup servers are federated under Veeam Backup Enterprise Manager — an optional component that provides centralized management and reporting for these servers through a web interface.



Veeam Backup Enterprise Manager collects data from backup servers and enables you to run backup and replication jobs across the entire backup infrastructure through a single web console, edit them and clone jobs using a single job as a template. It also provides reporting data for various areas (for example, all jobs performed within the last 24 hours or 7 days, all VMs engaged in these jobs and so on). Using indexing data consolidated on one server, Veeam Backup Enterprise Manager provides advanced capabilities to search for VM guest OS files in VM backups created on all backup servers (even if they are stored in backup repositories on different sites), and recover them in a single click. Search for VM guest OS files is enabled through Veeam Backup Enterprise Manager itself.

With flexible delegation options and security roles, IT administrators can delegate the necessary file restore or VM restore rights to authorized personnel in the organization – for example, allow database administrators to restore Oracle or Microsoft SQL Server VMs.

If you use Veeam Backup Enterprise Manager in your backup infrastructure, you do not need to install licenses on every backup server you deploy. Instead, you can install one license on the Veeam Backup Enterprise Manager server and it will be applied to all servers across your backup infrastructure. This approach simplifies tracking license usage and license updates across multiple backup servers.

In addition, VMware administrators will benefit from Veeam plug-in for vSphere Web Client that can be installed using Veeam Backup Enterprise Manager. They can analyze cumulative information on used and available storage space view and statistics on processed VMs, review success, warning, failure counts for all jobs, easily identify unprotected VMs and perform capacity planning for repositories, all directly from vSphere.

Resource Scheduling

Veeam Backup & Replication has the built-in mechanism of resource scheduling. Resource scheduling lets Veeam Backup & Replication automatically define what backup infrastructure resources are required for data protection and disaster recovery jobs and tasks, select optimal resources and assign them for the jobs and tasks.

Resource scheduling is performed by the Veeam Backup Service running on the backup server. When a job or task starts, it communicates with the service and informs it about the resources it needs. The service analyzes job settings, parameters specified for backup infrastructure components, current load on the components, and automatically allocates optimal resources to the job.

For resource scheduling, Veeam Backup Service uses the following settings and features:

- Limitation of Concurrent Tasks
- Limitation of Read and Write Data Rates for Backup Repositories
- Network Traffic Management
- Performance Bottlenecks

Limitation of Concurrent Tasks

When you start a data protection or disaster recovery job, Veeam Backup & Replication analyzes the list of VMs added to the job, and creates a separate task for every disk of every VM to be processed.
Veeam Backup & Replication then defines what backup infrastructure components must be used for the job, checks what backup infrastructure components are currently available, and assigns necessary components to process the created job tasks.

Backup infrastructure components typically process several tasks at the same time. You can limit the number of tasks that backup infrastructure components must process concurrently. Task limitations helps you balance the workload across the backup infrastructure and avoid performance bottlenecks.

Veeam Backup & Replication lets you limit the number of concurrent tasks for the following backup infrastructure components:

- Backup proxies
- Backup repositories

> **NOTE**
>
> Task limits set for backup infrastructure components influence the job performance. For example, you add a VM with 4 disks to a job and assign a backup proxy that can process maximum 2 tasks concurrently for the job. In this case, Veeam Backup & Replication will create 4 tasks (1 task per each VM disk) and start processing 2 tasks in parallel. The other 2 tasks will be pending.

How Task Limitation Works

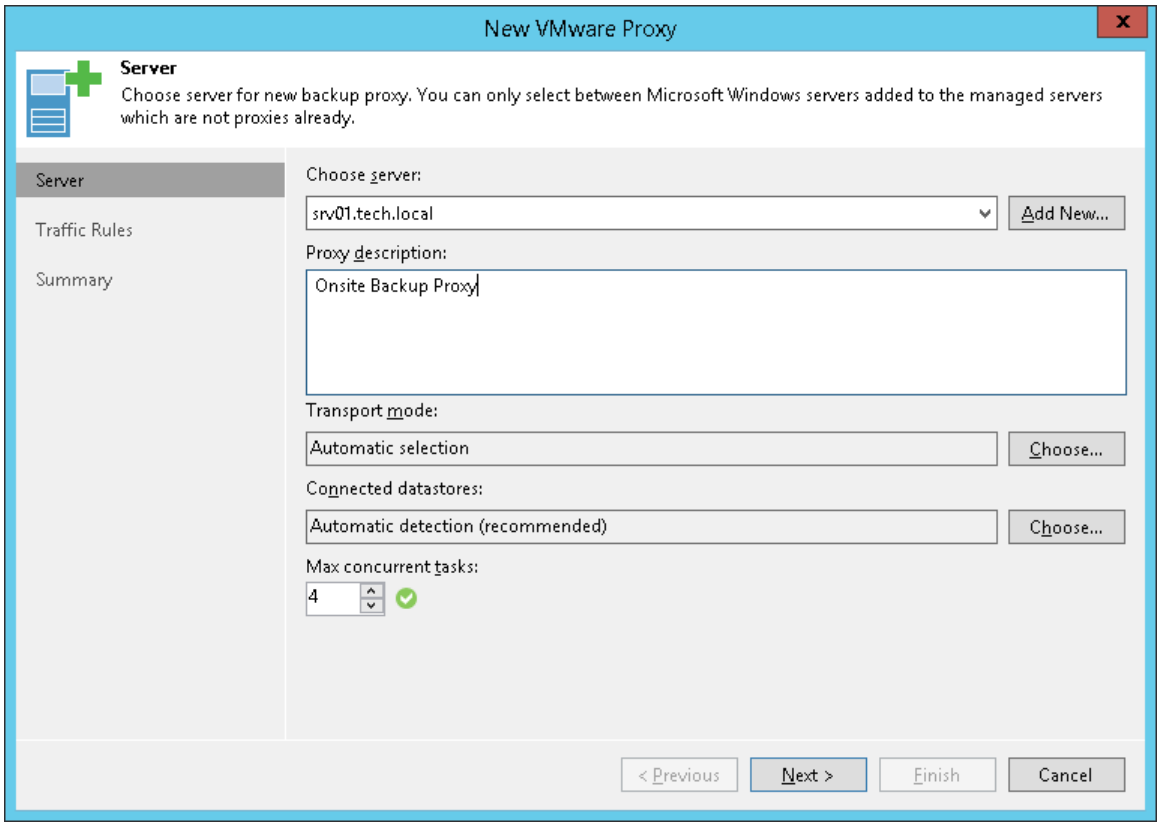
Task limiting is performed by the Veeam Backup Service. The Veeam Backup Service is aware of all backup proxies and backup repositories in the backup infrastructure, and task limitation settings configured for them.

When a job starts, it informs the Veeam Backup Service about the list of tasks created for the job, and backup infrastructure resources that must be used for the job. The Veeam Backup Service detects the number of tasks that required backup infrastructure components are currently processing, and analyzes the number of allowed tasks for these components. If the number of currently processed tasks has reached the allowed limit, the backup infrastructure component will not start processing a new task until one of the currently running tasks finishes.

Task Limitation for Backup Proxies

To limit the number of concurrent tasks on a backup proxy, you must define the **Max concurrent tasks** setting for the backup proxy.

The maximum number of concurrent tasks depends on the number of CPU cores available on the backup proxy. It is strongly recommended that you define task limitation settings using the following rule: 1 task = 1 CPU core. For example, if a backup proxy has 4 CPU cores, it is recommended that you limit the number of concurrent tasks for this backup proxy to 4.



Task Limitation for Backup Repositories

To limit the number of concurrent tasks in a backup repository, you must enable the **Limit maximum concurrent tasks to <N>** option on the backup repository and define the necessary task limit.

The maximum number of concurrent tasks depends on the number of CPU cores available in the backup repository. It is strongly recommended that you define task limitation settings using the following rule: 1 task = 1 CPU core.

It is recommended to configure 2 GB RAM per core. In case of shared folder backup repositories, the same amount of resources is required for gateway servers.

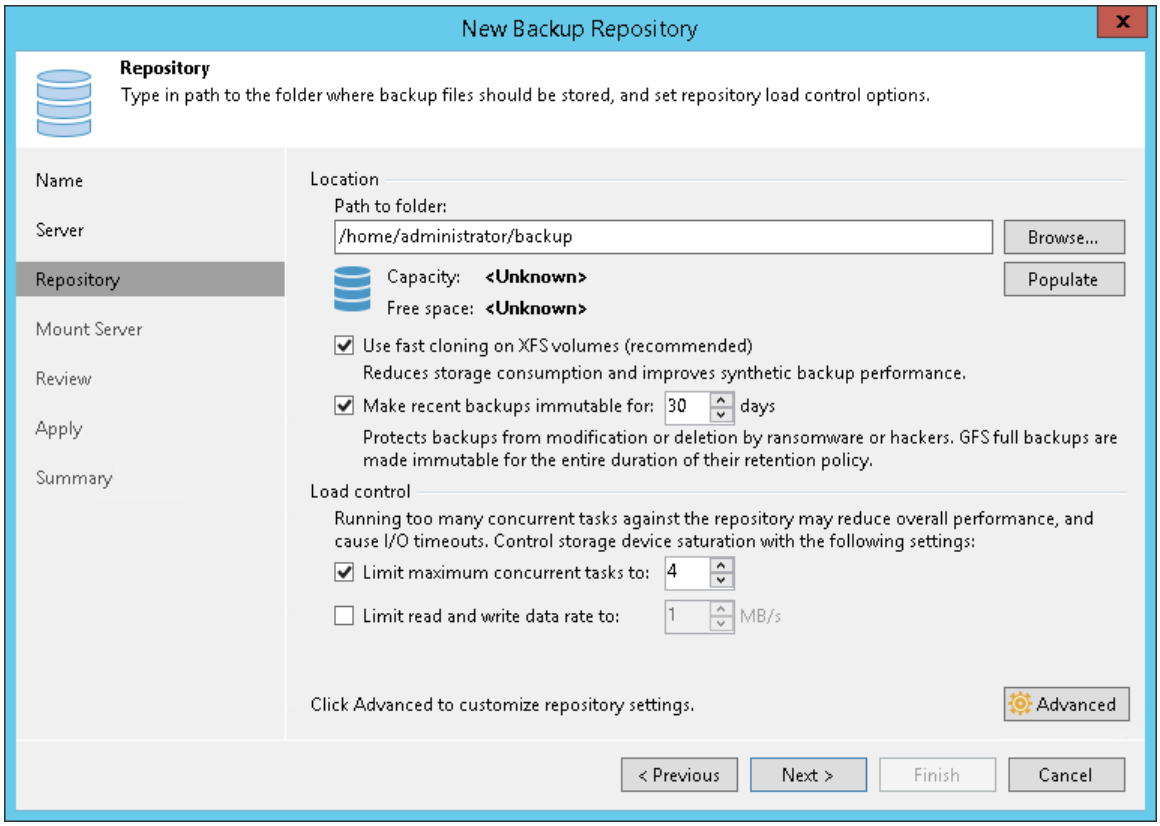
Synthetic operations performed in the backup repository (such as synthetic full backup, backup files merge and transformation) are also regarded as tasks. The number of tasks performed during these operations depends on the type of backup chains stored on the backup repository:

- For regular backup chains, Veeam Backup & Replication creates 1 task per job.
- For per-machine backup chains, Veeam Backup & Replication creates 1 task per every VM chain (that is, every VM added to the job).

If you use backup repositories for backup copy jobs, you must also consider tasks for read operations.

NOTE

When you limit the number of tasks for the backup repository, bear in mind the storage throughput. If the storage system is not able to keep up with the number of tasks that you have assigned, it will be the limiting factor. It is recommended that you test components and resources of the backup infrastructure to define the workload that they can handle.

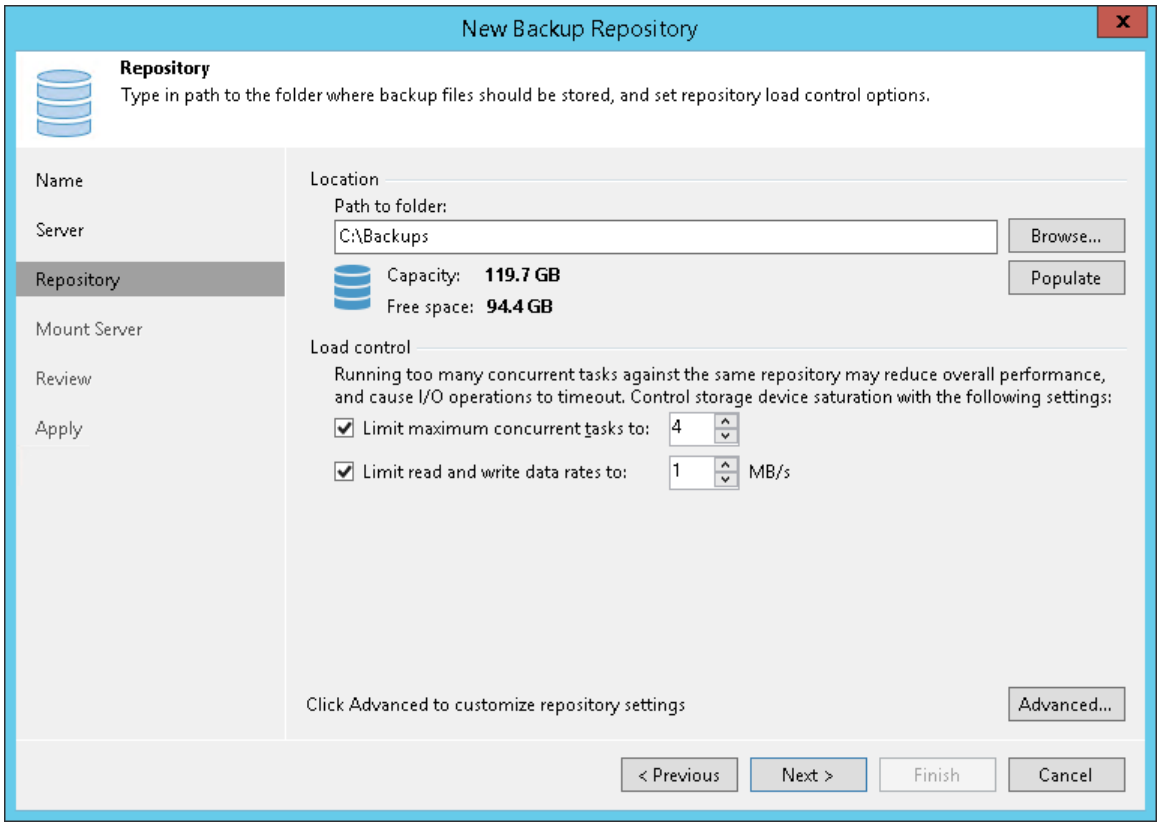


Task Limitation for Components with Several Roles

One machine can perform several roles. For example, you can assign roles of the backup proxy and backup repository to the same machine, or use a backup proxy as a gateway server for a shared folder backup repository. In such situation, you must make sure that the backup infrastructure component is able to process the cumulative number of tasks specified for different roles.

Limitation of Read and Write Data Rates for Backup Repositories

Veeam Backup & Replication can limit the speed with which Veeam Backup & Replication must read and write data to/from the backup repository. The data read and write speed is controlled with the **Limit read and write data rates to <N> MB/s** option that you can enable in backup repository settings.



The Veeam Backup Service is aware of read and write data rate settings configured for all backup repositories in the backup infrastructure. When a job targeted at a backup repository starts, the Veeam Backup Service informs the Veeam Data Mover running on this backup repository about the allowed read/write speed set for this repository so that the Veeam Data Mover can limit the read/write speed to the specified value.

If the backup repository is used by a number of tasks simultaneously, Veeam Backup & Replication splits the allowed read/write speed rate between these tasks equally. Note that the specified limit defines the allowed read speed and the allowed write speed at the same time.

For example, you set the **Limit read and write data rates to** option to 8 MB/s and start two backup jobs. Each job processes 1 VM with 1 VM disk. In this case, Veeam Backup & Replication will create 2 tasks and target them at the backup repository. The data write rate will be split between these 2 tasks equally: 4 MB/s for one task and 4 MB/s for the other task.

If at this moment you start some job reading data from the same backup repository, for example, a backup copy job processing 1 VM with 1 disk, Veeam Backup & Replication will assign the read speed rate equal to 8 MB/s to this job. If you start 2 backup copy jobs at the same time (each processing 1 VM with 1 disk), Veeam Backup & Replication will split the read speed rate between these 2 jobs equally: 4 MB/s for one backup copy job and 4 MB/s for the other backup copy job.

Network Traffic Management

Data protection requires large amount of data to be transferred through networks. This may cause heavy network loads.

Veeam Backup & Replication offers the following capabilities to reduce the load:

- Configuring Network Traffic Rules
- Managing Upload Streams
- Specifying Preferred Networks

Configuring Network Traffic Rules

Network traffic rules control traffic transferred between backup infrastructure components. These rules allow you to do the following:

- Throttle network traffic
- Encrypt transferred data

The rules apply only to the backup infrastructure components, so you do not have to change your network infrastructure.

How Network Rules Work

Each network rule contains IP address ranges for source and target components. When a job starts, Veeam Backup & Replication checks the rules against the components involved in the job. If the IP addresses of the components fall into the IP address ranges of a rule, the rule applies.

For example, 192.168.0.1-192.168.0.255 is the source range, and 172.16.0.1-172.16.0.255 is the target range. 192.168.0.12 is the IP address of one component, and 172.16.0.31 is the IP address of another component. Both IP addresses fall into the ranges, so the rule will apply.

Note that the rules are reversible. The rule from the example will also apply to the specified components if you swap the ranges: make 192.168.0.1-192.168.0.255 the target range and 172.16.0.1-172.16.0.255 the source range.

> **TIP**
>
> You can define a rule for specific components. For this, specify a single IP address in the source range and in the target range.

Creating Network Rules

You must create network rules at the backup server level. For details, see Enabling Traffic Throttling and Enabling Data Encryption.

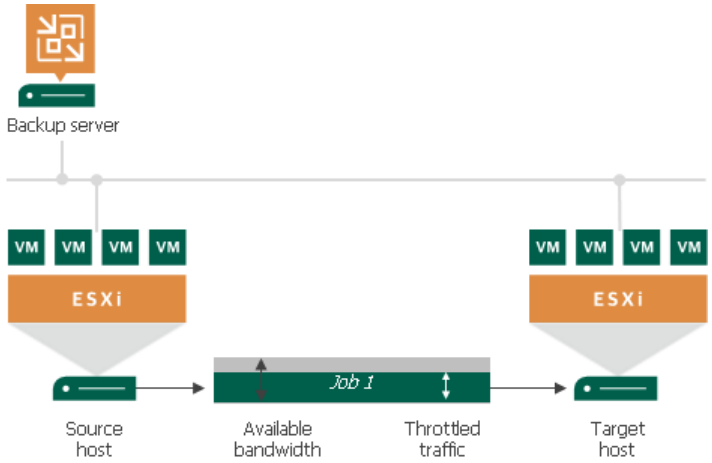
Veeam Backup & Replication has a predefined rule for traffic transferred between public networks. For more information, see Adjusting Internet Rule.

> **TIP**
>
> If you created a rule for a backup proxy, you can check whether it applies. For this, open the Traffic Rules step of the backup proxy wizard. The rule must be in the list of rules.

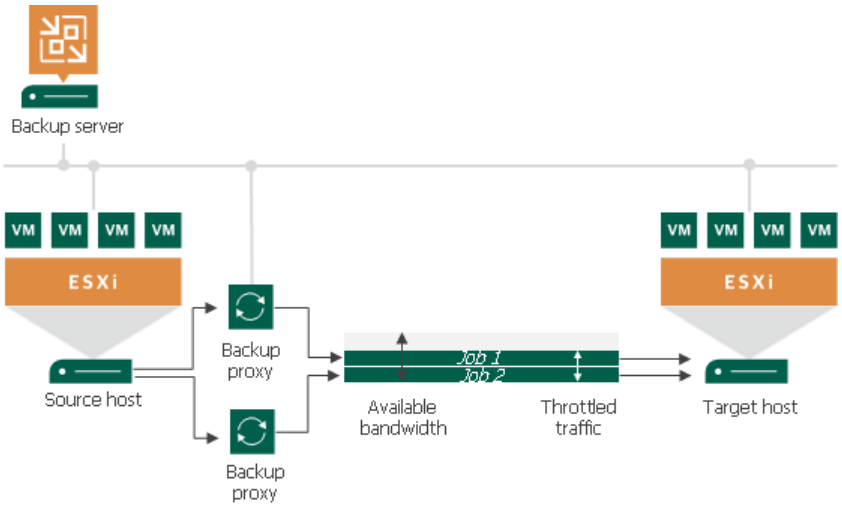
Enabling Traffic Throttling

Traffic throttling setting in a network rule allows you to limit the impact of Veeam Backup & Replication tasks on network performance. Traffic throttling prevents jobs from utilizing the entire bandwidth available in your environment and makes sure that other network operations get enough traffic.

> **IMPORTANT**
>
> Traffic throttling does not apply to restore tasks.



When several jobs that fall into the same network rule run simultaneously, Veeam Backup & Replication equally splits the throttled traffic between them. For example, if two jobs run at a time, each job gets half of the throttled traffic. Once one job finishes, the other gets the entire bandwidth allowed by the rule.



NOTE

Mind the following:

- It is recommended that you throttle network traffic if you perform off-site backup or replicate VMs to a DR site over slow WAN links.
- Traffic throttling rules apply to the transfer of virtual disks. Do not be surprised if you see bandwidth usage peaks at the beginning of the job: that may be caused by transferring other VM files.

Infrastructure Components

Traffic can be throttled only between the backup infrastructure components on which Veeam Data Movers are deployed. These components differ depending on a data protection scenario. The following table shows this dependency.

Scenario	Components
Backup to a Microsoft Windows or Linux backup repository*	Backup proxy and backup repository
Backup to an SMB share, Dell EMC Data Domain and HPE StoreOnce*	Backup proxy and gateway server.
CDP	Source and target VMware CDP proxies.
VM copy	Backup proxy and backup repository.
Backup copy*	Source and target backup repositories, gateway servers or WAN accelerators (if they are involved in the backup copy process).
Replication	Source and target backup proxies or WAN accelerators (if they are involved in the replication process).
Backup to tape*	Backup repository and tape server. For more information about backup to tape jobs, see the Machines Backup to Tape section in the Veeam Backup & Replication User Guide.
SOBR data offload	Gateway server and Amazon S3/Azure object storage.

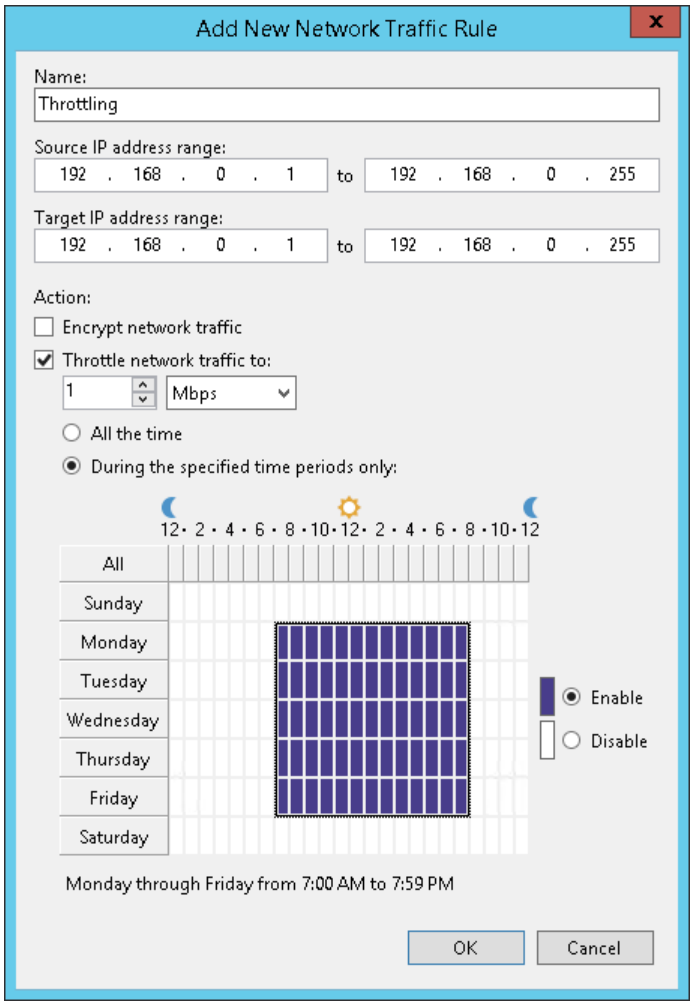
* Veeam Backup & Replication throttles traffic between the listed components also if backups are created with Veeam Agent for Microsoft Windows or Veeam Agent for Linux operating in the standalone or managed mode.

Configuring Traffic Throttling

To configure traffic throttling settings in a rule:

1. From the main menu, select **Network Traffic Rules**.
2. In the **Global Network Traffic Rules** window, click **Add**.

3. In the **Name** field, specify a name for the rule.
4. In the **Source IP address range** section, specify a range of IP addresses for the backup infrastructure components on the source side.
5. In the **Target IP address range** section, specify a range of IP addresses for the backup infrastructure components on the target side.
6. Select the **Throttle network traffic to** check box.
7. In the **Throttle network traffic to** field, specify the maximum speed that must be used to transfer data from source to target.
8. In the section below the **Throttle network traffic to** field, specify the time interval during which the speed must be limited. You can throttle the traffic all the time or at specific time intervals, for example, during business hours.



Several Rules with Traffic Throttling

If you create several rules with the same ranges of IP addresses, make sure that time intervals of the rules do not overlap. For example, to manage network traffic during business and non-business hours, you can create the rules as shown below. These rules have the same ranges of IP addresses.

- **Rule 1**. Speed limit: 1 Mbps; time interval: Monday through Friday from 7 AM to 7 PM.
- **Rule 2**. Speed limit: 10 Mbps; time interval: Saturday through Sunday from 7 AM to 7 PM.

With such rules, Veeam Backup & Replication will limit the speed up to 1 Mbps during business hours and up to 10 Mbps during non-business hours.

If several rules have the same target/source IP address range but different speed limits, the lowest limit is used. For example, if you configure the following rules:

- **Rule 1**. Source IP range: 192.168.0.1-192.168.0.30; target IP range: 192.168.0.1-192.168.0.255; speed limit: 4 Mbps.
- **Rule 2**. Source IP range: 192.168.0.1-192.168.0.255; target IP range: 192.168.0.1-192.168.0.255; speed limit: 1 Mbps.

In this case, Veeam Backup & Replication will use the lowest speed limit — 1 Mbps.

This rule also applies if traffic is throttled by Veeam Backup & Replication and also Veeam Agent for Microsoft Windows. For more information on how to throttle traffic by Veeam Agent for Microsoft Windows, see the Veeam Agent Management Guide and Veeam Agent for Microsoft Windows User Guide.

Enabling Data Encryption

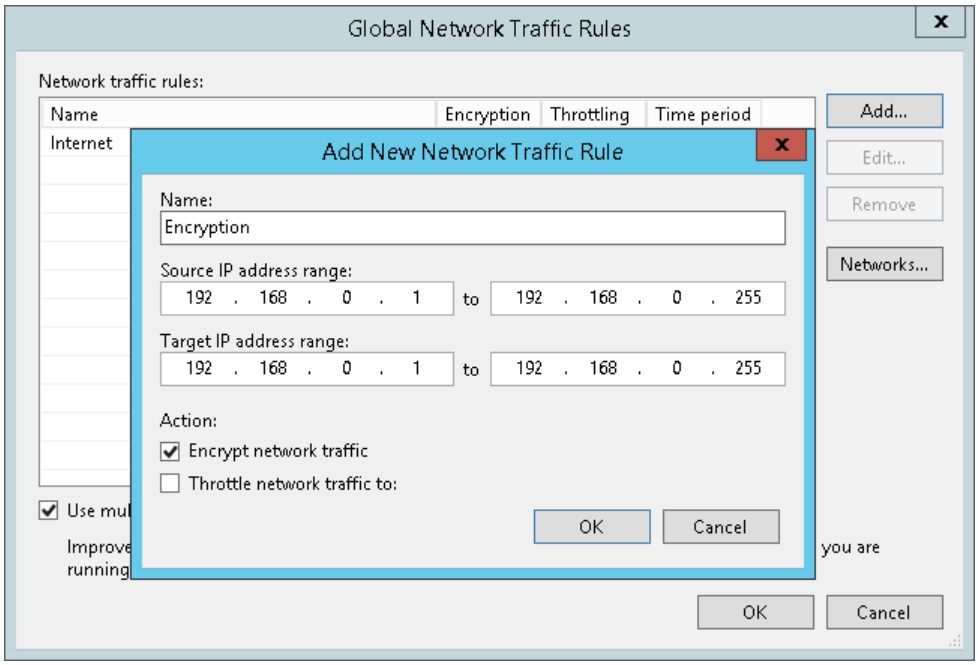
Network rules also allow you to encrypt transferred data. Data transferred through public networks is encrypted by default. For details, see Adjusting Internet Rule.

> **NOTE**
>
> Data encryption rules do not apply to traffic between the ESXi host and VMware CDP proxy.

To create a network rule with data encryption:

1. From the main menu, select **Network Traffic Rules**.
2. In the **Name** field, specify a name for the rule.
3. In the **Global Network Traffic Rules** window, click **Add**.
4. In the **Source IP address range** section, specify a range of IP addresses for backup infrastructure components on the source side.
5. In the **Target IP address range** section, specify a range of IP addresses for backup infrastructure components on the target side.

6. Select the **Encrypt network traffic** check box.



Related Topics

Data Encryption

Adjusting Internet Rule

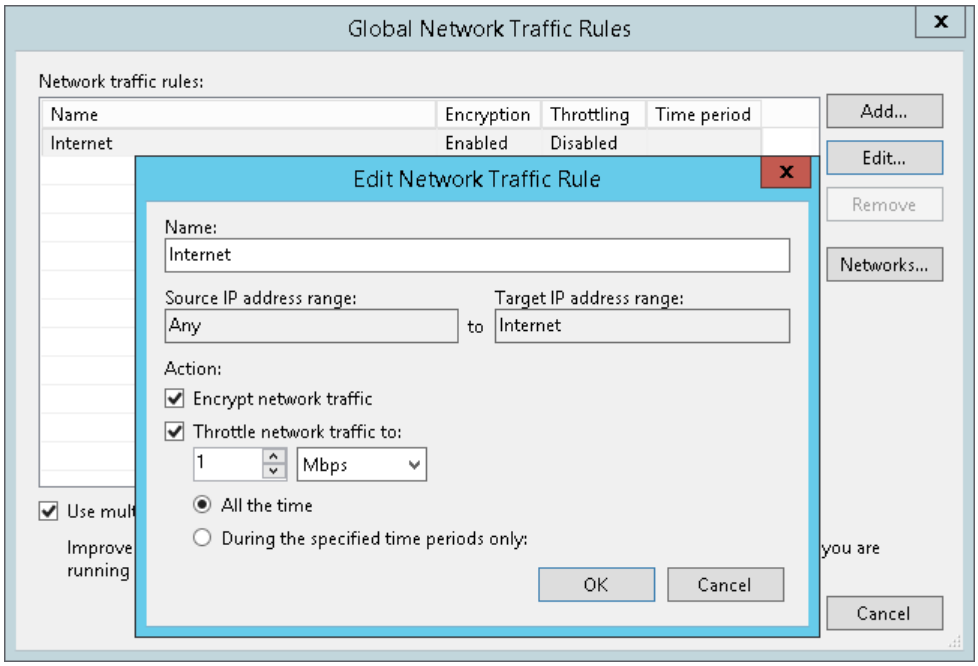
The Internet rule is a predefined network rule. This rule manages traffic transferred through public networks — all networks whose IP ranges differ from 10.0.0.0/8, 172.16.0.0/12 and 192.168.0.0/16. The Internet rule also encrypts traffic for such networks. You cannot delete the Internet rule — you can only adjust or turn it off.

To turn off the Internet rule, clear the **Throttle network traffic to** and **Encrypt network traffic** check boxes.

To adjust the rule:

1. From the main menu, select **Network Traffic Rules**.
2. In the **Global Network Traffic Rules** window, select **Internet** from the list and click **Edit**.
3. In the **Edit Network Traffic Rule** window:
    - To disable encryption, clear the **Encrypt network traffic** check box.

- To enable network traffic throttling, select the **Throttle network traffic to** check box. For details, see Enabling Traffic Throttling.



Managing Upload Streams

By default, Veeam Backup & Replication uses multithreaded data transfer for every job session. VM data going from source to target is transferred over 5 TCP/IP connections. However, if you schedule several jobs to run at the same time, load on the network may be heavy. If the network capacity is not sufficient to support multiple data transfer connections, you can disable multithreaded data transfer or change the number of TCP/IP connections.

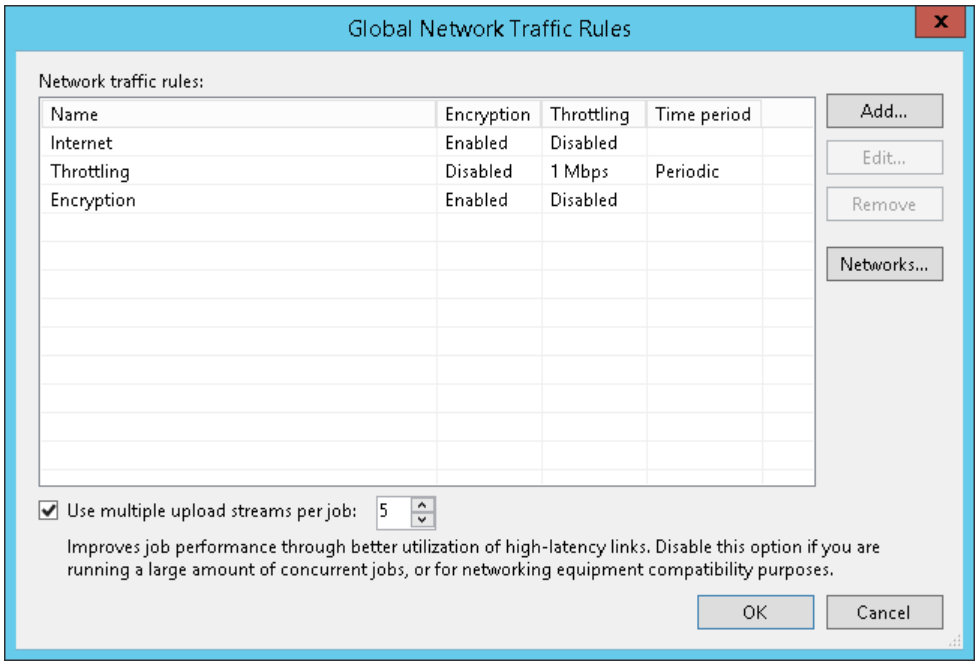
To change the number of connections:

1. From the main menu, select **Network Traffic Rules**.
2. In the **Global Network Traffic Rules** window, specify new data transfer settings:
   - To disable multithreaded data transfer, clear the **Use multiple upload streams per job** check box. Veeam Backup & Replication will use only one TCP/IP transfer connection for every job session.
   - To change the number of TCP/IP connections, leave the **Use multiple upload streams per job** check selected and specify the necessary number of connections in the field on the right.

NOTE

Mind the following:

- Veeam Backup & Replication performs a CRC check for the TCP traffic going between the source and the target. When you perform backup, replication or VM copy operations, Veeam Backup & Replication calculates checksums for data blocks going from the source. On the target, it recalculates checksums for received data blocks and compares them to the checksums created on the source. If the CRC check fails, Veeam Backup & Replication automatically re-sends data blocks without any impact on the job.
- [For Veeam Plug-ins for Enterprise Applications] The multithreaded data transfer setting does not affect backup jobs created by Veeam Plug-in for Oracle RMAN/SAP HANA/SAP on Oracle. To configure multiple channels for backup and restore operations for these plug-ins, see the Veeam Plug-ins for Enterprise Applications Guide.



Specifying Preferred Networks

You can choose networks over which Veeam Backup & Replication must transport data when you perform data protection and disaster recovery tasks. This option can be helpful if you have a non-production network and want to route data traffic over this network instead of the production one.

Preferred network rule applies only to traffic between the following backup infrastructure components:

- Backup server
- WAN accelerator[1]
- Gateway server[2]
- VMware CDP proxy
- Backup proxy
- Backup repository
- Log shipping server

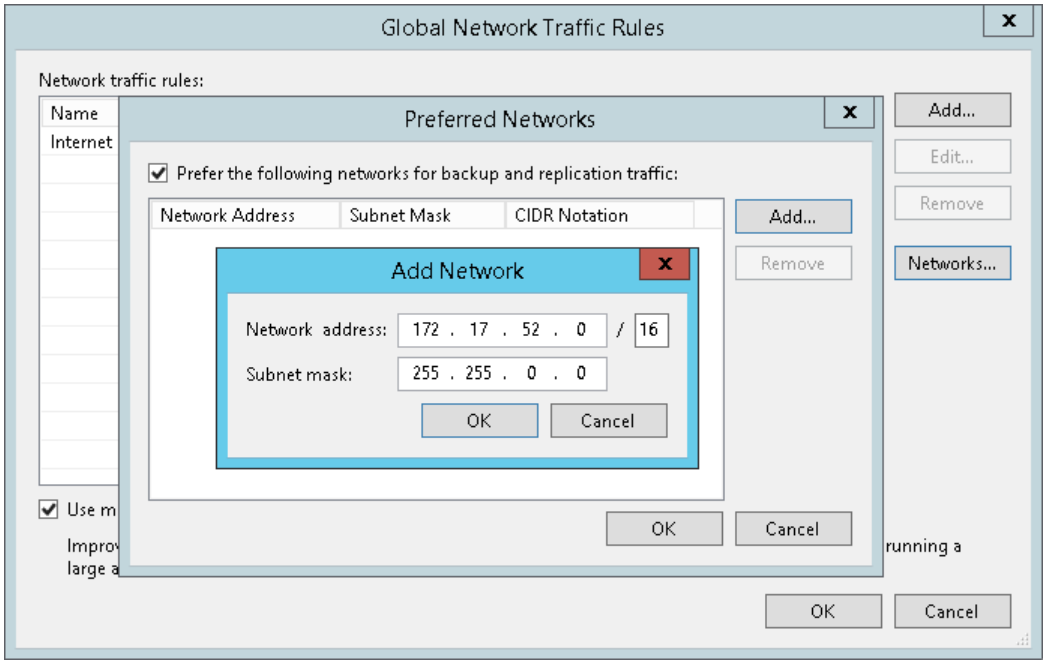
- Tape server
- Storage system (backup from storage snapshots scenario)[3]
- Veeam Agent

[1] The rule applies only to traffic between the source and target WAN accelerators.
[2] The rule does not apply to traffic between the gateway server and backup repository. For the list of backup repositories with which the gateway server communicates, see Gateway Server.
[3] The rule does not apply when backing up from Cisco HyperFlex systems or when the transport mode is Direct storage access.

To define networks for data transfer, you must create a list of preferred networks. When Veeam Backup & Replication needs to transfer data, it uses networks from this list. If a connection over preferred networks cannot be established for some reason, Veeam Backup & Replication will automatically fail over to the production network.

To set a network priority list:

1. From the main menu, select **Network Traffic Rules**.
2. In the **Global Network Traffic Rules** window, click **Networks**.
3. In the **Preferred Networks** window, select the **Prefer the following networks for backup and replication traffic** check box.
4. Click **Add**.
5. Specify a network address using a CIDR notation or a network mask and click **Add**.
6. Repeat steps 4-5 for all networks that you want to add.



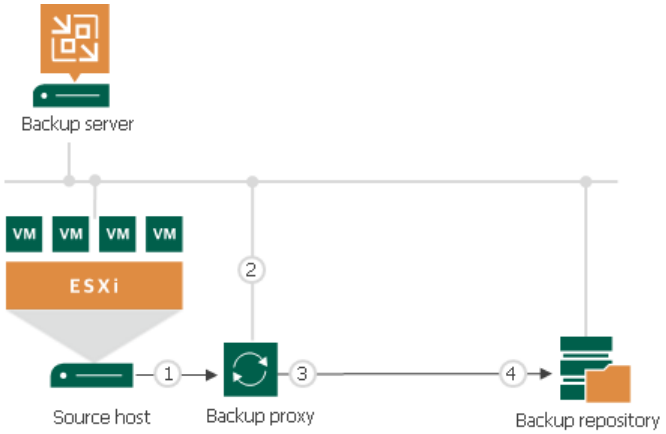
Performance Bottlenecks

As any backup application handles a great amount of data, it is important to make sure the data flow is efficient and all resources engaged in the backup process are optimally used. Veeam Backup & Replication provides advanced statistics about the data flow efficiency and lets you identify bottlenecks in the data transmission process.

Veeam Backup & Replication processes VM data in cycles. Every cycle includes a number of stages:

1. Reading VM data blocks from the source
2. Processing VM data on the backup proxy
3. Transporting data over the network
4. Writing data to the target

When one data processing cycle is over, the next cycle begins. VM data therefore goes over the "data pipe".



To evaluate the data pipe efficiency, Veeam Backup & Replication analyzes performance of all components in the data flow working as the cohesive system, and evaluates key factors on the source and target sides. Veeam Backup & Replication checks the following points in the data pipe:

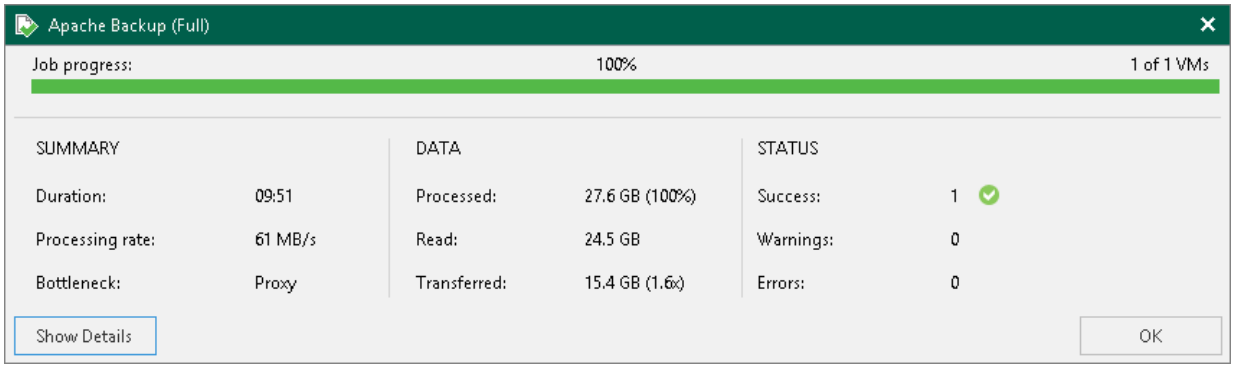
1. **Source** – source disk reader component responsible for retrieving data from the source storage.
2. **Proxy** – backup proxy component responsible for processing VM data.
3. **Source WAN accelerator** – WAN accelerator deployed on the source side. Used for backup copy and replication jobs working through WAN accelerators.
4. **Network** – network queue writer component responsible for getting processed VM data from the backup proxy and sending it over the network to the backup repository or another backup proxy.
5. **Target WAN Accelerator** – WAN accelerator deployed on the target side. Used for backup copy and replication jobs working through WAN accelerators.
6. **Target** – target disk writer component (backup storage or replica datastore).

The resource usage level for these points is evaluated in percent. This percent rate defines the amount of time for which components are busy during the job. An efficient data flow assumes that there is no latency at any point of the data pipe, and all its components work for approximately equal amount of time.

If any of the components operates inefficiently, there may appear a bottleneck in the data path. The insufficient component will work 100% of time while the others will be idling, waiting for data to be transferred. As a result, the whole data flow will slow down to the level of the slowest point in the data path, and the overall time of data processing will increase.

To identify a bottleneck in the data path, Veeam Backup & Replication detects the component with the maximum workload: that is, the component that works for the most time of the job. For example, you use a low-speed storage device as the backup repository. Even if VM data is retrieved from the SAN storage on the source side and transported over a high-speed link, VM data flow will still be impaired at the backup repository. The backup repository will be trying to consume transferred data at the rate that exceeds its capacity, and the other components will stay idle. As a result, the backup repository will be working 100% of job time, while other components may be employed, for example, for 60% only. In terms of Veeam Backup & Replication, such data path will be considered insufficient.

The bottleneck statistics for a job is displayed in the job session data. The bottleneck statistics does not necessarily mean that you have a problem in your backup infrastructure. It informs you about the weakest component in the data path. However, if you feel that the job performance is low, you may try taking some measures to get rid of the bottleneck. For example, in the case described above, you can limit the number of concurrent tasks for the backup repository.



Throttling as Bottleneck

In addition to main points in the data pipe, Veeam Backup & Replication may report throttling as a bottleneck. This can happen in the following cases:

- If you limit the read and write data rates for a backup repository, a backup repository may become a bottleneck. Veeam Backup & Replication will report *Throttling* in the bottleneck statistics.
- If you set up network throttling rules, network may become a bottleneck. Veeam Backup & Replication will report *Throttling* in the bottleneck statistics.

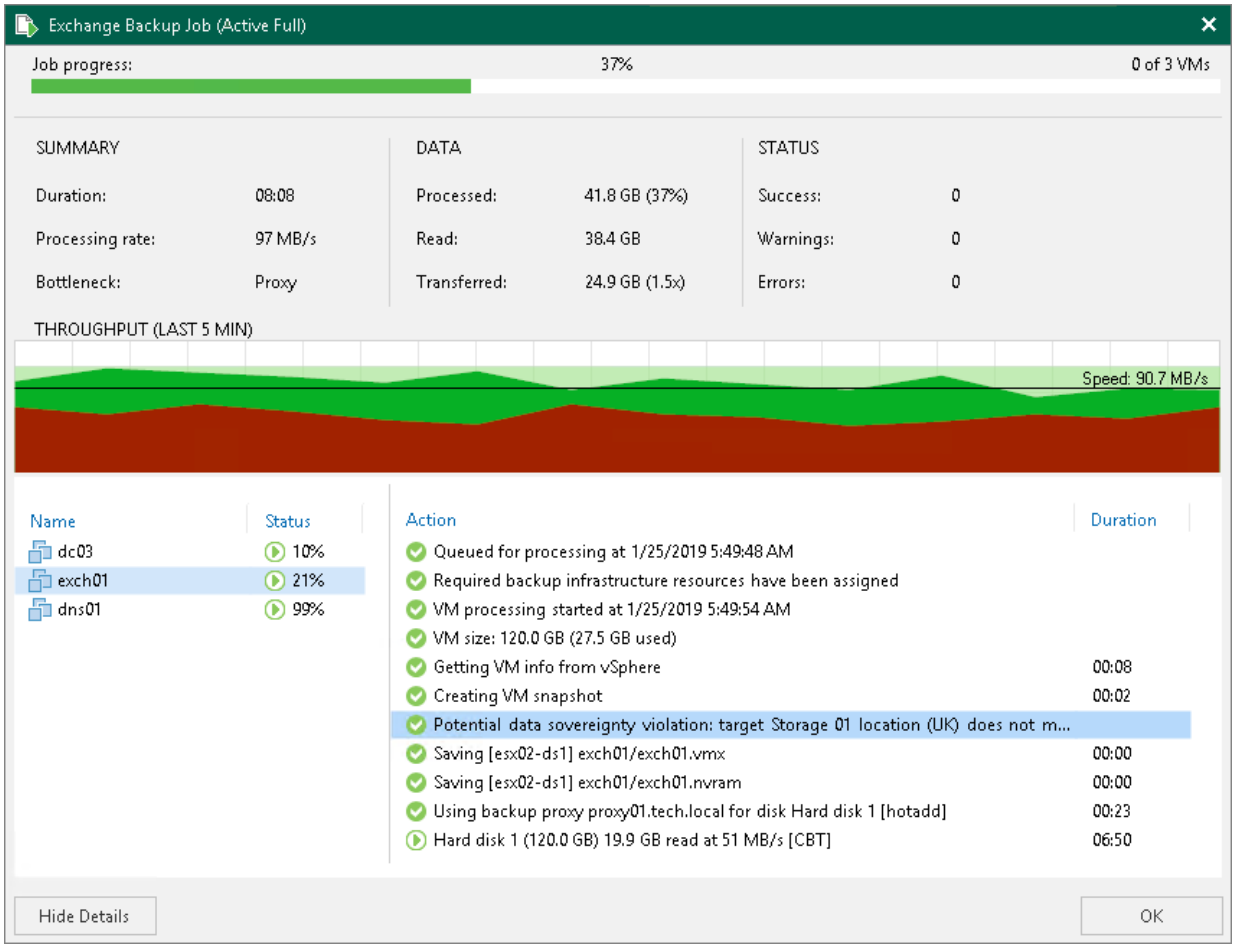
Locations

To control data migration in the virtual infrastructure, Veeam Backup & Replication introduces a notion of location. A location defines a geographic region, or country, in which an infrastructure object resides. You can create a list of locations, and assign to backup infrastructure objects information about locations to which they belong.

Veeam Backup & Replication allows you to assign information about locations to the following infrastructure objects:

- Virtual infrastructure objects: vCenter Servers, datacenters, clusters and hosts.
- Backup infrastructure objects: backup repositories, external repositories, scale-out backup repositories, tape libraries and tape vaults.
- Agent management objects: protection groups.
- Veeam Cloud Connect for service providers: cloud repositories and hardware plans.

Information about infrastructure objects location is stored in the Veeam Backup & Replication configuration database. When VM data in the virtual infrastructure migrate from one location to another, Veeam Backup & Replication displays a warning and stores a record about data migration to job or task session details. In addition to it, Veeam Backup & Replication logs this information to Microsoft Windows event logs. For example, if you back up VMs from a host that resides in Germany to a backup repository that resides in Australia, Veeam Backup & Replication will display a warning that VM data changes its location in the backup job wizard, display information about data migration in the backup job session details and log it to Microsoft Windows event logs.



Veeam Backup & Replication displays information about VM data migration in statistics for the following types of jobs:

- Backup jobs — Veeam Backup & Replication compares the location of the source host on which VMs are registered with the location of the target backup repository or cloud repository.
- Backup copy jobs — Veeam Backup & Replication compares the location of the source host with the location of the target host.
- VeeamZIP tasks (except the cases when you select to store the VeeamZIP file in a local or shared folder) — Veeam Backup & Replication compares the location of the source host on which VMs are registered with the location of the target backup repository.
- Replication jobs — Veeam Backup & Replication compares the location of the source host on which VMs are registered with the location of the target host.
- Replica failback tasks — Veeam Backup & Replication compares the location of the source host with the location of the host to which the VM is restored.
- VM copy jobs — Veeam Backup & Replication compares the location of the source host on which VMs are registered with the location of the target backup repository or target host.
- Quick migration tasks — Veeam Backup & Replication compares the location of the source host on which VMs are registered with the location of the target host.
- Entire VM restore tasks — Veeam Backup & Replication compares the location of the source host with the location of the host to which VMs are restored.
- External repository tasks:
    - Backup copy jobs: Veeam Backup & Replication compares the location of the source external repository with the location of the target backup repository.
    - Restore to Amazon EC2: Veeam Backup & Replication compares the geographic region of the backed-up Amazon EC2 instance with the geographic region of the target EC2 instance.
    - Restore to Microsoft Azure: Veeam Backup & Replication always displays a warning about VM data migration when restore to Microsoft Azure is performed from external repositories.
- SureBackup jobs — Veeam Backup & Replication compares the source location with the target location. The target location is always a host on which the virtual lab is registered. The source location may be one of the following:
    - If a VM is added to the application group, Veeam Backup & Replication compares the host on which the VM is registered (or was registered at the moment of backup) with the target location.
    - If a VM is added to the SureBackup job from the linked job, Veeam Backup & Replication compares the backup repository on which the backup file resides with the target location.
- Tape tasks:
    - Backup to tape jobs: In backup to tape jobs, Veeam Backup & Replication compares the location of the source job or repository with the location of the tape library in the target media pool. If the media pool spans multiple tape libraries, Veeam Backup & Replication analyzes locations of all libraries in the media pool.
    - Vaults: If a tape job exports offline backups to a vault, Veeam Backup & Replication compares the location of the source job or repository with the location of the vault. I a GFS tape job exports tapes to multiple vaults, Veeam Backup & Replication analyzes all vaults configured for target media pools of the GFS tape job.

- Media pools: Veeam Backup & Replication compares locations of all tape libraries added to the media pool. If the media pool exports tapes to a vault, Veeam Backup & Replication analyzes all vaults configured for the media pool.

Limitations for Locations

For SureReplica jobs, Veeam Backup & Replication does not compare information about source and target hosts location.

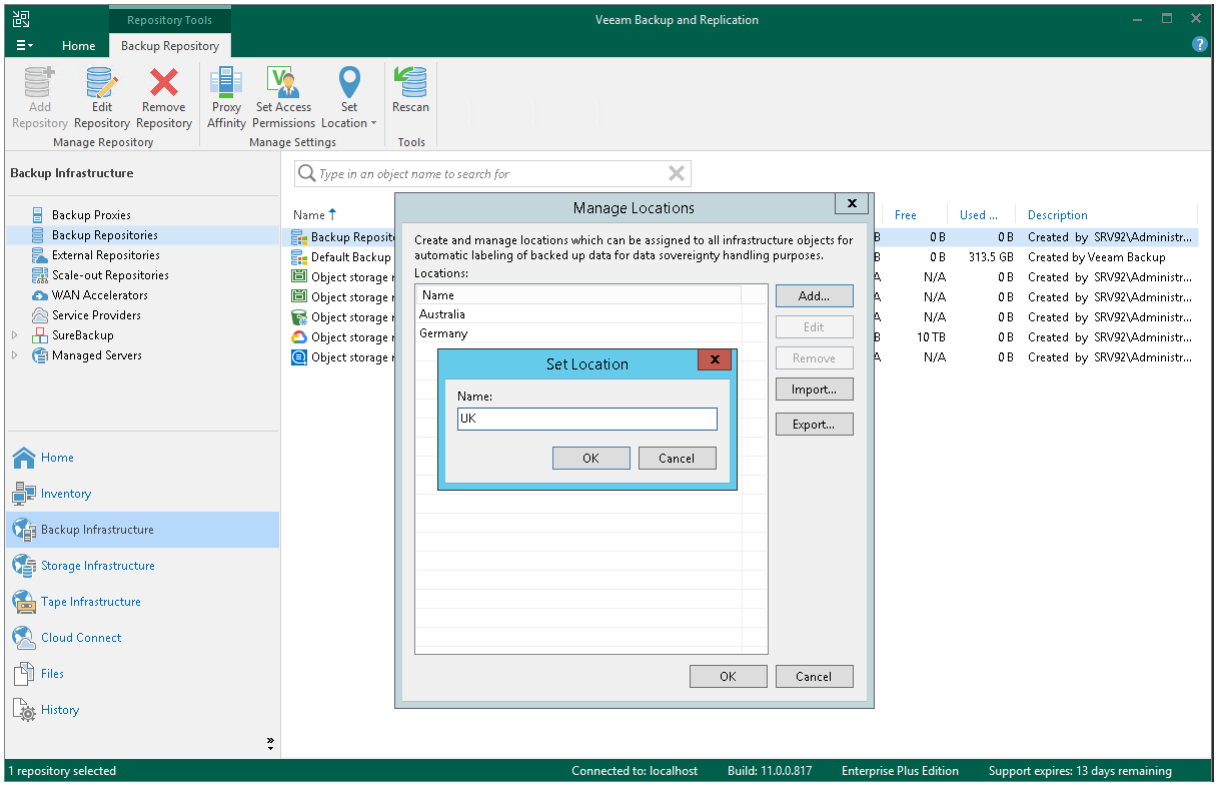
Veeam Backup & Replication does not display a warning about VM data migration for file copy jobs.

Creating and Assigning Locations to Infrastructure Objects

You can create a list of locations in Veeam Backup & Replication and assign locations to infrastructure objects. If you assign a location to a vCenter Server, it will be applied to all child hosts (clusters and ESXi hosts). You can also assign the location to a child host.

To create a location:

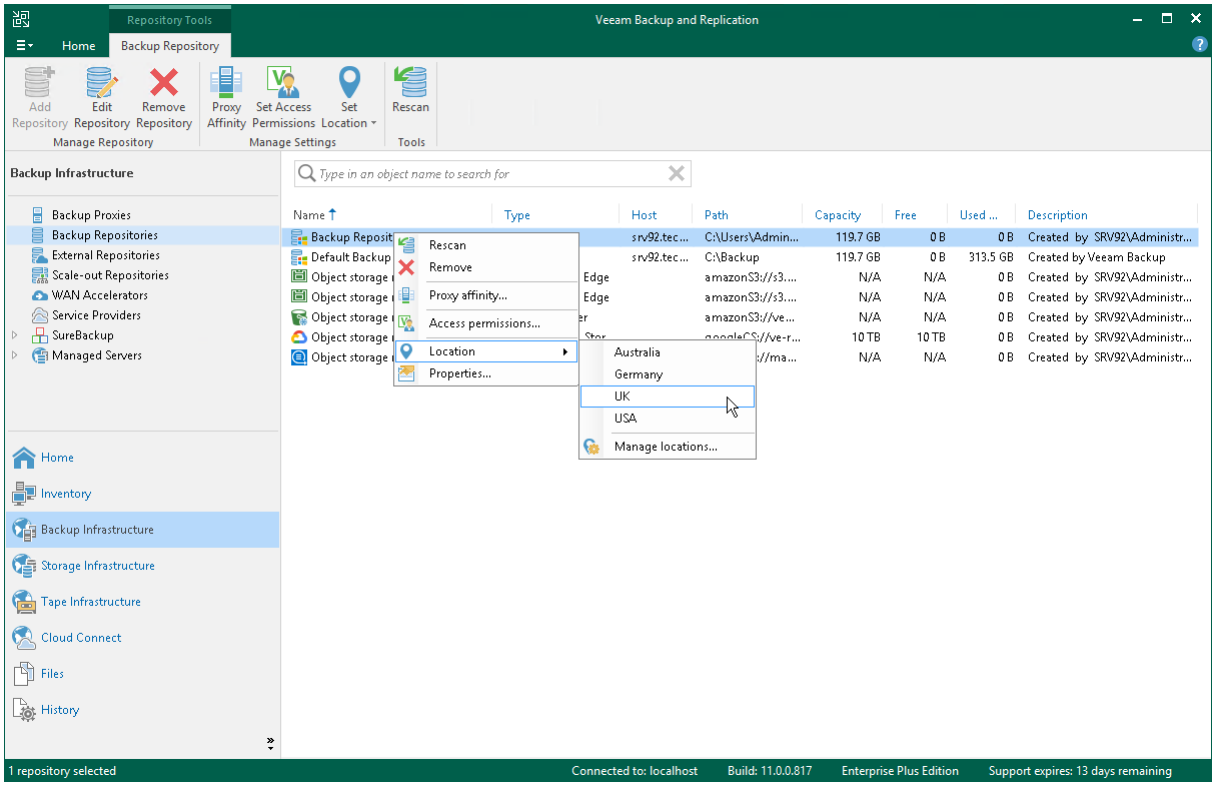
1. In the **Inventory** or **Backup Infrastructure** view, right-click the infrastructure object and select **Location** > **Manage locations**.
2. In the **Manage Locations** window, click **Add**.
3. In the **Name** field, enter a name of the location.



To assign a location to an infrastructure object, in the **Inventory** or **Backup Infrastructure** view, right-click the infrastructure object and select Location > *<Location name>*. If the location is not in the list, select **Location** > **Manage Locations** and add the location to the list.

NOTE

When assigning a location to a scale-out backup repository, the location will be global for all extents. If you add an extent whose location differs from the global location, it will be changed in favor of the location of the scale-out repository.

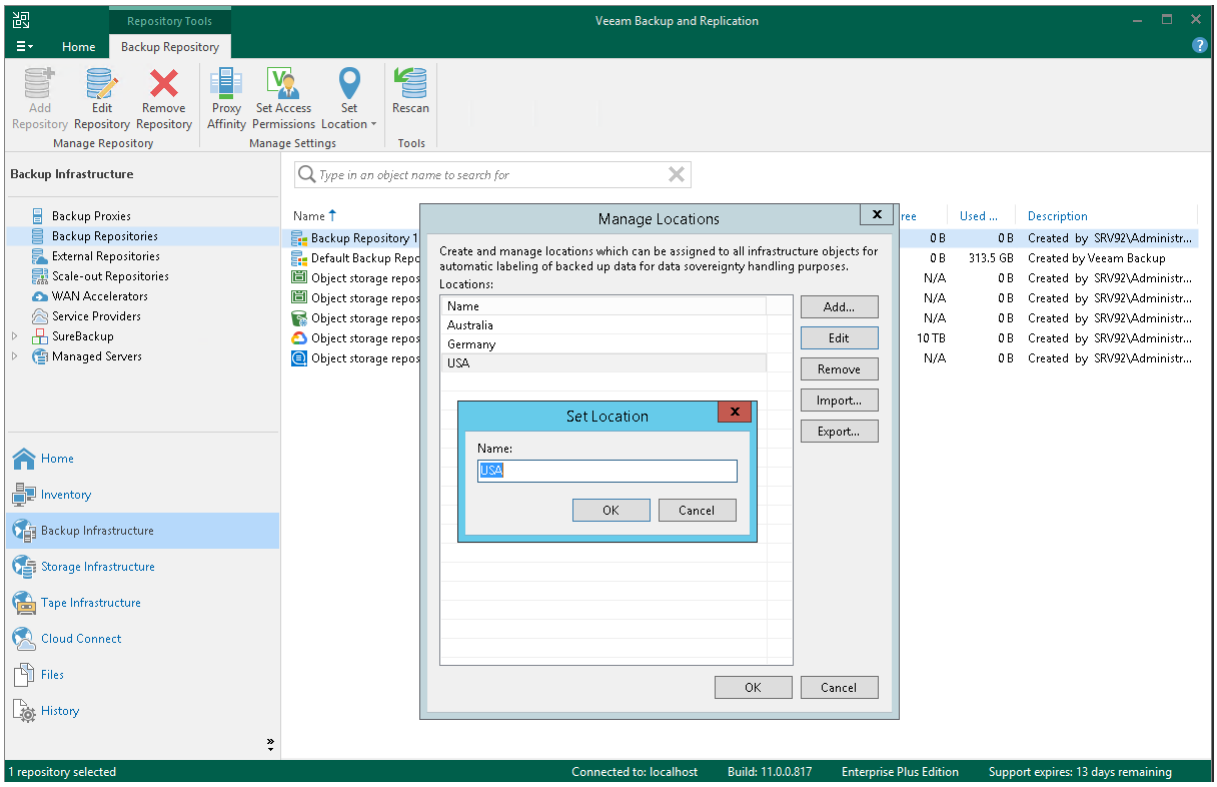


Editing Locations

You can edit a location in the locations list, for example, if you want to change the location name.

To edit a location:

1. In the **Inventory** or **Backup Infrastructure** view, right-click the infrastructure object and select **Location** > **Manage locations**.
2. In the **Manage Locations** window, select the location and click **Edit**.
3. In the **Name** field, change the location name as required.

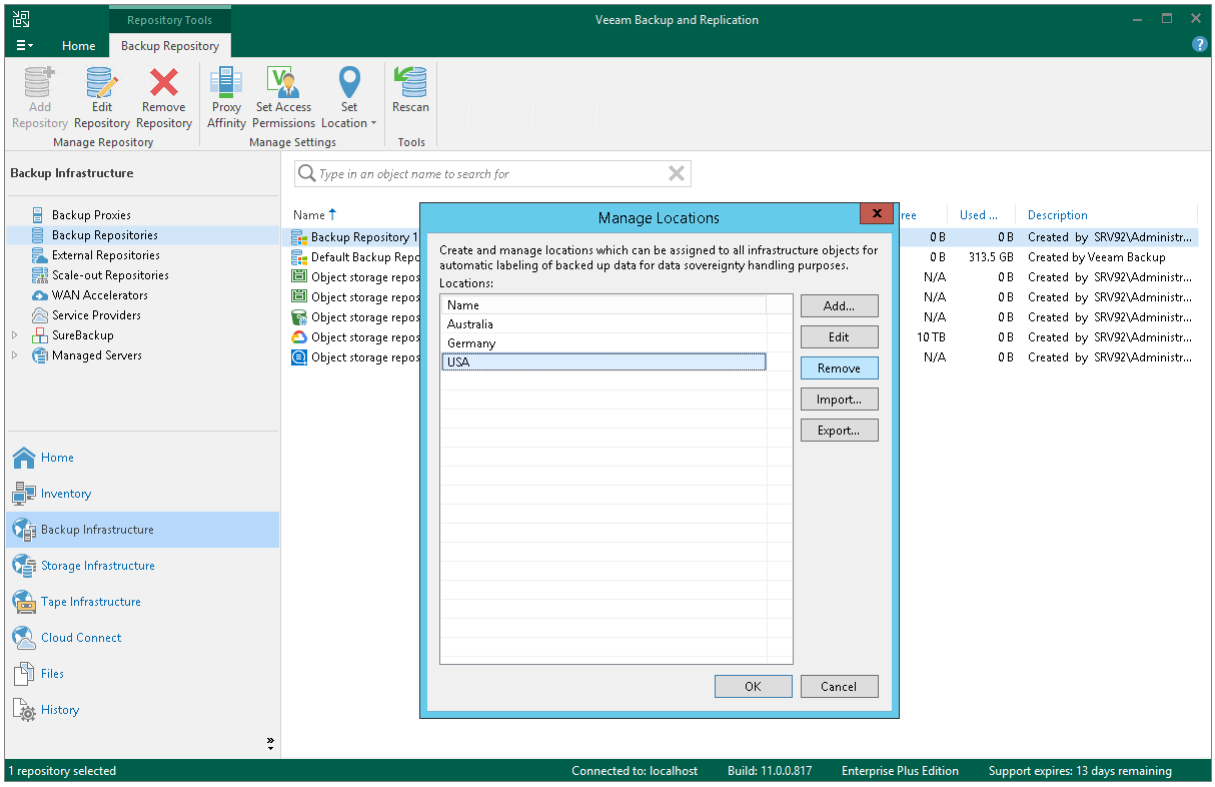


Deleting Locations

You can delete a location from the locations list, for example, if you no longer host infrastructure objects in this location.

To delete a location:

1. In the **Inventory** or **Backup Infrastructure** view, right-click the infrastructure object and select **Location** > **Manage locations**.
2. In the **Manage Locations** window, select the location and click **Delete**. If the location is currently assigned to some infrastructure objects, Veeam Backup & Replication will display a warning with the list of objects that belong to this location. Click **Yes** to confirm the location deletion.



Exporting and Importing Locations List

You can export and import the list of locations to/from a file of XML format.

The import and export functionality facilitates the process of locations creation and maintenance. For example, if you need to set up the same list of locations throughout the whole backup infrastructure, you can create a list of locations on one backup server manually, export this list to an XML file, and then import the list on other backup servers and machines running the Veeam Backup & Replication console.

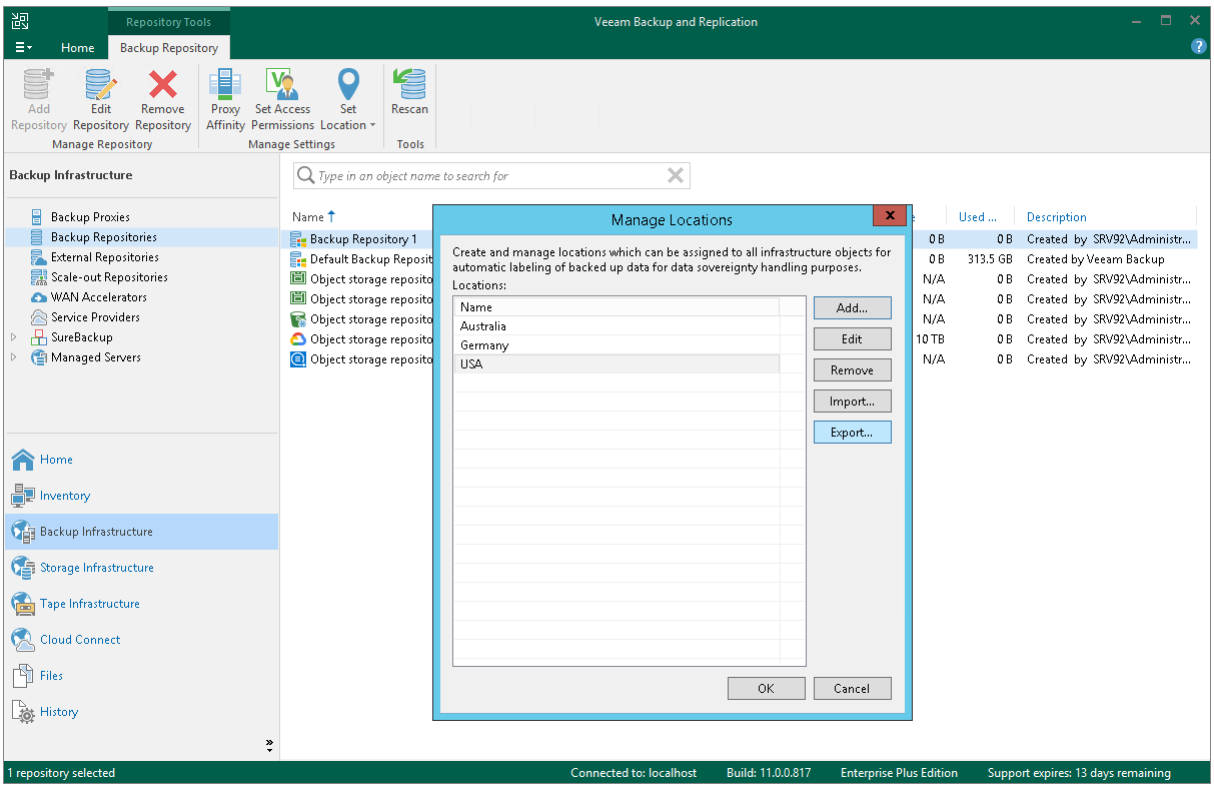
> **TIP**
>
> If you delete and recreate a location, Veeam Backup & Replication will create an object with a new ID in the database and consider it as a new location. Thus, to preserve the uniqueness of the location, use the location export/import operations.

To export the locations list:

1. In the **Inventory** or **Backup Infrastructure** view, right-click an infrastructure object and select **Location** > **Manage locations**.
2. In the **Manage Locations** window, click **Export** and specify a name of the XML file to which the locations list must be exported.

To import the locations list:

1. In the **Inventory** or **Backup Infrastructure** view, right-click an infrastructure object and select **Location** > **Manage locations**.
2. In the **Manage Locations** window, click **Import** and browse to the XML file from which the locations list must be imported.



Veeam Backup & Replication Settings

You can set up general settings for Veeam Backup & Replication. General settings are applied to all jobs, backup infrastructure components and other objects managed by the backup server.

Specifying I/O Settings

You can specify data processing settings.

Mind the following:

- The **Enable storage latency control** option is included in the Veeam Universal License. When using a legacy socket-based license, Enterprise or higher edition is required.
- The **Set custom thresholds on individual datastores** option is included in the Veeam Universal License. When using a legacy socket-based license, Enterprise Plus edition is required.
- The **Enable storage latency control** option is not supported for vVOLs/vSAN storage.

To specify data processing settings:

1. From the main menu, select **General Options**.
2. Click the **I/O Control** tab.
3. To control the I/O load on the production storage where VMs reside, select the **Enable storage latency control** check box. When you enable storage latency control, Veeam Backup & Replication monitors storage read latency on production datastores during data protection and disaster recovery activities. To monitor the storage latency, Veeam Backup & Replication uses real-time metrics from the hypervisor where VMs reside. By default, metrics from the hypervisor are collected every 20 seconds. These settings are inherited from VMware vSphere.

   Specify two thresholds:

   a. In the **Stop assigning new tasks to datastore at** field, specify the I/O latency limit at which Veeam Backup & Replication must not assign new tasks targeted at the datastore.

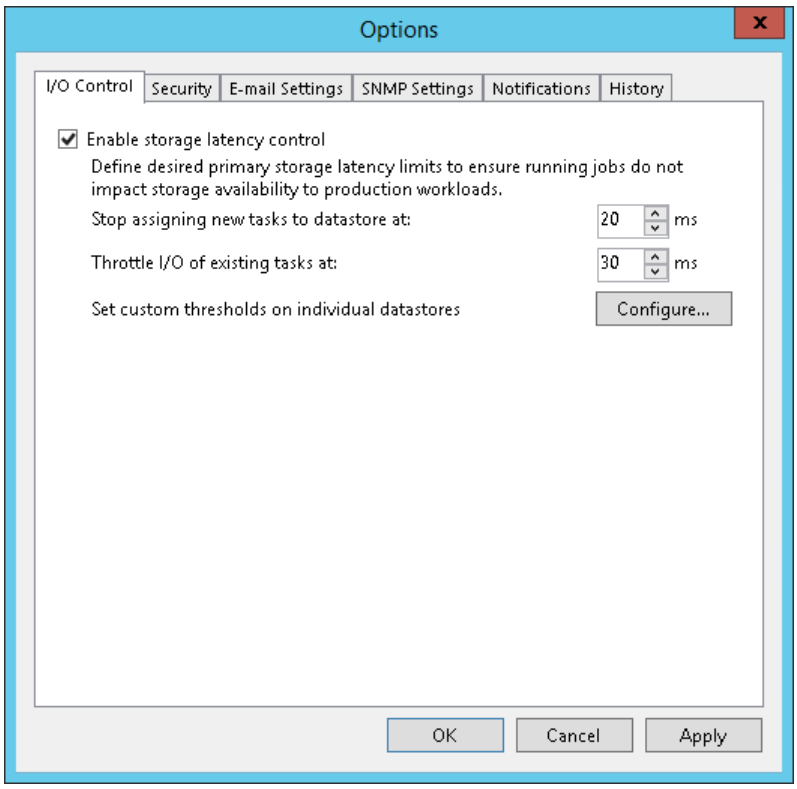
   b. In the **Throttle I/O of existing tasks at** field, specify the I/O latency limit at which Veeam Backup & Replication must decrease the speed of data retrieval or writing to/from the datastore. When the I/O latency for this datastore reaches this value, the Veeam Data Mover working with this datastore will slow down data retrieval or writing.

   The value in the **Stop assigning new tasks to datastore at** field cannot be greater than the value in the **Throttle I/O of existing tasks at** field.

NOTE

If you enable the storage latency control option, Veeam Backup & Replication starts processing VM disks residing on the same datastore with a 40-60 second time offset. This offset helps Veeam Backup & Replication evaluate the current I/O load on the datastore. For example, if you launch a job processing a VM with two disks, Veeam Backup & Replication will start processing the first VM disk, wait for 40-60 seconds to evaluate the I/O workload on the datastore, and then start processing the second VM disk.

Keep in mind this behavior. If you schedule jobs that process multiple VM disks residing on the same datastore to start at the same time, the jobs performance will degrade.

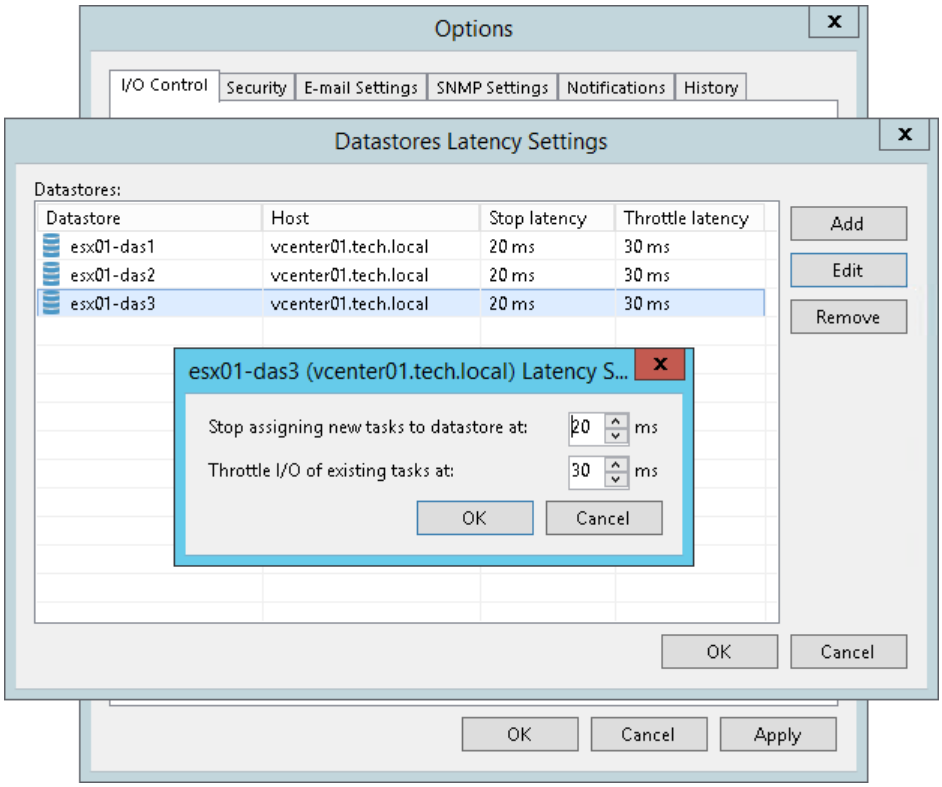


You can set the I/O latency limit for every storage in the virtual infrastructure separately.

To set the I/O latency limit for every storage separately:

1. From the main menu, select **General Options**.
2. Click the **I/O Control** tab.
3. Click **Configure**.
4. Click **Add** > **Datastore**, select the necessary datastore and click **OK** to add it to the storage list.
5. Select the added datastores in the list and click **Edit**.

6. Specify the I/O thresholds for the datastores as described above.



Configuring Security Settings

In the **Security** tab, you can configure the following:

- TLS Certificates

  Configure a TLS certificate to establish secure communication from backup infrastructure components to the backup server.

- Linux Hosts Authentication

  Enable the fingerprint check for Linux machines to protect connection from man-in-the-middle attacks.

- Cloud Connect

  Enable access to the cloud gateway for the Remote Access Console connected to an external network.

- FIPS Compliance

  Enable FIPS-compliant operation mode.

- Audit Logs Location

  Select a folder for storing audit logs.

TLS Certificates

When you configure the Veeam Backup & Replication infrastructure, you can specify what TLS certificate must be used to establish a secure connection from backup infrastructure components to the backup server. Veeam Backup & Replication offers the following options for TLS certificates:

- You can choose to keep the default self-signed TLS certificate generated by Veeam Backup & Replication at the process of upgrading to a new version of Veeam Backup & Replication.
- You can use Veeam Backup & Replication to generate a new self-signed TLS certificate. To learn more, see Generating Self-Signed Certificates.
- You can select an existing TLS certificate from the certificates store. To learn more, see Importing Certificates from Certificate Store.
- You can import a TLS certificate from a file in the PFX format. To learn more, see Importing Certificates from PFX Files.

> **NOTE**
>
> If you plan to use a certificate issued by your own CA, make sure that the certificate meets the following requirements:
>
> 1. The following Key Usage extensions are enabled in the certificate: *Digital Signature*, *Certificate Signing*, *Off-line CRL Signing*, *CRL Signing*.
> 2. The Key Type in the certificate is set to *Exchange*.

Generating Self-Signed Certificates

You can use Veeam Backup & Replication to generate a self-signed certificate for authenticating parties in the Veeam Backup & Replication infrastructure.

To generate TLS certificates, Veeam Backup & Replication employs the RSA Full cryptographic service provider by Microsoft Windows installed on the Veeam backup server. The created TLS certificate is saved to the *Shared* certificate store. The following types of users can access the generated TLS certificate:

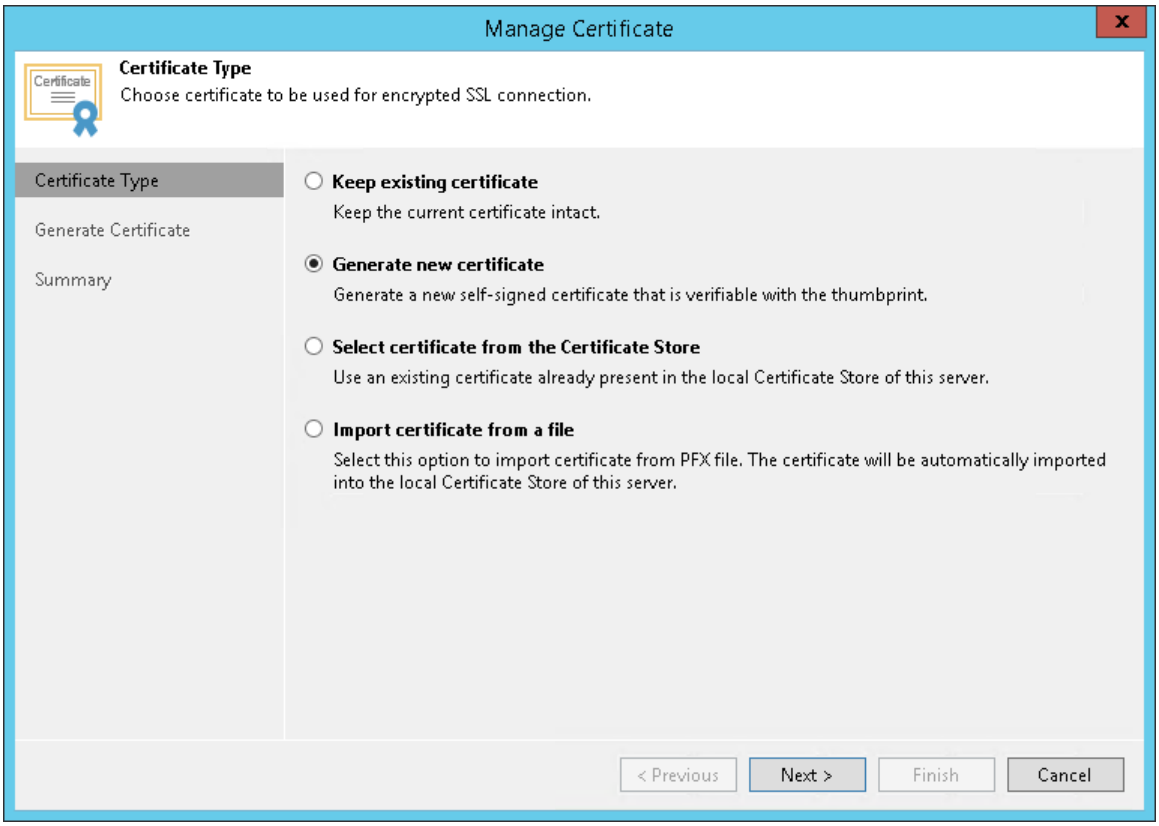
- User who created the TLS certificate
- LocalSystem user account
- Local Administrators group

If you use a self-signed TLS certificate generated by Veeam Backup & Replication, you do not need to take any additional actions to deploy the TLS certificate on a protected computer. When Veeam Backup & Replication discovers a protected computer, a matching TLS certificate with a public key is installed on the protected computer automatically. During discovery, Veeam Installer Service deployed on the protected computer retrieves the TLS certificate with a public key from the backup server and installs a TLS certificate with a public key on the protected computer.

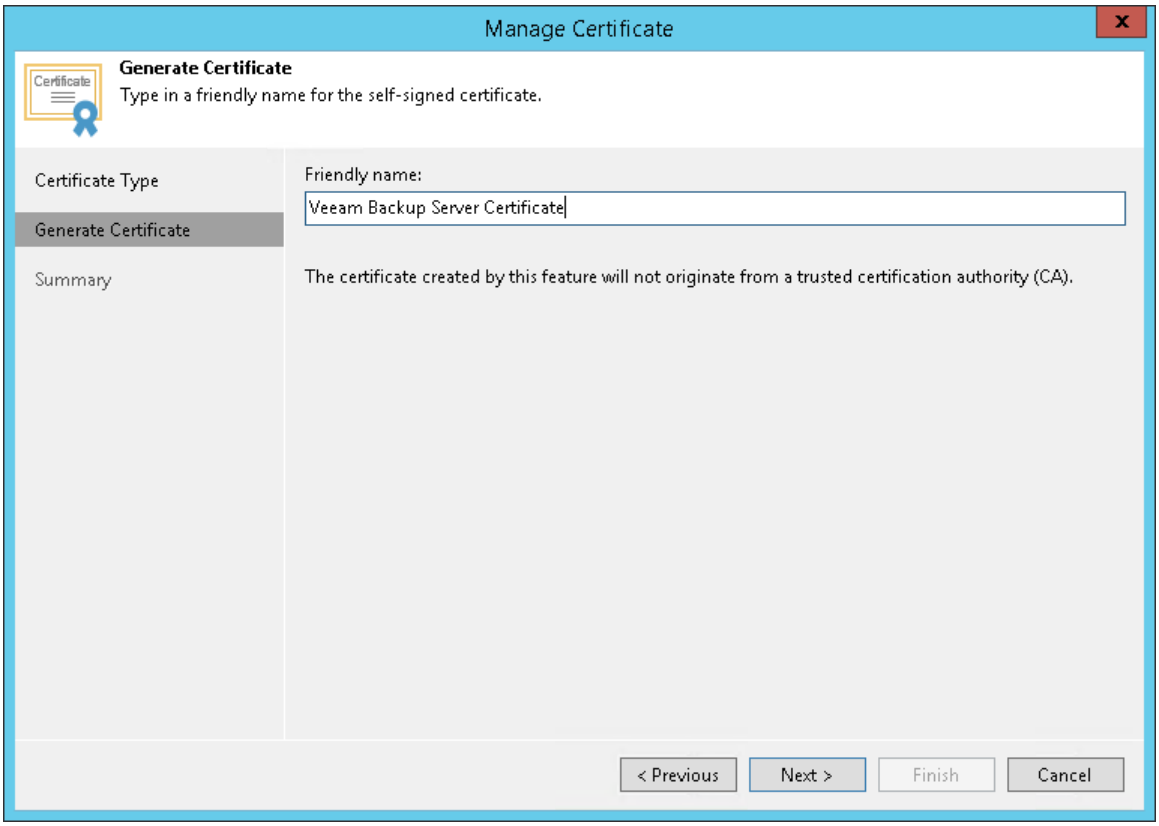
> **NOTE**
>
> When you generate a self-signed TLS certificate with Veeam Backup & Replication, you cannot include several aliases to the certificate and specify a custom value in the *Subject* field. The *Subject* field value is taken from the Veeam Backup & Replication license installed on the Veeam backup server.

To generate a self-signed TLS certificate:

1. From the main menu, select **General Options**.
2. Click the **Security** tab.
3. In the **Security** tab, click **Install**.
4. At the **Certificate Type** step of the wizard, select **Generate new certificate**.

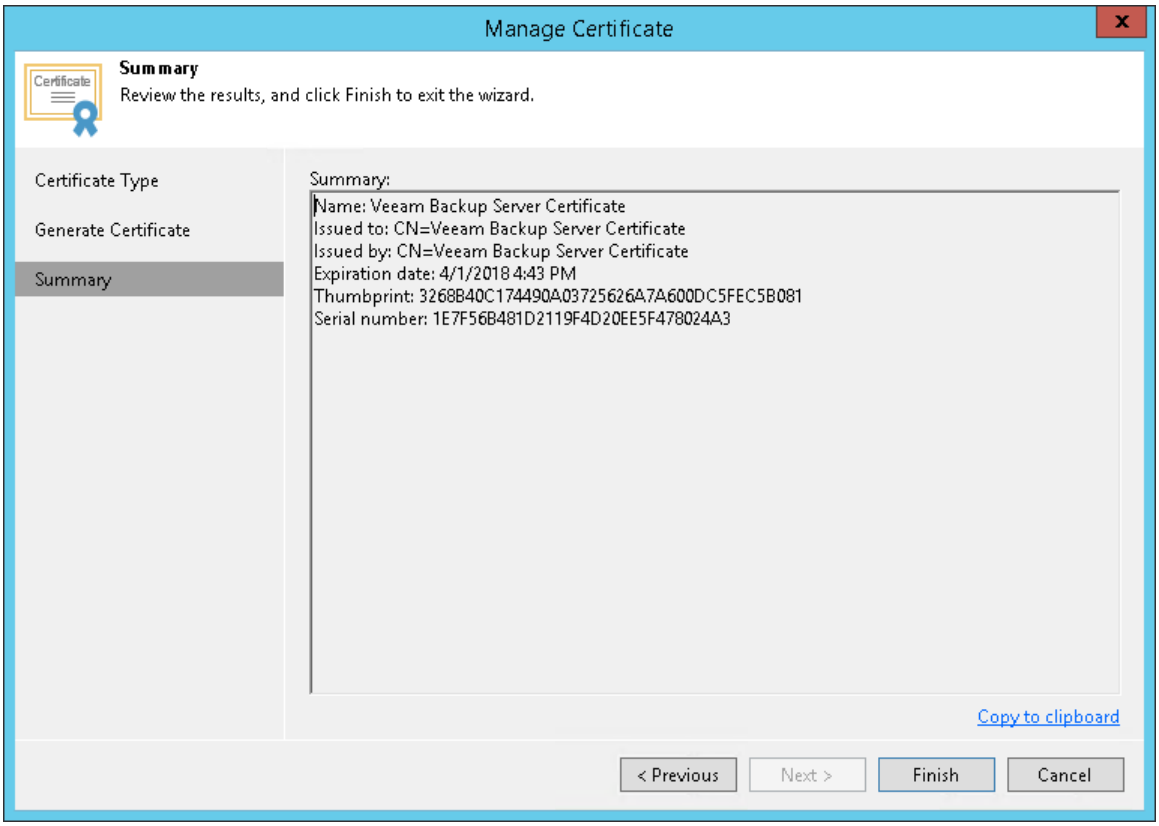


5. At the **Generate Certificate** step of the wizard, specify a friendly name for the created self-signed TLS certificate.



6. At the **Summary** step of the wizard, review the certificate properties. Use the **Copy to clipboard** link to copy and save information about the generated TLS certificate. You will be able to use the copied information to verify the TLS certificate with the certificate thumbprint.

7. Click **Finish**. Veeam Backup & Replication will save the generated certificate in the *Shared* certificate store on the Veeam backup server.



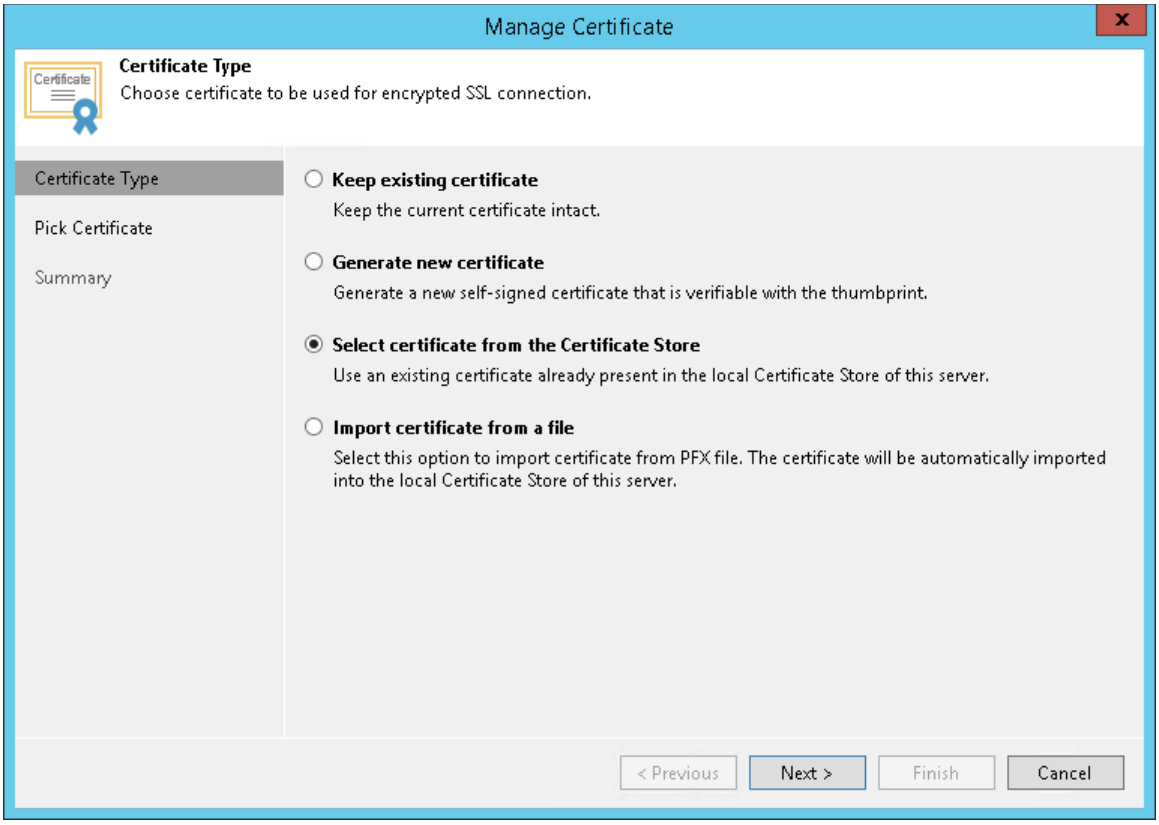
Importing Certificates from Certificate Store

If your organization has a TLS certificate signed by a CA and the TLS certificate is located in the Microsoft Windows certificate store, you can use this certificate for authenticating parties in the Veeam Backup & Replication infrastructure.

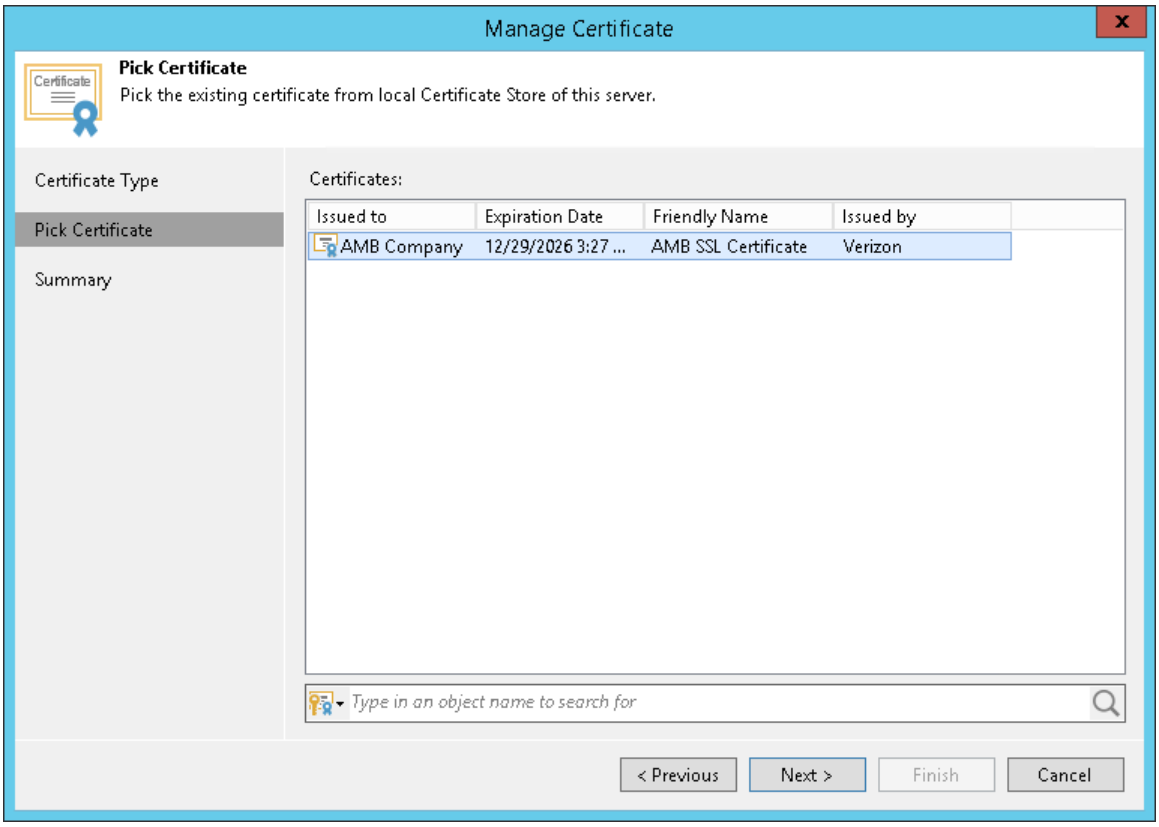
To select a certificate from the Microsoft Windows certificate store:

1. From the main menu, select **General Options**.
2. Click the **Security** tab.
3. In the **Security** tab, click **Install**.

4. At the **Certificate Type** step of the wizard, choose **Select certificate from the Certificate Store**.



5. At the **Pick Certificate** step of the wizard, select a TLS certificate that you want to use. You can select only certificates that contain both a public key and a private key. Certificates without private keys are not displayed in the list.



6. At the **Summary** step of the wizard, review the certificate properties.

7. Click **Finish** to apply the certificate.

Importing Certificates from PFX Files

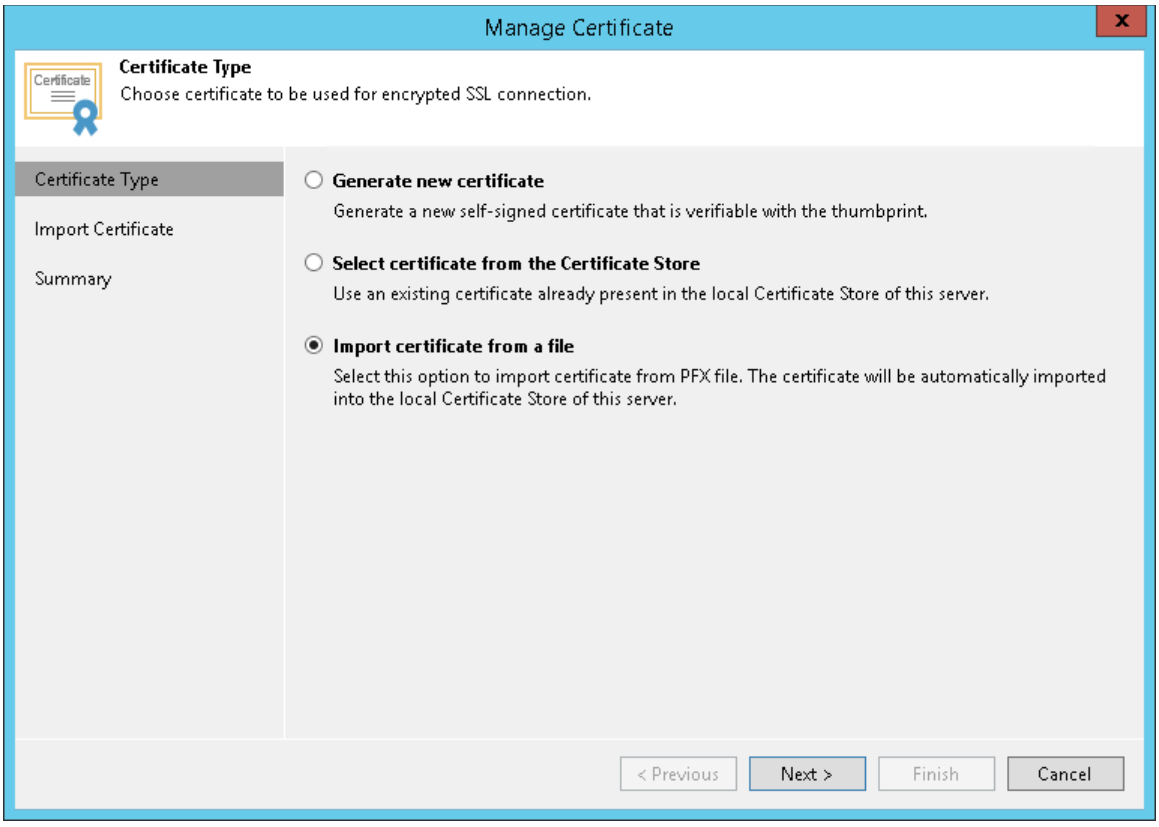
You can import a TLS certificate in the following situations:

- Your organization uses a TLS certificate signed by a CA and you have a copy of this certificate in a file of PFX format.
- You have generated a self-signed TLS certificate in the PFX format with a third-party tool and you want to import it to Veeam Backup & Replication.

> **IMPORTANT**
>
> Mind the following:
>
> - The TLS certificate must pass validation on the Veeam backup server. In the opposite case, you will not be able to import the TLS certificate.
> - If a PFX file contains a certificate chain, only the end entity certificate will be imported.

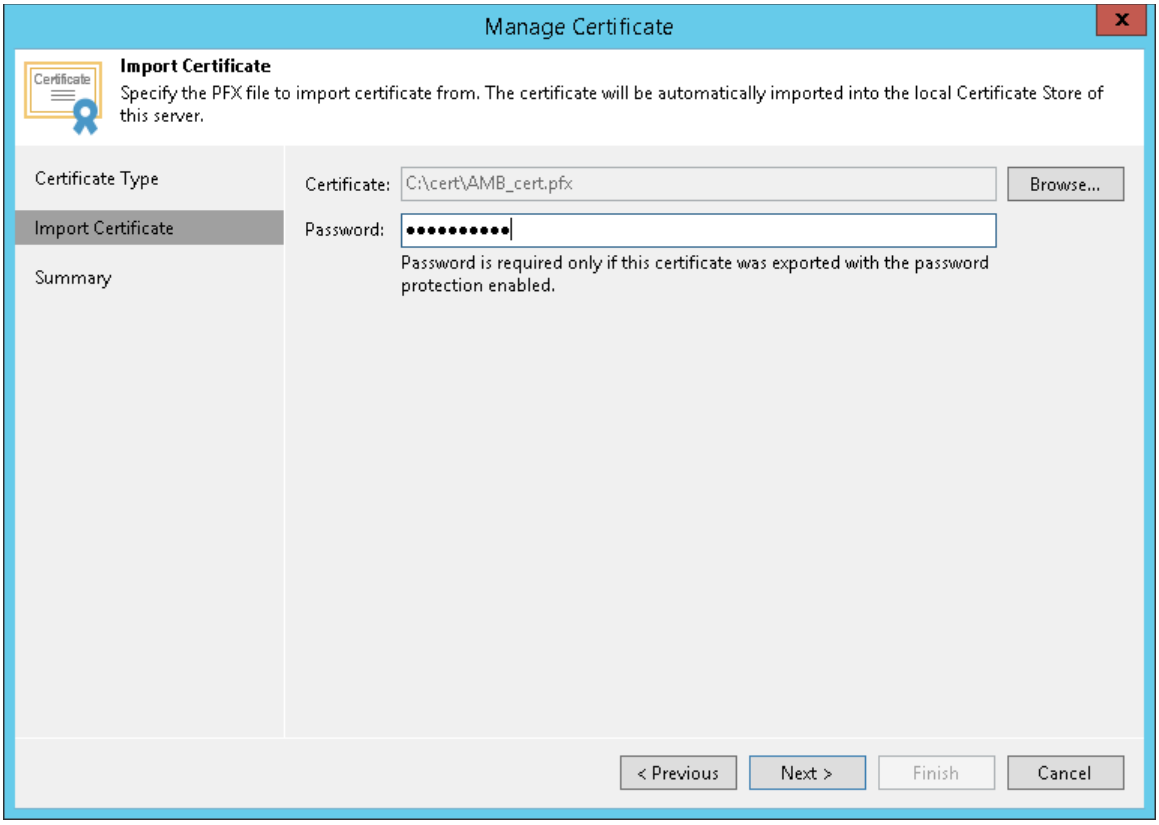
To import a TLS certificate from a PFX file:

1. From the main menu, select **General Options**.
2. Click the **Security** tab.
3. In the **Security** tab, click **Install**.
4. At the **Certificate Type** step of the wizard, choose **Import certificate from a file**.



5. At the **Import Certificate** step of the wizard, specify a path to the PXF file.

6. If the PFX file is protected with a password, specify the password in the field below.



7. At the **Summary** step of the wizard, review the certificate properties. Use the **Copy to clipboard** link to copy and save information about the TLS certificate. You can use the copied information on a protected computer to verify the TLS certificate with the certificate thumbprint.
8. Click **Finish** to apply the certificate.

Linux Host Authentication

In the **Linux hosts authentication** section of the Veeam Backup & Replication settings, you can specify SSH fingerprint verification settings for protected Linux machines.

You can select one of the following options:

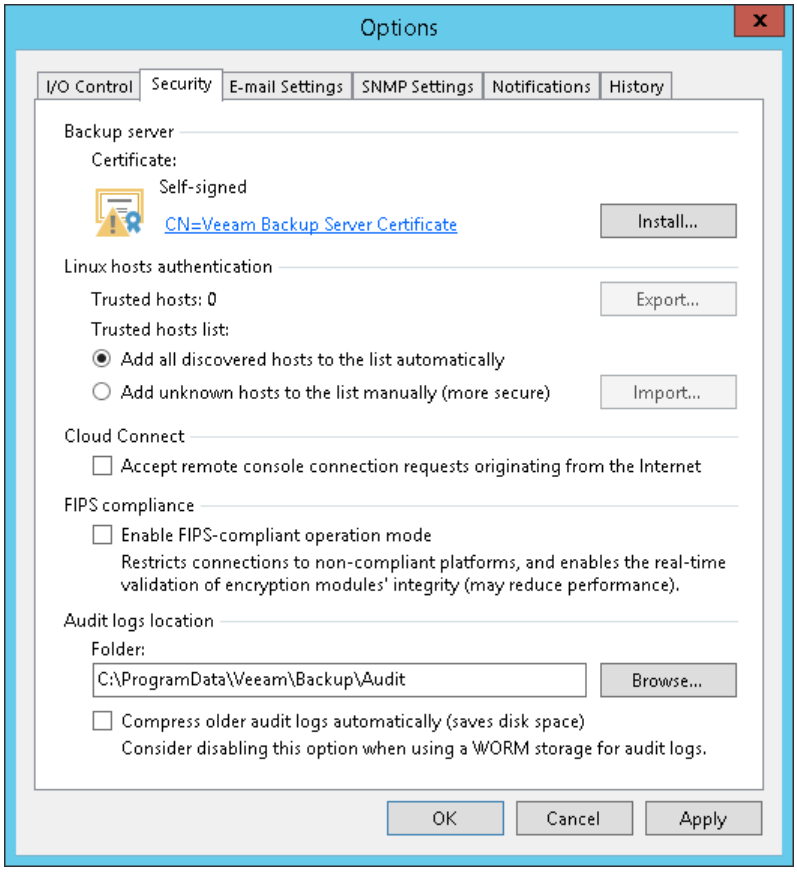
- **Add all discovered hosts to the list automatically** — with this option enabled, Veeam Backup & Replication allows all added Linux VMs and Linux servers to connect to the backup server.
- **Add unknown hosts to the list manually (more secure)** — with this option enabled, only the following Linux machines can connect to the backup server:
  - Protected machines that have already established a connection to the backup server and have their fingerprints stored in the Veeam Backup & Replication database.

    Veeam Backup & Replication displays the number of trusted machines in the **Trusted hosts** field. To export the list of trusted machines to the `known_hosts` file, click **Export** and specify a path to the folder to save the file.

- Protected machines specified in the `known_hosts` file imported to Veeam Backup & Replication. To import the `known_hosts` file, click **Import** and specify a path to the folder where the file resides.

  Machines that do not meet the above-mentioned conditions cannot connect to the Veeam backup server and download Veeam Agent for Linux installation packages during discovery. Also, guest OS processing of untrusted VMs will fail.

  Veeam Backup & Replication displays these computers under the **Untrusted** node in the inventory. To start managing an untrusted computer, you need to validate its fingerprint manually in the Veeam Backup & Replication console. To learn more, see Validating SSH Fingerprints.



Validating SSH Fingerprints

Veeam Backup & Replication treats fingerprints of Linux VMs and Linux hosts differently:

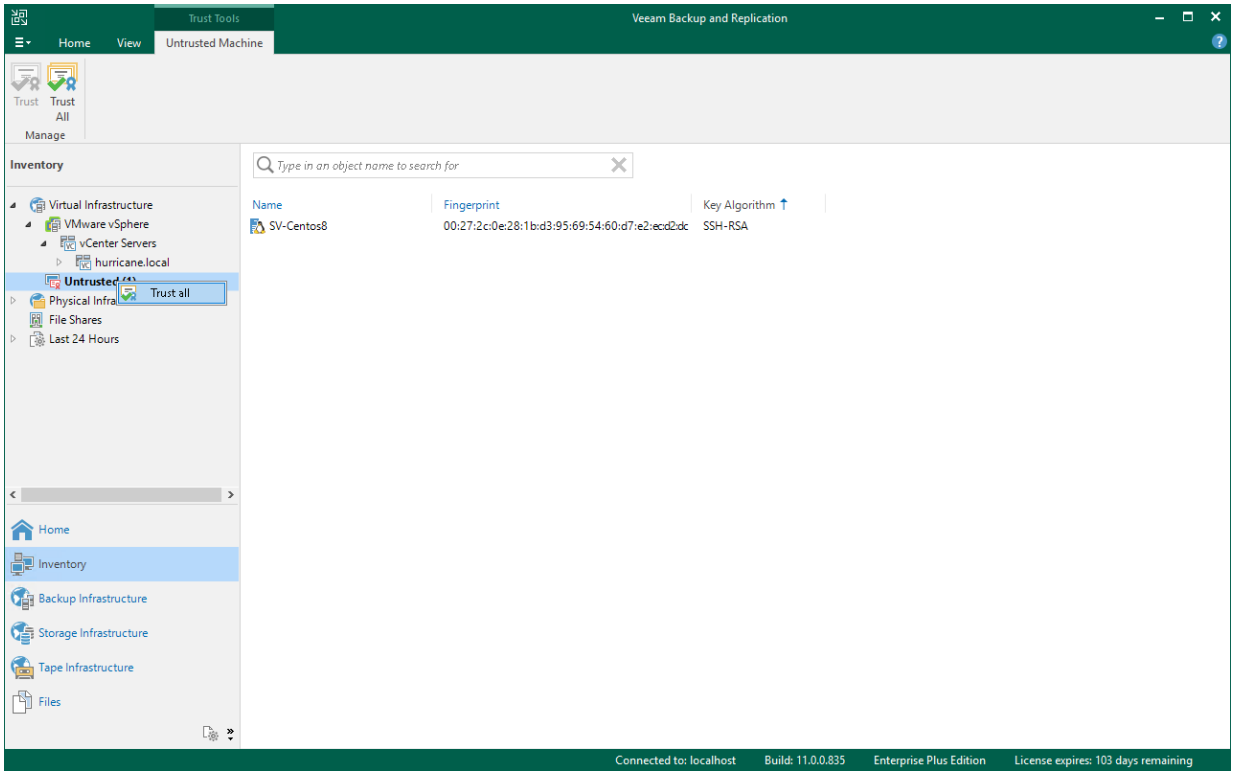
- Validating SSH Fingerprints of Linux VMs
- Validating SSH Fingerprints of Linux Hosts

Validating SSH Fingerprints of Linux VMs

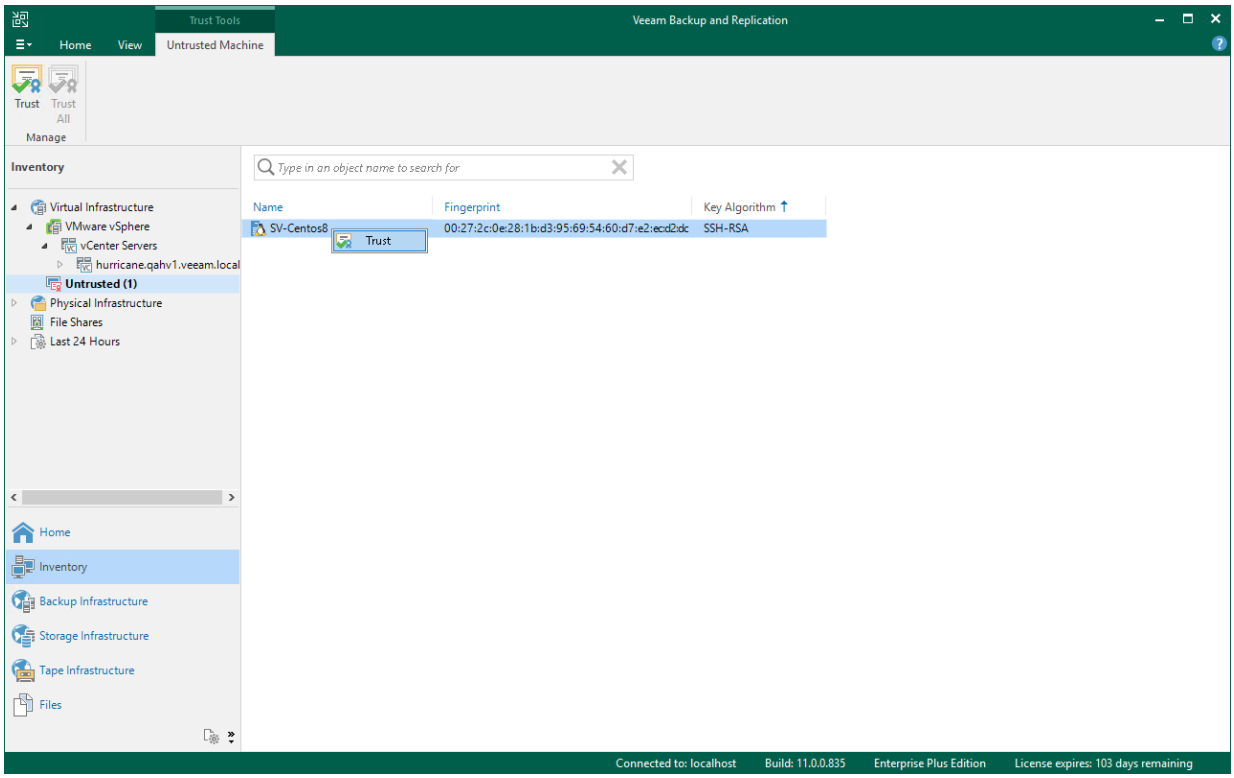
When you enable the **Add unknown hosts to the list manually (more secure)** option in Veeam Backup & Replication settings, Linux-based machines whose fingerprints are not stored in the Veeam Backup & Replication database or the `known_hosts` file become unable to communicate to the Veeam backup server. During discovery, Veeam Backup & Replication puts such machines to the *Untrusted* node. To start managing an untrusted computer, you need to validate its fingerprint manually in the Veeam Backup & Replication console.

To validate the SSH fingerprint:

1. Open the **Inventory** view.
2. In the inventory pane, expand the **Virtual Infrastructure** node and click **Untrusted**.
3. In the working area, Veeam Backup & Replication will display a list of machines whose fingerprints need to be validated. Check fingerprints of the machines in the list and validate them in one of the following ways:
   - To validate fingerprints of all untrusted machines at once, select the **Untrusted** node in the inventory pane and click **Trust All** on the ribbon. Alternatively, you can right-click the **Untrusted** node and select **Trust all**.



- To validate a fingerprint of a specific computer in the list, select the necessary computer in the working area and click **Trust** on the ribbon. Alternatively, you can right-click the computer and select **Trust**.

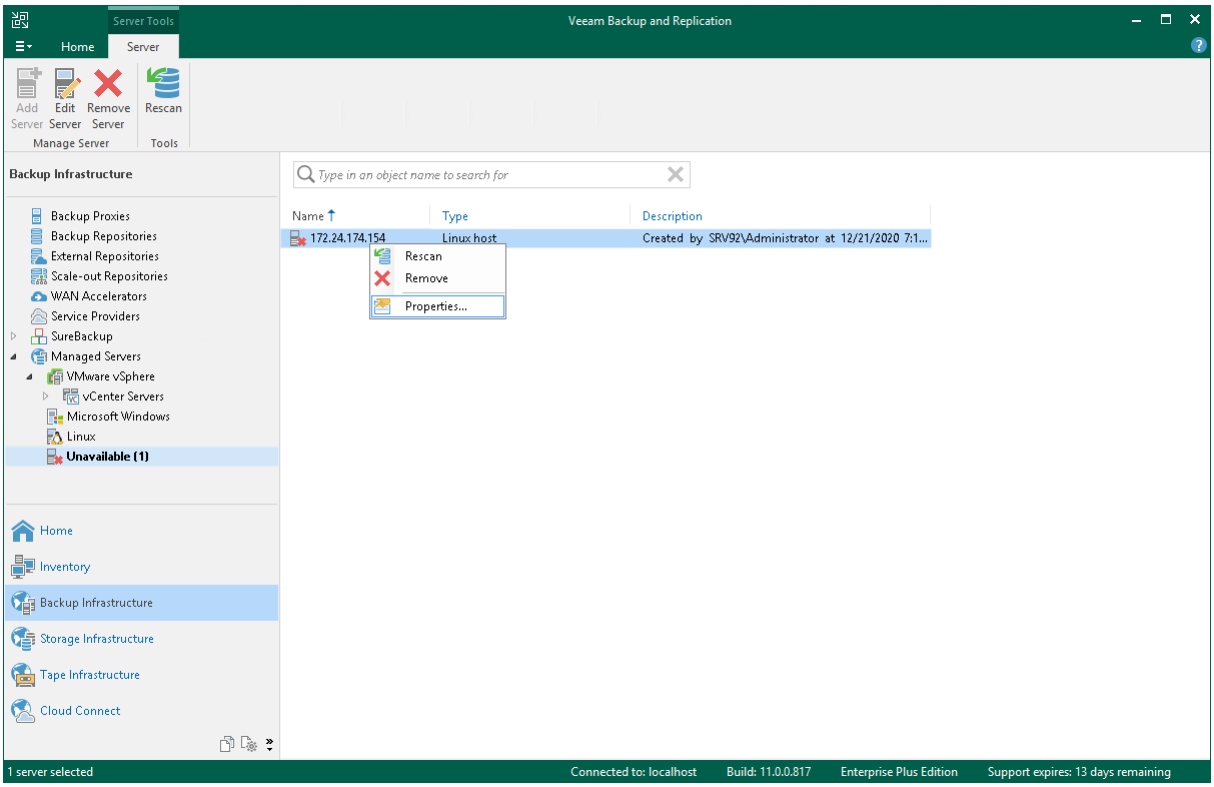


Validating SSH Fingerprints of Linux Hosts

If the SSH public key fingerprint of a Linux host is changed, Veeam Backup & Replication places this host in the **Unavailable** folder. To be able to use this server, do the following:

1. In the **Backup Infrastructure** view, expand the **Managed Servers** node and select **Unavailable.**
2. Right-click the Linux server and select **Properties.**
3. In the **SSH Connection** step of the **Edit Linux Server** wizard, click **Apply.**
4. In the pop-up dialogue window, click **Yes** to confirm that you trust this server.

5. Click **Finish** to close the wizard.



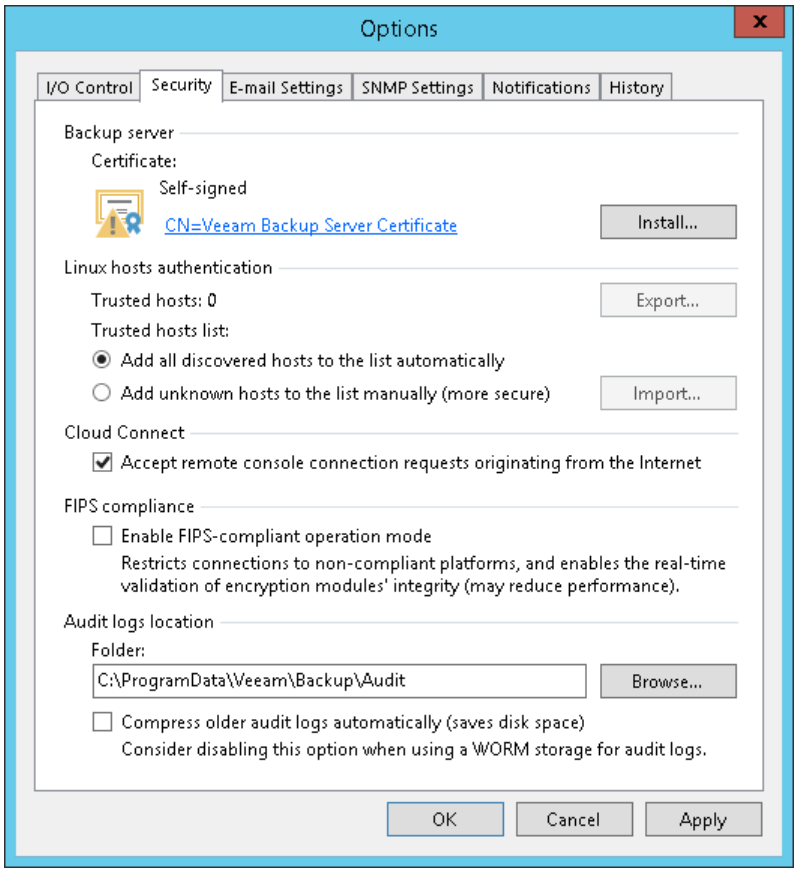
Cloud Connect

If you use the Veeam Cloud Connect functionality, you can remotely access the tenant backup server to manage Veeam Backup & Replication deployed on the tenant side. One of the ways to do this is to use the Remote Access Console. In case it is installed on a remote machine connected to an external network, you will need to enable access to the cloud gateway for the Remote Access Console over the internet. For more information, see Veeam Cloud Connect Guide.

To enable access to the cloud gateway for the Remote Access Console:

1. From the main menu, select **General Options**.
2. Open the **Security** tab.
3. In the **Cloud Connect** section, select the **Accept remote console connection requests originating from the Internet** check box.

4. Click **OK**.



FIPS Compliance

You can enable the FIPS policy to be used on all backup infrastructure components. For more information on FIPS, see this Microsoft article.

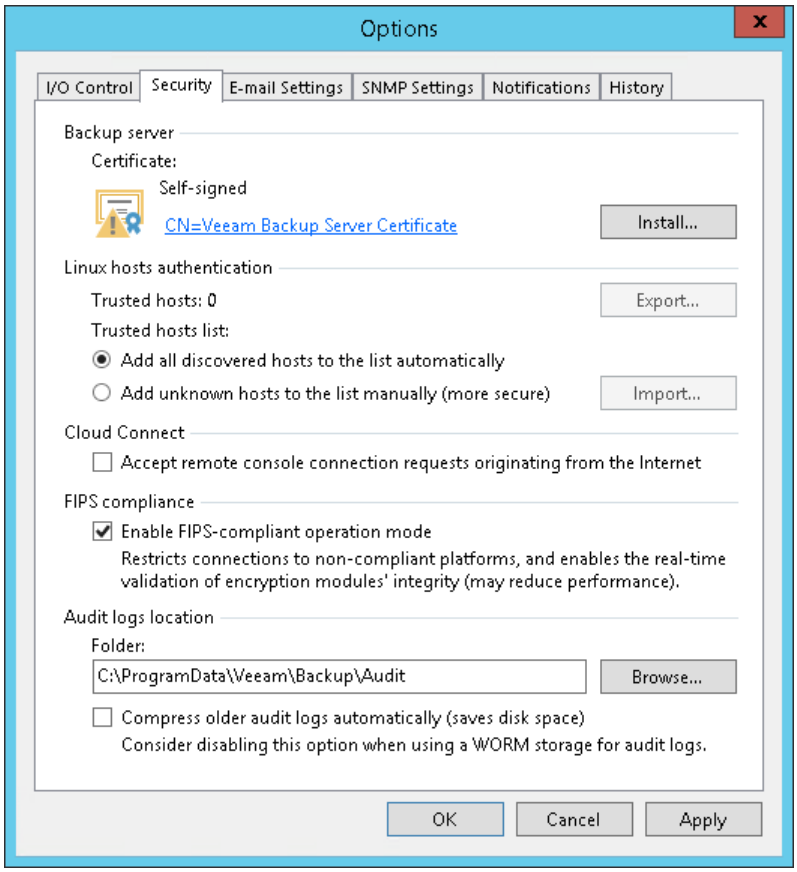
To enable the FIPS-compliant operation mode:

1. From the main menu on the backup server, select **General Options**.
2. Open the **Security** tab.
3. In the **FIPS compliance** section, select the **Enable FIPS-compliant operation mode** check box.
4. Click **OK**.

IMPORTANT

Mind the following for backup infrastructure components based on Linux servers with persistent Veeam Data Movers:

- After you select or clear the selection of the **Enable FIPS-compliant operation mode** check box, you must open the Edit Linux Server wizard for each Linux server with the persistent Veeam Data Mover and proceed to the end of the wizard. This will update server settings. If you do not update the settings, the servers will be unavailable.
- In Veeam Backup & Replication version 11, the component of persistent Linux Veeam Data Mover that is responsible for data processing launch cannot operate in the FIPS-compliant mode. However, components responsible for data processing itself can operate in this mode. Starting from Veeam Backup & Replication version 11a, whole Data Mover can become FIPS-compliant.

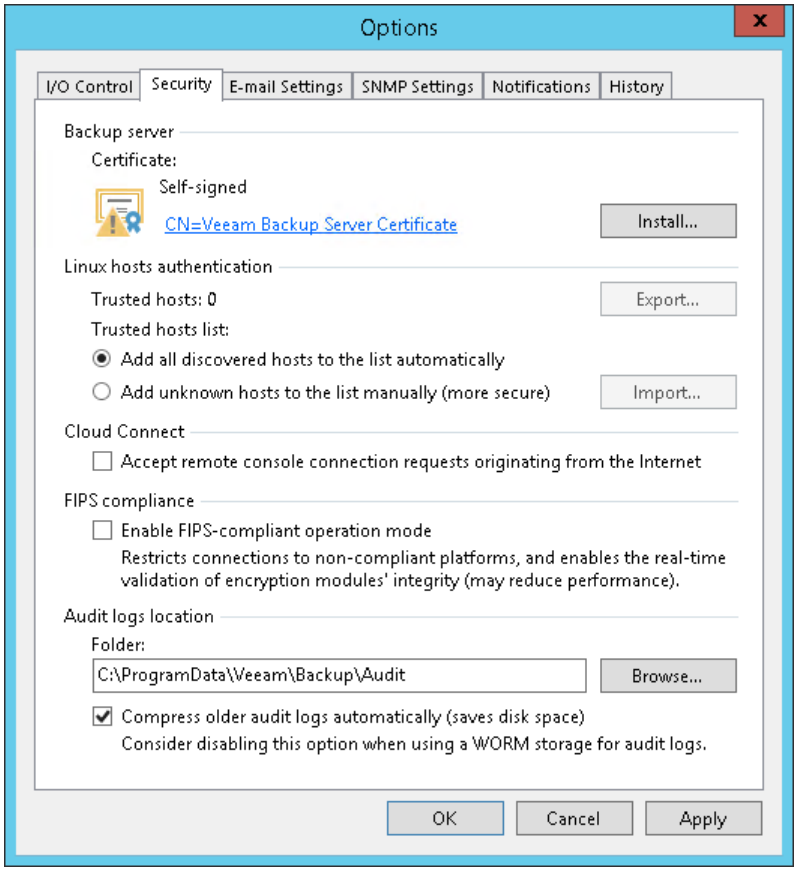


Audit Logs Location

Veeam Backup & Replication provides logging of performed activities, such as data protection and disaster recovery tasks: for example, the list of files restored during **File-Level Restore** sessions. The results of the audit of such activities are stored in form of `.csv` files that are called audit logs. For more information about log files, see Logging.

In **Audit Logs Location** field, you can specify folder where the audit logs will be stored. By default, log files are stored in the following folder: `%ProgramData%\Veeam\Backup\Audit`. You can also specify an SMB (CIFS) folder.

If you use an SMB (CIFS) folder, the service account that is used for Veeam Backup Service on the machine with Veeam Backup & Replication must have access to that SMB (CIFS) folder. By default, this is the **Local System** account, so you will need to grant write access to the **VBR Server Active Directory** computer account.



Specifying Email Notification Settings

You can receive email notifications with results on jobs performed on the backup server.

To receive email notifications, you must perform the following tasks:

- Configure global email notification settings in Veeam Backup & Replication
- Configure job notification settings

> **TIP**
>
> To receive email notification about all jobs performed on the backup server in one email, configure email notification settings in Veeam Backup Enterprise Manager.

Configuring Global Email Notification Settings

To configure global email notification settings:

1. From the main menu, select **General Options**.
2. Open the **E-mail Settings** tab.
3. Select the **Enable e-mail notifications** check box.
4. In the **SMTP server** field, enter a full DNS name or IP address of the SMTP server that will be used for sending email notifications.
5. Click the **Advanced** button to specify user credentials and connection options:
   a. Specify the port number and connection timeout for the SMTP server.
   b. To use a secure connection for email operations, select the **Connect using SSL** check box.
   c. If you need to connect to the SMTP server using a specific account, select the **This SMTP server requires authentication** check box and select the necessary credentials from the **Log on as** list. If you have not set up credentials beforehand, click the **Manage accounts** link or click **Add** on the right to add credentials. For more information, see Managing Credentials.
6. In the **From** field, specify an email from which email notifications must be sent.
7. In the **To** field, specify the recipient addresses. Use a semicolon to separate multiple addresses. Recipients specified in this field will receive notification about every job managed by the backup server. You can leave the field empty if required.

   For every particular job, you can specify additional recipients. For more information, see Configuring Job Notification Settings.

   > **NOTE**
   >
   > If you specify the same email recipient in both job notification and global notification settings, Veeam Backup & Replication will send two separate notifications only if a subject for the email message specified in job notification is different from the subject specified in global notification settings. Otherwise, Veeam Backup & Replication will suppress global notification settings and will send job notifications only.

8. In the **Subject** field, specify a subject for the sent message. You can use the following variables in the subject:
   a. *%Time%* — completion time
   b. *%JobName%*
   c. *%JobResult%*
   d. *%ObjectCount%* — number of VMs in the job
   e. *%Issues%* — number of VMs in the job that have been processed with the *Warning* or *Failed* status
9. In the **Send daily reports at** field, specify at what time Veeam Backup & Replication will send daily email reports.

   Daily reports are generated for different purposes throughout Veeam Backup & Replication:

   - Reports about processing results of scale-out repository data.

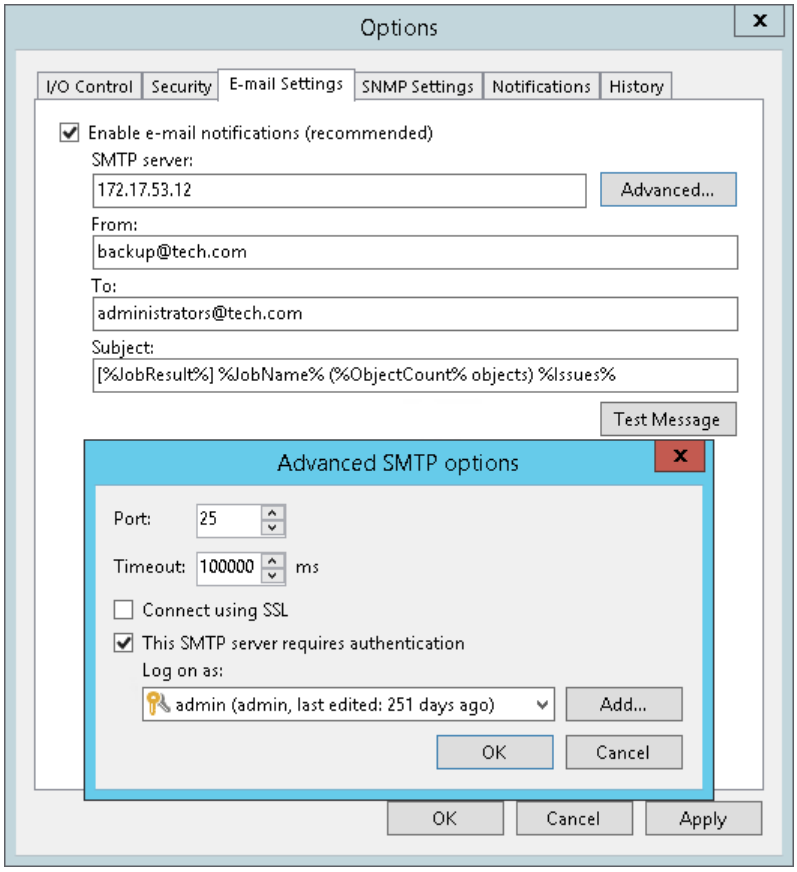
     For more information, see Receiving Scale-Out Backup Repository Reports.
   - Reports about processing results of backup copy jobs.

     For more information, see Notification Settings in **Creating Backup Copy Jobs for VMs and Physical Machines**.
   - Reports about processing results of backup copy jobs for transaction log backups.

     For more information about transaction log backups, see Microsoft SQL Server Logs Backup.
   - Reports about backups of virtual and physical machines created with Veeam Agent for Microsoft Windows or Veeam Agent for Linux in the Managed by Agent mode.
   - Reports with statistics for rescan job sessions performed for protection groups of virtual and physical machines created with Veeam Agent for Microsoft Windows or Veeam Agent for Linux.
   - Reports about processing results of backup copy jobs for backups of Oracle RMAN and SAP HANA databases created with Veeam Plug-ins for Enterprise Applications.

   > **NOTE**
   >
   > Settings configured for a certain report override global notification settings.
10. In the **Notify on** group, select the **Success**, **Warning** and/or **Failure** check boxes to receive email notification if a job is run successfully, not successfully or with a warning.
11. Select the **Suppress notifications until the last retry** check box to receive a notification about the final job status. If you do not enable this option, Veeam Backup & Replication will send one notification per every job retry.

12. Veeam Backup & Replication allows sending a test email to check if all settings have been configured correctly. To send a test email, click **Test Message**.



Configuring Job Notification Settings

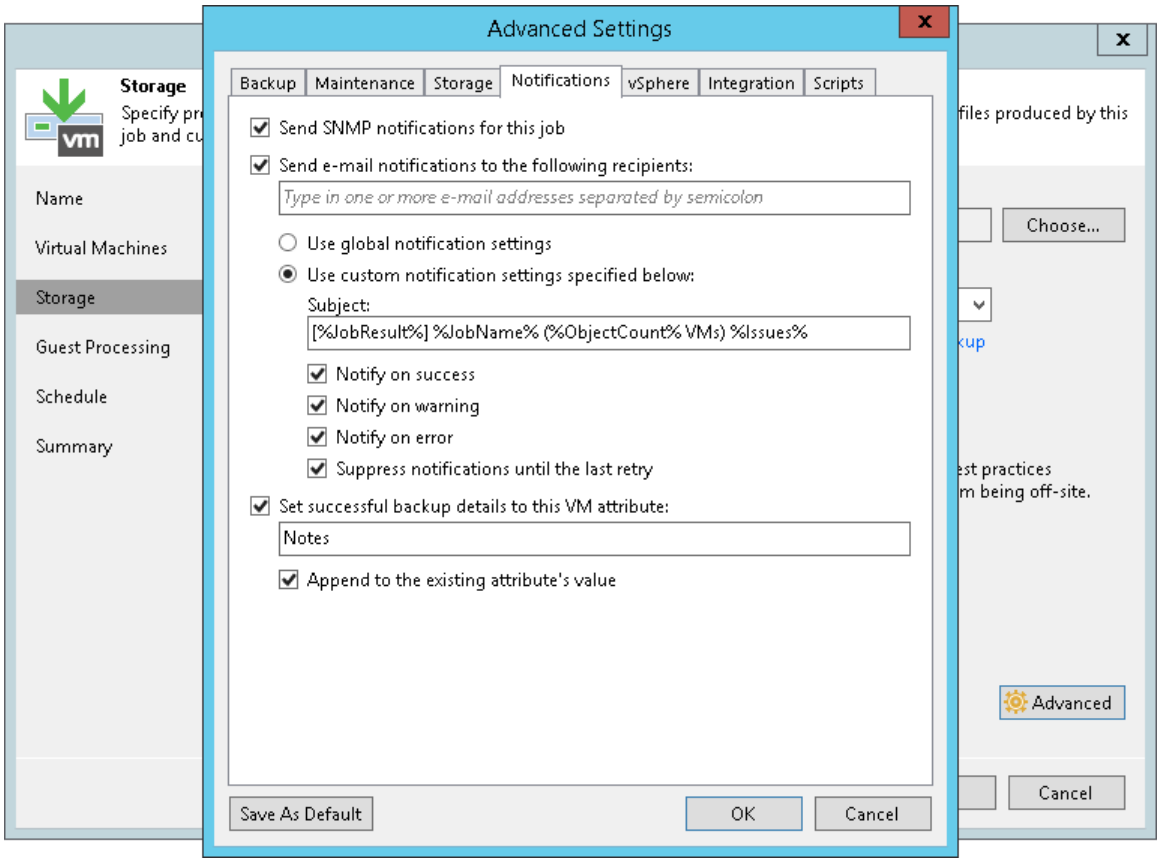
To configure job notification settings:

1. Open advanced settings of the job.
2. On the **Notifications** tab, select the **Send email notifications to the following recipients** check box.
3. In the field below, enter an email address to which a notification must be sent. You can enter several email addresses separated with a semicolon.

> **NOTE**
>
> If you specify the same email recipient in both job notification and global notification settings, Veeam Backup & Replication will send two separate notifications to this recipient in the following cases:
>
> - If a subject for the email message specified in job notification and global notification settings is different.
> - If a list of email recipients specified in job notification and global notification settings is different.

4. You can choose to use global notification settings for the job or specify custom notification settings.
    - To receive a typical notification for the job, select **Use global notification settings**. In this case, Veeam Backup & Replication will apply to the job global email notification settings specified for the backup server. For more information, see Configuring Global Email Notification Settings.

- To configure a custom notification for the job, select **Use custom notification settings** and specify notification settings as required.



Specifying SNMP Settings

You can receive SNMP traps with results on jobs performed on the backup server. You can use SNMP traps to feed data to other monitoring systems such as CA Unicenter, BMC Patrol, IBM Tivoli or HPE OpenView. SNMP traps can be sent to 5 different destinations.

Veeam Backup & Replication supports SNMP versions 1 and 2.

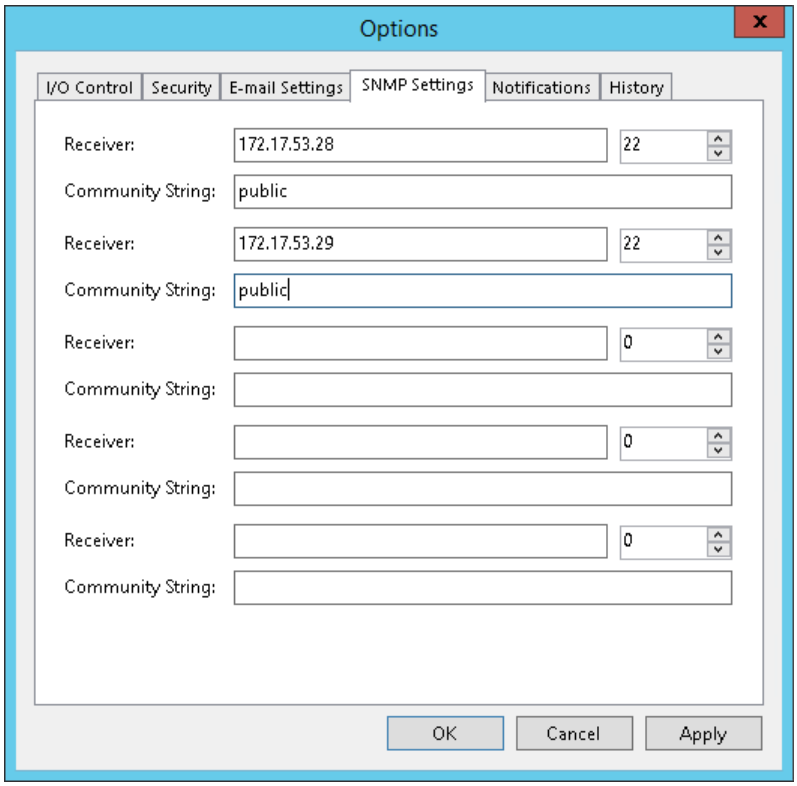
To receive SNMP traps, you must perform the following tasks:

- Configure global SNMP settings
- Configure SNMP service properties
- Configure SNMP settings for jobs

Configuring Global SNMP Settings

To configure global SNMP settings:

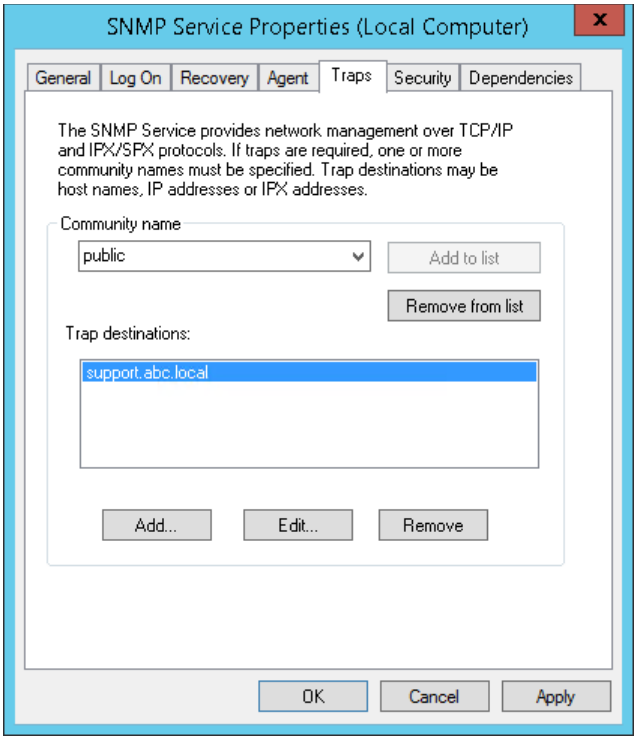
1. From the main menu, select **General Options**.
2. Click the **SNMP Settings** tab.
3. In the **Receiver** field, specify an IP address of the SNMP recipient.
4. In the field on the right, enter the port number to be used.
5. In the **Community String** field, enter the community identifier.



Configuring SNMP Service Properties

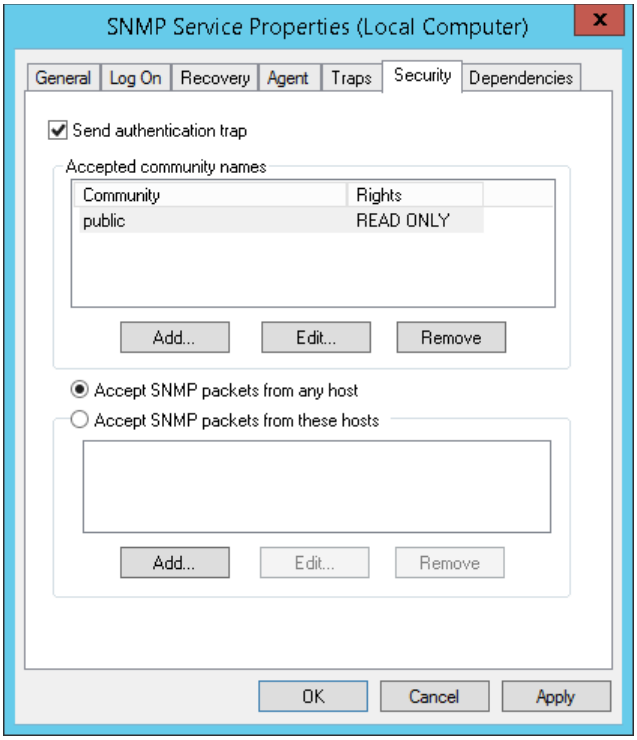
To configure SNMP service properties on recipient systems:

1. Install a standard Microsoft SNMP agent from the Microsoft Windows distribution on the computer.
2. From the **Start** menu, select **Control Panel** > **Administrative Tools** > **Services**.
3. Double-click **SNMP Service** to open the **SNMP Service Properties** window.
4. Click the **Traps** tab.
5. Add the public string to the **Community name** list and name of the necessary host to the **Trap destinations** list.



6. Click the **Security** tab.
7. Make sure the **Send authentication trap** check box is selected.
8. Add the public string to the **Accepted community names** list.
9. Select the **Accept SNMP packets from any host** check box.

10. Click **OK** to save changes.



TIP

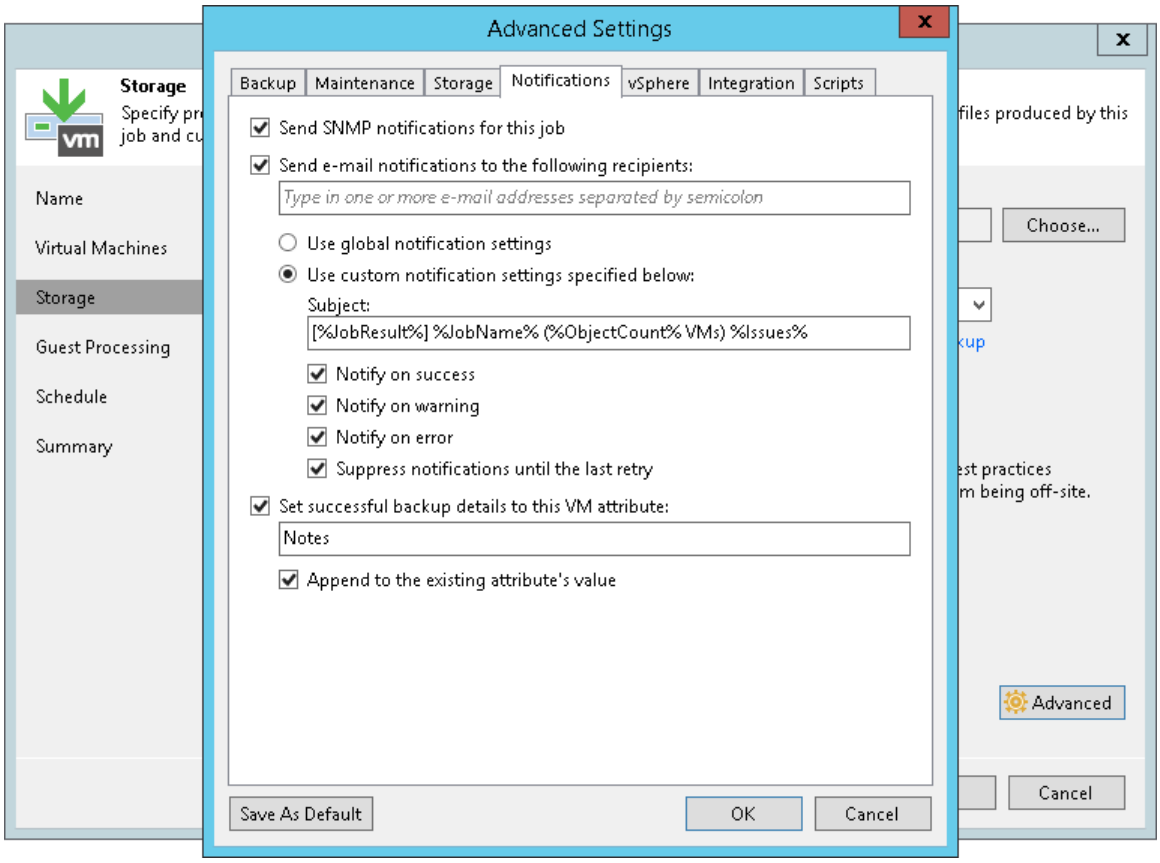
To interpret the traps incoming from the backup server, import the `VeeamBackup.mib` file to your recipient systems. The backup server stores this file in the `<vbr_installation_folder>\Backup\` folder. The installation folder is specified at the Program features step of the installation wizard.

Configuring SNMP Settings for Jobs

To receive SNMP traps with results of a specific job:

1. Open advanced settings of the job.

2. On the **Notifications** tab, select the **Send SNMP notifications for this job** check box.



Specifying Other Notification Settings

You can configure Veeam Backup & Replication to automatically notify you about the following events:

- Low disk space
- Support contract expiration
- New product versions, available updates

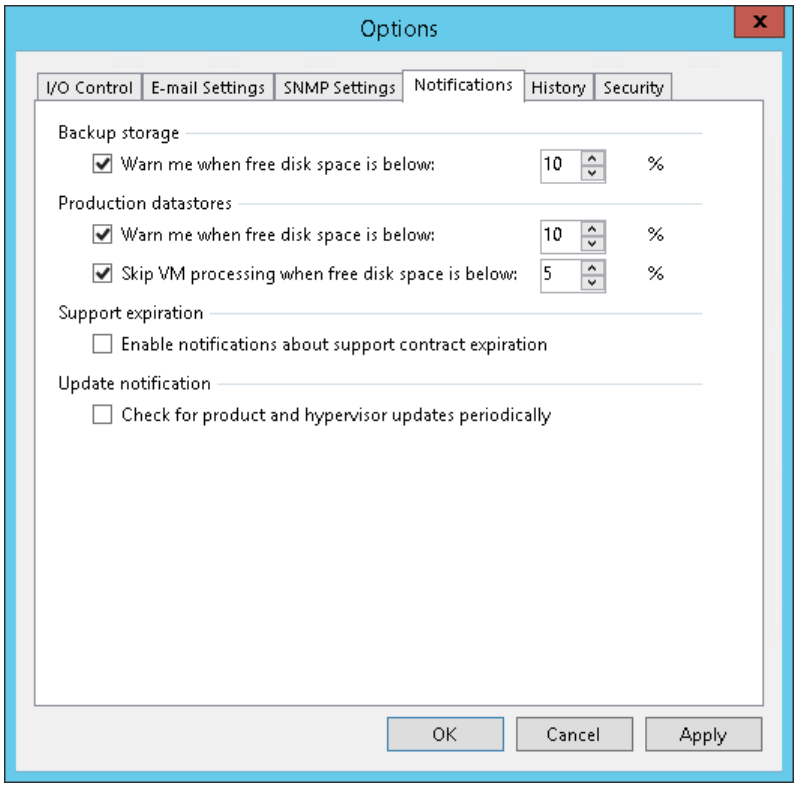
Low Disk Space Notification

When you run a job, Veeam Backup & Replication checks disk space in the target backup repository and production storage. If the disk space is below a specific value, Veeam Backup & Replication will display a warning message in the job session details.

To specify the disk space threshold:

1. From the main menu, select **General Options**.
2. Click the **Notifications** tab.
3. In the **Backup storage** and **Production datastores** sections, select the **Warn me when free disk space is below <N> %** options and specify a desired disk space threshold.
4. In the **Production datastores** section, select the **Skip VMs when free disk is below <N> %** option and specify a desired disk space threshold. When the threshold is reached, Veeam Backup & Replication will terminate backup and replication jobs working with production datastores before VM snapshots are taken. Such behaviour helps ensure that production datastores do not run out of space.

   Veeam Backup & Replication also terminates jobs if the amount of free space on the datastore is below 2 GB. You can change this threshold limit with registry keys. For more information, contact Veeam Customer Support.

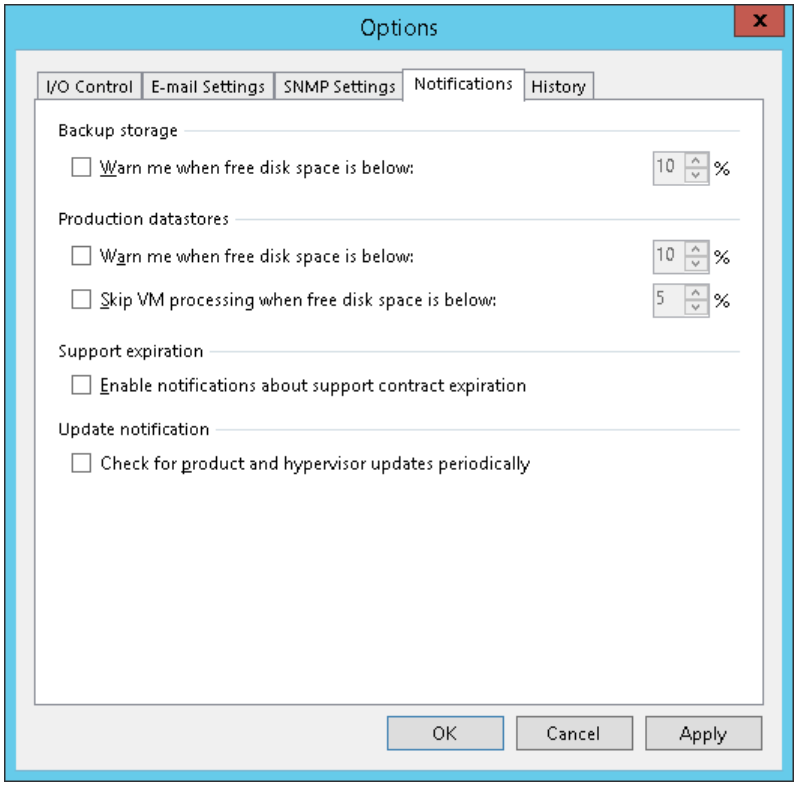


Support Contract Expiration Notification

By default, Veeam Backup & Replication informs email recipients specified in global notification settings about the support expiration date in every email notification. Veeam Backup & Replication starts sending such notifications 14 days before the expiration date. Expiration information is also shown on the splash screen and on the **License Information** window (to display the **License Information** window, select **Help** > **License** from the main menu).

To stop receiving notifications about support contract expiration:

1. From the main menu, select **General Options**.
2. Click the **Notifications** tab.
3. Clear the **Enable notifications about support contract expiration** check box.



New Product Versions

Veeam Backup & Replication automatically checks and notifies you about new product versions and updates available on the Veeam website. For more information, see Update Notification.

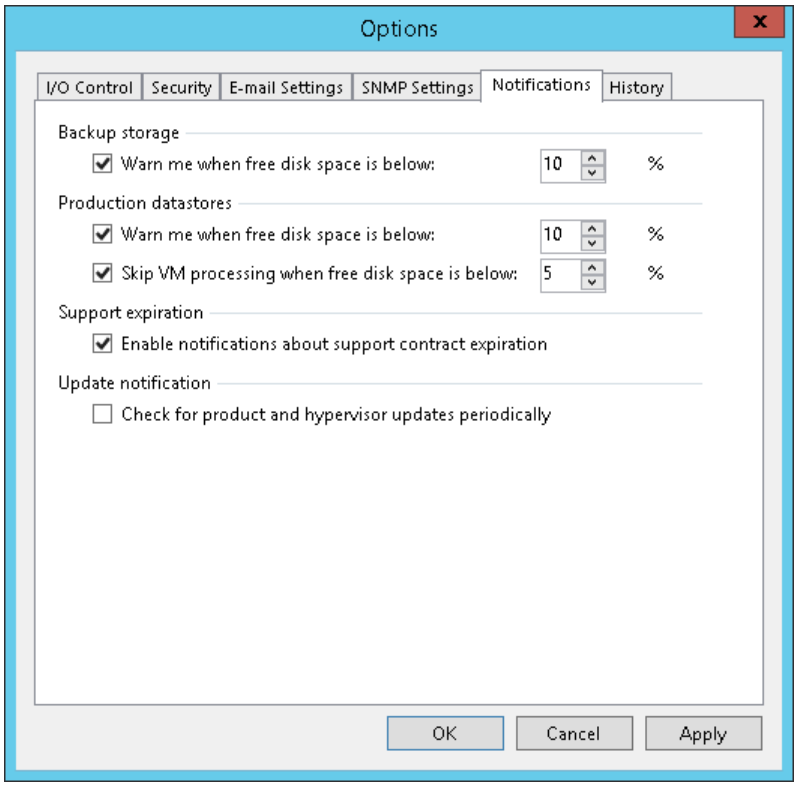
You can disable update notifications. However, it is recommended that you leave the update notifications enabled not to miss critical updates.

To stop receiving the notifications:

1. From the main menu, select **General Options**.
2. Click the **Notifications** tab.
3. Clear the **Check for product and hypervisor updates periodically** check box.

IMPORTANT

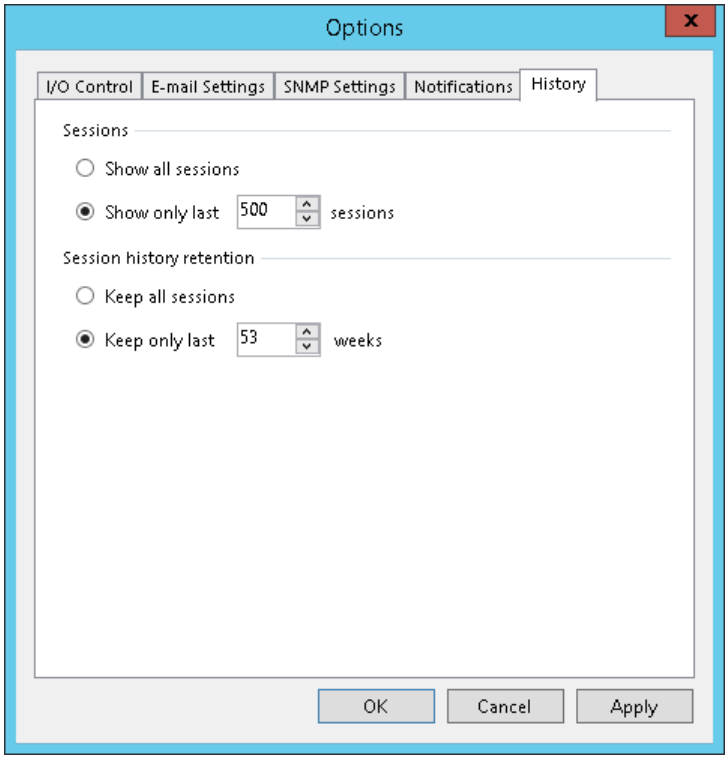
Make sure that the backup server is connected to the Internet. In the opposite case, you will not be able to receive the notifications about updates.



Specifying Session History Settings

You can specify session history settings for jobs performed on the backup server.

1. From the main menu, select **General Options**.
2. Click the **History** tab.
3. In the **Sessions** section, configure how many sessions you want to display in the **Sessions** list of the **History** view:
   - Select **Show all sessions** if you want to show all sessions.
   - Select **Show only last sessions** if you want to show a limited number of sessions. In the field, also specify the total number of sessions that you want to display.
4. In the **Session history retention** section, configure for how long you want to keep session information in the database:
   - Select **Keep all sessions** if you do not want to delete sessions.
   - Select **Keep only last weeks** if you want to keep sessions for a limited period of time. In the field, also specify this period in weeks.



Roles and Users

You can assign one of the following roles to users or groups of users who plan to work with Veeam Backup & Replication:

- Veeam Restore Operator
- Veeam Backup Viewer
- Veeam Backup Operator
- Veeam Backup Administrator
- Veeam Tape Operator

A role assigned to the user defines the user activity scope: what operations in Veeam Backup & Replication the user can perform. Role security settings affect the following operations:

- Starting and stopping jobs
- Performing restore operations

Users having different roles can perform a different set of operations:

Role	Operations
Veeam Restore Operator	Can perform restore operations using existing backups and replicas. Mind that during restore, Veeam Restore Operator can overwrite existing instances: VMs during VM restore, disks during disk restore and files during file-level restore.
Veeam Backup Viewer	Has the "read-only" access to Veeam Backup & Replication. Can view existing jobs and review the job session details.
Veeam Backup Operator	Can start and stop existing jobs, export backups and create VeeamZip backups.
Veeam Backup Administrator	Can perform all administrative activities in Veeam Backup & Replication. Mind that with the Veeam Backup & Replication console, Veeam Backup Administrator has full access to all files on servers and hosts added to the backup infrastructure.
Veeam Tape Operator	Can manage tapes and perform the following operations: tape inventory, tape export, tape eject, tape catalog, inventory library, catalog library, rescan library, import tapes, eject tape from drive.

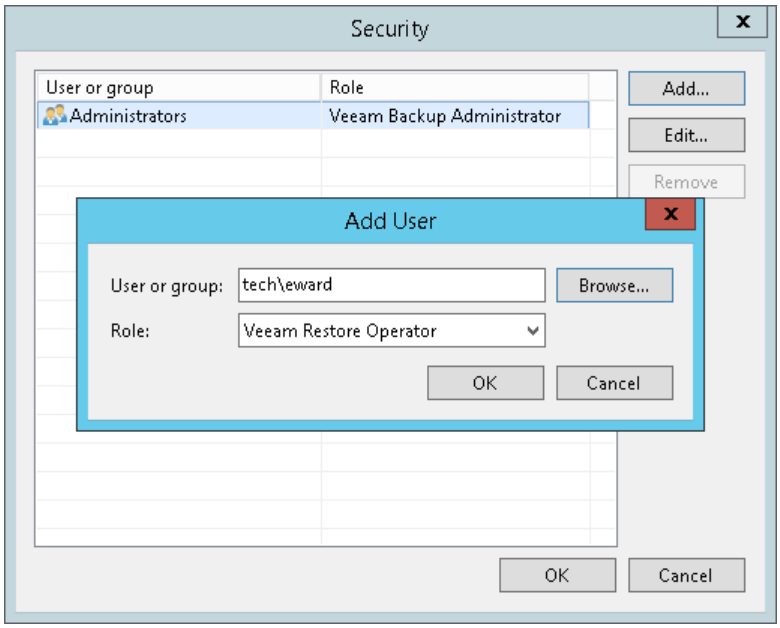
You can assign several roles to the same user. For example, if the user must be able to start jobs and perform restore operations, you can assign the **Veeam Backup Operator** and **Veeam Restore Operator** roles to this user.

Mind the following:

- The user account under which the Veeam Backup Service runs must have the Veeam Backup Administrator role. By default, during installation the Veeam Backup Administrator role is assigned to users in the *Administrators* group. If you change the default settings, make sure that you assign the Veeam Backup Administrator role to the necessary user account. It is recommended to assign the Veeam Backup Administrator role to the user account explicitly rather than the group to which the user belongs.
- Built-in administrator accounts (Domain\Administrator and Machine\Administrator) always have full access to Veeam Backup & Replication, even if you exclude them from all Veeam Backup & Replication roles. If you delete the *Administrators* group from the Veeam Backup & Replication roles, the users who are added to this group will still have access to Veeam Backup & Replication.

To assign a role to the user or user group:

1. From the main menu, select **Users and Roles**.
2. Click **Add**.
3. In the **User or group** field, enter a name of a user or user group in the *DOMAIN|USERNAME* format.
4. From the **Role** list, select the necessary role to be assigned.

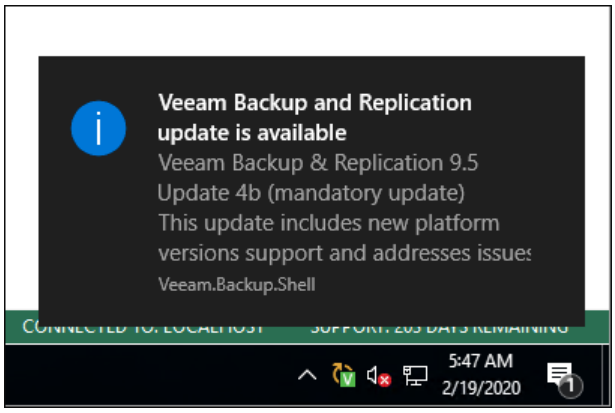


Update Notification

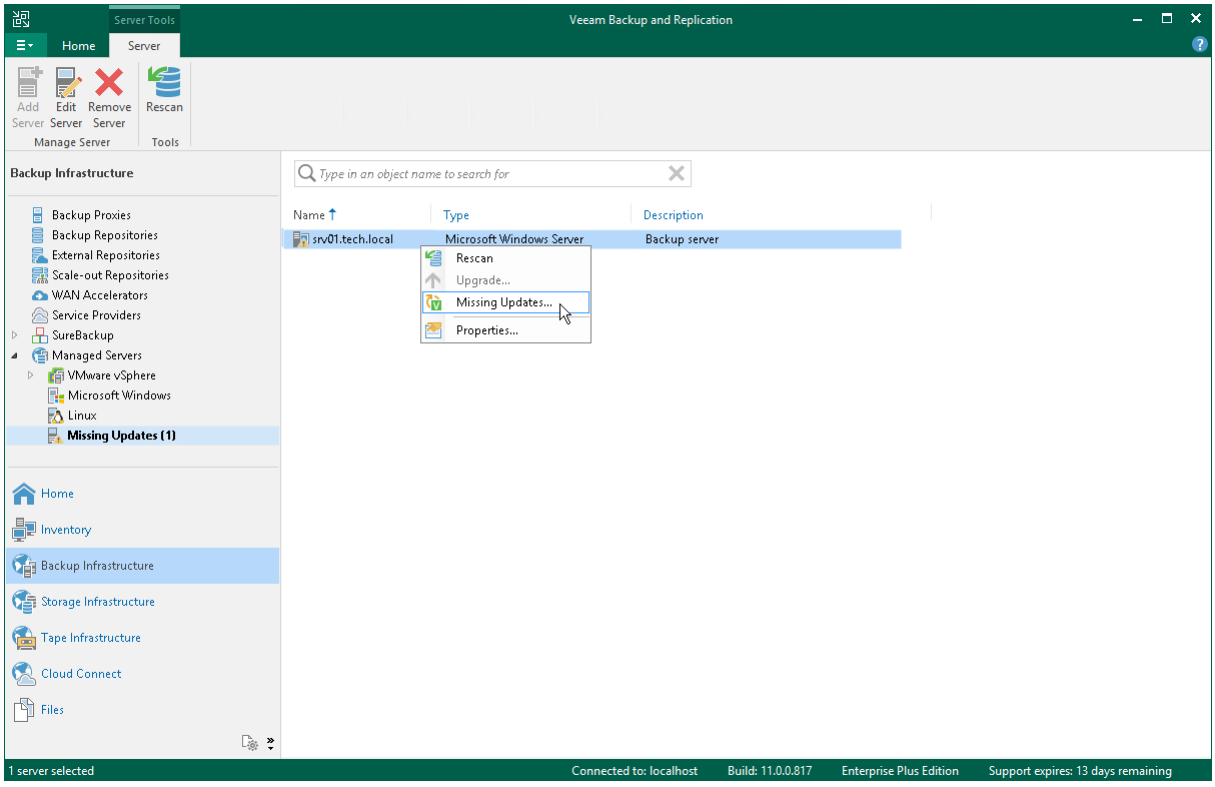
Automatic Update Notification

Veeam Backup & Replication may automatically notify you about updates that must or can be installed to enhance your work experience with the product. Update notifications eliminate the risk of using out-of-date components in the backup infrastructure or missing critical updates that can have a negative impact on data protection and disaster recovery tasks.

After a new build of Veeam Backup & Replication is published on the Veeam update server, the backup console will display a notification in the Windows Action Center (or an icon in the system tray for earlier Windows versions). If the update is not installed, this notification will keep appearing once a week as a reminder.



You can also see available updates in the **Managed Servers** > **Missing Updates** node in the **Backup Infrastructure** view.



The update notifications are enabled by default. You can disable notifications by clearing the **Check for product and hypervisor updates periodically** check box on the **General Options** > **Notifications** tab, as described in Specifying Other Notification Settings. However, it is recommended that you leave update notifications enabled not to miss critical updates.

Veeam Backup & Replication notifies about new product versions and new product updates.

How Update Notification Works

To check for updates, Veeam Backup & Replication uses a special XML file on the Veeam Update Notification Server (dev.veeam.com). The XML file contains information about the most up-to-date product version and updates.

Veeam Backup & Replication downloads an XML file from the Veeam Update Notification Server once a week. It also collects information about the installed product. The collected information is compared with the information in the downloaded file. If new product versions and updates are available, Veeam Backup & Replication informs you about them.

> **NOTE**
>
> Make sure that the backup server is connected to the Internet and update notification is enabled in Veeam Backup & Replication options. In the opposite case, update notification will not function.

Installing Updates

To install a product update, double-click the Veeam Backup & Replication notification in the Windows Action Center (or an icon in the system tray for earlier Windows versions). Veeam Backup & Replication will open a KB webpage with the update description and links to the installation archive of the new product version or new update.

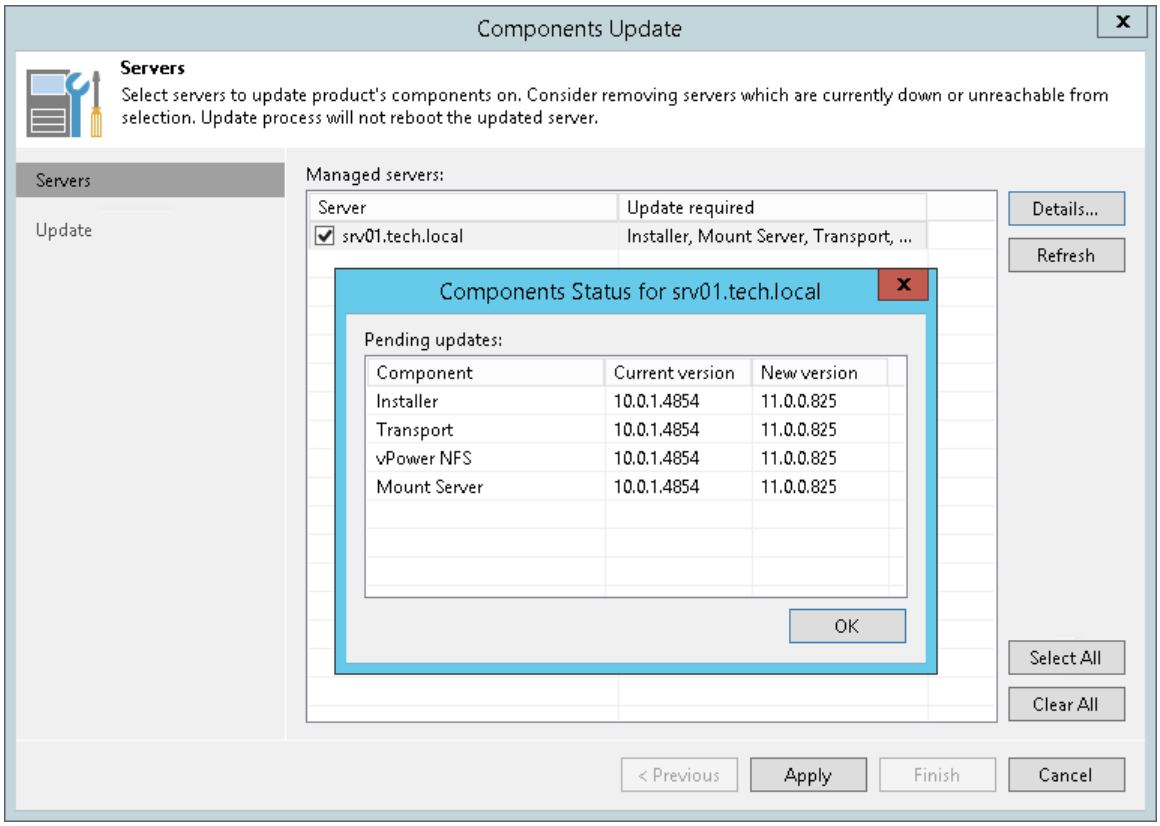
Server Components Upgrade

Every time you launch the Veeam Backup & Replication console, Veeam Backup & Replication automatically checks if Veeam Backup & Replication components installed on managed servers are up to date. If a later version of components is available, Veeam Backup & Replication displays the **Components Update** window and prompts you to upgrade components on managed servers. Components upgrade may be necessary, for example, after you have upgraded Veeam Backup & Replication.

You can manually check if components upgrade is required. To do this, select **Upgrade** from the main menu. If components on all managed servers are up to date, the menu item will be disabled.

To upgrade components on managed servers:

1. In the **Components Update** window, select a server and click **Details**. Veeam Backup & Replication will display the current and latest available versions for installed components.
2. In the **Components Update** window, select check boxes next to servers for which you want to upgrade components and click **Next**.



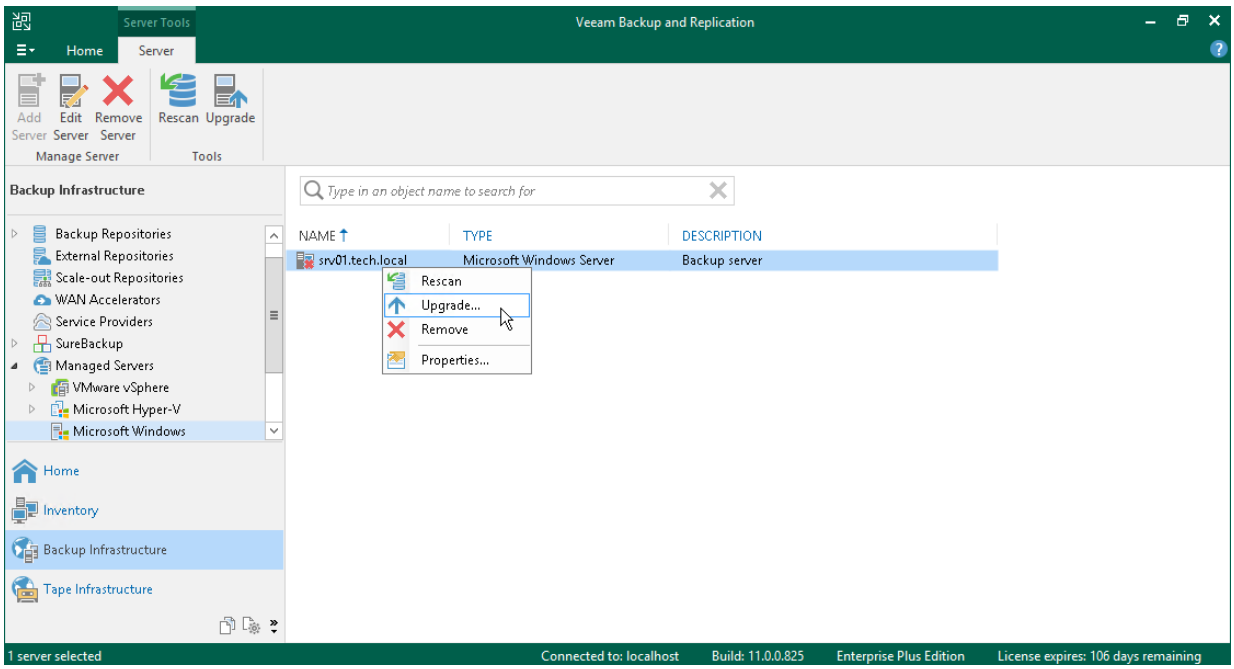
You can update components on every managed server separately. If components installed on the server require upgrade, Veeam Backup & Replication displays a warning icon next to the server.

To update components for a managed server:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, click **Managed servers**.

3. In the working area, select the server and click **Upgrade** on the ribbon.

   Alternatively, you can open the **Infrastructure** view, in the inventory pane select **Managed servers**, in the working area right-click the server and select **Upgrade**.



Logging

Veeam Backup & Replication provides detailed logging of performed activities data protection and disaster recovery tasks.

On the backup server, log files are stored in the following folder: `%ProgramData%\Veeam\Backup`. If you run Veeam Backup & Replication Console as a non-privileged user, log files are stored in the user's AppData folder: `\<username>\AppData\Local\Veeam\Backup`.

Veeam Backup & Replication keeps a separate log file for each of its components: Veeam Shell, Veeam Backup Service, Veeam Guest Catalog Service, Veeam vPower NFS Service, Veeam Installer, Veeam Data Mover and performed jobs.

In addition to logs stored on the backup server, log files are also stored on all servers added to the backup infrastructure:

- On Linux servers and ESXi hosts, logs are stored in the following directory: `/var/log/VeeamBackup/` or `/tmp/VeeamBackup`
- On Microsoft Windows servers, logs are stored in the following directory: `%ProgramData%\Veeam\Backup`

You can collect log files from the backup server and servers managed by Veeam Backup & Replication using the **Export Logs** wizard.

> **TIP**
>
> You can change default log files settings. For more information, see this Veeam KB article.

Exporting Logs

You can use log files to submit a support ticket. It is recommended that you send all log files when submitting a support ticket to ensure that overall and comprehensive information is provided to Veeam Support Team.

To aggregate all log files in the same location, use the **Export Logs** wizard. To launch the wizard, from the main menu select **Help** > **Support Information**.

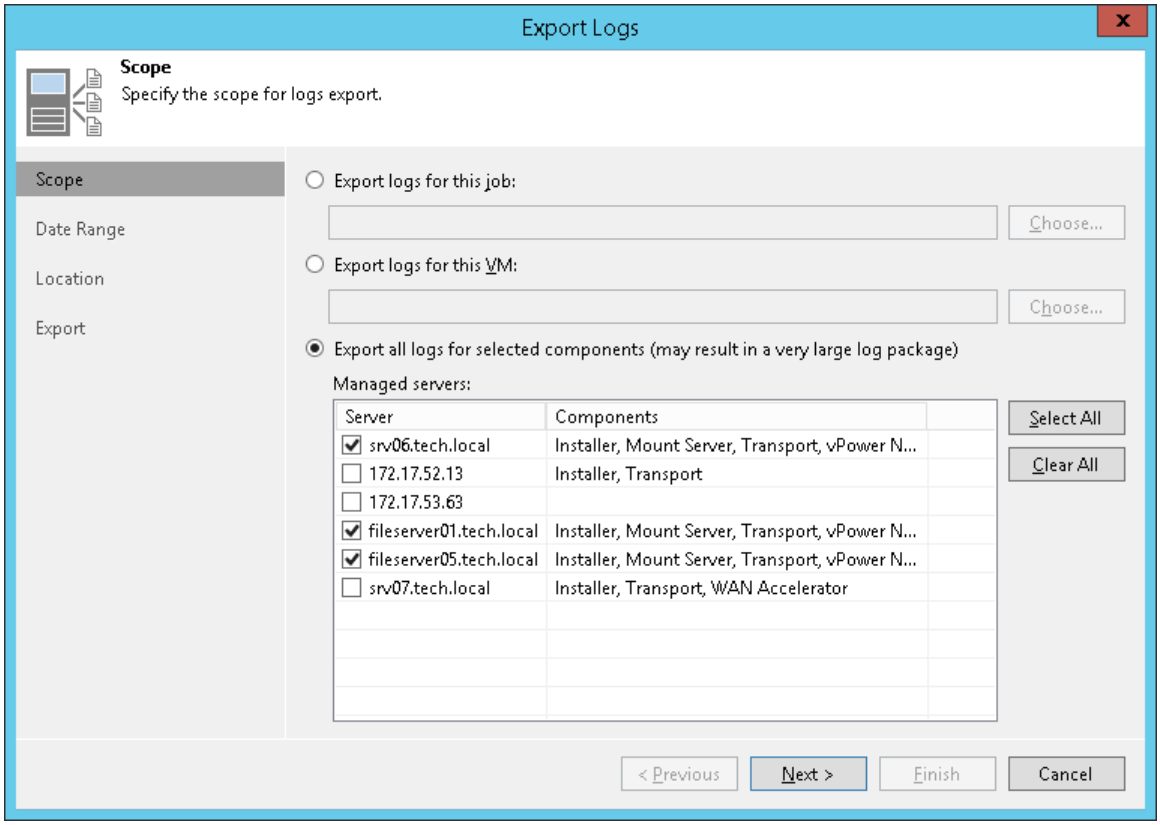
Step 1. Select Virtual Infrastructure Scope

At the **Scope** step of the wizard, define the scope for logs export. You can export logs for the following objects:

- Specific jobs on the backup server
- Specific VMs in the virtual environment
- Specific components in the backup infrastructure

NOTE

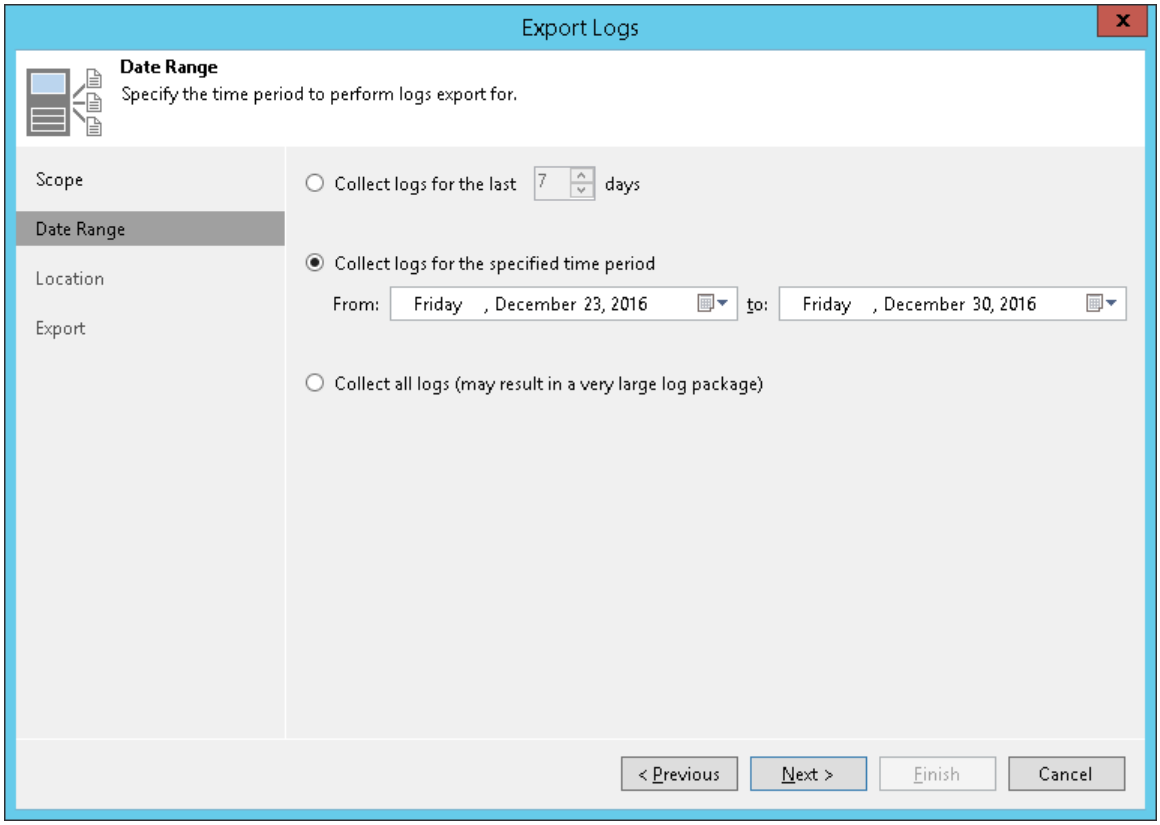
If you export logs from the Veeam Backup & Replication console, the exported logs will be copied to the machine where the console is installed. The log archive will also contain logs from the console machine.



Step 2. Specify Time Interval

At the **Date Range** step of the wizard, define the time interval for which logs must be collected. You can select one of the following options:

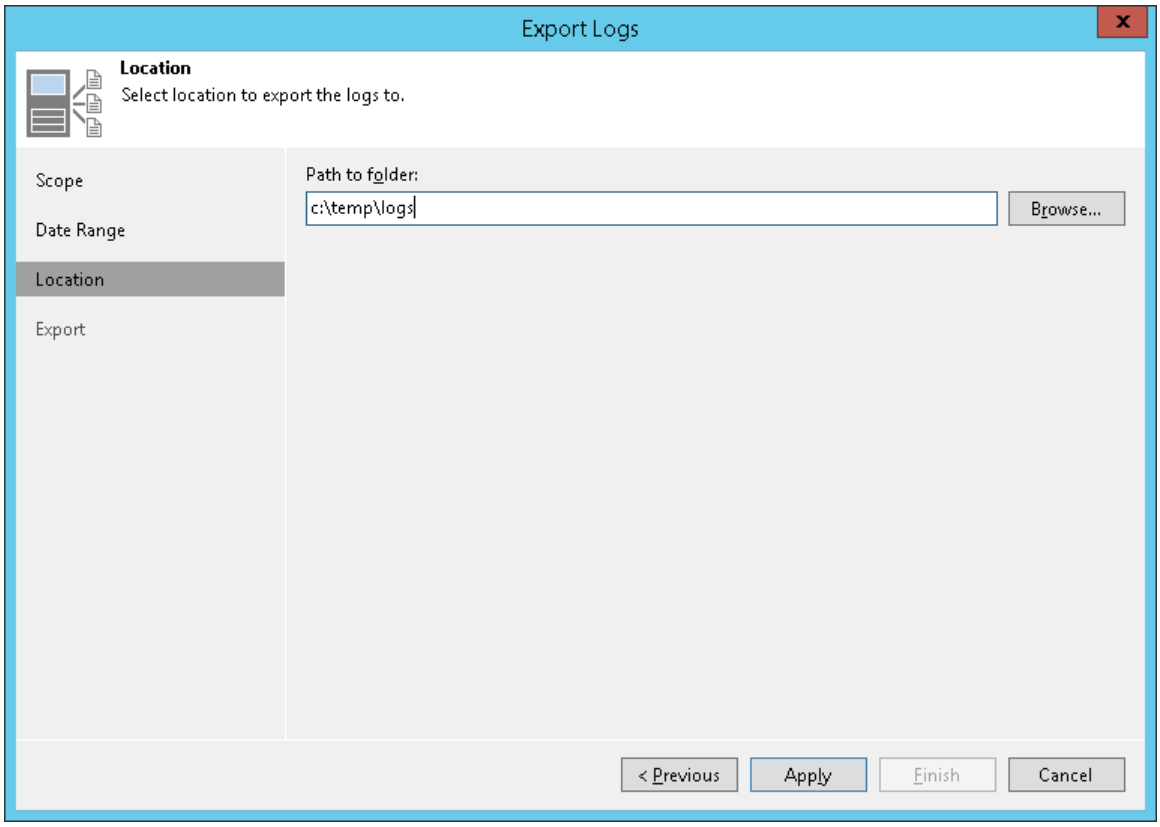
- Collect logs for the last N days
- Collect logs for a specific period of time
- Collect all available logs



Step 3. Specify Destination Folder

At the **Location** step of the wizard, specify the destination folder to which the logs will be exported.

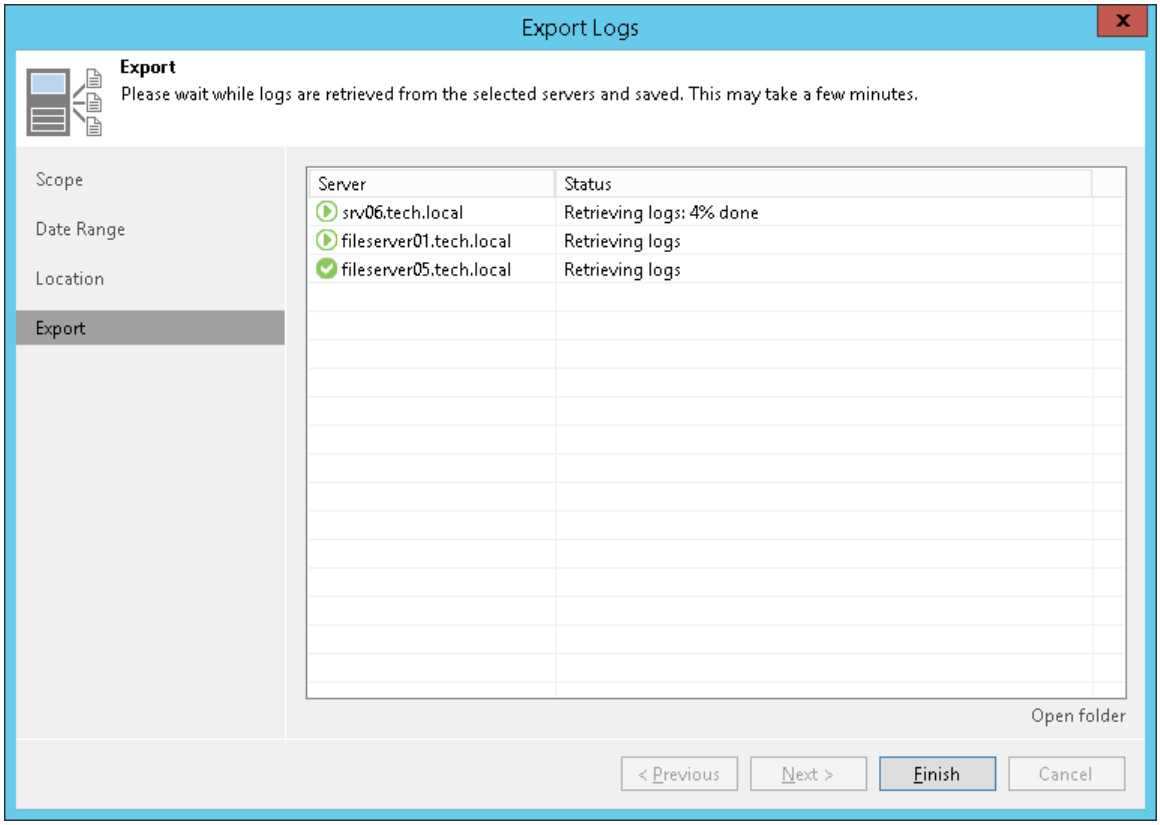
In the **Path to folder** field, specify a path to an archive with log files that will be created. By default, the archive is placed to the `C:\temp\logs` folder on the backup server.



Step 4. Review Results

At the **Export** step of the wizard, Veeam Backup & Replication will collect specified logs and create a log archive. Wait for the export process to complete, review the results and click the **Open folder** link to browse to exported log files and log package.

> **TIP**
>
> If you do not want to wait until the log export process completes, click **Finish** and Veeam Backup & Replication will continue exporting logs in the background. To check the status of the log export session, open the **History** view. In the inventory pane, click **System** and in the working area select the necessary log export session.



Configuration Backup and Restore

You can back up and restore the configuration database that Veeam Backup & Replication uses. If the backup server fails for some reason, you can re-install the backup server and quickly restore its configuration from the configuration backup. You can also use configuration backups to apply the configuration of one backup server to another backup server in the backup infrastructure. During configuration backup, Veeam Backup & Replication exports data from the configuration database and saves it to a backup file in the backup repository.

It is recommended that you regularly perform configuration backup for every backup server in the backup infrastructure. Periodic configuration backups reduce the risk of data loss and minimize the administrative overhead if any problem with backup servers occurs.

It is not recommended to back up the backup server configuration using backup jobs in Veeam Backup & Replication. For backup, Veeam Backup & Replication uses VM snapshots. During snapshot creation and commit, the VM freezes for some time, which can potentially lead to the following consequences:

- Disconnection from the configuration database. For more information, see the Veeam KB1681 article.
- Disconnection from remote Veeam Backup & Replication agents.
- Disconnection from network storage (for example, storage presented through iSCSI) and so on.

For this reason, you must always use the configuration backup functionality to back up and restore configuration of the backup server.

Configuration Backup

By default, Veeam Backup & Replication is configured to create a configuration backup daily. You can change the schedule or run the backup manually.

See the following topics:

- Scheduling Configuration Backups
- Configuring Notification Settings for Configuration Backups
- Running Configuration Backups Manually
- Creating Encrypted Configuration Backups

Configuration Backup Files

When you perform configuration backup, Veeam Backup & Replication retrieves data for the backup server from the configuration database, writes this data into a set of XML files and archives these XML files to a backup file of the BCO format.

Veeam Backup & Replication exports information about the following objects:

- **Backup infrastructure components and objects**: hosts, servers, backup proxies, repositories, WAN accelerators and jobs, global settings configured on the backup server and so on.
- **Backups**: backups, replicas and backup copies created on the backup server.
- **Sessions**: job sessions performed on the backup server.
- **Tapes**: tape libraries connected to the backup server.

> **NOTE**
>
> If you use custom configuration registry keys, keep in mind that configuration backup will not apply to them. You might want to back them up manually.

Configuration backup is job-driven. You can schedule it to run regularly or start it manually. You can choose the backup repository in which the configuration backup must be stored and specify the necessary retention settings.

> **NOTE**
>
> The configuration backup job creates a snapshot of the configuration database and retrieves data required for successful restore from it. If the database size is large, the job may produce significant load on the Microsoft SQL Server. Make sure that you schedule the configuration backup job for a period of low operation intensity on the backup server.

Backup Repository Target

The resulting configuration backup file is stored in the `\VeeamConfigBackup\%BackupServer%` folder on the default backup repository. However, for security's sake, it is recommended that you do not store configuration backups on the default backup repository or in any other folder on the backup server. In this case, if the backup server fails, its configuration data will remain, and you will be able to recover the failed backup server.

When you configure a new backup repository, Veeam Backup & Replication offers you to change the configuration backup file location from the default backup repository to the new backup repository. Click **Yes**, and Veeam Backup & Replication will automatically change the backup target in the configuration backup job settings and will use this target in future.

Configuration backups that were created before the target change will remain in the default backup repository. You can manually copy them to the new backup repository to have all restore points of the configuration backup in one place.

> **IMPORTANT**
>
> You cannot store configuration backups in scale-out backup repositories.

Scheduling Configuration Backups

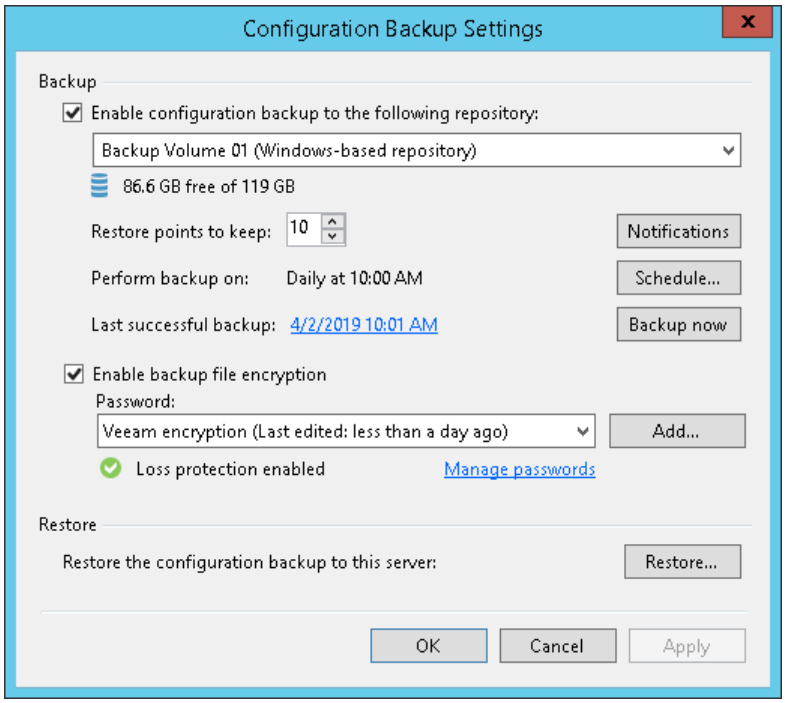
You can instruct Veeam Backup & Replication to perform configuration backup automatically by schedule.

> **IMPORTANT**
>
> If you plan to migrate configuration data to the database used by another backup server, stop all running jobs and disable scheduled jobs before creating the configuration backup. In the opposite case, job sessions may be failing after configuration restore. For more information, see Migrating Configuration Database.

To schedule a configuration backup:

1. From the main menu, select **Configuration Backup**.
2. Make sure that the **Enable configuration backup to the following repository** check box is selected.
3. From the **Backup repository** list, choose a backup repository on which the configuration backup must be stored.
4. In the **Restore points to keep** field, specify the number of restore points that you want to maintain in the backup repository.
5. Click **Schedule** next to the **Perform backup on** field and specify the time schedule according to which the configuration backup must be created.

6. To create an encrypted backup, select the **Enable backup file encryption** check box. From the **Password** drop-down list, select a password you want to use for encryption. If you have not created a password beforehand, click **Add** or use the **Manage passwords** link to specify a new password. For more information, see Creating Encrypted Configuration Backups.



Configuring Notification Settings for Configuration Backups

You can configure notifications for the configuration backup:

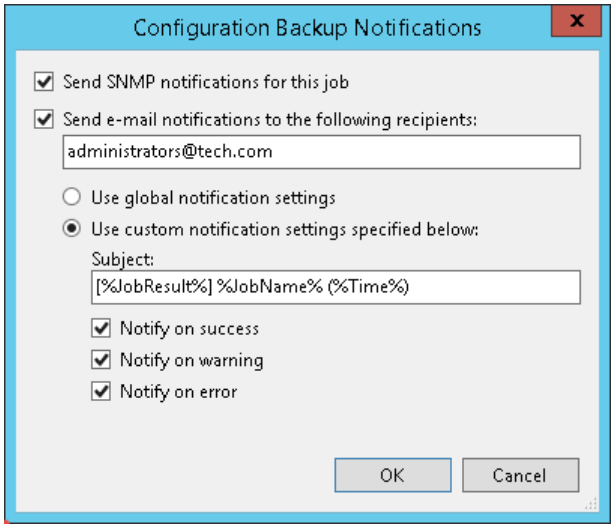
1. From the main menu, select **Configuration Backup**.
2. Click **Notifications**.
3. Select the **Send SNMP notifications for this job** check box if you want to receive SNMP traps when the job completes successfully.

   SNMP traps will be sent if you specify global SNMP settings in Veeam Backup & Replication and configure software on recipient's machine to receive SNMP traps. For more information, see Specifying SNMP Settings.
4. Select the **Send email notifications to the following recipients** check box if you want to receive notifications about the job completion status by email. In the field below, specify recipient's email address. You can enter several addresses separated by a semicolon.

   Email notifications will be sent if you configure global email notification settings in Veeam Backup & Replication. For more information, see Configuring Global Email Notification Settings.
5. You can choose to use global notification settings or specify custom notification settings.
   - To receive a typical notification for the job, select **Use global notification settings**. In this case, Veeam Backup & Replication will apply to the job global email notification settings specified for the backup server. For more information, see Configuring Global Email Notification Settings.
   - To configure a custom notification for the job, select **Use custom notification settings specified below** check box. You can specify the following notification settings:

i. In the **Subject** field, specify a notification subject. You can use the following variables in the subject: *%JobResult%*, *%JobName%*, *%Time%* (completion time).

ii. Select the **Notify on success**, **Notify on warning** and/or **Notify on error** check boxes to receive email notification if the job completes successfully, fails or completes with a warning.



Running Configuration Backups Manually

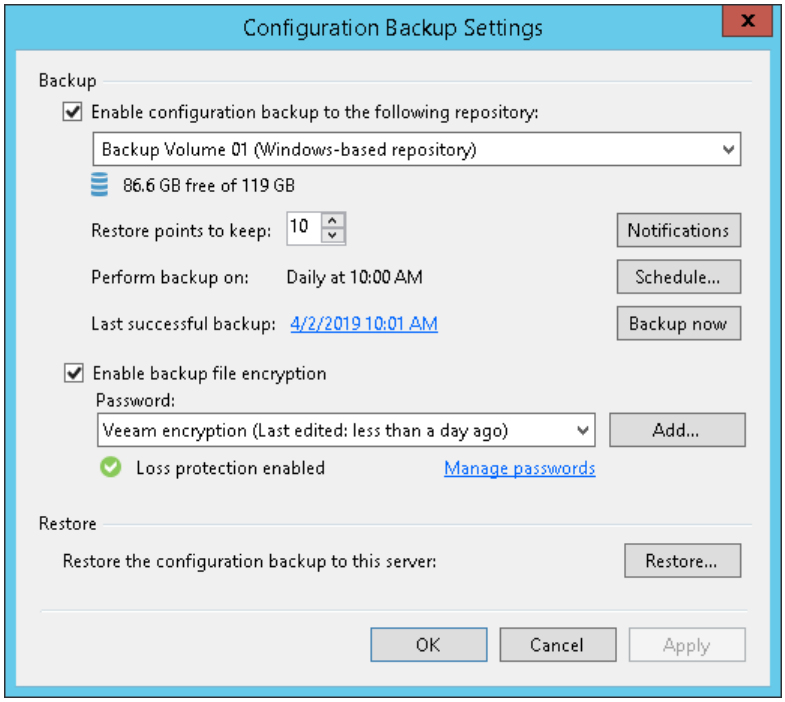
You can create a configuration backup manually when you need it, for example, if you want to capture a state of the configuration database at a specific point in time.

> **IMPORTANT**
>
> If you plan to migrate configuration data to the database used by another backup server, stop all running jobs and disable scheduled jobs before creating the configuration backup. In the opposite case, job sessions may be failing after configuration restore. For more information, see Migrating Configuration Database.

To create a configuration backup manually:

1. From the main menu, select **Configuration Backup**.
2. Make sure that the **Enable configuration backup to the following repository** check box is selected.
3. From the **Backup repository** list, choose a backup repository on which the configuration backup must be stored.
4. In the **Restore points to keep** field, specify the number of restore points that you want to maintain in the backup repository.
5. To create an encrypted backup, select the **Enable backup file encryption** check box. From the **Password** field, select a password you want to use for encryption. If you have not created a password beforehand, click **Add** or use the **Manage passwords** link to specify a new password. For more information, see Creating Encrypted Configuration Backups.
6. Click **Backup now**.

Veeam Backup & Replication will back up the configuration database and store a new restore point to the selected backup repository.

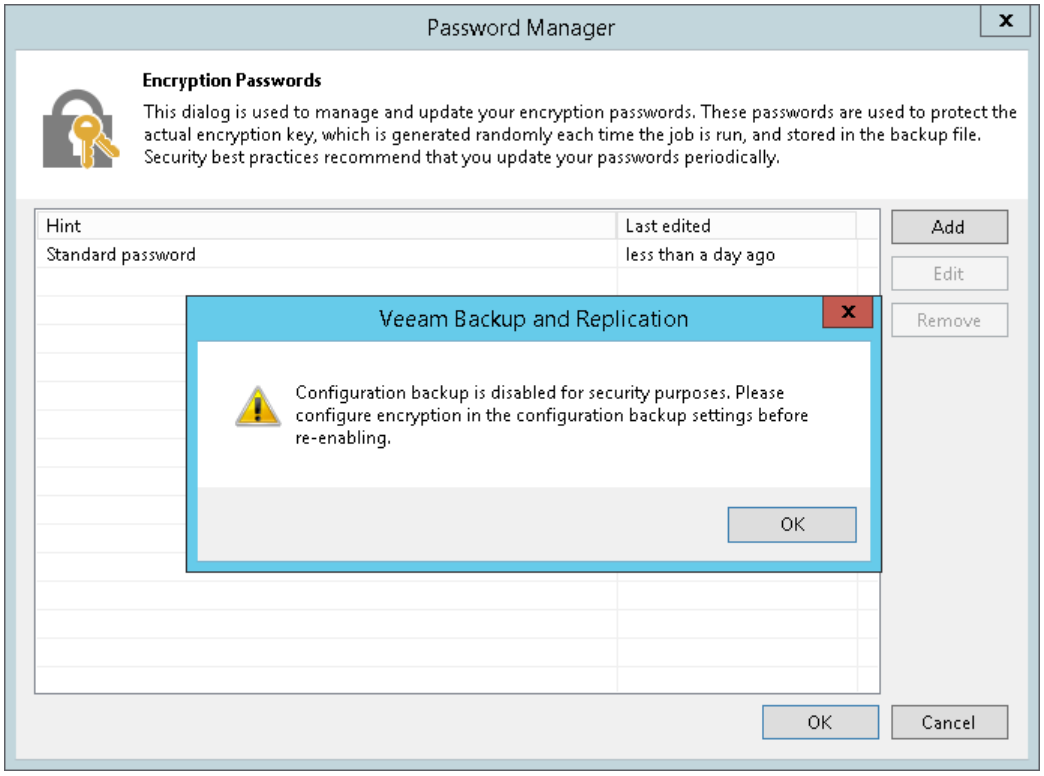


Creating Encrypted Configuration Backups

Veeam Backup & Replication requires that you encrypt the configuration backup if you have created at least one password in the Password Manager on the backup server.

When you encrypt jobs or tapes with passwords, Veeam Backup & Replication creates a set of keys that are employed in the encryption process. Some encryption keys, for example, storage keys and metakeys, are stored in the configuration database. If a configuration backup was non-encrypted, data from it could be freely restored on any backup server. Encryption keys saved to the configuration database and the content of encrypted files might become accessible for unintended audience.

If the Password Manager contains at least one password, and you do not enable encryption for the configuration backup, Veeam Backup & Replication disables configuration backup. To enable the configuration backup, you must enable encryption in the configuration backup job settings.

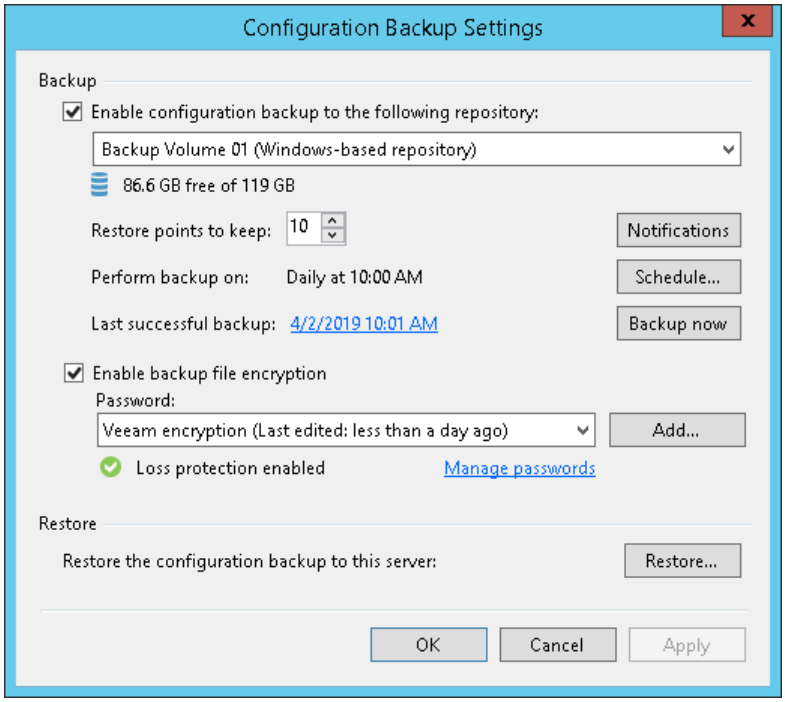


After you enable the encryption option, Veeam Backup & Replication will create encrypted configuration backups. Beside encryption keys, the created backups capture credential records specified in the Credentials Manager. When you restore data from such backup, you will not have to enter passwords for credentials records again (unless the passwords for credentials records have changed by the time of restore).

To encrypt the configuration backup:

1. From the main menu, select **Configuration Backup**.
2. Select the **Enable backup file encryption** check box.

3. From the **Password** field, select a password you want to use for encryption. If you have not created a password beforehand, click **Add** or use the **Manage passwords** link to specify a new password. For more information, see Managing Passwords for Data Encryption.



Restoring Configuration Database

Restore of the configuration database is helpful in the following situations:

- The configuration database got corrupted and you want to recover data from the configuration backup.
- The Microsoft SQL Server in which the configuration database resides got corrupted, and you want to deploy the configuration database on a new Microsoft SQL Server, and restore data from the configuration backup to it.
- You want to roll back the configuration database to a specific point in time.
- You want to restore data to a new configuration database on the same Microsoft SQL Server instance, for example, for testing purposes.

You can restore a configuration backup on the same backup server where the backup was created or on another backup server.

Before you start the restore process, check prerequisites. Then use the **Veeam Backup & Replication Configuration Restore** wizard to restore the configuration database.

Before You Begin

Before you start the restore process, check the following prerequisites:

- Stop all jobs that are currently running. During restore of configuration, Veeam Backup & Replication temporary stops the Veeam Backup Service and jobs.
- Save registry keys that you changed or created on the backup server. After restore, you will need to recreate or change the keys manually because the configuration database does not store them.
- Check the version of the backup server. On the backup server running Veeam Backup & Replication 11, you can restore configuration backups created with the following product versions: 10 and 9.5 Update 4b.
- Make sure that the certificate chain restored from a configuration backup will successfully pass validation on the target backup server. This precaution is required if the following conditions are met:
  a. You want to restore configuration database of a backup server used in the Veeam Agent management scenario.
  b. The backup server whose configuration database you want to restore uses a custom certificate issued by a Certificate Authority instead of the default self-signed certificate to ensure a secure connection in the Veeam Agent management infrastructure.
- If you plan to restore configuration data to the database on another Microsoft SQL Server, make sure the account for using Veeam Backup & Replication has sufficient permissions. For more information, see Required Permissions.
- [for Veeam Backup for Nutanix AHV users] After you restore the configuration database of Veeam Backup & Replication, you must remove existing AHV backup proxies from Veeam Backup & Replication infrastructure and connect to them again. Otherwise, these AHV backup proxies will not be able to perform backup and restore due to missing authentication certificates. For instructions on how to connect to an existing AHV backup proxy, see the Adding Nutanix AHV Backup Proxy section in the Veeam Backup for Nutanix AHV User Guide.
- After you run configuration restore for capacity tier, Veeam Backup & Replication will put the capacity tier extents into the seal mode. Rescan the backup infrastructure and then remove the extents from the seal mode.

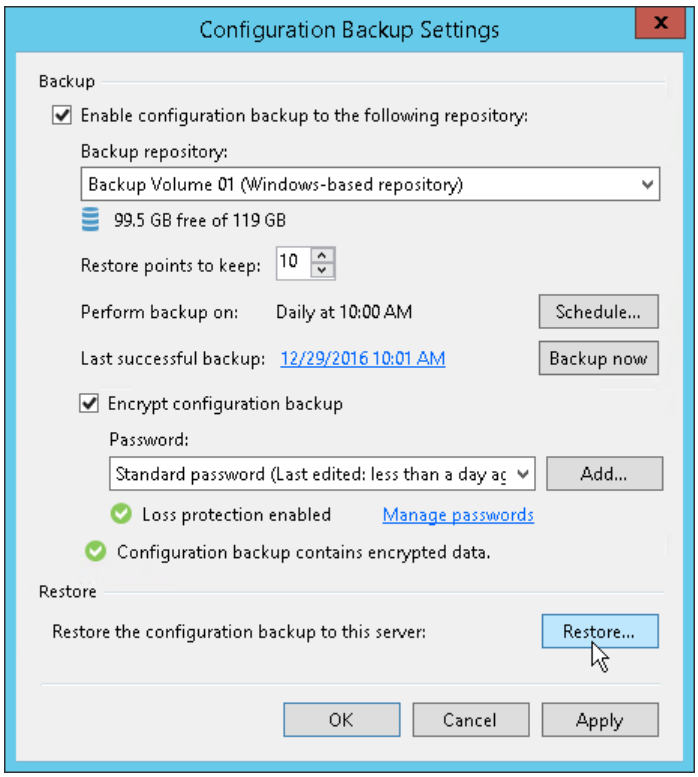
- If you use a Veeam Backup & Replication server as a Veeam repository, after restore to a new host this repository will be located in the same file path as it was before the migration

> **IMPORTANT**
>
> You can start configuration restore only from the Veeam Backup & Replication console installed locally on the backup server. You cannot start configuration restore from the console installed on a remote machine.

Step 1. Launch Configuration Database Restore Wizard

To launch the **Veeam Backup and Replication Configuration Restore** wizard, do either one of the following:

- From the main menu of the Veeam Backup & Replication console, select **Configuration Backup**. In the **Restore** section, click **Restore**.
- In the **Start** menu of the backup server, click **Configuration Restore**.
- Use the `Veeam.Backup.Configuration.Restore.exe` file located in the installation folder on the backup server. By default, the path to the folder is the following: `%PROGRAMFILES%\Veeam\Backup and Replication\Backup`.
- [If the configuration backup is stored on the backup server] In Microsoft Windows Explorer, open the folder where configuration backups are stored (by default, `Backup\VeeamConfigBackup\<BackupServerName>` on the volume with most disk space on the backup server) and double-click the necessary configuration backup file.



Step 2. Select Restore Mode

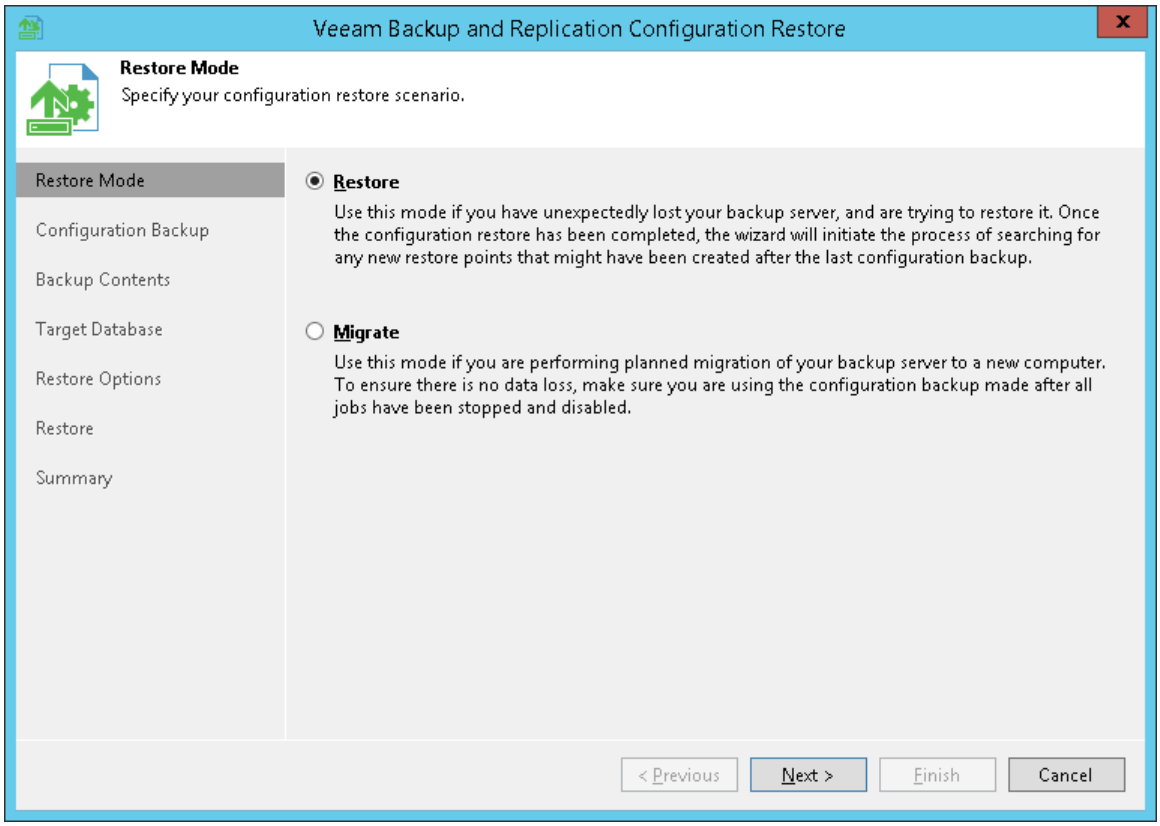
At the **Restore Mode** step of the wizard, choose a restore mode that you want to use.

- Select **Restore** if you want to restore data from the configuration backup to the database used by the initial backup server.

  In the Restore mode, Veeam Backup & Replication retrieves configuration data from the backup and stores it to the target database. After that, Veeam Backup & Replication performs additional rescan of VM replicas, backup repositories and tape libraries connected to the backup server. Rescan helps synchronize potential changes between the backup infrastructure and restored database that took place from the moment when the configuration backup was created till the present time. As a result, the target configuration database will contain information about restore points that were created after the configuration backup was taken, and this information is displayed in the Veeam Backup & Replication console.

- Select **Migrate** if you want to restore data from the configuration backup to the database used by another backup server.

  In the Migrate mode, Veeam Backup & Replication retrieves configuration data from the backup and stores it to the target database. No rescan operation is performed.

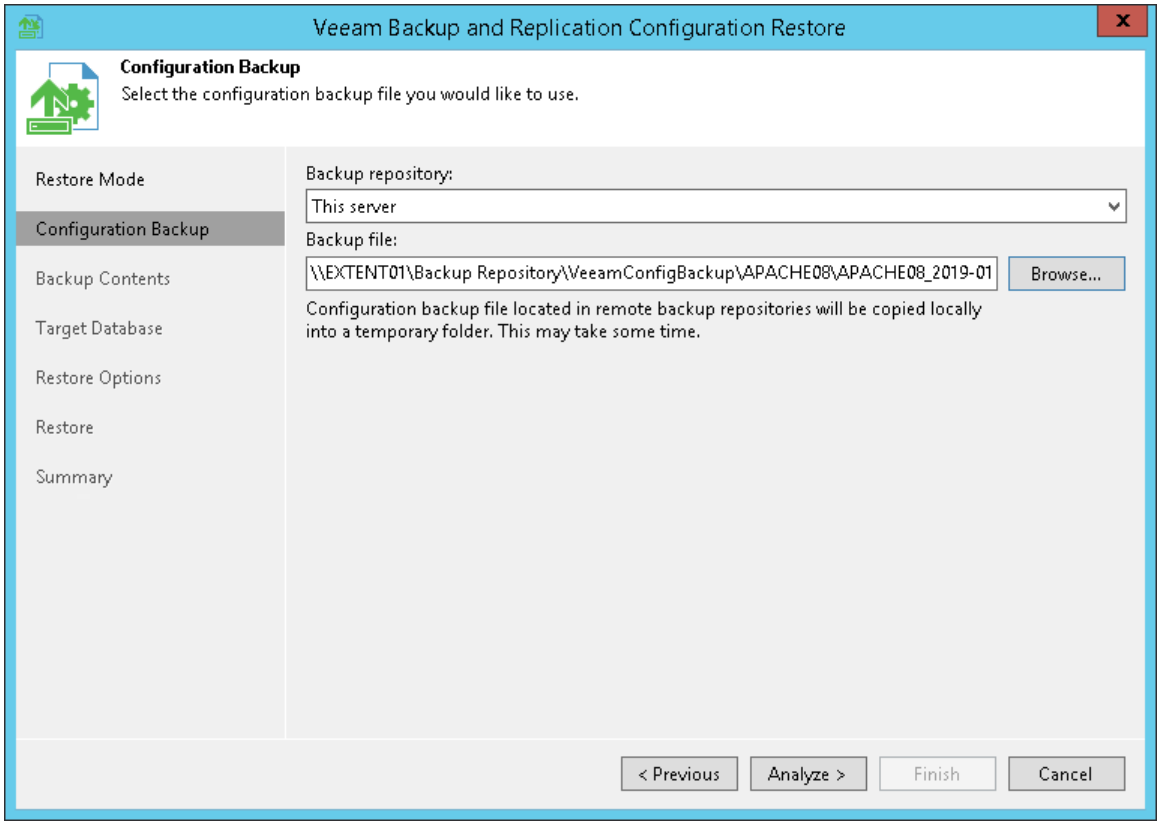


Step 3. Select Configuration Backup

At the **Configuration Backup** step of the wizard, select a configuration backup from which you want to restore data.

1. From the **Backup repository** list, select a server or backup repository on which the configuration backup file is located.
2. Click **Browse** next to the **Backup file** field and select the backup file.

If you select to restore configuration data from a backup in a remote backup repository, during restore Veeam Backup & Replication will first copy the backup file to a temporary folder on the backup server. After you finish the restore process and close the wizard, Veeam Backup & Replication will automatically delete the configuration file from the temporary folder.

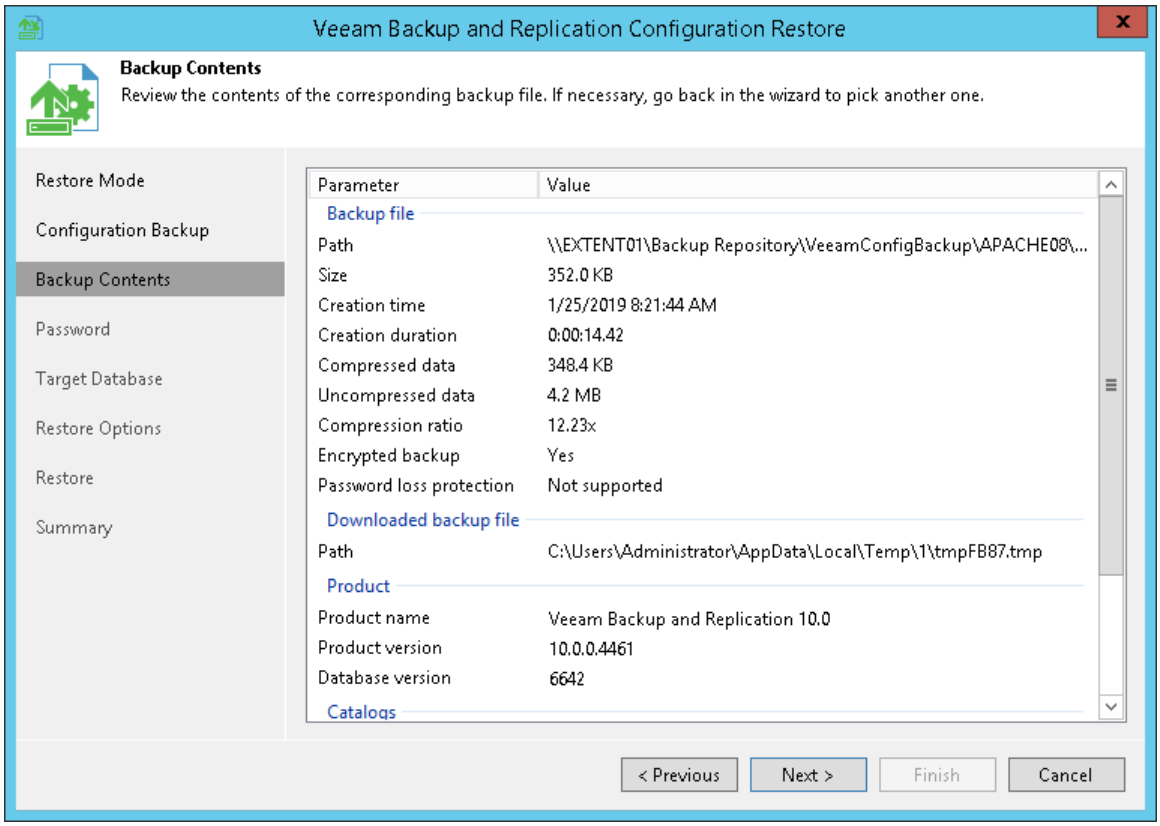


Step 4. Review Configuration Backup Parameters

At the **Backup Contents** step of the wizard, Veeam Backup & Replication will analyze the content of the selected backup file and display the following settings:

- **Backup file settings**: settings of the configuration backup file itself.
- **Product settings**: version of Veeam Backup & Replication installed on the initial backup server and configuration database version.
- **List of catalogs**: catalogs storing backup configuration data.

Review the displayed settings and click **Next**.



Step 5. Specify Password

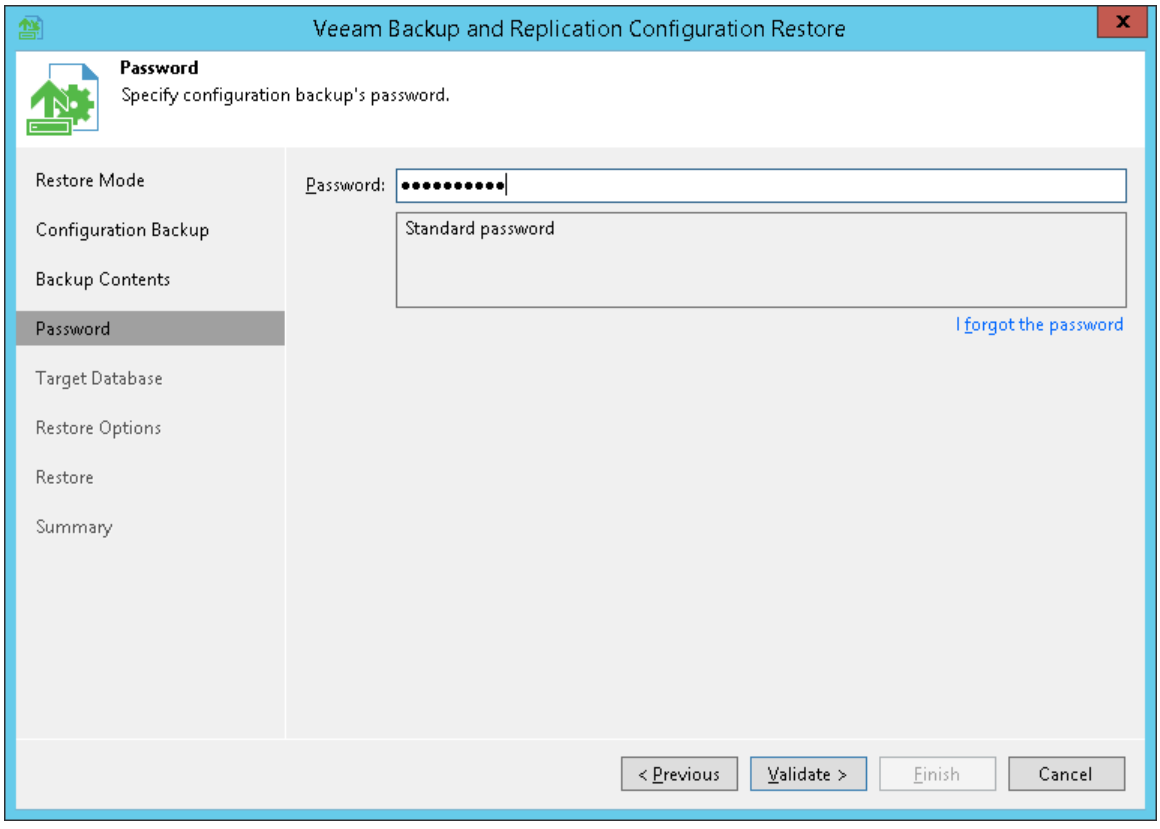
The **Password** step of the wizard is available if you have enabled the encryption option in the configuration backup properties.

Enter the password to decrypt configuration backup data:

1. Check the password hint to recall the password.
2. In the **Password** field, enter the password to decrypt the configuration backup file.

If you forgot or lost the password, click the **I forgot the password** link. For more information, see Decrypting Data Without Password.

> **NOTE**
>
> Restoring configuration data without a password is included in the Veeam Universal License. When using a legacy socket-based license, Enterprise or higher edition is required.
>
> Also, your backup server must be connected to Veeam Backup Enterprise Manager. Otherwise, you will not see the **I forgot the password** link.



Step 6. Specify Target Database

At the **Target Database** step of the wizard, specify the target Microsoft SQL Server instance to which configuration data must be restored.

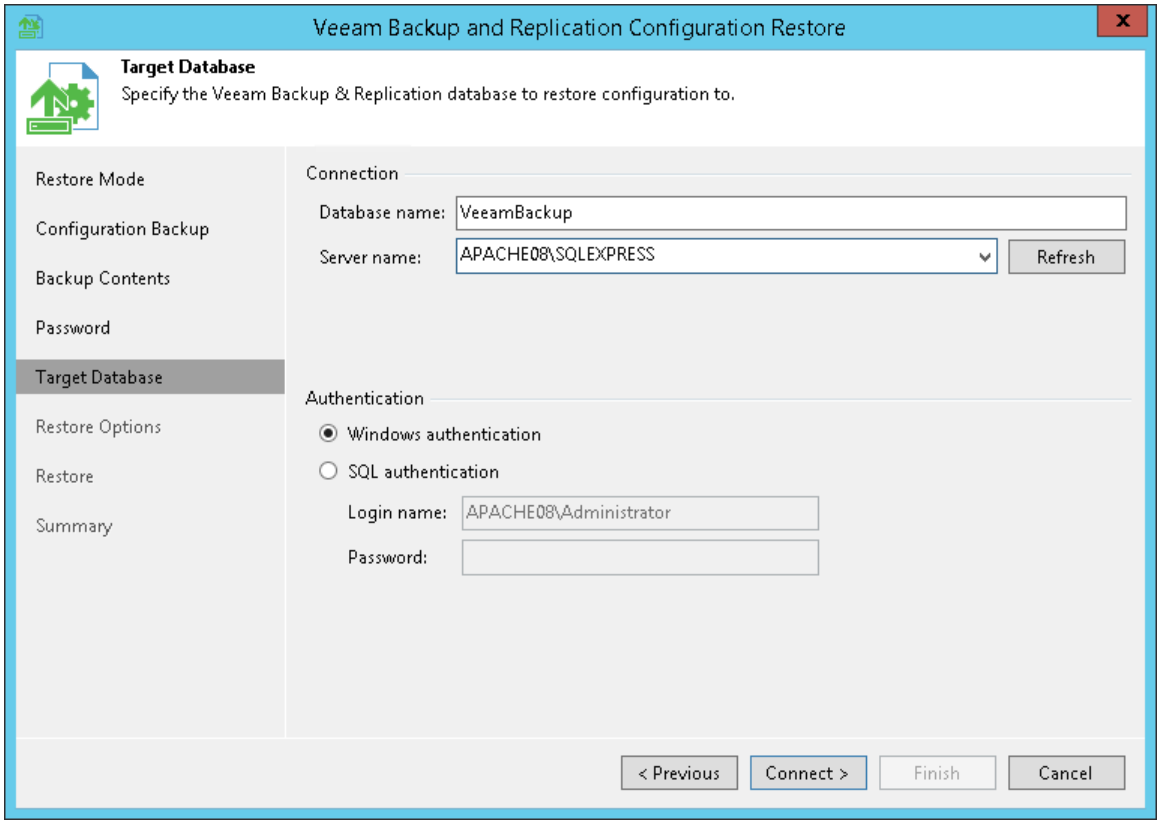
1. In the **Database name** field, specify a name of the database to which configuration data must be restored. By default, Veeam Backup & Replication uses the default name for the target database.

   If you specify a name of an existing target database, Veeam Backup & Replication will overwrite this database. If you specify a name of the database that does not exist, Veeam Backup & Replication will create it on the specified Microsoft SQL Server instance.

   > **NOTE**
   >
   > If a backup repository is located in the backup server, after configuration restore, this repository will point to the same path as it was before the migration but in a new host. For example, if you keep backed-up data on the D disk, after migration Veeam Backup & Replication will keep new backups on the D disk of the new host.

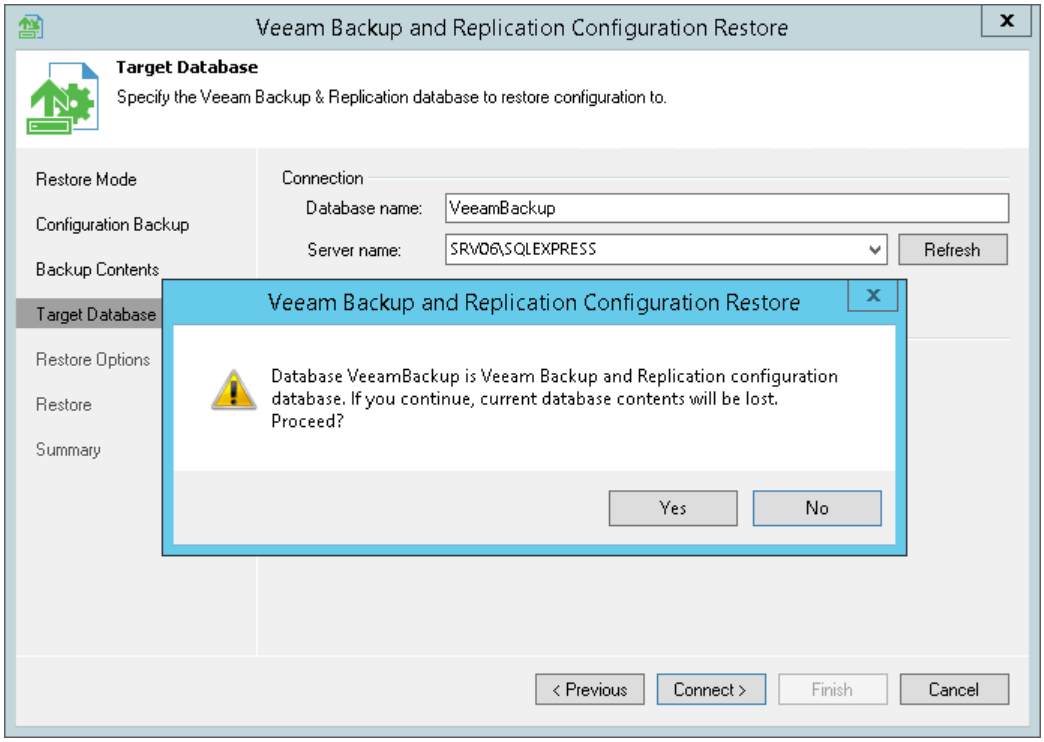
2. From the **Server name** list, select a Microsoft SQL Server instance on which the database is deployed or must be deployed. In the list of Microsoft SQL Server instances, Veeam Backup & Replication displays all servers from the network where the backup server resides. To update the list of servers, click **Refresh** on the right.

3. In the **Authentication** section, select the authentication mode to connect to the target Microsoft SQL Server instance: **Windows Authentication** or **SQL authentication**. If you select the SQL authentication mode, specify credentials of the Microsoft SQL Server account that you want to use.



When you restore configuration to an existing database, the configuration restore process will delete the current state of the database contents and replace it with the restored data. Veeam Backup & Replication will display a warning. If you want to replace the contents, click **Yes** to confirm.

If you do not want to lose the current data, restore the configuration to a new database. To do this, click **No** to the warning and specify a non-existing database name in the **Database name** field.

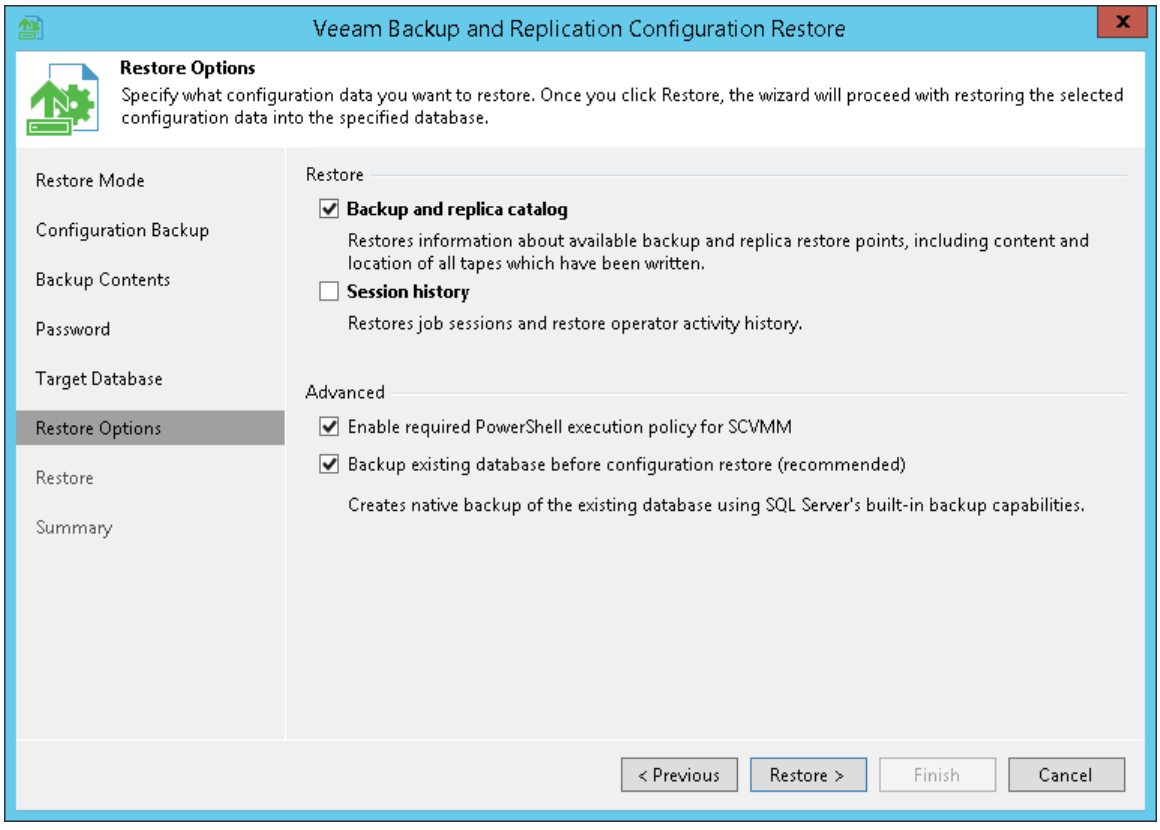


Step 7. Specify Restore Options

At the **Restore Options** step of the wizard, specify additional restore options.

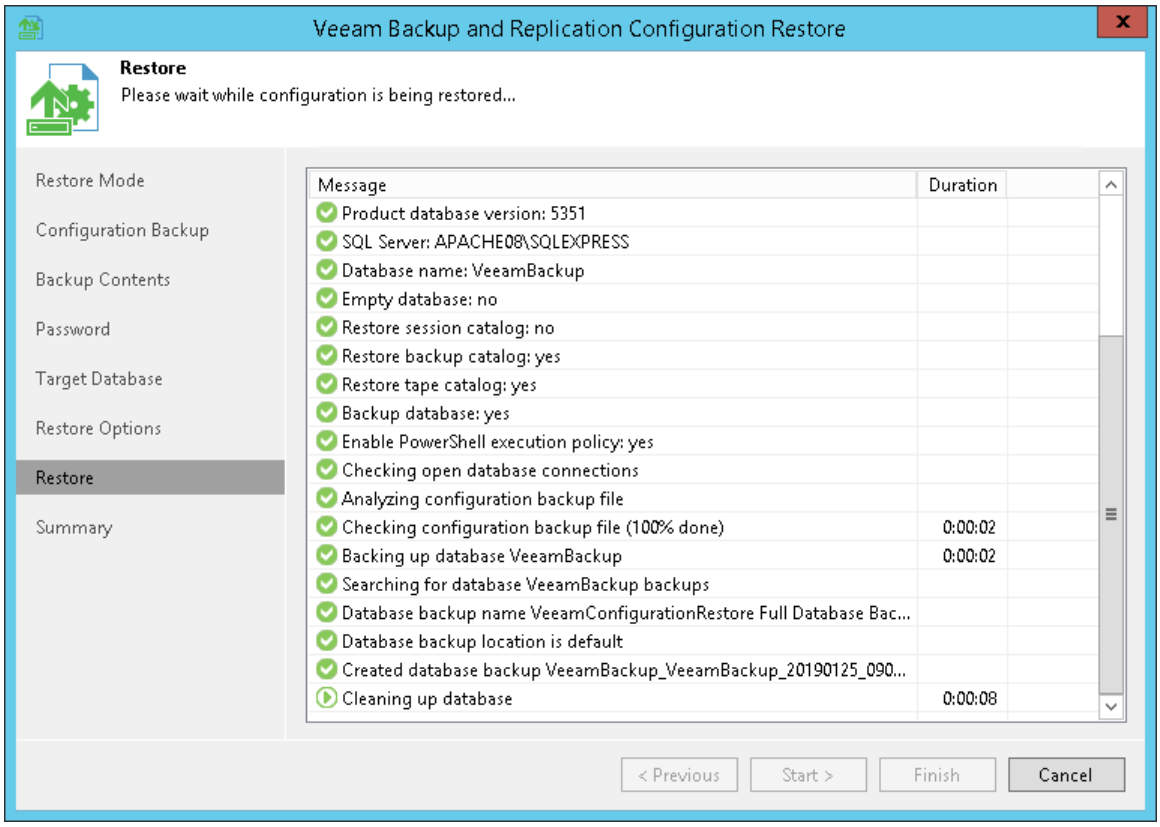
1. In the **Restore** section, select what data you want to restore from the configuration backup. Veeam Backup & Replication always restores configuration data for backup infrastructure components, jobs and global settings specified at the level of the backup server. You can additionally restore the following data:
   - **Backup and replica catalog**: data about all backups and replicas registered on the backup server and information about tapes to which backups were written and location of these tapes.
   - **Session history**: data about all sessions performed on the backup server.
2. If you plan to use PowerShell on the restored backup server, select the **Enable required PowerShell policy for SCVMM** check box. During restore, Veeam Backup & Replication will enable the PowerShell execution policy and you will not have to enable it manually afterwards. Enabling this option is identical to running the '*Set-ExecutionPolicy RemoteSigned*' command on the backup server.
3. If you are restoring configuration data to the same database, select the **Backup existing database before configuration restore** check box. This option will help you protect the current database from accidental errors during the restore process. During restore, Veeam Backup & Replication will first back up the current database using the native tools of Microsoft SQL Server. After that, Veeam Backup & Replication will purge the current database and import data from the configuration backup to it. In such scenario, if an error occurs during the restore process, you will be able to restore the current database from the Microsoft SQL backup using Microsoft SQL Management Studio or SQL scripts.

   The created Microsoft SQL database backup is named by the following pattern: *VeeamBackup<DatabaseName><date>.bak* and stored to the default Microsoft SQL backups location, for example: `%ProgramFiles%\Microsoft SQL Server\MSSQL11.MSSQLSERVER\MSSQL\Backup\`.
4. Click **Restore**. Veeam Backup & Replication will stop currently running jobs and Veeam Backup & Replication services and will restore the database to the specified location.



Step 8. Review Restore Settings

At the **Restore** step of the wizard, Veeam Backup & Replication will display the progress on the restore process. Wait for the restore process to complete and click **Next**.



If you have chosen to restore data in the Migrate mode and the configuration backup file does not meet the Migrate mode requirements, Veeam Backup & Replication will display a warning and offer you to switch to the Restore mode. The Restore mode requires more time but guarantees that information about all new restore points will be available in the restored database.

- To switch to the **Restore** mode, in the warning window click **Yes**.
- To carry on data restore in the **Migrate** mode, in the warning window click **No**.
- To stop the restore process, in the warning window click **Cancel**.

For more information, see Migrating Configuration Database.

Step 9. Finalize Restore Process

After the restore process has finished, you may need to perform the following actions to finalize the configuration database restore:

1. Specify credentials for backup infrastructure objects.
2. Specifying credentials for cloud services.
3. Perform components upgrade.

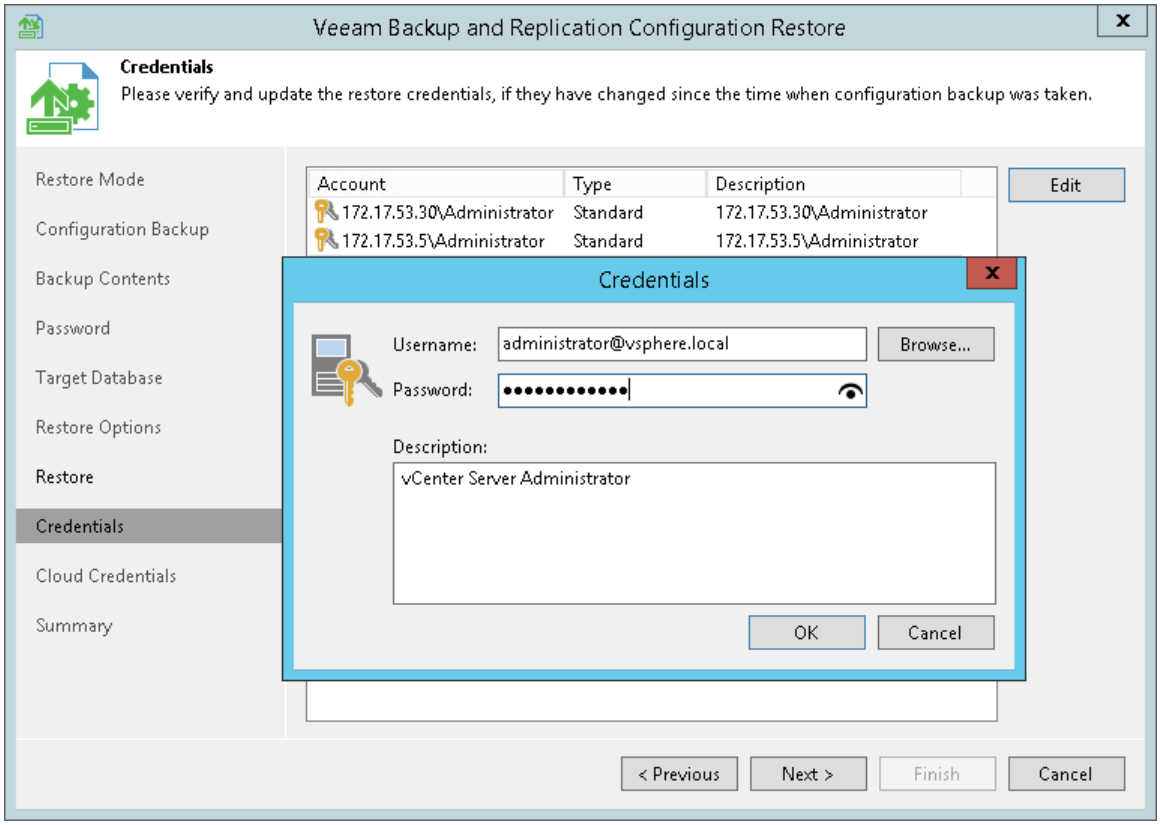
Specifying Credentials

At the **Credentials** step of the wizard, Veeam Backup & Replication displays a list of credentials records that existed on the backup server at the time when the configuration backup was created. If by the time of restore passwords for credentials records have changed, you can specify new values for these records.

> **IMPORTANT**
>
> If you have not enabled encryption for configuration backups, Veeam Backup & Replication will not restore passwords for credentials records. You need to re-enter passwords for all credentials records to make sure that backup infrastructure components and jobs work in a proper way after you complete configuration restore.

To edit credentials records:

1. Select a record in the list and click **Edit**.
2. Edit settings of the record as required.
3. Repeat the procedure for all records in the list.



Specifying Cloud Credentials

At the **Cloud Credentials** step of the wizard, Veeam Backup & Replication displays a list of cloud credentials records that existed on the backup server at the time when the configuration backup was created. If by the time of restore passwords for cloud credentials records have changed, you can specify new values for these records.

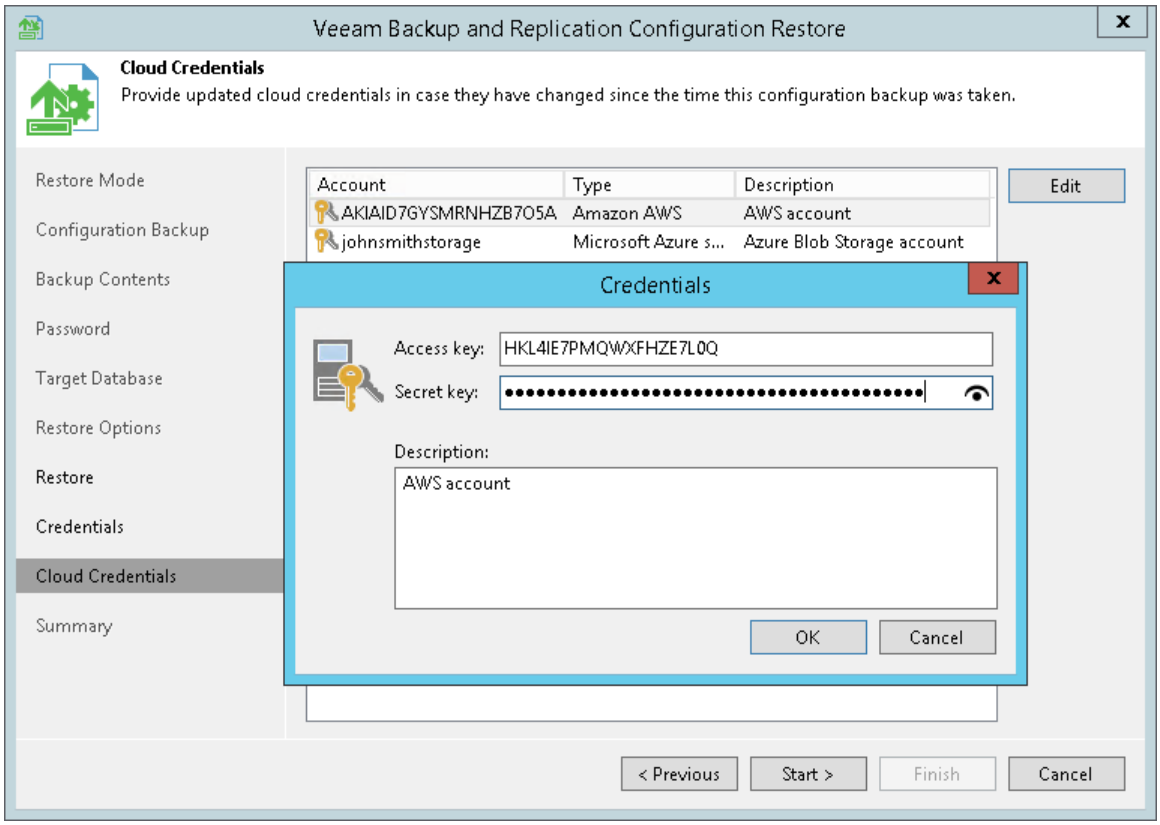
IMPORTANT

Consider the following:

- If you have not enabled encryption for configuration backups, Veeam Backup & Replication will not restore passwords for cloud credentials records. You need to re-enter passwords for all cloud credentials records to make sure that cloud services and jobs work in a proper way after you complete configuration restore.
- You cannot edit credentials of Microsoft Azure compute accounts and Google Cloud service account in the configuration restore wizard. You can edit these credentials only after configuration restore in Cloud Credentials Manager. For details, see Editing and Deleting Credentials Records.
- Veeam Backup & Replication does not restore credentials records of Google Cloud service accounts that were created automatically using the Cloud Credentials Manager. For details, see Google Cloud Service Account.

To edit cloud credentials records:

1. Select a record in the list and click **Edit**.
2. Edit settings of the record as required.
3. Repeat the procedure for all records in the list.



Performing Components Upgrade

After the restore process is complete, Veeam Backup & Replication will check if services on backup infrastructure components must be upgraded and display a list of outdated components.

To upgrade backup infrastructure components, select check boxes next to the necessary components and click **Next**. If some component fails to upgrade, you can get back to a previous step of the wizard and repeat the procedure or close the wizard and upgrade the components manually. For more information, see Server Components Upgrade.

Step 10. Synchronize Backups and Tape Libraries

After the configuration database is restored, Veeam Backup & Replication can perform a synchronization operation for backups and replicas created on the backup server and tape libraries connected to the backup server.

- The synchronization operation for backups and replicas is performed if you are restoring a database from a backup created with Veeam Backup & Replication in the Restore mode and you have selected to restore data from the backup and replica catalog.
- The synchronization operation for tape libraries is performed if you are restoring a database from a backup created with Veeam Backup & Replication in the Restore mode and you have selected to restore data from the backup and replica catalog.

Wait for the synchronization operation to complete.

Step 11. Finish Working with Wizard

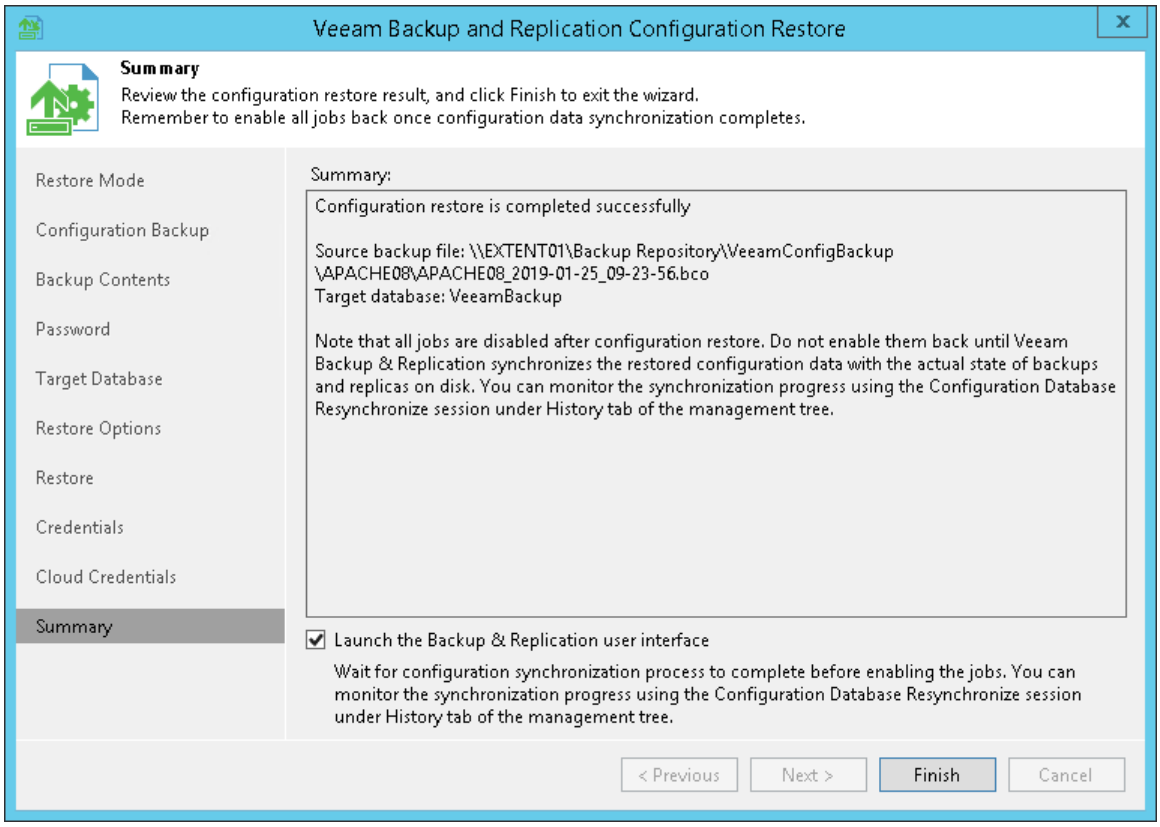
At the **Summary** step of the wizard, finalize the process of configuration data restore.

1. Review the restore process results.
2. If you want to start Veeam Backup & Replication after you finish working with the wizard, select the **Launch the Backup & Replication user interface** check box.
3. Click **Finish** to exit the wizard.

NOTE

Mind the following:

- If you created custom registry keys or changed the existing ones on the backup server, you must recreate or change the registry keys again.
- If you restore data from the configuration backup in the Restore mode, all jobs on the backup server will be disabled after the restore process is complete. You need to enable them manually.



Migrating Veeam Backup & Replication to Another Server

If you need to migrate Veeam Backup & Replication to another server, you can back up its configuration database, install Veeam Backup & Replication on the target server and restore the configuration data from the backup. As a result, you will have a new Veeam Backup & Replication server with all settings, jobs and backup infrastructure elements from the old server.

> **NOTE**
>
> Mind the following:
>
> - It is recommended that you use Veeam Backup & Replication tools to create configuration backups and migrate the configuration database. If you use native Microsoft SQL Server tools, after migration, some information, such as secure configuration data, may not be accessible.
> - This section gives instructions on how to migrate Veeam Backup & Replication together with its configuration database to another server. If you need to migrate only the configuration database, see Migrating Configuration Database to Another SQL Server.

If you want to migrate Veeam Backup & Replication to another server, perform the following steps:

1. Stop running jobs and disable scheduled jobs.
2. Save registry keys that you changed or created.
3. Back up the configuration database of Veeam Backup & Replication.
4. Install Veeam Backup & Replication on the target machine.
5. Restore the configuration database from the backup.
6. Finish the configuration.

Step 1. Stop and Disable Jobs

Before you start the migration process, stop all running jobs and disable all scheduled jobs on the source backup server before you create the configuration backup.

> **NOTE**
>
> Do not start or enable any jobs until the migration of Veeam Backup & Replication is finished. If you start a job before migration is completed, Veeam Backup & Replication will produce a new restore point in the chain and update the chain metadata. The created configuration backup will not contain information about this new restore point. When you migrate data from the configuration backup to the database and start the job again, Veeam Backup & Replication will fail to synchronize the chain metadata with data in the database. As a result, the job will fail.

Step 2. [Optional] Save Registry Keys That You Changed or Created

If you have created new registry keys or changed the existing keys on the backup server, you will need to recreate the keys or change the keys manually after the migration. The configuration database does not store registry keys.

To save registry keys, you can use the `reg export` command or you can use Registry Editor to save keys manually. For details, see Windows Documentation.

Step 3. Back Up Configuration Database

Back up the configuration database of Veeam Backup & Replication. For instructions, see Running Configuration Backups Manually.

Step 4. Install Veeam Backup & Replication on Target Machine

Install Veeam Backup & Replication on the machine on which you plan to move your source backup server. The machine must meet system requirements for a backup server.

During installation, you need to specify a new database name to store the configuration data. You can select a Microsoft SQL Server instance on the local or remote server. It does not matter whether you select an existing instance or create a new one. In the next steps of this guide, you will restore the previous configuration data to the database selected at this step.

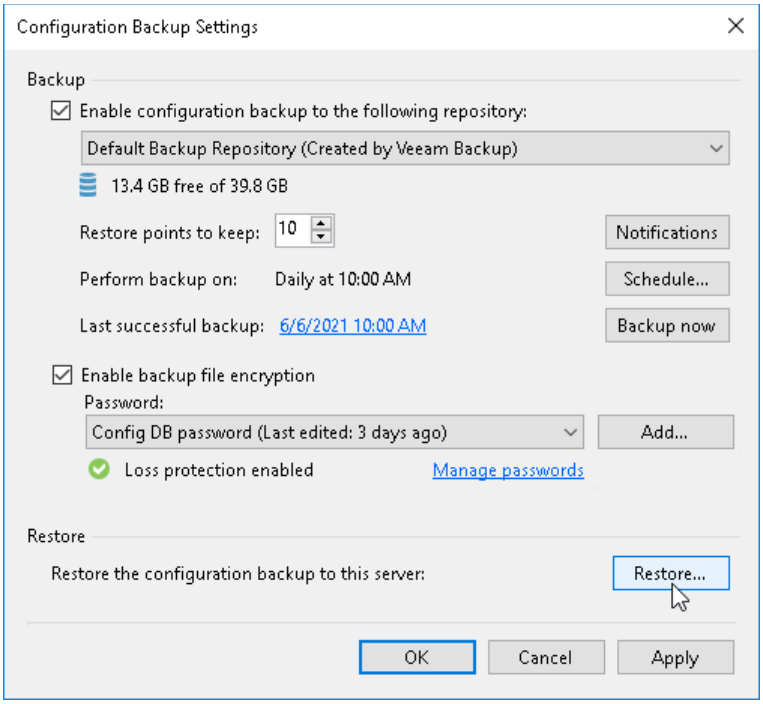
For instructions on how to install Veeam Backup & Replication, see Installing Veeam Backup & Replication.

Step 5. Restore Configuration Database from Backup

Use the native feature of Veeam Backup & Replication to restore the configuration database from the backup created in Step 3.

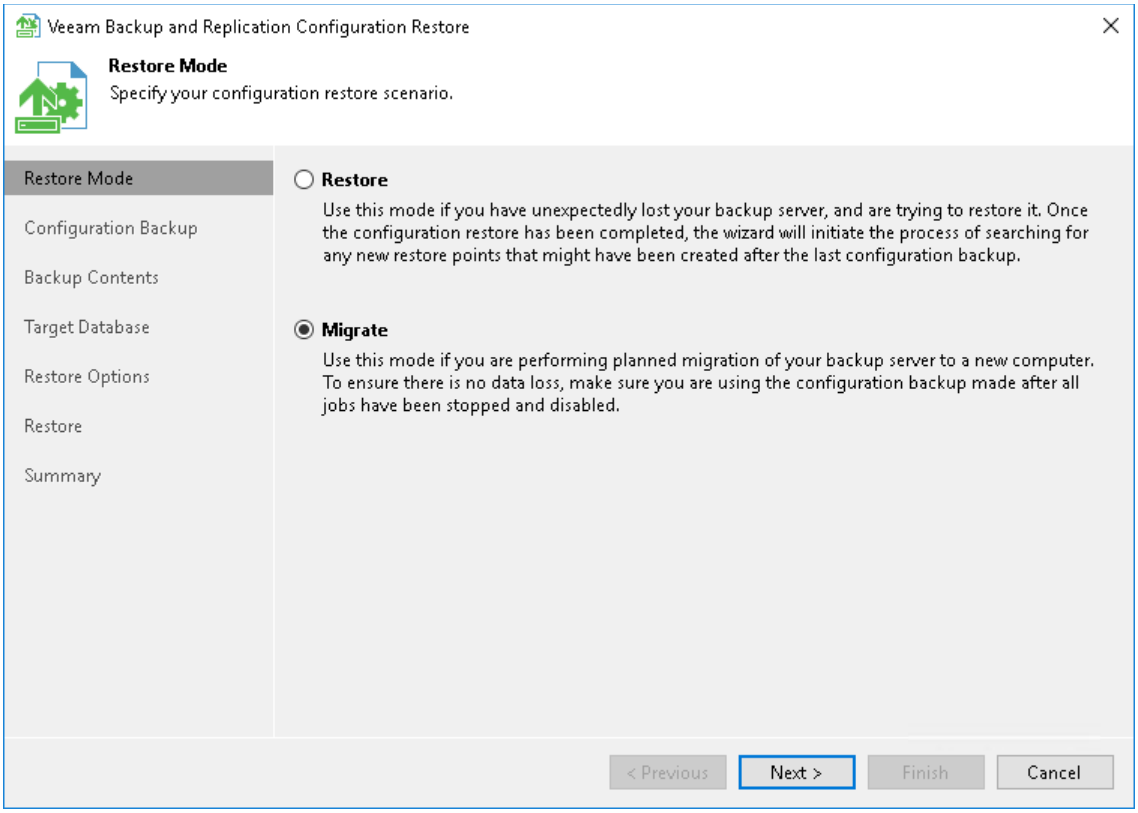
1. On the target backup server, log in to the Veeam Backup & Replication console. For details, see Logging in to Veeam Backup & Replication.
2. From the main menu of the Veeam Backup & Replication console, select **Configuration Backup**.

3. In the **Restore** section, click **Restore** to launch the **Configuration Database Restore** wizard.



4. At the **Restore Mode** step of the wizard, select **Migrate**.

   Before migrating the configuration backup, Veeam Backup & Replication performs an additional check. If the configuration backup does not meet the requirements, Veeam Backup & Replication will offer you an option to switch to the **Restore** mode.

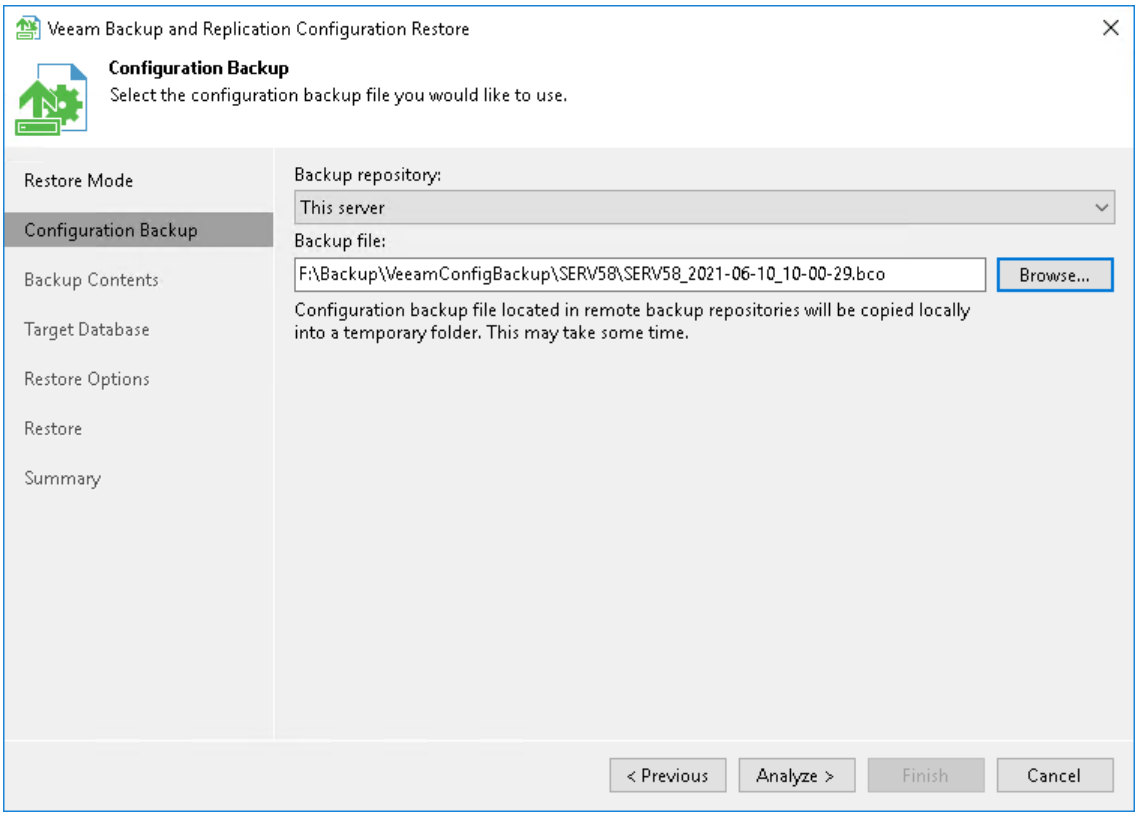


5. At the **Configuration Backup** step of the wizard, specify the backup file of the configuration database

   a. From the **Backup repository** drop-down list, select **This server**.

b. Click **Browse** to specify the backup file location.

> **NOTE**
>
> This backup file has been created in Step 3. You must copy the backup file to the target Veeam Backup & Replication server beforehand.



6. Complete the wizard as described in the Restoring Configuration Database section.

   Veeam Backup & Replication will rescan VM replicas, backup repositories and tape libraries connected to the backup server. The database will be updated to include information about new restore points, and subsequent job sessions will work in a proper way.

Step 6. Finish Configuration

After you restore the configuration backup, finalize the configuration:

- If you created custom registry keys or changed the existing ones on the previous backup server, you must recreate or change the registry keys again manually on the target backup server. You can import saved keys using the `reg import` command or Registry Editor.
- If you have local repositories, after migration to another VM they may be displayed as empty. In this case, add them again and remap the jobs.
- Enable your backup jobs and backup copy jobs. Take a closer look at your backup infrastructure to ensure that everything is working as expected.
- If you use Veeam Backup & Replication to back up storage systems, after migration they will not be added to the backup infrastructure. In this case, you must re-add them after migration completes. Integration with Storage Systems section in the Veeam Backup & Replication User Guide.

- If you use a hardened repository with immutability, after migration this server will not be available. In this case, you must specify single-use credentials for this repository again. For more information, see Editing Settings of Backup Repositories.

> **NOTE**
>
> If the Veeam Backup & Replication server was added to the Veeam Backup Enterprise Manager or Veeam ONE infrastructure, you must re-add the backup server after you migrate it to another server.

Migrating Configuration Database to Another SQL Server

It is the best practice to keep the Veeam Backup & Replication application and its configuration database on the same server to maintain lowest latency and highest performance. However, in some scenarios a remote Microsoft SQL Server instance can be the better choice:

- High Availability. SQL Clustering and AlwaysOn Availability Group on external SQL Servers can be used for high availability of the configuration database. To learn about the configuration details, see this Veeam KB article.
- Licensing. Some enterprises have dedicated virtual clusters for SQL Servers due to licensing constraints. In such cases, you can place the Veeam configuration database on an existing instance to lower the total cost of ownership.

If you need to migrate the Veeam Backup & Replication configuration database to another server, you can connect the configuration database to a Microsoft SQL Server instance deployed on another server and restore the configuration settings from the backup. As a result, you will be able to continue using the same Veeam Backup & Replication server but it will be connected to a configuration database on another server.

> **NOTE**
>
> Mind the following:
>
> - This section gives instructions on how to migrate a configuration database to another server. If you need to migrate the Veeam Backup & Replication application itself, see Migrating Veeam Backup & Replication to Another Server.
> - It is recommended that you use Veeam Backup & Replication tools to create configuration backups and migrate the configuration database. If you use native Microsoft SQL Server tools or others, after migration, some information, such as secure configuration data, may not be accessible.

If you want to migrate the configuration database of Veeam Backup & Replication to another server, perform the following steps:

1. Stop and disable jobs.
2. Back up the configuration database.
3. Connect Veeam Backup & Replication to another SQL Server instance.
4. Restore the configuration database from the backup.
5. [Optional] Reactivate the Enterprise Manager Keyset.
6. Finish the configuration.

Step 1. Stop and Disable Jobs

Before you start the database migration, you must finish all jobs and restore sessions. If the job is scheduled, you must disable the job. For instructions on how to stop and disable jobs, see Managing Jobs.

NOTE

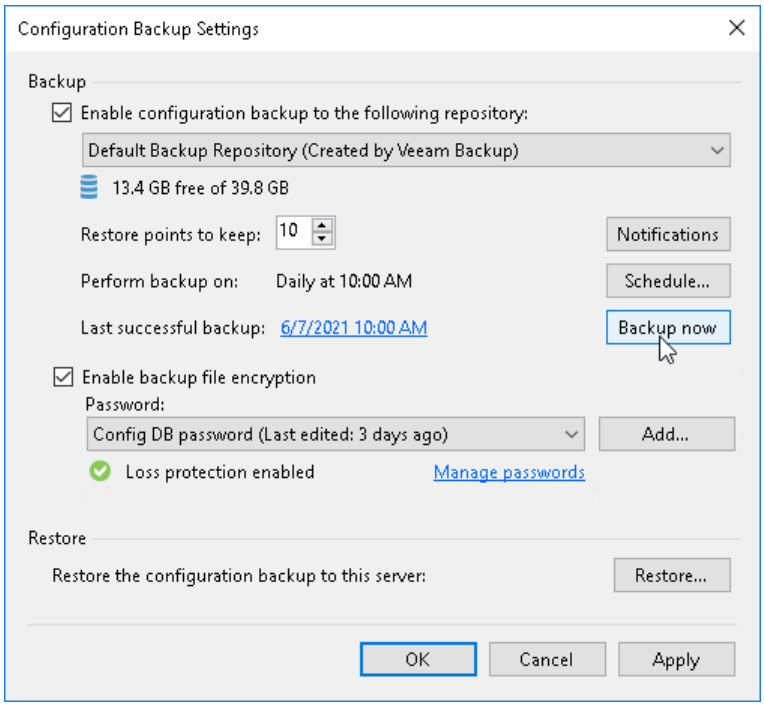
Do not start or enable any jobs until the migration of Veeam Backup & Replication is finished. If you start a job before migration is completed, Veeam Backup & Replication will produce a new restore point in the chain and update the chain metadata. The created configuration backup will not contain information about this new restore point. When you migrate data from the configuration backup to the database and start the job again, Veeam Backup & Replication will fail to synchronize the metadata of the backup chain with data in the database. As a result, the job will fail.

Step 2. Create Configuration Database Backup

To create a configuration database backup manually, perform the following steps:

1. From the main menu of the Veeam Backup & Replication console, select **Configuration Backup**.
2. Make sure that the **Enable configuration backup to the following repository** check box is selected.
3. From the list of repositories, select a backup repository in which the configuration backup must be stored.

4. Click **Backup now**.



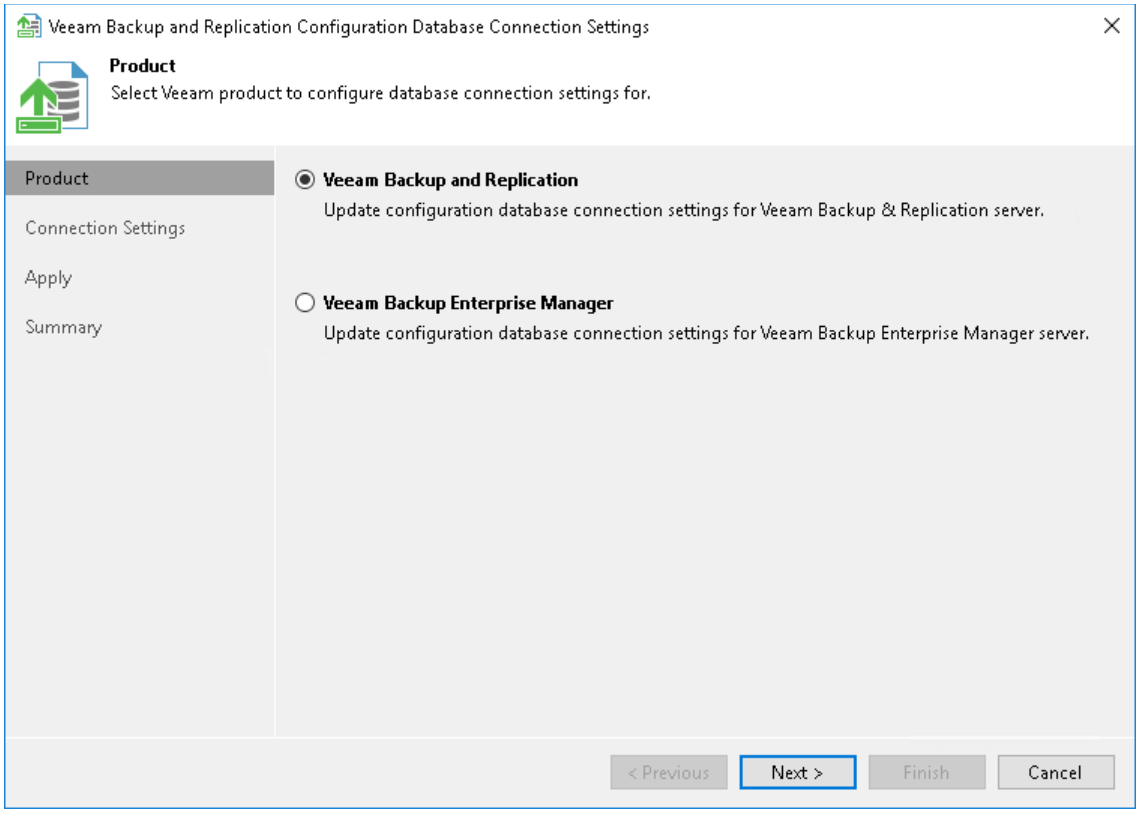
Step 3. Connect Veeam Backup & Replication to Another MS SQL Server Instance

Veeam Backup & Replication comes with the **Configuration Database Connection Settings** utility (Veeam.Backup.DBConfig.exe) that allows you to manage connection settings for Veeam Backup & Replication configuration database. You can use the utility to connect Veeam Backup & Replication to another SQL Server database. To do this, perform the following steps:

1. In the **Start** menu, click **Configuration Database Connection Settings** to launch the `Veeam.Backup.DBConfig.exe` utility. To run the utility, you must have administrative rights on the local machine, as long as the utility makes changes to the registry. If prompted at the launch, choose Run as administrator.

   You can also start the utility from its installation folder. By default, the utility is in the following folder: `%PROGRAMFILES%\Common Files\Veeam\Backup and Replication\DBConfig`.

2. In the **Product** step of the wizard, select **Veeam Backup and Replication**.



3. In the **Connection Settings** step of the wizard, specify the target database settings:

   a. At the **Database name** field, specify a name for a new database that will be created on the target server.

   b. At the **Server name** field, specify the server and Microsoft SQL Server instance name where you want to create the configuration database. Use the following format: `SERVER_NAME\INSTANCE_NAME`. Microsoft SQL Server instances available on the network are shown in the **Server name** list.

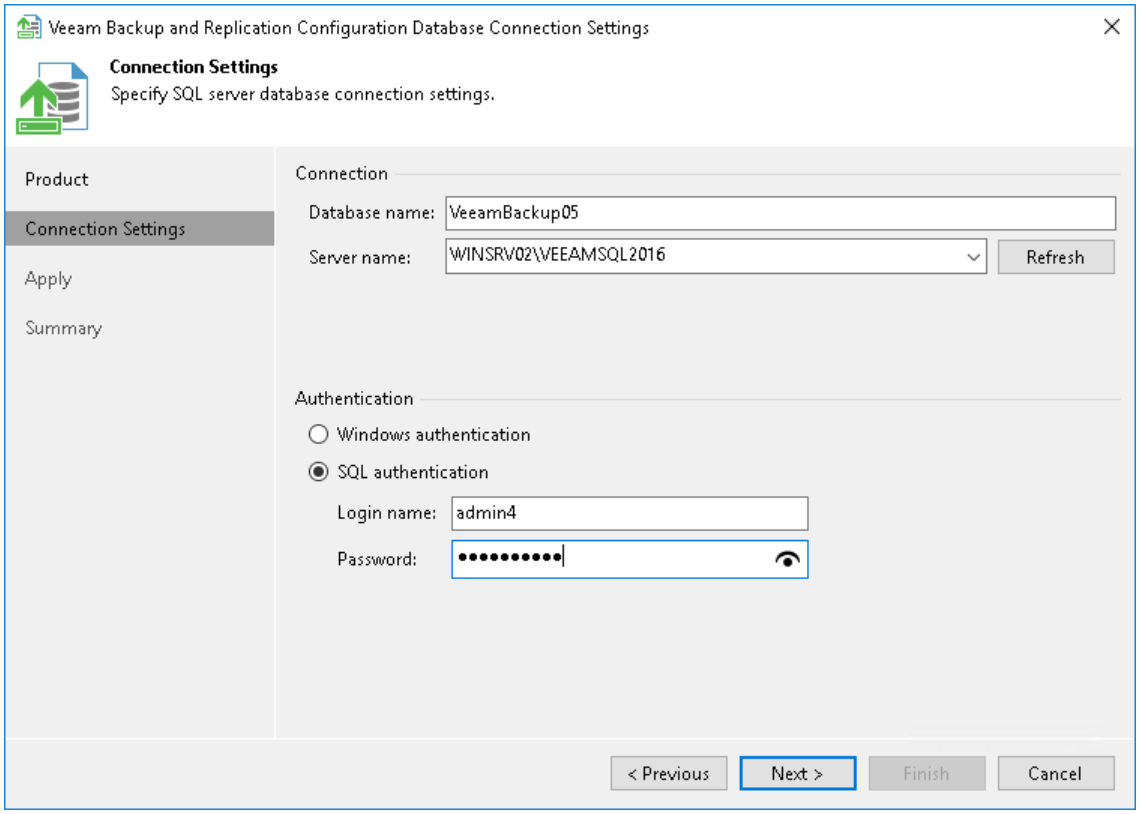
c. Select the **SQL authentication** method to use the Microsoft SQL Server authentication.

d. At the **Login name** and **Password** fields, specify credentials that will be used to connect to the target Microsoft SQL Server instance and create a new database.

IMPORTANT

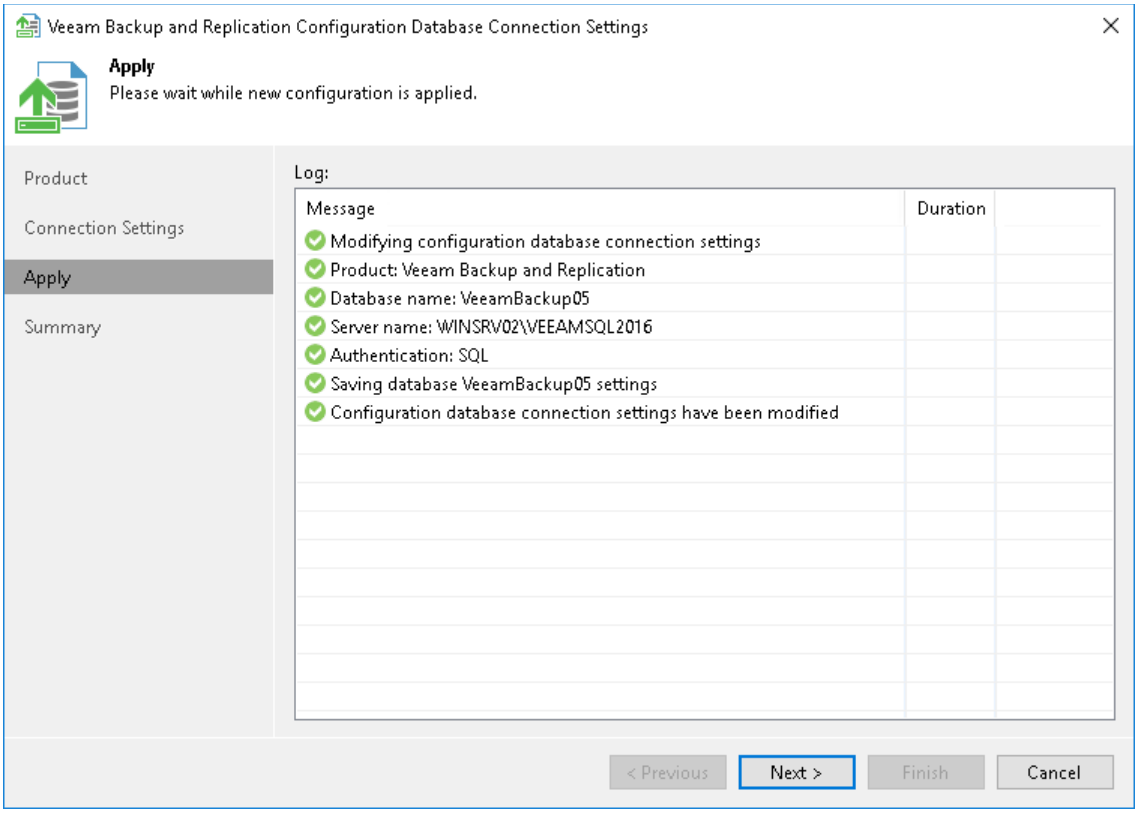
When you migrate the configuration database to another server, you must use the Microsoft SQL Server credentials that have CREATE ANY DATABASE permission on the target Microsoft SQL Server. For details, see Microsoft Docs.

After database creation this account automatically gets a *db_owner* role and can perform all operations with the database. If the current account does not have this permission, a Database Administrator may create an empty database in advance and grant the *db_owner* role to the account that will be used for migration of the configuration database.

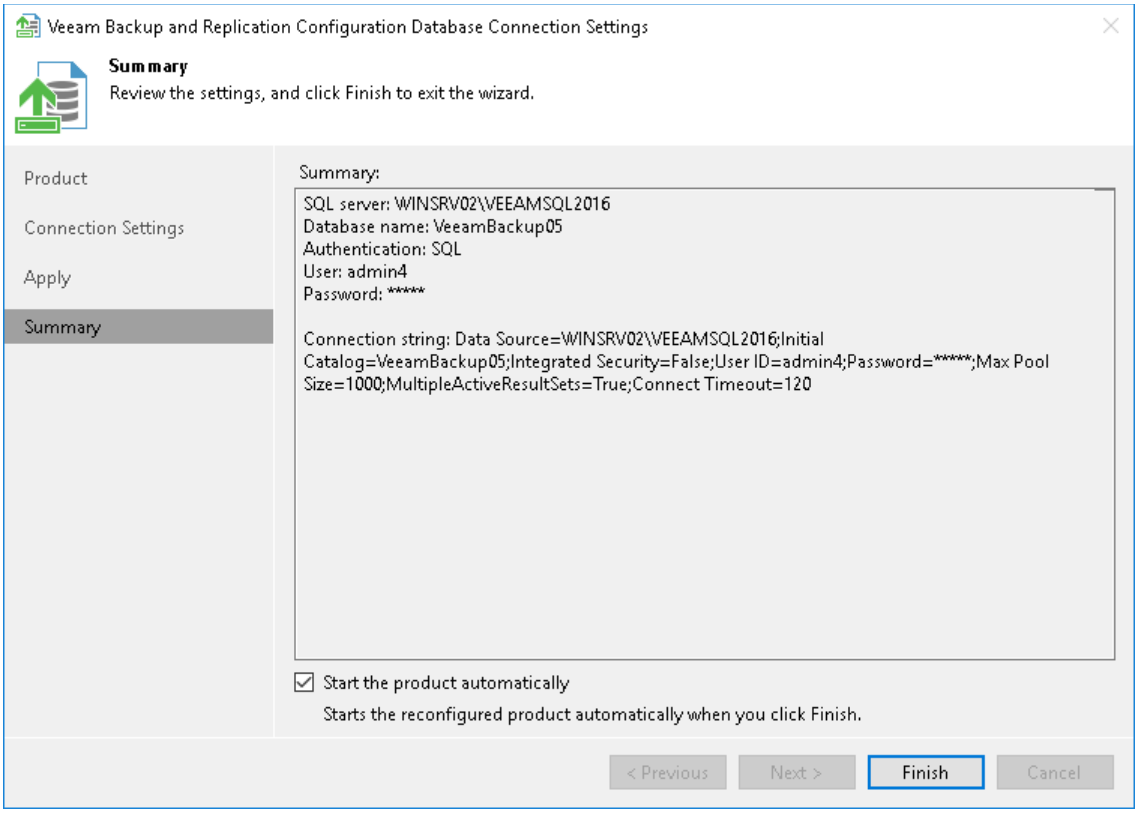


4. Before proceeding, the utility validates the specified settings to make sure that the user account has enough privileges to access the database. For the new settings to be applied, the utility needs to stop Veeam Backup & Replication services and jobs that are currently running. Click **Yes** to confirm the operation.

   Wait for the operation to complete and click **Next** to proceed to the **Summary** step of the wizard.



5. At the **Summary** step of the wizard, select the **Start the product automatically** check box and click **Finish**.

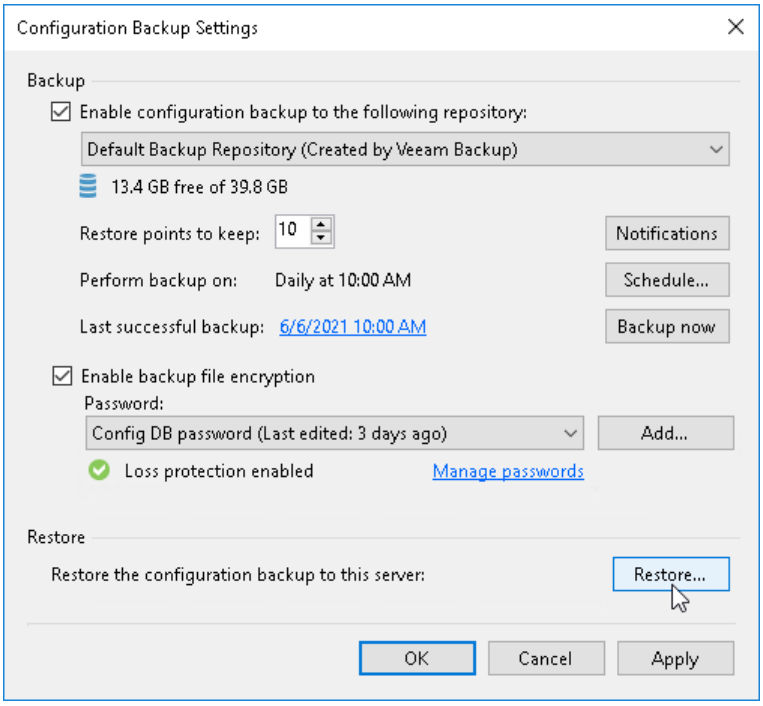


Step 4. Restore Configuration Database from Backup

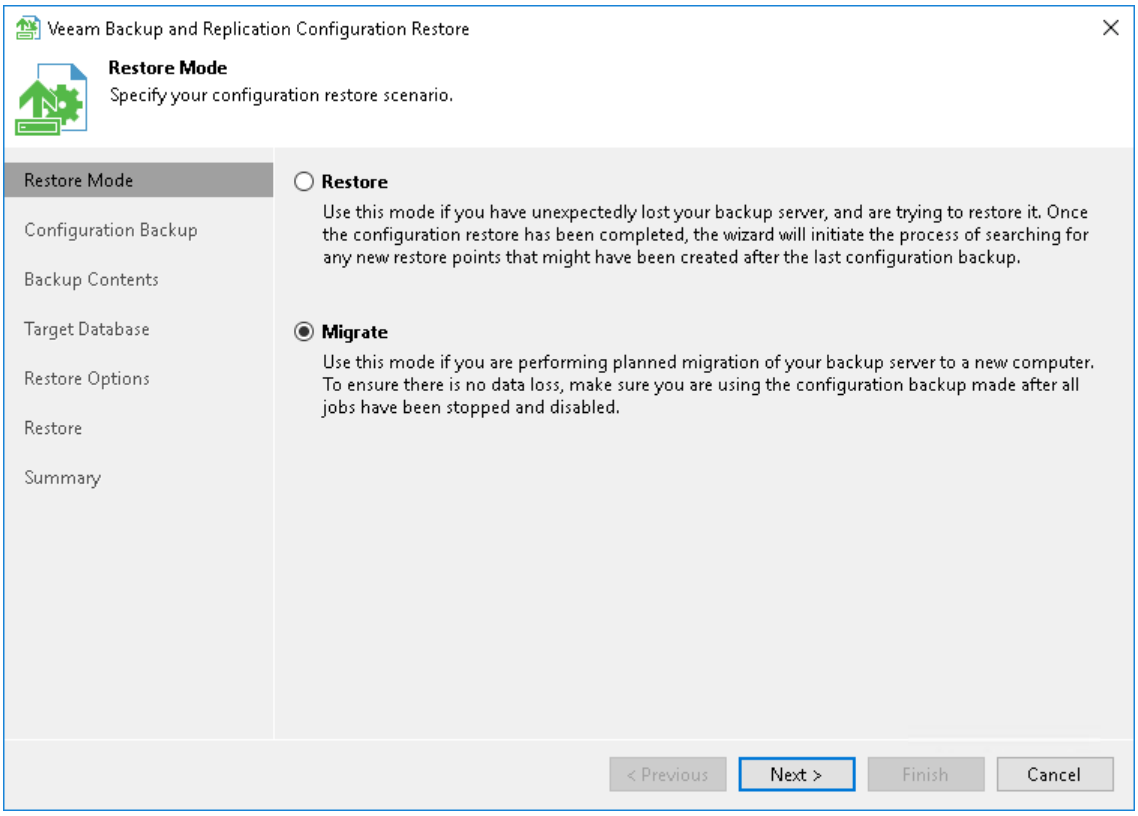
After you connect Veeam Backup & Replication to another Microsoft SQL Server instance, the configuration database will be empty. You must restore the configuration database from the backup created in Step 2.

To restore the configuration database, perform the following:

1. From the main menu of the Veeam Backup & Replication console, select **Configuration Backup**.
2. In the **Restore** section, click **Restore**.



3. At the **Restore Mode** step of the **Veeam Backup & Replication Configuration Restore** wizard, select **Migrate**.



4. Complete the wizard as described in the Restoring Configuration Database section.

Step 5. [Optional] Reactivate Enterprise Manager Keyset

After you migrate the Veeam Backup & Replication configuration database to another server, Veeam Backup Enterprise Manager still sees the Veeam Backup & Replication server. However, you might need to reactivate encryption keys.

If you use the Data Encryption feature to encrypt backups and your Veeam Backup & Replication server is added to the Veeam Backup Enterprise Manager infrastructure, then you must reactivate the Enterprise Manager keyset.

To reactivate the Enterprise Manager key, perform the following steps:

1. In the Veeam Backup Enterprise Manager web console, open the **Settings** section of the **Configuration** view.
2. Open the **Key Management** tab.
3. In the **Managed keys** section, select the necessary keyset and click **Activate**.

For detailed instructions, see the Activating Enterprise Manager Keyset section in the Veeam Backup Enterprise Manager Guide.

Step 6. Finish Configuration

After restoring the configuration database from the backup, finalize the configuration:

- Configure all necessary settings to ensure that you have a working configuration database backup. You can now perform a backup of your new configuration database in the **Configuration Backup Settings** window.

  Reschedule your configuration database backup. Also, check if you can see the *Loss protection enabled* label.

- If you have local repositories, after migration to another VM they may be displayed as empty. In this case, add them again and remap the jobs.

- Enable your backup jobs and backup copy jobs. Take a closer look at your backup infrastructure to ensure that everything is working as expected.

Backup

Veeam Backup & Replication produces image-level backups of VMs. It treats VMs as objects, not as a set of files. When you back up VMs, Veeam Backup & Replication copies a VM image as a whole, at a block level. Image-level backups can be used for different types of restore, including Instant Recovery, entire VM restore, VM file recovery, file-level recovery and so on.

The backup technology is typically used for VMs with lower RTOs. When the primary VM fails, you need some time to restore VM data from a compressed and deduplicated backup file.

About Backup

Veeam Backup & Replication is built for virtual environments. It operates at the virtualization layer and uses an image-based approach for VM backup.

Veeam Backup & Replication does not install agent software inside the VM guest OS to retrieve VM data. To back up VMs, it leverages VMware vSphere snapshot capabilities. When you back up a VM, Veeam Backup & Replication requests VMware vSphere to create a VM snapshot. The VM snapshot can be thought of as a point-in-time copy of a VM that includes virtual disks, system state, configuration and so on. Veeam Backup & Replication uses this point-in-time copy as a source of data for backup.

Veeam Backup & Replication copies VM data from the source datastore at a block level. It retrieves VM data, compresses and deduplicates it, and stores in backup files in the backup repository in Veeam proprietary format.

In Veeam Backup & Replication, backup is a job-driven process. To perform backup, you need to configure backup jobs. A backup job is a configuration unit of the backup activity. The backup job defines when, what, how and where to back up. One backup job can be used to process one or several VMs. You can instruct Veeam Backup & Replication to run jobs automatically by schedule or start them manually.

The first backup job session always produces a full backup of the VM image. Subsequent backup job sessions are incremental — Veeam Backup & Replication copies only those data blocks that have changed since the last backup job session. To keep track of changed data blocks, Veeam Backup & Replication uses different approaches. For more information, see Changed Block Tracking.

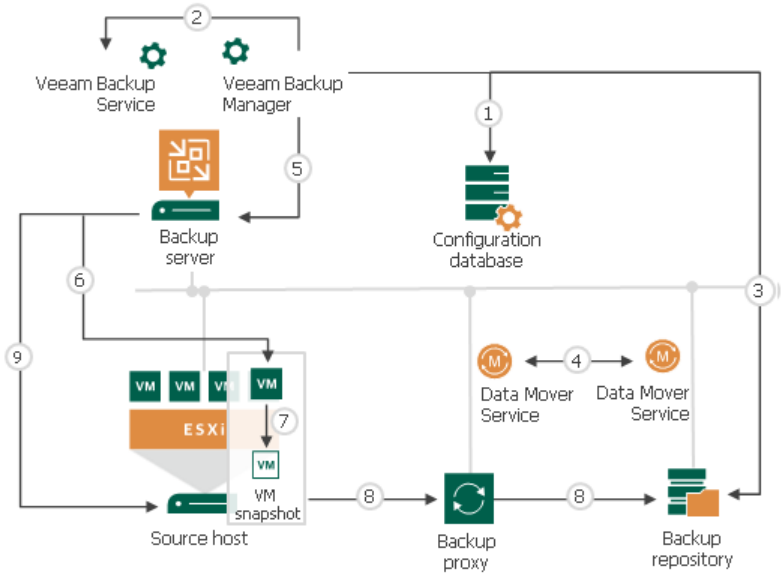
How Backup Works

Veeam Backup & Replication performs VM backup in the following way:

1. When a new backup job session starts, Veeam Backup & Replication starts the Veeam Backup Manager process on the backup server. Veeam Backup Manager reads job settings from the configuration database and creates a list of VM tasks to process. For every disk of VMs added to the job, Veeam Backup & Replication creates a new task.
2. Veeam Backup Manager connects to the Veeam Backup Service. The Veeam Backup Service includes a resource scheduling component that manages all tasks and resources in the backup infrastructure. The resource scheduler checks what backup infrastructure resources are available, and assigns backup proxies and backup repositories to process job tasks.
3. Veeam Backup Manager establishes a connection with Veeam Data Movers on the target backup repository and backup proxy, and sets a number of rules for data transfer, such as network traffic throttling rules and so on.
4. Veeam Data Movers on the backup proxy and backup repository establish a connection with each other for data transfer.
5. Veeam Backup Manager queries information about VMs and virtualization hosts from the Veeam Broker Service.
6. If application-aware image processing is enabled for the job, Veeam Backup & Replication connects to VM guest OSes, deploys non-persistent runtime components or, if necessary, persistent agent components on VM guest OSes and performs in-guest processing tasks.
7. Veeam Backup & Replication requests vCenter Server or ESXi host to create a VM snapshot. VM disks are put to the read-only state, and every virtual disk receives a delta file. All changes that the user makes to the VM during backup are written to delta files.
8. The source Veeam Data Mover reads the VM data from the read-only VM disk and transfers the data to the backup repository in one of transport modes. During incremental job sessions, the source Veeam Data Mover uses CBT to retrieve only those data blocks that have changed since the previous job session. If CBT is not available, the source Veeam Data Mover interacts with the target Veeam Data Mover on the backup repository to obtain backup metadata, and uses this metadata to detect blocks that have changed since the previous job session.

   While transporting VM data, the source Veeam Data Mover performs additional processing. It filters out zero data blocks, blocks of swap files and blocks of excluded VM guest OS files. The source Veeam Data Mover compresses VM data and transports it to the target Veeam Data Mover.

9. After the backup proxy finishes reading VM data, Veeam Backup & Replication requests the vCenter Server or ESXi host to commit the VM snapshot.



Backup Architecture

Veeam Backup & Replication uses the following components for the backup process:

- One or more source hosts with associated datastores
- One or more backup proxies
- Backup repository
- [Optional] One or more guest interaction proxies
- [For shared folder backup repository] Gateway server

All backup infrastructure components that are engaged in the job make up a data pipe. The source host and backup repository produce two terminal points for the data flow. Veeam Backup & Replication processes VM data in multiple cycles, moving VM data over the data pipe block by block.

Veeam Backup & Replication collects VM data, transforms and transports it to target with the help of Veeam Data Movers. Veeam Backup & Replication uses two-service architecture — one Veeam Data Mover controls interaction with the source host, and the other one controls interaction with the backup repository. Veeam Data Movers communicate with each other and maintain a stable connection.

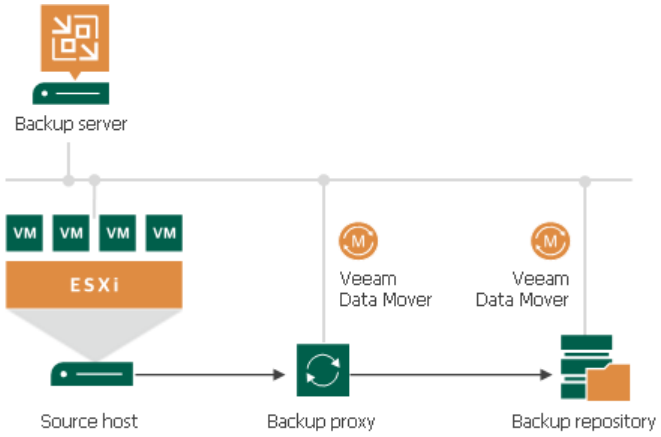
When a new backup session starts, Veeam Backup & Replication performs the following actions:

1. Veeam Backup & Replication deploys non-persistent runtime components or, if necessary, persistent agent components on VM guest OSes using the guest interaction proxy (for Microsoft Windows VMs) or backup server (for VMs with other OSes).
2. The target-side Veeam Data Mover obtains job instructions and communicates with the source-side Veeam Data Mover to begin data collection.
3. The source-side Veeam Data Mover copies VM data from the source storage in one of transport modes. During incremental job runs, the source-side Veeam Data Mover retrieves only those data blocks that have changed since the previous job session.

   While copying, the source-side Veeam Data Mover performs additional data processing. It filters out zero data blocks, blocks of swap files and blocks of excluded VM guest OS files, compresses and deduplicates VM data blocks and moves them to the target-side Veeam Data Mover.
4. The target-side Veeam Data Mover deduplicates similar blocks of data on the target side and writes the result to the backup file in the backup repository.

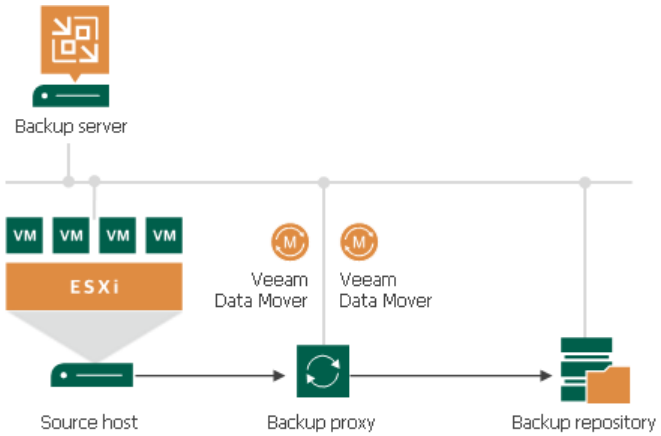
On-Site Backup

To back up to a Microsoft Windows or Linux backup repository in the local site, you need to deploy a backup proxy on a machine that has access to the source datastore, and point the backup job to this backup proxy. In this scenario, the source-side Veeam Data Mover is started on the backup proxy, and the target-side Veeam Data Mover is started on the Microsoft Windows or Linux repository. VM data is sent from the backup proxy to the backup repository over the LAN.



To back up to a shared folder in the local site, you need to deploy a gateway server that has access to the shared folder backup repository. You can assign the role of a gateway server to the backup server itself or any Microsoft Windows machine added to the backup infrastructure.

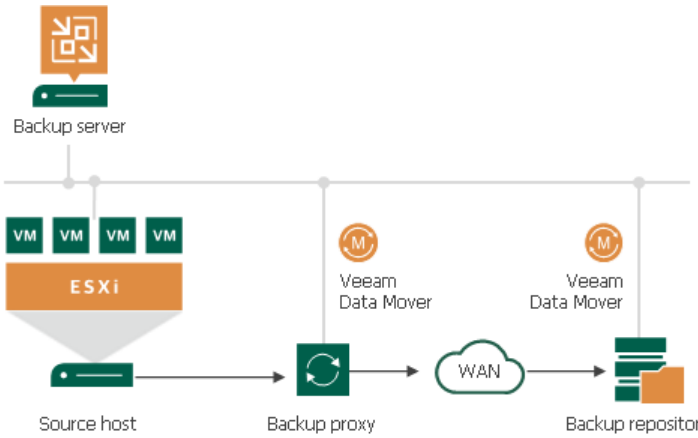
You can use the same Microsoft Windows machine as the backup proxy and gateway server for SMB. In this scenario, Veeam Backup & Replication starts the source-side and target-side Veeam Data Movers on the same machine, and sends VM data from the backup proxy to the shared folder backup repository over the LAN.



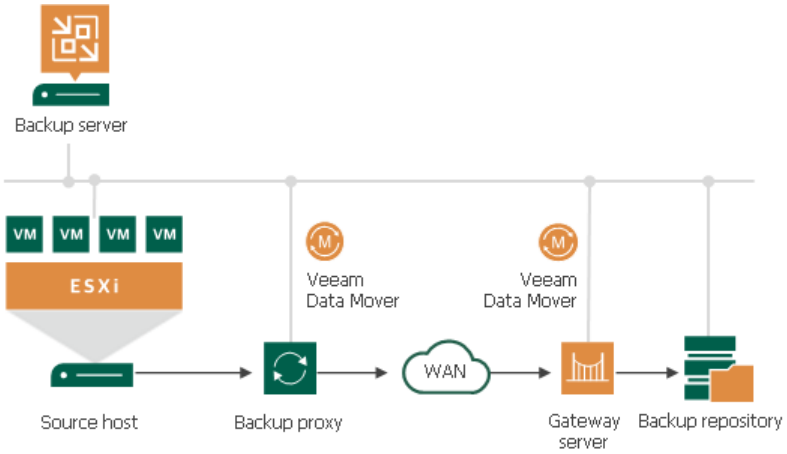
Off-Site Backup

The common requirement for off-site backup is that one Veeam Data Mover runs in the production site (closer to the source datastore), and the other Veeam Data Mover runs in the remote site, closer to the backup repository. During backup, Veeam Data Movers maintain a stable connection, which allows for uninterrupted operation over the WAN or slow links.

To back up to a Microsoft Windows or Linux repository in the remote site, you need to deploy a backup proxy in the production site, closer to the source datastore. In this scenario, the source-side Veeam Data Mover is started on the backup proxy, and the target-side Veeam Data Mover is started on the Microsoft Windows or Linux repository. VM data is sent from the backup proxy to the backup repository over the WAN.



To back up VMs to shared folder backup repository in the remote site, you must deploy a backup proxy in the source site and a gateway server in the remote site. The shared folder backup repository must be pointed at the target-side gateway server. During backup, the source-side Veeam Data Mover is started on the source backup proxy in the production site, and the target-side Veeam Data Mover is started on the target gateway server in the remote site. VM data is transferred between the backup proxy and gateway server over the WAN.



Backup Chain

Veeam Backup & Replication creates and maintains the following types of backup files:

- VBK — full backup files that store copies of full VM images.
- VIB or VRB — incremental backup files that store incremental changes of VM images.
- VBM — backup metadata files that store information about the backup job, VMs processed by the backup job, number and structure of backup files, restore points, and so on. Metadata files facilitate import of backups, backup mapping and other operations.

In addition to these file types, Veeam Backup & Replication can create the following files in the backup repository:

- VSB — virtual synthetic backup files used for generation of virtual full backups on tapes. For more information, see the Virtual Full Backup section in the Veeam Backup & Replication User Guide.
- VLB, VSM and VLM — files that store Microsoft SQL Server transaction log data. For more information, see Microsoft SQL Server Logs Backup.
- VLB, VOM and VLM — files that store Oracle archived log data. For more information, see Oracle Logs Backup.

All backup files created by the backup job reside in a dedicated job folder in the backup repository. For example, if you create a backup job with the *DC Backup* name, Veeam Backup & Replication will create the `DC Backup` folder on the target backup repository and store all backup files produced with this job in this folder.

Backup files make up a backup chain. The backup chain consists of first full backup file, incremental backup files and, additionally, backup metadata file. The number of backup files and how Veeam Backup & Replication orders them in the chain depend on the chosen backup method. For details, see Backup Methods.

Full and incremental backup files correspond to restore points of backed-up VMs. You can think of restore points as of "snapshots" of VM data at specific points in time. Restore points let you roll back VMs to the necessary state.

To roll back a VM to a specific point in time, you need a chain of backup files: a full backup file plus a set of incremental backup files dependent on this full backup file. If some file in the backup chain is missing, you will not be able to roll back to the necessary state. For this reason, you must not delete separate backup files from the backup repository manually. Instead, you must specify retention policy settings that will let you maintain the desired number of backup files in the backup repository.

By default, during every backup job session Veeam Backup & Replication writes data of all VMs to the same backup file. If necessary, you can instruct Veeam Backup & Replication to create per-machine backup chains — that is, produce a separate backup chain for every VM added to the backup job. For details, see Per-Machine Backup Files.

Backup Methods

Veeam Backup & Replication provides the following methods for creating backup chains:

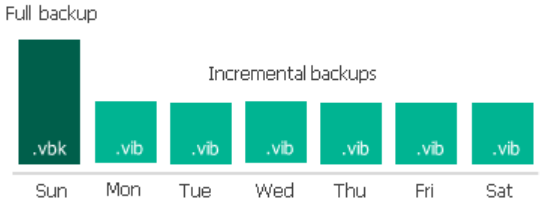
- **Forever forward incremental (FFI)**

  When the forever forward incremental (FFI) backup method is used, Veeam Backup & Replication creates a backup chain that consists of the first full backup file (VBK) and a set of forward incremental backup files (VIBs) following it.

This backup method helps you save space on the backup storage because Veeam Backup & Replication stores only one full backup file and removes incremental backup files once the retention period is exceeded. To meet retention policy settings, Veeam Backup & Replication injects data of an incremental file into the full backup file before deleting the increment. Such transformations can lead to the fragmentation of the full backup file, and you have to schedule the compact of full backup file operation. This operation produces additional I/O load on the backup storage. Overall, the FFI method produces the medium I/O impact on the storage compared to other backup methods.

Restore to the earliest restore point from backup files created using the FFI method is the fastest compared to other methods because the first available restore point is always a full backup. Restore to other restore points can be compared by speed with FI method.

For more information on the FFI backup method and how it works, see Forever Forward Incremental Backup.



- **Forward incremental (FI)**

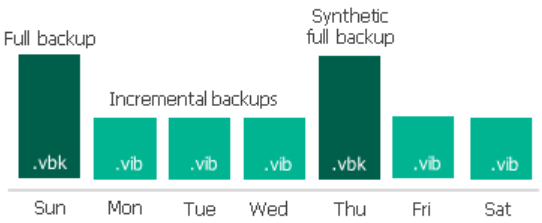
  When the forward incremental (FI) backup method is used, Veeam Backup & Replication creates a backup chain that consists of multiple full backup files (VBKs) and sets of forward incremental backup files (VIBs) following each full backup file. Full backups can be created using synthetic full and/or active full method. By doing regular full backups, the backup chain is split into shorter series. This lowers the chances of losing the backup chain completely and makes this backup method the most reliable.

  This backup method requires more storage space than other methods because the backup chains contains multiple full backup files and sometimes Veeam Backup & Replication stores more restore points than specified in the retention policy settings due to the specifics of the FI retention policy. For more information on how backups are retained, see Forward Incremental Backup Retention Policy.

  The FI backup method produces the lowest I/O impact on the backup storage. However, the impact on the backup storage increases on days when synthetic full backups are scheduled, and impact on the production storage increases on days when active full backups are scheduled.

  Restore from backup files created using the FI method is the most optimal in time compared to other methods (in cases when you you do not restore to the earliest or latest restore point). That is because the backup chain is usually divided into short series of full backup files and incremental files and the aggregation of the desired restore point does not take long time.

  For more information on the FI backup method and how it works, see Forward Incremental Backup.



- **Reverse incremental (RI)**

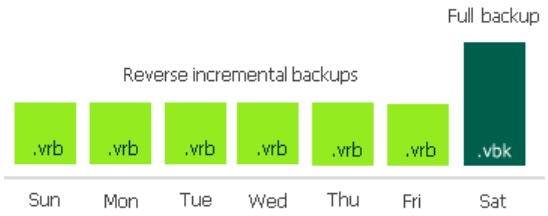
  When the reverse incremental (RI) backup method is used, Veeam Backup & Replication creates a backup chain that consists of the full backup file (VBK) and a set of reverse incremental backup files (VRBs) preceding it.

This backup method helps you save space on the backup storage because Veeam Backup & Replication stores only one full backup file (if you do not schedule periodic full backups) and removes incremental backup files once the retention period is exceeded. For more information on how backups are retained, see Reverse Incremental Backup.

The RI method produces the heaviest I/O impact on the backup storage compared to other backup methods. That is because during backup, Veeam Backup & Replication injects changed data blocks into the full backup file and also creates reverse incremental backup files. Such transformations can lead to the fragmentation of the full backup file, and you have to schedule the compact of full backup file operation. This operation produces additional I/O load on the backup storage.

Restore to the latest restore point from backup files created using the RI method is the fastest compared to other methods because the most recent restore point is always a full backup, and it gets updated after every backup cycle. Restore to earlier restore points is slower than for other backup methods.

For more information on the RI backup method and how it works, see Reverse Incremental Backup.

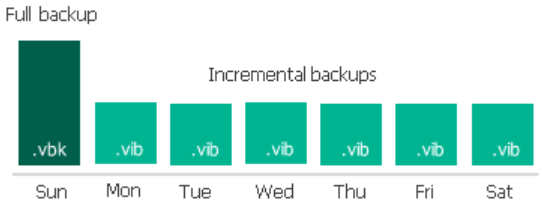


Forever Forward Incremental Backup

The forever forward incremental backup method produces a backup chain that consists of the first full backup file (VBK) and a set of forward incremental backup files (VIB) following it.

Veeam Backup & Replication creates a forever forward incremental backup chain in the following way:

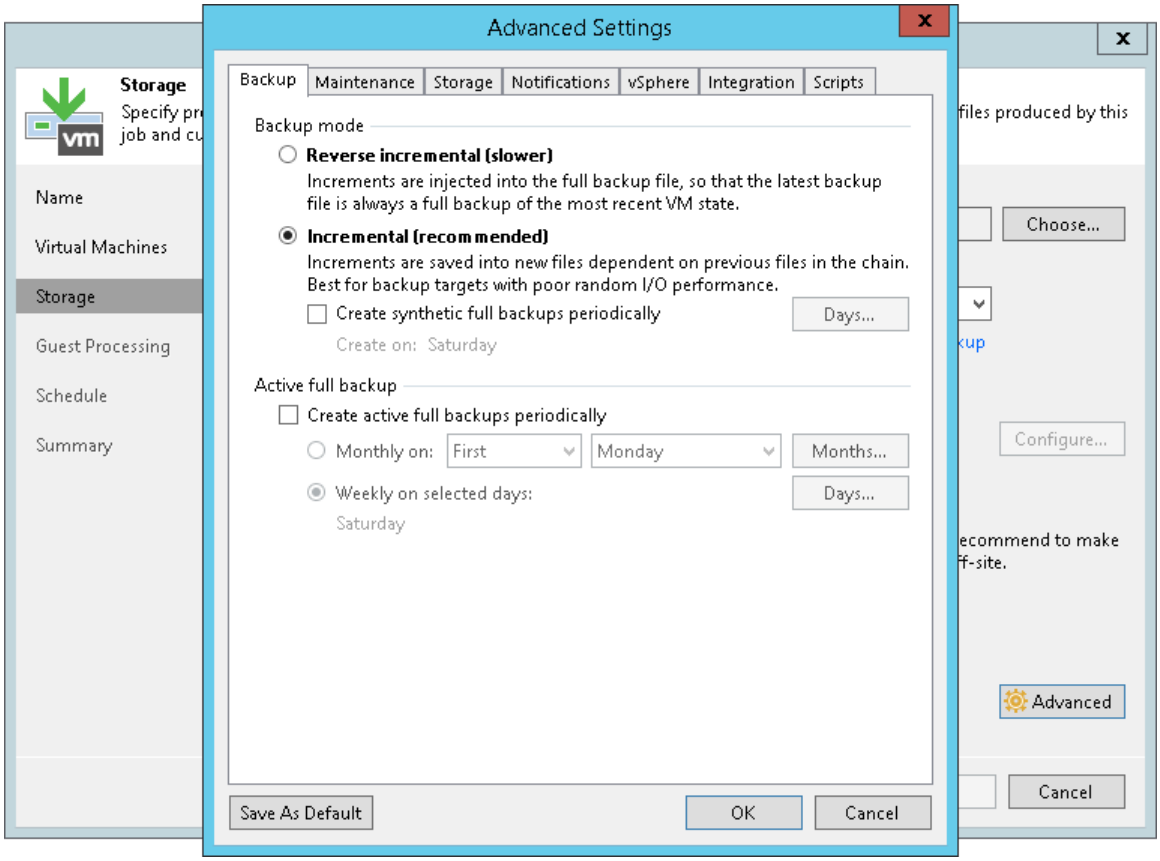
1. During the first session of a backup job, Veeam Backup & Replication creates a full backup file in the backup repository.
2. During subsequent backup job sessions, Veeam Backup & Replication copies only VM data blocks that have changed since the last backup job session (full or incremental) and saves these blocks as an incremental backup file in the backup chain.
3. After adding a new restore point to the backup chain, Veeam Backup & Replication checks the retention policy for the job. If Veeam Backup & Replication detects an outdated restore point, it transforms the backup chain to make room for the most recent restore point. For more information, see Forever Forward Incremental Backup Retention Policy.



To use the forever forward incremental backup method, you must select the following options in the backup job settings:

1. Select the **Incremental** backup mode.

2. Do not enable synthetic full backups and/or active full backups. If you enable synthetic and/or active full backups, Veeam Backup & Replication will produce a forward incremental backup chain.



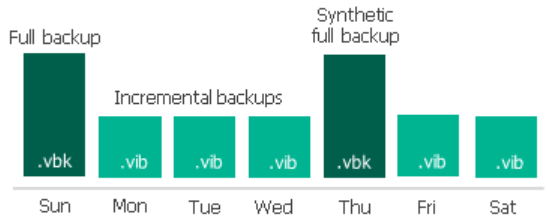
Forward Incremental Backup

The forward incremental backup method produces a backup chain that consists of the first full backup file (VBK) and a set of forward incremental backup files (VIB) following it. Additionally, the forward incremental backup chain contains synthetic full and/or active full backup files that "split" the backup chain into shorter series.

Veeam Backup & Replication creates a forward incremental backup chain in the following way:

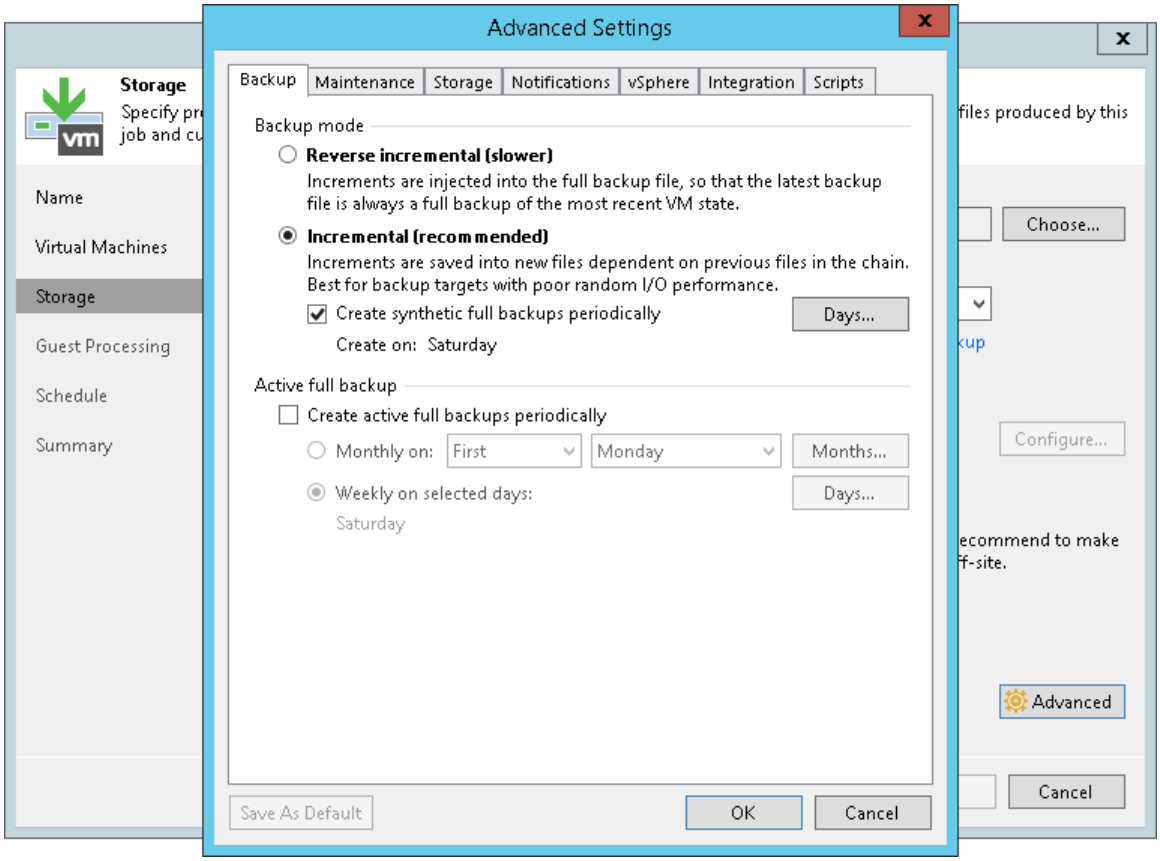
1. During the first backup job session, Veeam Backup & Replication creates a full backup file in the backup repository.
2. During subsequent backup job sessions, Veeam Backup & Replication copies only VM data blocks that have changed since the last backup job session (full or incremental) and saves these blocks as an incremental backup file in the backup chain.
3. On a day when the synthetic full or active full backup is scheduled, Veeam Backup & Replication creates a full backup file and adds it to the backup chain. Incremental restore points produced after this full backup file use it as a new starting point.

4. After adding a new restore point to the backup chain, Veeam Backup & Replication checks the retention policy set for the job. If Veeam Backup & Replication detects an outdated restore point, it attempts to remove this point from the backup chain. For more information, see Retention for Forward Incremental Backup.



The forward incremental backup with synthetic full backup enabled is a default method for backup chain creation. To use the forward incremental backup method, you can leave the default settings or select the following options in the backup job settings:

1. Select the **Incremental** backup mode.
2. Enable synthetic full backups and/or active full backups. If the synthetic full backup and/or active full backups are not enabled, Veeam Backup & Replication will produce a forever forward incremental backup chain.



Reverse Incremental Backup

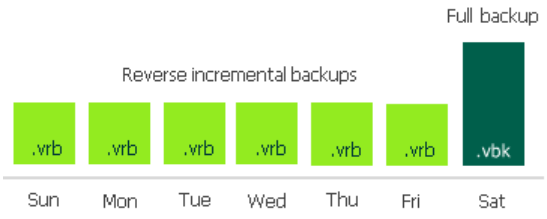
The reverse incremental backup method produces a backup chain that consists of the last full backup file (VBK) and a set of reverse incremental backup files (VRB) preceding it.

Veeam Backup & Replication creates a reverse incremental backup chain in the following way:

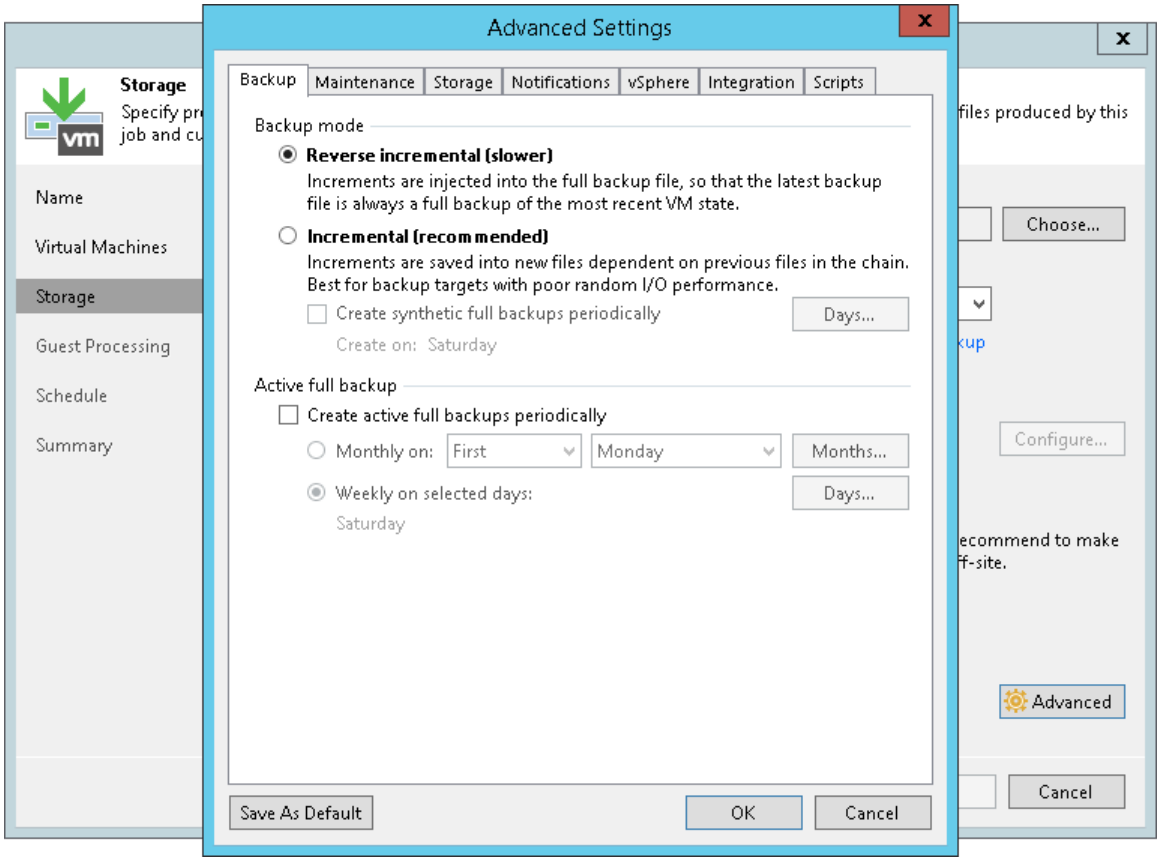
1. During the first backup job session, Veeam Backup & Replication creates a full backup file in the backup repository.
2. During subsequent backup job sessions, Veeam Backup & Replication copies only VM data blocks that have changed since the last backup job session. Veeam Backup & Replication "injects" copied data blocks into the full backup file to rebuild it to the most recent state of the VM. Additionally, Veeam Backup & Replication creates a reverse incremental backup file containing data blocks that are replaced when the full backup file is rebuilt, and adds this reverse incremental backup file before the full backup file in the backup chain.
3. After adding a new restore point to the backup chain, Veeam Backup & Replication checks the retention policy set for the job. If Veeam Backup & Replication detects an outdated restore point, it removes this point from the backup chain. For more information, see Retention for Reverse Incremental Backup.

As a result, the most recent restore point in the backup chain is always a full backup, and it gets updated after every successful backup job session.

The reverse incremental backup method lets you immediately restore a VM to the most recent state without extra processing because the most recent restore point is a full backup file. If you need to restore a VM to a particular point in time, Veeam Backup & Replication applies the required VRB files to the VBK file to get to the required restore point.



To use the reverse incremental backup method, you must select the **Reverse incremental** option in the backup job settings.



Switching Between Backup Methods

You can easily switch between backup methods. Veeam Backup & Replication does not transform the previously created chain. Instead, it creates a new backup chain next to the existing one in the following manner:

- If you switch from the reverse incremental method to the forever forward incremental or forward incremental method, Veeam Backup & Replication creates a set of incremental backup files next to the reverse incremental chain. The full backup file in the reverse incremental chain is used as a starting point for incremental backup files.
- If you switch from the forever forward incremental or forward incremental method to the reverse incremental method, Veeam Backup & Replication first creates a full backup file next to incremental backup files. During every new job session, Veeam Backup & Replication transforms this full backup file and adds reverse incremental backup files to the backup chain.
- If you switch from the forever forward incremental method to the forward incremental method, Veeam Backup & Replication creates synthetic full backups according to the specified schedule. Old backup chain is deleted when the number of restore points in the new chain reach the retention limit.
- If you switch from the forward incremental method to the forever forward incremental method, synthetic full backups are no longer created. When the number of restore points created since the last full backup reaches the retention limit, the old backup chain is deleted. Thereafter, with each restore point creation the earliest increment file will merge with the full backup file.

Full Backup Methods

Veeam Backup & Replication provides the following methods for creation of full backup files:

- Active Full Backup

  When you perform active full backup, Veeam Backup & Replication retrieves VM data from the source datastore where the VM resides, compresses and deduplicates it and writes it to the VBK file in the backup repository.

- Synthetic Full Backup

  When you perform synthetic full backup, Veeam Backup & Replication does not retrieve VM data from the source datastore. Instead, it synthesizes a full backup from data you already have in the backup repository. Veeam Backup & Replication accesses the previous full backup file and a chain of subsequent incremental backup files on the backup repository, consolidates VM data from these files and writes consolidated data into a new full backup file. As a result, the created synthetic full backup file contains the same data you would have if you created an active full backup.

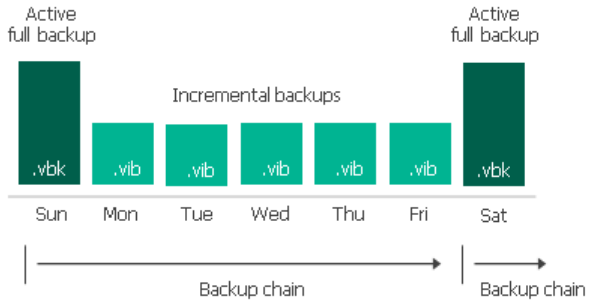
> **TIP**
>
> You can perform both active and synthetic full backups. For more information on how to do that, see Backup Settings.

Active Full Backup

In some cases, you need to regularly create a full backup. For example, your corporate backup policy may require that you create a full backup on weekend and run incremental backup on work days. To let you conform to these requirements, Veeam Backup & Replication lets you create active full backups.

The active full backup produces a full backup of a VM, just as if you run the backup job for the first time. Veeam Backup & Replication retrieves data for the whole VM from the source, compresses and deduplicates it and stores it to the full backup file — VBK.

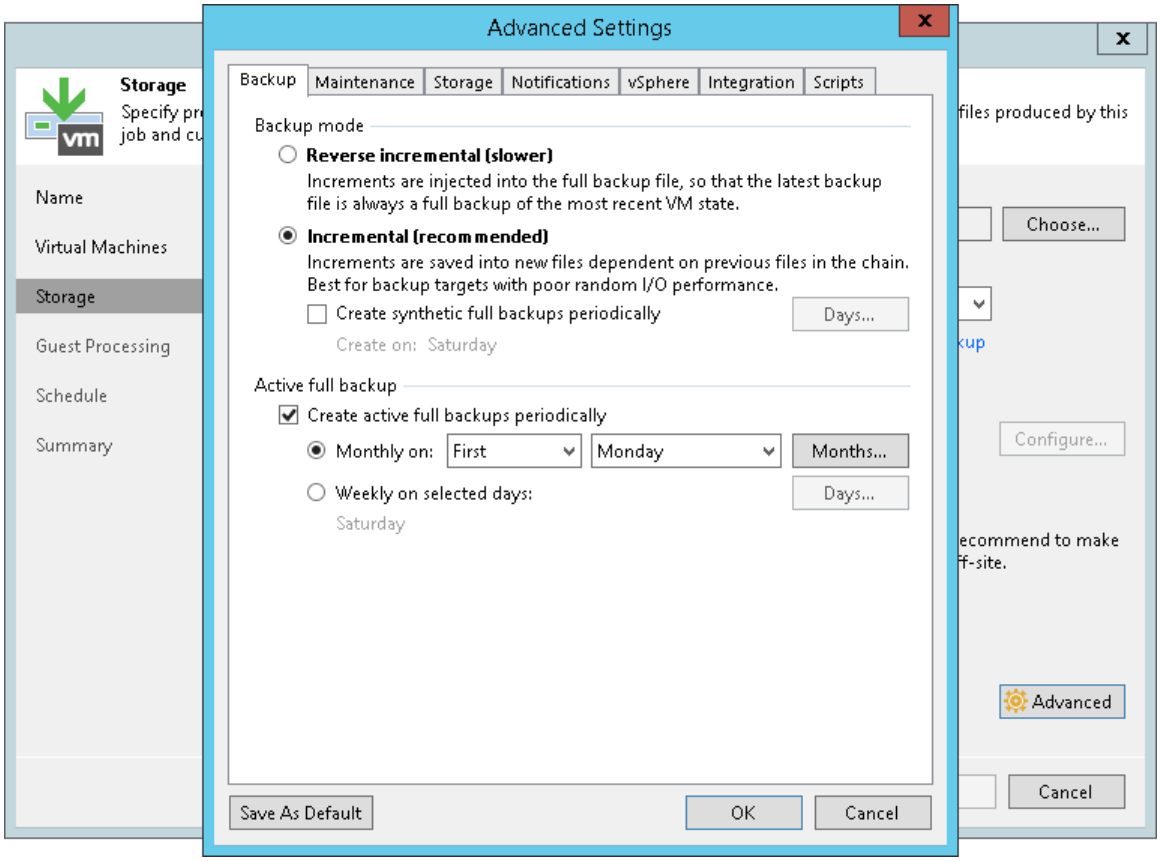
The active full backup resets a backup chain. All incremental backup files use the latest active full backup file as a new starting point. A previously used full backup file remains on disk until it is automatically deleted according to the retention policy.



You can create active full backups manually or schedule a backup job to create active full backups periodically.

- To create an active full backup manually, use the **Active Full** command from the shortcut menu of a corresponding backup job.

- To schedule active full backups, specify scheduling settings in the **Advanced** section of a corresponding backup job. You can schedule active full backups to run weekly, for example, every Saturday, or monthly, for example, every fourth Sunday of a month.



Active Full Backup Schedule

Veeam Backup & Replication automatically triggers a backup job to create an active full backup, even if a regular backup job session is not scheduled on this day. The job session is started at the same time when the parent backup job is scheduled. For example, if you schedule the parent backup job at 12:00 AM Sunday through Friday, and schedule active full backup on Saturday, Veeam Backup & Replication will start a backup job session that will produce an active full backup at 12:00 AM on Saturday.

If the parent backup job is not scheduled to run automatically or disabled, Veeam Backup & Replication will not perform active full backup.

If a regular backup job is scheduled together with active full backup, Veeam Backup & Replication will produce only one backup file — an active full backup that will contain the latest state of the source VM. An incremental backup file that should have been created by the backup job schedule will not be added to the backup chain.

Veeam Backup & Replication creates an active full backup only once a day on which active full backup is scheduled (unless you create a full backup manually). If you run the backup job again on the same day, Veeam Backup & Replication will perform incremental backup in a regular manner.

Synthetic Full Backup

In some situations, running active full backups periodically may not be an option. Active full backups are resource-intensive and consume considerable amount of network bandwidth. As an alternative, you can create synthetic full backups.

In terms of data, the synthetic full backup is identical to a regular full backup. Synthetic full backup produces a VBK file that contains data of the whole VM. The difference between active and synthetic full backup lies in the way how VM data is retrieved:

- When you perform active full backup, Veeam Backup & Replication retrieves VM data from the source datastore where the VM resides, compresses and deduplicates it and writes it to the VBK file in the backup repository.
- When you perform synthetic full backup, Veeam Backup & Replication does not retrieve VM data from the source datastore. Instead, it synthesizes a full backup from data you already have in the backup repository. Veeam Backup & Replication accesses the previous full backup file and a chain of subsequent incremental backup files on the backup repository, consolidates VM data from these files and writes consolidated data into a new full backup file. As a result, the created synthetic full backup file contains the same data you would have if you created an active full backup.

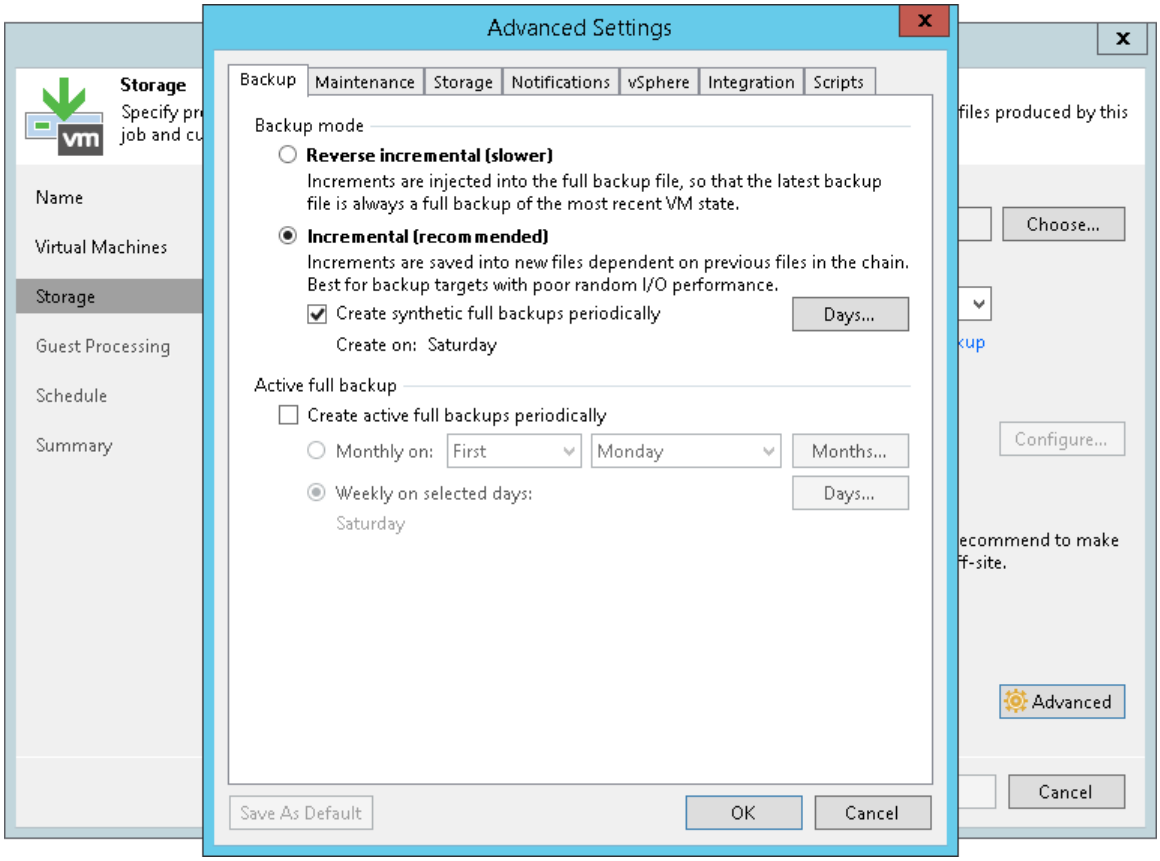
The synthetic full backup has a number of advantages:

- The synthetic full backup does not use network resources: it is created from backup files you already have on disk.
- The synthetic full backup produces less load on the production environment: it is synthesized right on the backup repository.

Veeam Backup & Replication treats synthetic full backups as regular full backups. As well as any other full backup file, the synthetic full backup file resets the backup chain. All subsequent incremental backup files use the synthetic full backup file as a new starting point. A previously used full backup file remains on disk until it is automatically deleted according to the retention policy.

> **NOTE**
>
> If you enable both synthetic and active full backups and schedule their creation on the same day, the synthetic full backup is not created.

To create synthetic full backups, you must enable the **Create synthetic full backups periodically** option and schedule creation of synthetic full backups on specific days in the backup job settings.

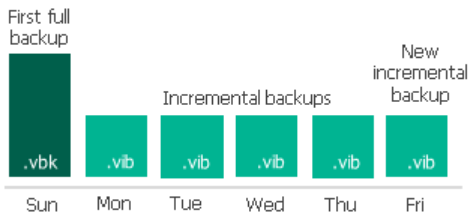


How Synthetic Full Backup Works

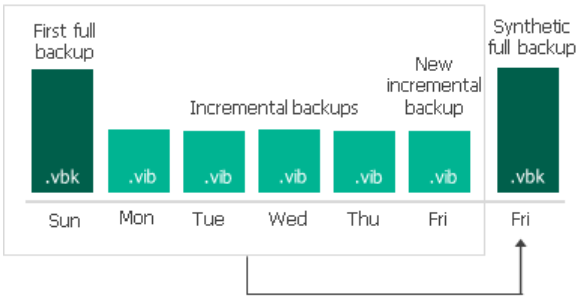
To create a synthetic full backup, Veeam Backup & Replication performs the following steps:

1. On a day when synthetic full backup is scheduled, Veeam Backup & Replication triggers a new backup job session. During this session, Veeam Backup & Replication first performs incremental backup in a regular manner and adds a new incremental backup file to the backup chain.

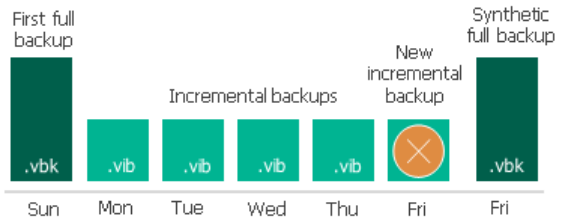
   Veeam Backup & Replication retrieves VM data for this incremental backup file from the production storage. Incremental backup helps Veeam Backup & Replication ensure that the synthetic full backup includes the latest changes of the source VM in the production environment.



2. At the end of the backup job session, Veeam Data Mover on the backup repository builds a new synthetic full backup using backup files that are already available in the backup chain, including the newly created incremental backup file.



3. When the synthetic full backup is created, Veeam Data Mover on the backup repository deletes the incremental backup file created at the beginning of the job session. As a result, you have a backup chain that consists of a full backup file, set of incremental backup files and synthetic full backup file.



4. Every next job session creates a new incremental restore point starting from the synthetic full backup until the day on which synthetic full backup is scheduled. On this day, Veeam Backup & Replication creates a new synthetic full backup.

Synthetic Full Backup Schedule

Veeam Backup & Replication automatically triggers a backup job session to create a synthetic full backup, even if a regular backup job session is not scheduled on this day. The job session is started at the same time when the parent backup job is scheduled. For example, if you schedule the parent backup job at 12:00 AM Sunday through Friday, and schedule synthetic full backup on Saturday, Veeam Backup & Replication will start a backup job session that will produce a synthetic full backup at 12:00 AM on Saturday.

If a regular backup job is scheduled together with a synthetic full backup, Veeam Backup & Replication will produce only one backup file — a synthetic full backup that will contain the latest state of the source VM. An incremental backup file that should have been created by the backup job schedule will not be added to the backup chain.

If an active full backup is scheduled together with a synthetic full backup, Veeam Backup & Replication will create only the active full backup.

Veeam Backup & Replication creates a synthetic full backup only once a day on which synthetic full backup is scheduled. If you run the backup job again on the same day, Veeam Backup & Replication will perform incremental backup in a regular manner.

Short-Term Retention Policy

Every successful backup job session creates a new restore point that lets you roll back VM data to an earlier point in time. To control the number of restore points in the backup chain, you must specify retention policy settings. The retention policy defines how many restore points you want to retain on disk and thus how 'far' you are able to roll back. After the allowed number of restore points is exceeded, Veeam Backup & Replication applies the retention policy — so it removes the earliest restore point from the backup chain.

To define the retention policy for a backup job, you must specify the necessary number of restore points or days in the **Retention policy** field in the backup job settings. By default, Veeam Backup & Replication keeps restore points for the last 7 days. In the **Storage** step of the **New Backup Job** wizard you can select the following units of retention policy.

- **Restore points**: Veeam Backup & Replication keeps the last *N* restore points, where *N* is the specified number of restore points.
- **Days**: Veeam Backup & Replication keeps restore points created during the last *N* days, where *N* is the specified number of days.

  The daily retention policy can be helpful if you periodically create off-schedule backups. For example, you create a daily backup job and you want to store retention points for 14 calendar days. If you select 14 restore points and manually create off-schedule backups, there will be several restore points in one day. So, you will have retention points for less than 14 days. In this case, you can use the **days** option.

  Mind the following for the daily retention policy:

  - The minimum number of retained restore points is 3. This number does not depend on the number of days set in the retention policy. For example, the retention policy is set to 5 days. You launch the job after it was stopped for 10 days. Normally, Veeam Backup & Replication would delete all previous restore points. However, due to the minimum number of retained restore points, you will still have at least 3 restore points: the newly created restore point and the two previous ones.

    You can change the minimum number of retained restore points with a registry key. For more information, contact Veeam Customer Support.
  - If the backup job starts at the end of the day and finishes the next day, Veeam Backup & Replication assumes that the restore point is created at the moment when the backup job started. However, Veeam Backup & Replication starts counting retention policy days only after the backup job finishes processing VMs.
  - When determining whether the number of allowed days is exceeded, Veeam Backup & Replication ignores the day when the daily retention policy is applied.

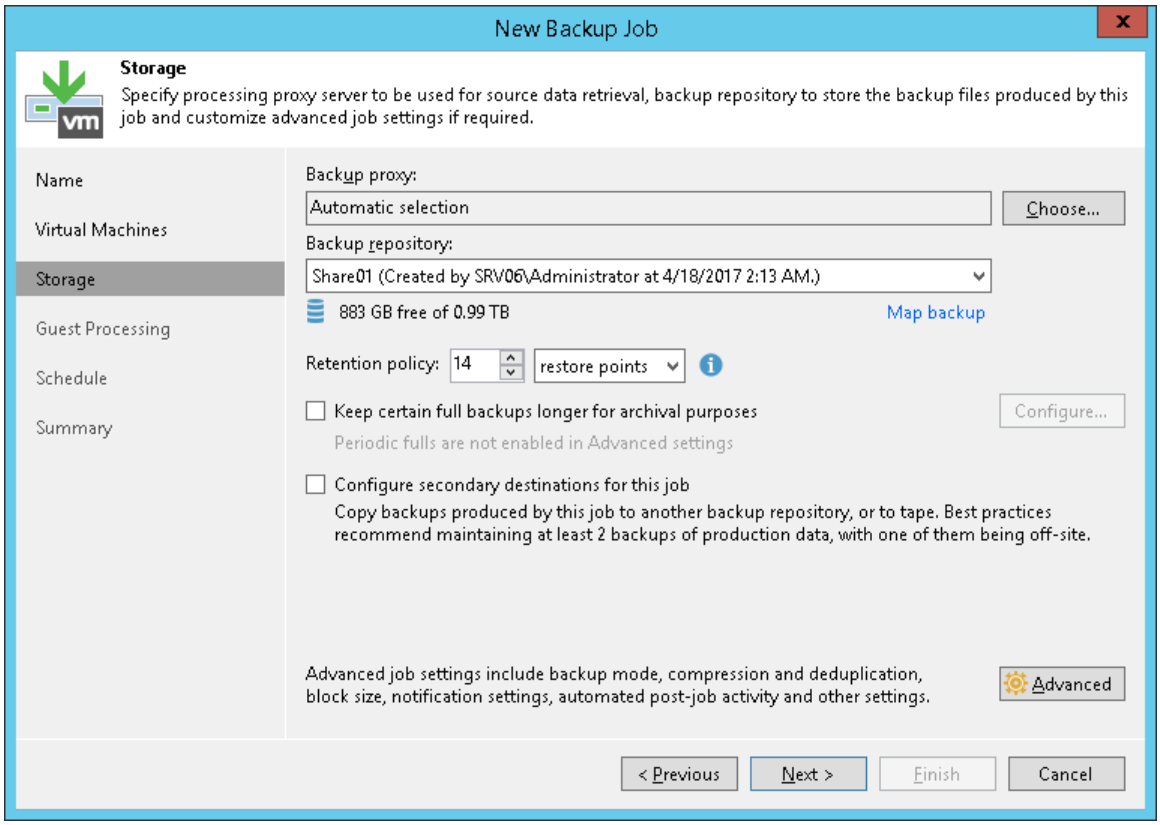
    In fact, Veeam Backup & Replication keeps restore points for the *N + 1* days, where *N* is the number of days that you specify in the job settings. In this case, Veeam Backup & Replication applies retention policy after the *N+1* days are passed. For example, if you set retention policy to keep restore points for 6 days, Veeam Backup & Replication will keep restore points for 7 days and will apply the retention policy on the 8th day. Note that the retention period may be longer depending on the specified backup method.
  - When determining whether the number of allowed days is exceeded, Veeam Backup & Replication also counts days when the backup job did not create any backups.

When the specified number is exceeded, the earliest restore points will be removed from the backup chain or merged with the next closest restore point. Veeam Backup & Replication handles restore points in different ways for forever forward incremental, forward incremental and reverse incremental backup chains:

- Forever Forward Incremental Backup Retention Policy
- Retention for Forward Incremental Backup
- Retention for Reverse Incremental Backup

NOTE

When the specified number of restore points or days is exceeded, Veeam Backup & Replication deletes the whole backup file, not separate VMs from it. For more information, see Removal of Restore Points.

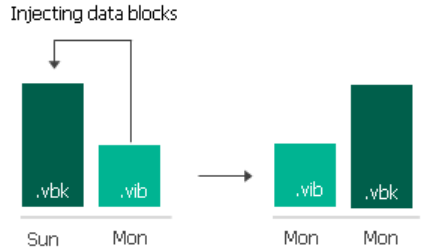


Forever Forward Incremental Backup Retention Policy

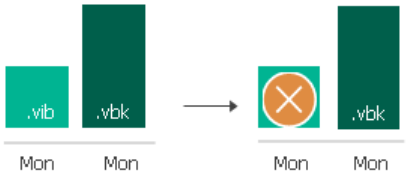
If the number of days or restore points in forever forward incremental backup chains exceeds retention policy settings, Veeam Backup & Replication transforms the backup chain to make room for the most recent restore point. The transformation process is performed in the following way:

1. Veeam Backup & Replication adds a new incremental backup file to the backup chain and detects that the number of allowed restore points or days is exceeded.
2. Veeam Backup & Replication reuses empty data blocks in the full backup file to include changes of the incremental backup file that follows the full backup. To do that, Veeam Backup & Replication injects data blocks from the first incremental backup file in the chain into the full backup file. As a result, the full backup file 'moves' one step forward in the backup chain.

> **NOTE**
>
> If the forever forward incremental backup chain resides on a deduplicating storage appliance, Veeam Backup & Replication does not reuse empty data blocks of the full backup file. Instead, Veeam Backup & Replication appends data from the first incremental backup file in the chain to the full backup file. As a result, the backup chain may consume more disk space on the appliance.

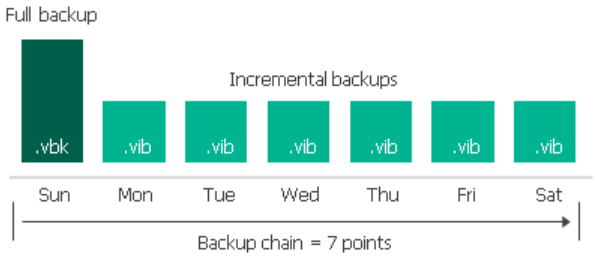


3. The first incremental backup file is removed from the backup chain as redundant. Its data has already been injected into the full backup file, and the full backup file contains the same data as this incremental backup file.



For example, you want to keep 7 restore points in the backup chain. The backup job starts on Sunday and runs daily. In this case, Veeam Backup & Replication will create the backup chain in the following way:

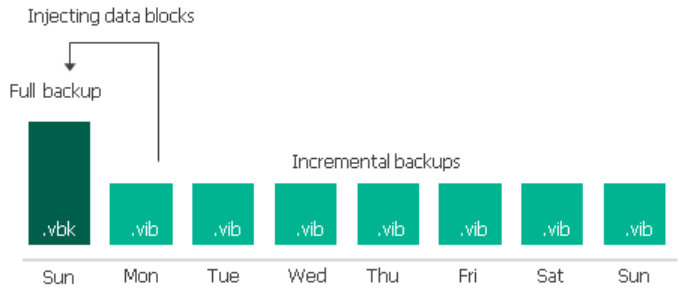
1. During the first backup job session on Sunday, Veeam Backup & Replication creates the first restore point — a full backup file.
2. Monday through Saturday Veeam Backup & Replication adds six incremental backup files to the backup chain.



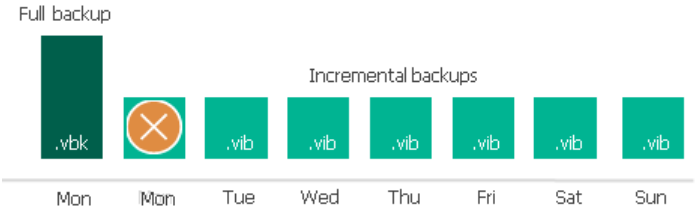
3. The next Sunday, Veeam Backup & Replication adds a new incremental backup file to the backup chain.

4. Veeam Backup & Replication detects that the number of allowed restore points is exceeded, and starts the transformation process:

   a. Veeam Backup & Replication merges data blocks from the incremental backup file created on Monday into the full backup file created on Sunday. This way, the full backup file 'moves' one step forward — from Sunday to Monday.



   b. The incremental backup created on Monday becomes redundant and is removed from the backup chain.



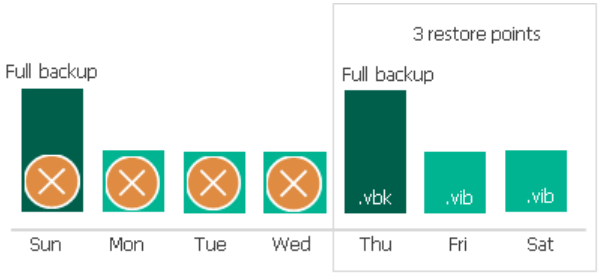
As a result, you have a chain of a full backup file as of Monday and six incremental backup files Tuesday through Sunday.

Forward Incremental Backup Retention Policy

To be able to restore from a forward incremental backup, you need to have a full backup file and a chain of subsequent incremental backup files on disk. If you delete a full backup file, the whole chain of incremental backup files will become useless. In a similar manner, if you delete any incremental backup file before the point to which you want to roll back, you won't be able to restore VM data (since later incremental backup files depend on earlier incremental backup files).

For this reason, if you select forward incremental backup method, in some days there will be more restore points on disk than specified by retention policy settings. Veeam Backup & Replication will remove the full backup chain only after the last incremental backup file in the chain becomes outdated.

For example, the retention policy is set to 3 restore points. The first full backup file is created on Sunday, incremental backup files are created Monday through Saturday, and the second full backup is created on Thursday. Although the retention policy is already breached on Wednesday, the first full backup is not deleted. Without the full backup, backup chain would be useless, leaving you without any restore point at all. Veeam Backup & Replication will wait for the second full backup file and 2 incremental backup files to be created, and only then will delete the whole previous chain, which will happen on Saturday.



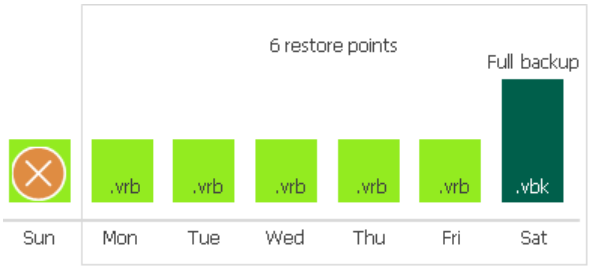
Reverse Incremental Backup Retention Policy

In case of reverse incremental backup, Veeam Backup & Replication immediately deletes the earliest reverse incremental backup file as soon as it becomes outdated.

For example, you configure a backup job in the following way:

- The backup job starts on Sunday.
- The backup method is reverse incremental.
- Retention policy is set to 6 restore points.

Veeam Backup & Replication will start the backup job on Sunday. Monday through Friday, it will add new restore points to the backup chain and rebuild the full backup file. On Saturday, Veeam Backup & Replication will add a new restore point and remove the earliest reverse incremental backup file (VRB) from the backup chain.



Retention Policy for Deleted Items

In some situations, after you configure and run backup jobs in Veeam Backup & Replication, you may want to change something in the virtual infrastructure or in the backup strategy. For example, you may remove some machines from the virtual infrastructure or move them to another location. You may also exclude some machines from jobs that have already run for some time.

Retention policy for deleted items functions differently depending on the per-machine backup file option. For details, see Per-Machine Backup Files.

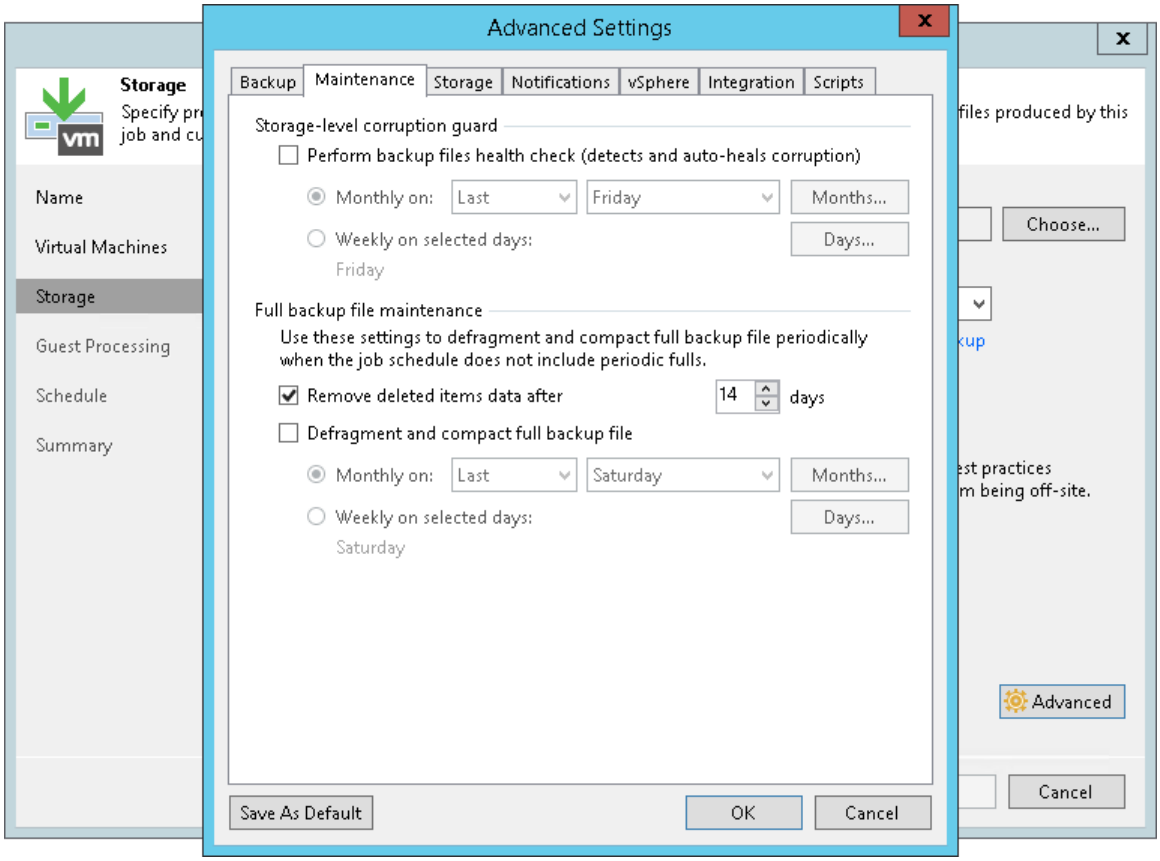
- [If per-machine is enabled] When you enable retention policy for deleted items, Veeam Backup & Replication will remove data for machines that are no longer processed by the backup job from the backup repository.
- [If per-machine is disabled] When you enable retention policy for deleted items, Veeam Backup & Replication will remove the data about deleted items from the backup job and Veeam Backup & Replication database. The stored blocks of deleted machines will remain in the repository. The stored blocks of deleted machines will be removed only when the restore point retention limit is reached or by the compact full backup file option.

Retention policy for deleted items data is set at the level of the backup job. You must enable the **Remove deleted items data after** option in backup job settings and specify the period of time for which data for deleted items must be retained in the backup repository.

Mind the following:

- You must use retention policy for deleted items data carefully. It is strongly recommended that you set the retention policy to 3 days or more to prevent unwanted data loss.

- The **Remove deleted items data after** option lets you control data of deleted or excluded items. In addition to it, Veeam Backup & Replication applies general retention policy rules to maintain the necessary number of restore points in the backup chain. For more information, see Retention Policy.



How Retention Policy for Deleted Items Works

If you enable retention policy for deleted items data in backup job settings, Veeam Backup & Replication performs the following actions:

1. If all machines in the job are processed with the *Success* status, at the end of the backup job session Veeam Backup & Replication gets a list of machines in the backup.
2. For every machine in the backup, Veeam Backup & Replication checks the configuration database and gets the date of the latest backup job session completed with the *Success* status.
3. Veeam Backup & Replication checks if any machine in the backup meets the following conditions:
   - There are no successful backups for the machine for the last N days.
   - There are no corrupted backups for the machine for the last N days.

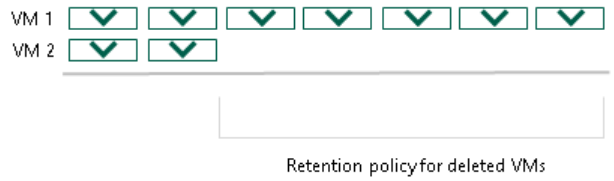
   Where N is the number of days specified in the **Remove deleted items data after N days** setting.
4. If both conditions are true for some machine, Veeam Backup & Replication removes data for this machine from the backup. Note that if per-machine is disabled, it does not free up space in the backup repository. It marks the space as available to be overwritten, and this space is removed during subsequent job sessions or the full backup file compact operation.

Example 1

You create a backup job for 2 VMs and set the retention policy for deleted items to 5 days. The backup job runs once a day for 7 times and processes VMs in the following way:

- *VM 1* is successfully processed during all job sessions.
- *VM 2* is successfully processed during the 1st and 2nd backup job sessions. Before the 3rd job session, *VM 2* is excluded from the job and is not processed by subsequent job sessions.

During the 8th job session, Veeam Backup & Replication will remove data for *VM 2* from backups in the backup repository since there are no successful and corrupted backups for *VM 2* for the last 5 days.

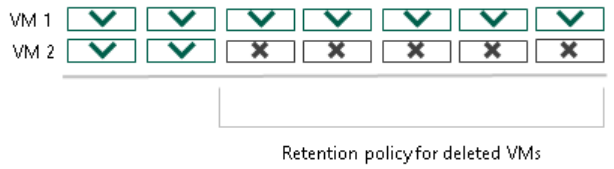


Example 2

You create a backup job for 2 machines and set the retention policy for deleted machines to 5 days. The backup job runs once a day for 7 times and processes machines in the following way:

- *VM 1* is successfully processed during all job sessions.
- *VM 2* is successfully processed during the 1st and 2nd backup job sessions. Starting from the 3rd job session, *VM 2* fails to be processed, for example, due power loss while machine data is transported.

During the 8th job session, Veeam Backup & Replication will not remove data for *VM 2* from backups in the backup repository. Even though there are no successfully created backups for *VM 2* for the last 5 days, Veeam Backup & Replication will detect that the configuration database contains information about corrupted backups for *VM 2* for the last 5 days.



Limitations for Retention Policy for Deleted Items

- [Per-machine is disabled] Retention policy for deleted items does not function if you enable synthetic full backups and/or active full backup.

  [Per-machine is enabled] Retention policy for deleted items functions without limitations.
- [For vCD backup jobs] To apply retention policy for deleted items, Veeam Backup & Replication checks backups created for the vApp itself, not for a machine in this vApp. Thus, the retention policy is applied only if the job stops creating backups for the entire vApp.

Removal of Restore Points

To keep up with the retention policy, Veeam Backup & Replication deletes the whole backup file from the backup chain, not data for separate VMs from the backup file. In some situations a certain VM may have fewer restore points than it is specified in retention policy settings. This can happen if a backup job processes a number of VMs or VM containers (hosts, clusters, folders, resource pools, VirtualApps, datastores or tags), and some VMs or VM containers fail to be processed during some job sessions.

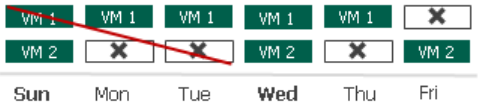
Removal of Restore Points from Forward Incremental Chains

In case of a forward incremental backup chain, Veeam Backup & Replication does not remove a restore point immediately. Instead, Veeam Backup & Replication waits for a new full backup (synthetic or active) to be created and a new backup chain to be started. As soon as the last incremental restore point in the "old" backup chain is marked as redundant, Veeam Backup & Replication removes the whole "old" backup chain from the backup repository. For more information, see Retention for Incremental Backup.

For example, a backup job processes 2 VMs: *VM 1* and *VM 2*. According to the retention policy settings, the backup chain must contain 3 restore points. The backup job has already had 5 job sessions and VMs have been processed in the following way:

- *VM 1* has been successfully backed up 3 times and has 3 restore points
- *VM 2* has failed to be processed in 2 job sessions and has 1 valid restore point

When Veeam Backup & Replication adds a new restore point to the backup chain, it will not remove the earliest restore point. Veeam Backup & Replication will wait until a new full backup file and 2 incremental backup files are added to the backup chain. After that, it will remove the whole outdated backup chain from the backup repository. Restore points in the new backup chain, at the same time, may contain data for both VMs or for one VM only: Veeam Backup & Replication regards backup files as restore points, not separate VMs in these files.

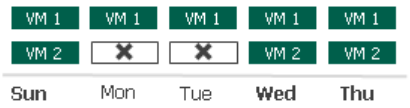


Removal of Restore Points from Reverse Incremental Chains

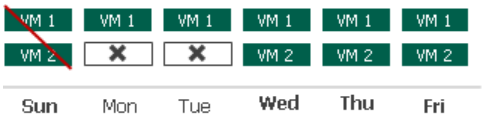
In case of a reverse incremental backup chain, Veeam Backup & Replication immediately deletes a redundant restore point when the allowed number of restore points is exceeded. For more information, see Retention for Reverse Incremental Backup.

For example, a backup job processes two VMs: *VM 1* and *VM 2*. According to the retention policy settings, the backup chain must contain 5 restore points. The backup job has already had 5 job sessions and VMs have been processed in the following way:

- *VM 1* has been successfully backed up 5 times and has 5 valid restore points
- *VM 2* has failed to be processed in 2 job sessions and has 3 valid restore points



After that, Veeam Backup & Replication runs a new backup job session in which *VM 1* and *VM 2* are successfully processed. When a new restore point is added to the chain, Veeam Backup & Replication removes the earliest restore point because the number of restore points in the backup chain has exceeded 5. As a result, you will have 5 restore points for *VM 1* and 3 restore points for *VM 2*.



Long-Term Retention Policy (GFS)

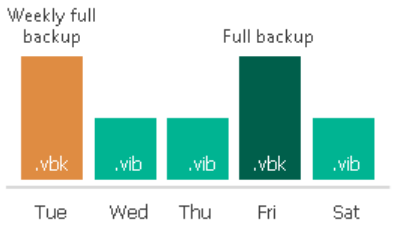
The long-term or Grandfather-Father-Son (GFS) retention policy allows you to store backup files for long periods of time — for weeks, months and even years. For this purpose, Veeam Backup & Replication does not create any special new backup files — it uses backup files created while backup job runs and marks these backups with specific GFS flags.

To mark a backup file for long-term retention, Veeam Backup & Replication can assign to the file the following types of GFS flags: weekly (W), monthly (M) and yearly (Y). The types of GFS flags that Veeam Backup & Replication assigns depend on the configured GFS retention policy settings. Depending on which flag is assigned to a backup, it will be stored for specified number of weeks, months or years.

> **IMPORTANT**
>
> GFS flags can be assigned only to full backup files created during the time period specified in GFS policy settings.

As soon as Veeam Backup & Replication assigns a GFS flag to a full backup file, this backup file can no longer be deleted or modified. Also, Veeam Backup & Replication does not apply short-term retention policy settings to the full backup file — that is, Veeam Backup & Replication ignores the backup file when determining whether the number of allowed backup files is exceeded. For more information on how GFS flags are assigned, see Assignment of GFS Flags.

When the specified retention period ends, Veeam Backup & Replication unassigns the GFS flag from the full backup file. If the backup file does not have any other GFS flags assigned, it can be modified and deleted according to the short-term retention policy. For more information on when Veeam Backup & Replication removes GFS flags, see Removal of GFS Flags.



Limitations

When planning to use GFS retention policy, mind the following limitations:

- GFS retention policy does not apply to reverse incremental backup chains.
- GFS retention policy applies to forever forward incremental backup chain only if you periodically create full backups manually or using scheduled scripts. For more information on cmdlets that you can use in scripts, see the Veeam PowerShell Reference.
- GFS retention policy does not apply to immutable backups in capacity tier.
- As Veeam Backup & Replication does not create new full backup files while applying the GFS retention policy, you must configure your backup jobs in a way you do not lose any essential data due to an insufficient number of full backup files. For example, if you configure monthly GFS retention, you need at least one full backup file per month.
- If a GFS flag is assigned to a full backup file in an active backup chain, Veeam Backup & Replication is not able to merge data from incremental backup files into the full backup file. For forever forward incremental backup chain, this means, that the short-term retention policy does not apply.

- Veeam Backup & Replication assigns GFS flags only after you save GFS retention policy settings. This means that GFS flags are assigned only to those backup files created after the configuration, while backup files created earlier are not affected and previously assigned flags are not modified.
- You cannot store full backups to which GFS flags are assigned in backup repositories with rotated drives.
- Retention policy for deleted items does not apply to full backup files to which GFS flags are assigned.

Assignment of GFS Flags

When configuring GFS retention policy settings, you can choose a number of GFS flag types that Veeam Backup & Replication will use to mark backup files for long-term retention. Depending on this number, Veeam Backup & Replication will apply slightly different algorithms when assigning GFS flags:

- Algorithm for One Flag Type
- Algorithm for Multiple Flag Types

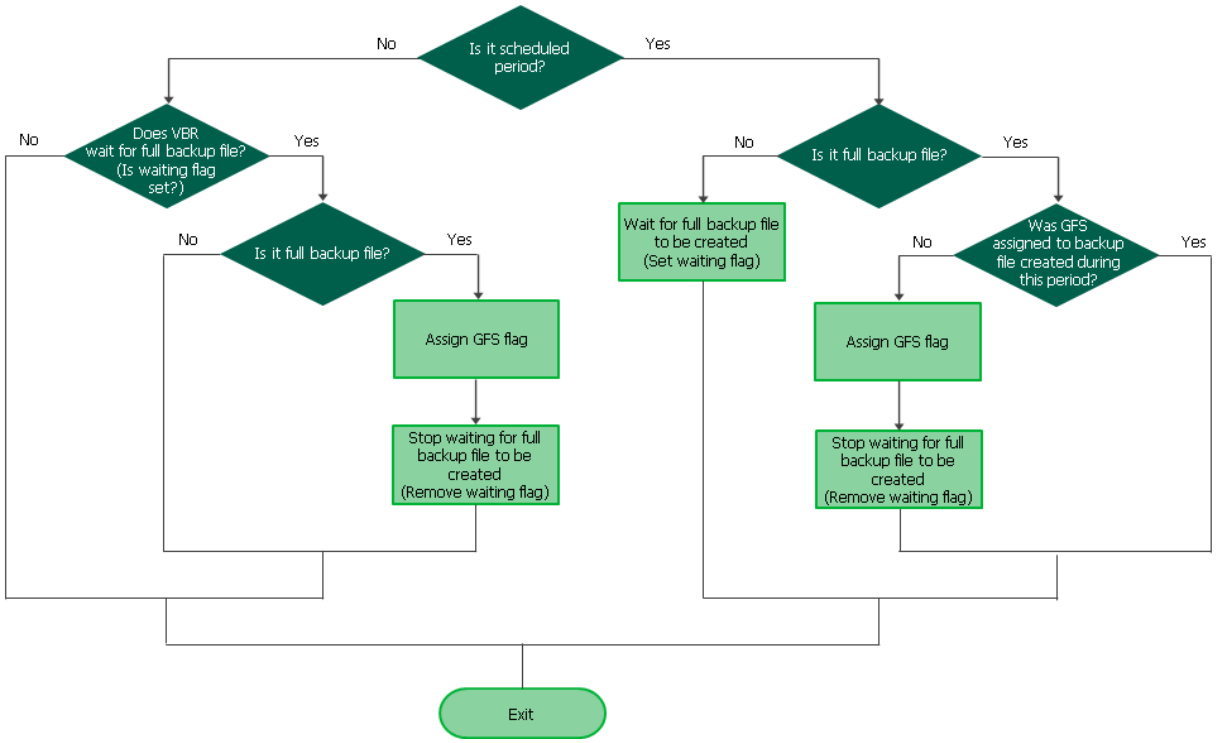
> **IMPORTANT**
>
> When you reconfigure GFS retention policy and save it, the following applies:
>
> - If you change the amount of time for which backups with GFS flags must be kept, the previously created GFS backups will also be kept according to the new settings.
> - If you change the time period when GFS flags must be assigned, the previously assigned GFS flags are not considered when determining whether new GFS flags must be assigned.

Algorithm for One Flag Type

If you select only one type of GFS flags when configuring retention policy settings for a backup job, Veeam Backup & Replication performs the following steps as soon as the job finishes. In the described algorithm, the "scheduled period" is the following: for weekly GFS — the day selected in the GFS retention policy settings; for monthly GFS — the selected week; for yearly GFS — the selected month.

- If the job finishes within the period when a new GFS flag must be assigned (the scheduled period), Veeam Backup & Replication checks whether the backup job has created a full backup file.
  - Yes, the full backup has been created: If the GFS flag has already been assigned to another backup file during the scheduled period, Veeam Backup & Replication does not assign another GFS flag to the backup file. If the GFS flag has not been assigned, Veeam Backup & Replication assigns it and stops waiting for the full backup file (removes the waiting flag).
  - No, the full backup has not been created: Veeam Backup & Replication waits for a full backup file to be created (sets the waiting flag).
- If the job finishes outside the scheduled period, Veeam Backup & Replication checks whether the full backup file is being waited for (check whether the waiting flag is set).
  - Yes, the waiting flag is set: Veeam Backup & Replication checks whether the backup job has created a full backup file.
    - Yes: Veeam Backup & Replication assigns the GFS flag and stops waiting for the full backup file (removes the waiting flag).
    - No: Veeam Backup & Replication does not assign a new GFS flag.

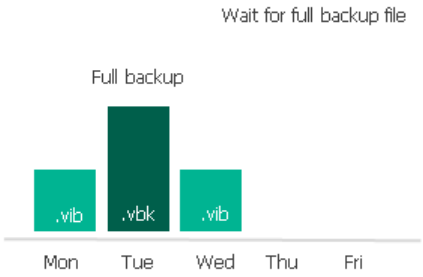
- No, the waiting flag is not set: Veeam Backup & Replication does not assign a new GFS flag.



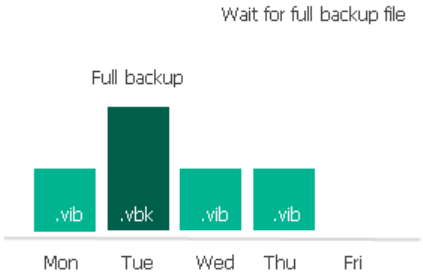
Consider the following example. On Monday, you configure GFS policy settings of a backup job in a way weekly GFS flags must be assigned every Wednesday. In this example, Veeam Backup & Replication will take the following steps.

Example Description

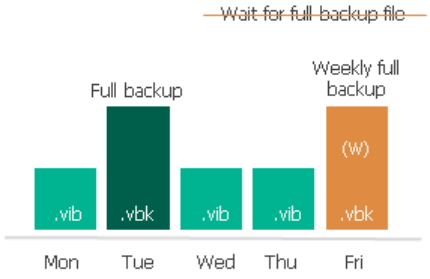
1. Until Wednesday, Veeam Backup & Replication will not assign any GFS flags because the scheduled period has not started yet.
2. On Wednesday, the backup job will produce an incremental backup file. Veeam Backup & Replication will start waiting for a full backup file to be created since the scheduled period is now started.



3. On Thursday, the backup job will produce another incremental backup file, while Veeam Backup & Replication will still be waiting for a full backup file.



4. On Friday, the backup job will produce a full backup file, and Veeam Backup & Replication will immediately assign the weekly GFS flag to the backup file.



Algorithm for Multiple Flag Types

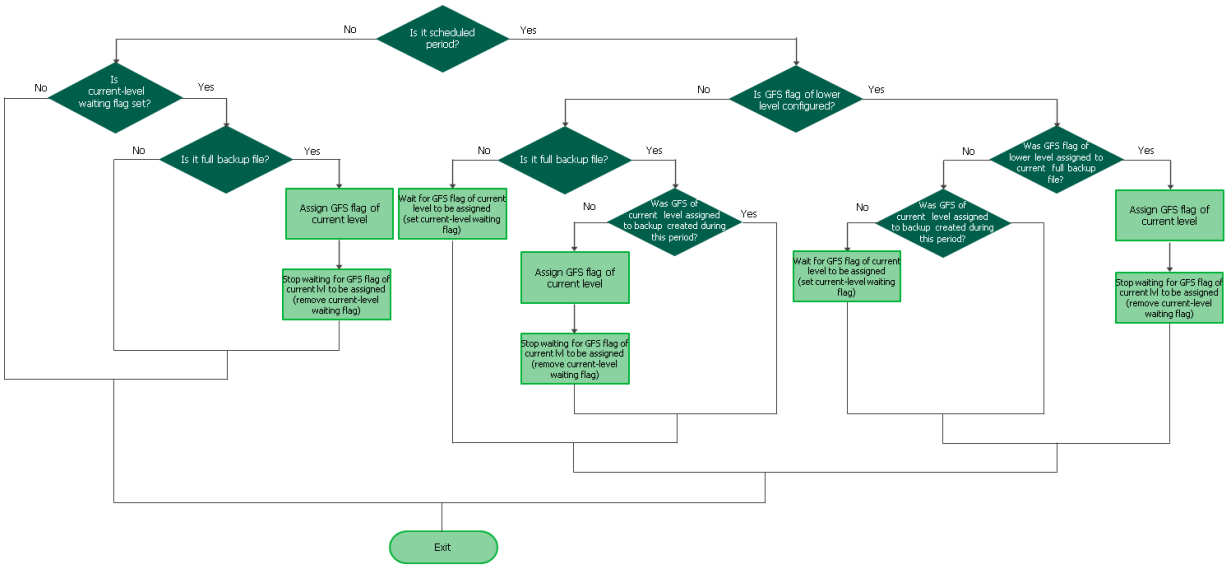
If you select multiple types of GFS flags when configuring retention policy settings for a backup job, GFS flags depend on each other. Yearly flags (high level) depend on monthly flags, monthly flags depend on weekly flags, weekly flags (low level) do not depend on any flags. This means that Veeam Backup & Replication can only assign GFS flags of a higher level to backup files with GFS flags of a lower level — this mechanism helps you save space in the backup repository.

> **IMPORTANT**
>
> The described mechanism does not apply to a situation where you select only two types of GFS flags, yearly and weekly. If you select these two types and do not select the monthly type, Veeam Backup & Replication treats the case as if you select only one flag type. This means that flags of each type are assigned according to the algorithm for one flag type.

Veeam Backup & Replication performs the following steps as soon as the job finishes and performs them for GFS flags of each level. Note that GFS flags of lower levels are processed before flags of higher levels. In the described algorithm, the "scheduled period" is the following: for weekly GFS — the day selected in the GFS retention policy settings; for monthly GFS — the selected week; for yearly GFS — the selected month.

- If the job finishes within the period when a new GFS flag must be assigned, Veeam Backup & Replication checks whether the flag of a lower level is configured in the job settings.
  - Yes, the flag of a lower level is configured: If the GFS flag of a lower level has already been assigned to the backup file created by the job, Veeam Backup & Replication assigns the GFS flag of the current level and stops waiting for the flag of the current level to be assigned (removes the current-level waiting flag). If the GFS flag of a lower level has not been assigned, Veeam Backup & Replication checks whether the GFS flag of current level has already been assigned to another backup file during the scheduled period:
    - Yes, the GFS flag of the current level has already been assigned: Veeam Backup & Replication does not assign any GFS flags.
    - No, the flag has not been assigned: Veeam Backup & Replication waits for the flag of the current level to be assigned (sets the current-level waiting flag)
  - No, the flag of a lower level is not configured: Veeam Backup & Replication checks whether the backup job has created a full backup file.
    - Yes: If the GFS flag of current level has already been assigned to another backup file during the scheduled period, Veeam Backup & Replication does not assign the GFS flag of the current level to the backup file. If the GFS flag has not been assigned, assigns it and stops waiting for the GFS flag of the current level to be assigned (removes the current-level waiting flag).
    - No: Veeam Backup & Replication waits for a full backup file to be created (sets the current-level waiting flag).
- If the job finishes outside the scheduled period, Veeam Backup & Replication checks whether the GFS flag of current level is being waited for (checks whether the current-level waiting flag is set).
  - Yes, the flag of the current level has already been assigned: Veeam Backup & Replication does not assign any GFS flags.
  - No, the flag has not been assigned: Veeam Backup & Replication checks whether the backup job has created a full backup file.
    - Yes: Veeam Backup & Replication assigns the flag of the current level and stops waiting for the GFS flag of the current level to be assigned (removes the current-level waiting flag).

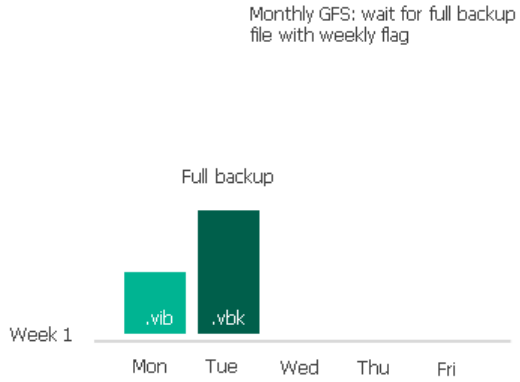
- No: Veeam Backup & Replication does not assign any GFS flags.



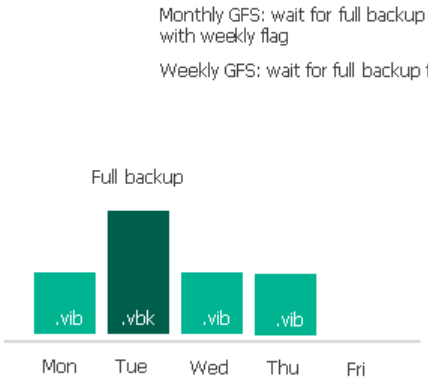
Consider the following example. On Monday, you configure GFS policy settings of a backup job in a way weekly GFS flags must be assigned every Wednesday and monthly GFS flags must be assigned every first week of a month. In this example, Veeam Backup & Replication will take the following steps.

Example Description

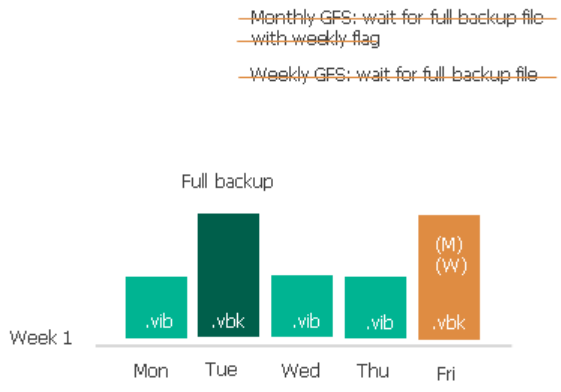
1. On Monday, the backup job will produce an incremental backup file. Veeam Backup & Replication will start waiting for a full backup file since the monthly scheduled period is now started.
2. On Tuesday, the backup job will produce a full backup file. Although Veeam Backup & Replication will still be waiting for a full backup file, the full backup file with weekly GFS flag is required.



3. On Wednesday, the backup job will produce an incremental backup file.
4. On Thursday, the backup job will produce an incremental backup file. Veeam Backup & Replication will start waiting for a full backup file to assign the weekly GFS flag to it since the weekly scheduled period is now started.



5. On Friday, the backup job will produce a full backup file. Veeam Backup & Replication will immediately assign the weekly GFS flag to the backup file. As the weekly GFS flag will have been assigned, Veeam Backup & Replication will also assign the monthly GFS flag to the backup file.



Removal of GFS Flags

When configuring GFS retention policy settings, you can specify the retention period for each type of GFS flags. After the specified retention period exceeds, Veeam Backup & Replication removes GFS flags.

The date when Veeam Backup & Replication can remove GFS flag is calculated by the following formulas:

- Weekly: `date of GFS flag assignment + N * 7 days`
- Monthly: `date of GFS flag assignment + N months + 1 day`
  When calculating the `date of GFS flag assignment + N months`, Veeam Backup & Replication increases the month ordinal number by `N`. If the calculated date does not exist, Veeam Backup & Replication uses the last date of the calculated month.
- Yearly: `date of GFS flag assignment + N years + 1 day`

Where, `N` is the value specified in the **Keep weekly/monthly/yearly full backups for** field. For more information, see Configure Long-Term Retention

> **NOTE**
>
> Mind the following:
>
> - Veeam Backup & Replication removes GFS flags only during running backup job sessions. This means that if the backup job does not run on the calculated date, Veeam Backup & Replication will remove the GFS flag later during the next job session.
> - After you change GFS retention policy settings, the date of GFS flag removal is recalculated for already created restore points.

Consider the following example. At the beginning of January, you create a backup job whose GFS retention policy settings are configured to assign monthly GFS flags. You want to keep backup files with monthly flags for 1 month and set the value of the **Keep monthly full backups for** field to 1. Veeam Backup & Replication will perform the following steps to assign and remove the flags.

1. Veeam Backup & Replication will assign the monthly GFS flag on `1/31/2019`.
2. To calculate the date when the monthly flag must be removed, the following formula is used: `date of GFS flag assignment + 1 month`. This means that the flag must be removed on `2/31/2019`. However, this date does not exist since the last date of February is `2/28/2019`. That is why Veeam Backup & Replication will remove the GFS flag on `3/1/2019` (which is `2/28/2019 + 1 day`).

Background Retention Job

Background retention job is a process that starts automatically every 24 hours and runs in the background. It consists of two parts: deleted agent retention and background GFS retention.

Deleted Agent Retention

Deleted agent retention is a process that detects and deletes outdated backup files according to Veeam Agent retention policy for outdated backups.

An outdated backup file is a backup file for which no new restore points were created and no new backup job sessions were started within the last N days, where N is the retention period specified in the Veeam Agent backup job settings.

For more information, see Retention Policy for Outdated Backups article in the Veeam Agent for Microsoft Windows User Guide and Maintenance Settings article in the Veeam Agent Management Guide.

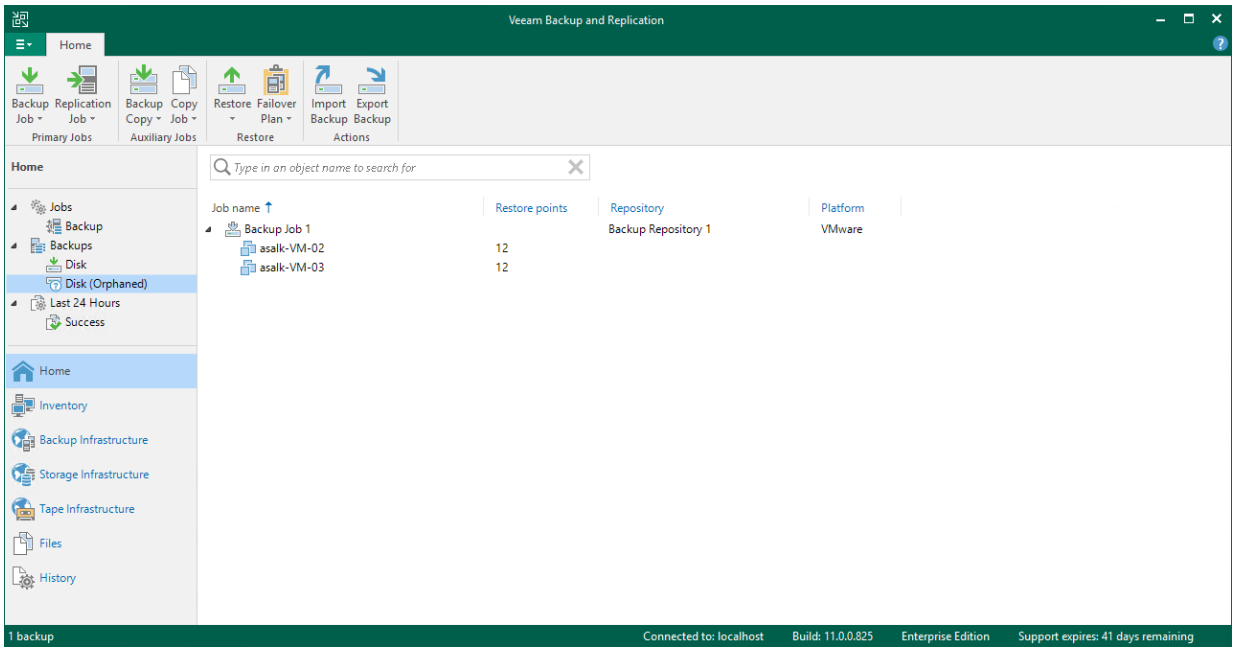
Background GFS Retention

Background GFS retention applies to standard and orphaned backup files that are no longer processed by any jobs. The orphaned backup files are shown under the **Home > Backups > Disk (Orphaned)** node.

Background GFS retention process detects backup files that are no longer processed by any jobs and analyzes their retention period. If the retention period for GFS flags of such backup files has expired, Veeam Backup & Replication removes GFS flags. Then Veeam Backup & Replication deletes these backup files according to the short-term retention policy, if the following conditions are met:

- The backup file does not have any other GFS flags assigned.
- The backup file does not have any dependent files.

For more information about GFS retention policy, see Long-Term Retention Policy (GFS).



Job and Backup File Types

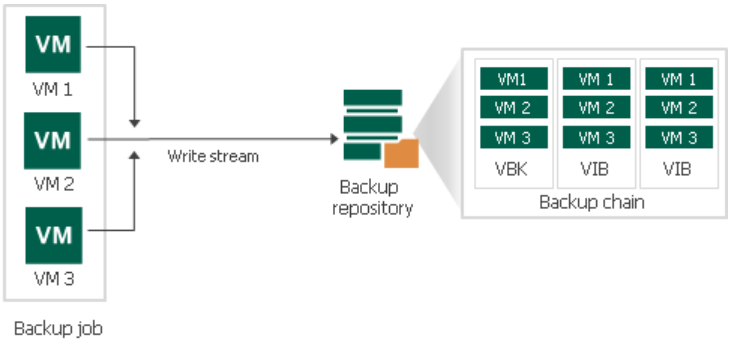
Background retention job applies to backup files created with all the types of jobs that use GFS:

- Backup
- Backup copy
- Veeam Agent
- Veeam Agent management

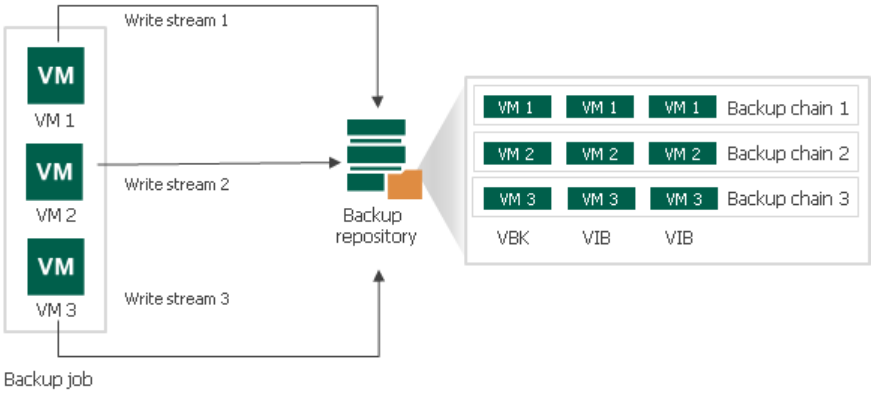
Background retention job applies both to standard and orphaned backup files. It does not apply to imported backup files.

Per-Machine Backup Files

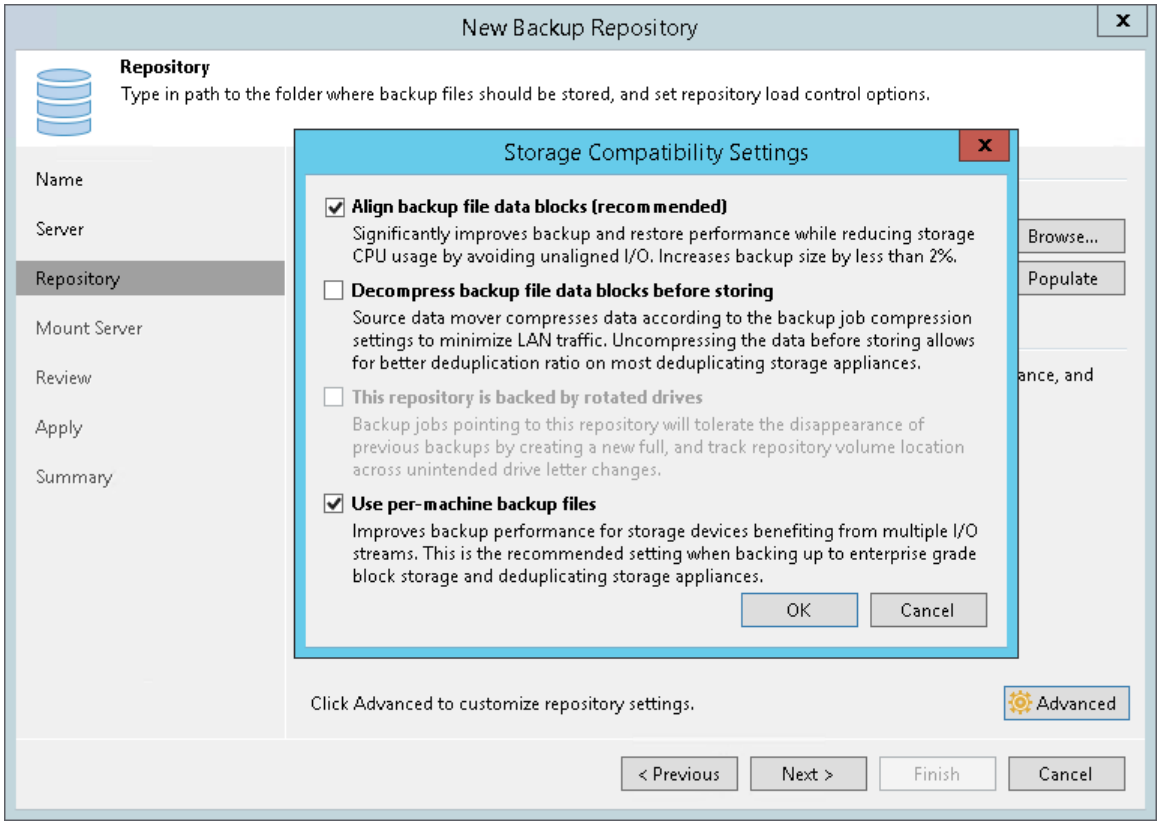
By default, backup jobs write VM data to the backup repository in one write stream, and store data of all VMs to the same backup file. Such behavior can be non-optimal if the target storage device is able to write data in multiple streams simultaneously. In this situation, the backup repository may become the bottleneck for the data transfer, even though its resources will not be fully utilized.



You can instruct Veeam Backup & Replication to create per-machine backup files in the backup repository. In this case, the backup job will use a separate write stream for every VM in the job, and store data of every VM to a separate backup file. Resources of the storage device will be used more efficiently, and the job performance may increase.



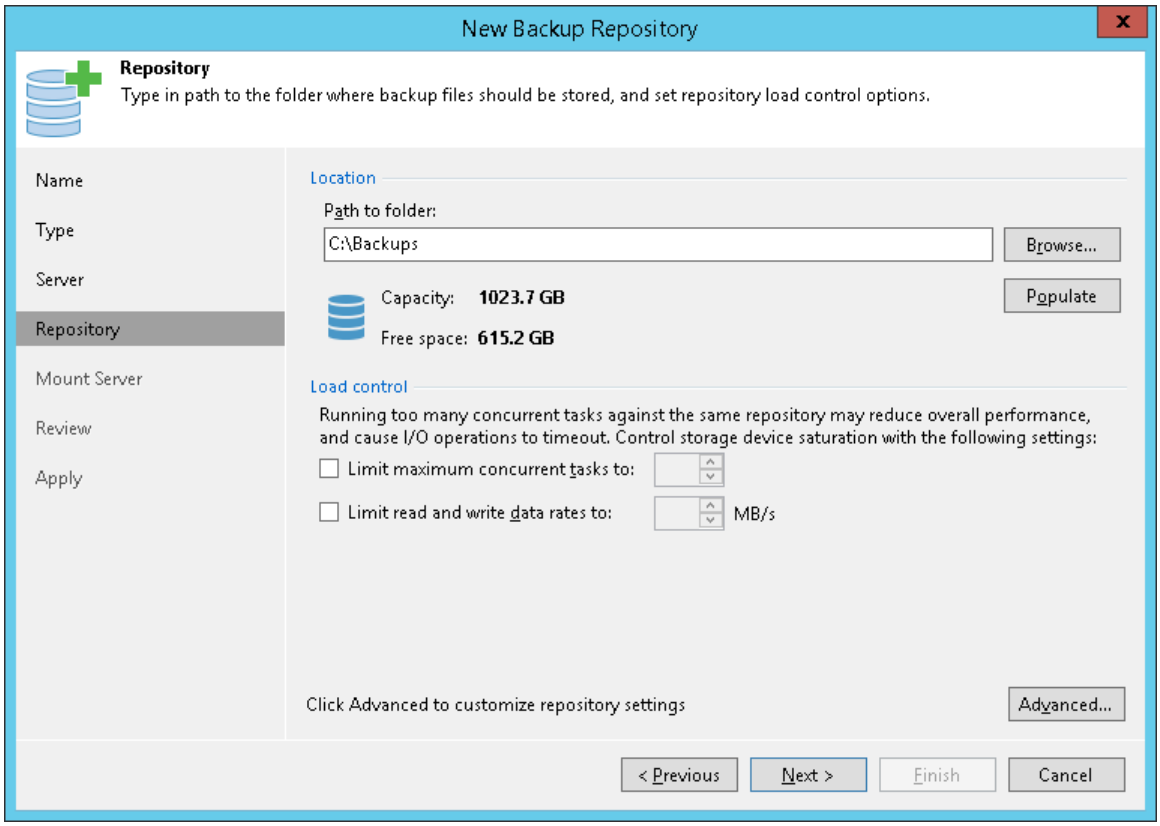
To create per-machine backup files, you must enable the **Use per-machine backup files** option at the level of the backup repository. It is recommended that you enable this option for deduplicating storage appliances that support multiple write streams. The option is also enabled for scale-out backup repositories by default.



It is recommended that you balance the number of tasks on backup proxies and backup repository to avoid the situation where some backup infrastructure resources remain idle while others are overloaded.

NOTE

It is not recommended that you disable the **Limit maximum concurrent tasks to N** option for backup repositories with per-machine backup chains. In case of per-machine backup chains, synthetic operations (synthetic full backup, backup files merge and transformation) work in parallel for every VM in the backup. The number of parallel operations is limited by the number of concurrent tasks that can be performed on the backup repository. If you disable the **Limit maximum concurrent tasks to N** option (which results in using an unlimited number of slots), the load on the backup repository may be high.



Limitations for Per-Machine Backup Files

When planning to use per-machine backup files, mind the following limitations:

- The **Use per-machine backup files** option cannot be enabled for backup repositories with rotated drives.
- If you enable the **Use per-machine backup files** option, data deduplication between VMs will not work. For more information, see Data Compression and Deduplication.
- The per-machine backup chains functionality is available in Veeam Backup & Replication Enterprise edition and higher. If you configure backup repositories to produce per-machine backup chains and then install a license that does not support this functionality, you must manually disable the **Use per-machine backup files** option for backup repositories. Otherwise backup jobs targeted at these backup repositories will be failing.
- [For Veeam Agent backup jobs] Veeam Backup & Replication ignores the **Use per-machine backup files** option. The way how Veeam Backup & Replication creates backup files depends on the objects included in the backup job:
  - If several Veeam Agent machines are included in the backup job, Veeam Backup & Replication creates a separate backup file for each machine.

- If failover clusters are included in the backup job, Veeam Backup & Replication creates a separate backup file for each cluster.

Per-Machine Backup Files Option for Existing Backup Repositories

You can enable or disable the **Use per-machine backup files** option for existing backup repositories at which backup jobs are already targeted. The new setting will not have any effect on previously created backup files in the backup repository. It will affect new backup files created after the setting is changed.

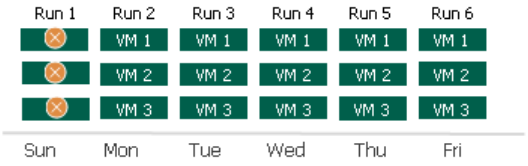
Veeam Backup & Replication applies the new setting starting from the next active full backup. You can create an active full backup manually or wait for Veeam Backup & Replication to automatically create active full backup (if active full backups are scheduled). Synthetic full backups do not affect the **Use per-machine backup files** setting.

Retention Policy for Per-Machine Backup Files

If you enable the **Use per-machine backup files** option for the backup repository, Veeam Backup & Replication creates a separate backup chain for every VM added to the job. Backup files are stored together in the folder of the backup job in the backup repository. The job produces one metadata file. This file stores information about all created backup files and backup chains.

Veeam Backup & Replication regards all backup files that are created during one backup job run as one restore point. When Veeam Backup & Replication needs to remove earlier restore points by retention policy, it removes backup files for all VMs that were created during one job run.

For example, you have added 3 VMs to the job, set the retention setting to 5 restore points and run the job 5 times. The job will produce 15 backup files, 5 per each VM in the job. On the 6th job run, the job will remove from the backup chain 3 backup files — the earliest restore points for every VM.



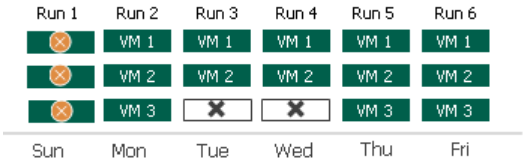
If the job backs up some VMs during the job run and does not manage to back up others, Veeam Backup & Replication will still regard that the restore point is valid. When the earliest restore point gets outdated, Veeam Backup & Replication will remove backup files for all VMs at once, even though backup chains for some VMs may contain fewer backup files than you expect.

The rules of restore points deletion for regular backup chains also apply to per-machine backup chains. For more information, see Removing Restore Points from the Backup Chain.

For example, you have added 3 VMs to the job and set retention policy to 5. The backup job worked in the following way:

1. During the first two job runs, Veeam Backup & Replication backed up all VMs.
2. During the 3rd and 4th job runs, *VM 1* and *VM 2* were successfully backed up, and *VM 3* failed.

During the 6th job run, Veeam Backup & Replication will delete the earliest restore point for all VMs. As a result, the *VM 1* and *VM 2* will have 5 restore points and *VM3* will have 3 restore points.



Changed Block Tracking

To perform incremental backup, Veeam Backup & Replication needs to know what data blocks have changed since the previous job session.

For VMware VMs with hardware version 7 and later, Veeam Backup & Replication employs a native VMware vSphere feature — VMware vSphere Changed Block Tracking (CBT). Instead of scanning VMware VMFS (Virtual Machine File System), Veeam Backup & Replication queries CBT through VMware VADP (VMware vStorage API for Data Protection) and gets a list of blocks that have changed since the last job session. Use of CBT increases the speed and efficiency of block-level incremental backups.



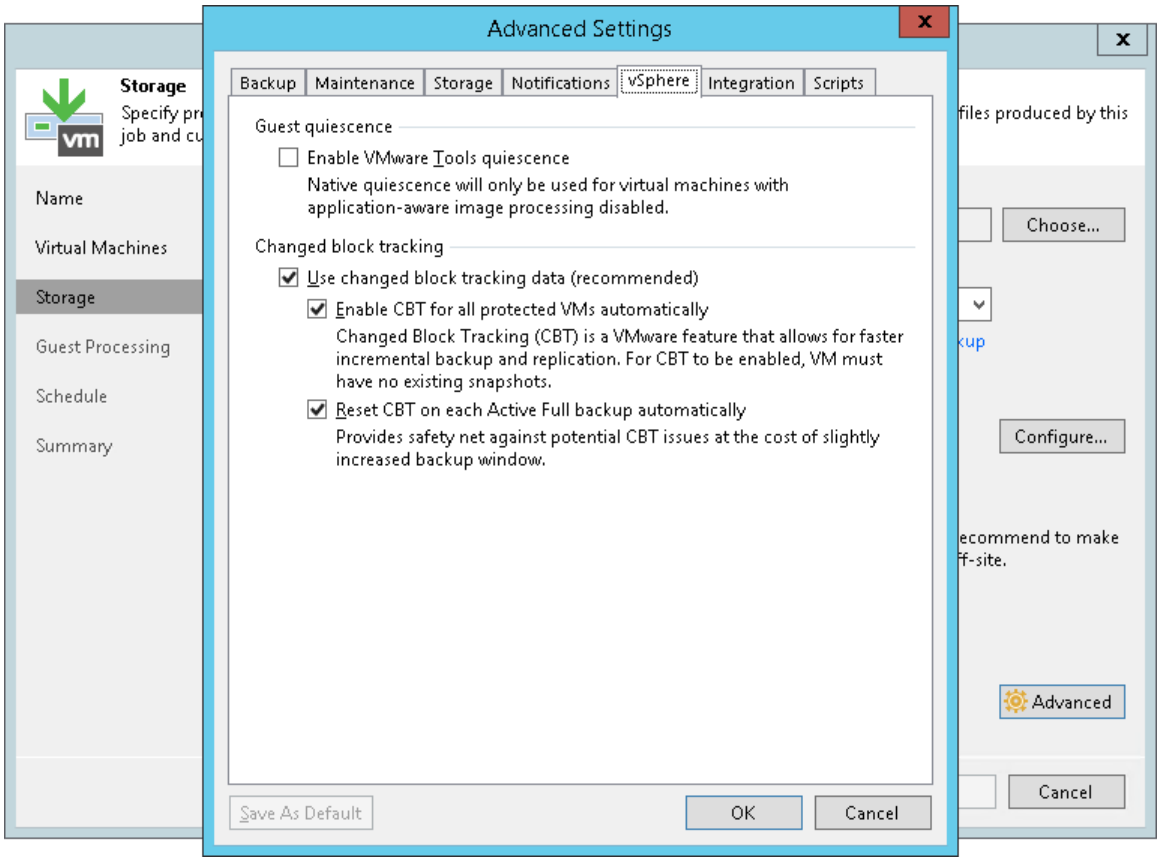
Veeam Backup & Replication uses CBT for the following operations:

- Backup
- Replication
- Entire VM restore
- VM disk restore

Veeam Backup & Replication enables CBT by default. If necessary, you can disable it in job settings.

NOTE

If you back up a VM to which the backup proxy role is assigned and that uses Virtual appliance (HotAdd) transport mode, CBT for this VM is disabled and cannot be enabled.



NOTE

For VMs with virtual disks in thin format, Veeam Backup & Replication also uses CBT during active full backup sessions to detect unallocated regions of virtual disks and skip them. For VMs with virtual disks on NFS datastore, Veeam Backup & Replication uses CBT as well, but cannot leverage CBT on first full run (for more information, see this KB article).

In some situations, Veeam Backup & Replication cannot leverage VMware vSphere CBT, for example, if VMs run an earlier version of virtual hardware. If Veeam Backup & Replication cannot leverage VMware vSphere CBT, it fails over to the Veeam proprietary filtering mechanism. Instead of tracking changed blocks of data, Veeam Backup & Replication filters out unchanged data blocks.

During VM processing, Veeam Backup & Replication consolidates virtual disk content, scans through the VM image and calculates a checksum for every data block. Checksums are stored as metadata to backup files next to VM data. When incremental backup is run, Veeam Backup & Replication opens all backup files in the chain of previous full and incremental backups, reads metadata from these files and compares it with checksums calculated for a VM in its current state. If a match is found (which means the block already exists in the backup), the corresponding block is filtered out.

Data Compression and Deduplication

Veeam Backup & Replication provides mechanisms of data compression and deduplication. Data compression and deduplication let you decrease traffic going over the network and disk space required for storing backup files and VM replica files.

Data Compression

Data compression decreases the size of created files but affects duration of the backup or replication procedure. Veeam Backup & Replication allows you to select one of the following compression levels:

- **None** compression level is recommended if you plan to store backup files and VM replica files on storage devices that support hardware compression and deduplication.
- **Dedupe-friendly** is an optimized compression level for very low CPU usage. Select this compression level if you want to decrease the load on the backup proxy.
- **Optimal** is the recommended compression level. It provides the best ratio between size of the file and time of the procedure.
- **High** compression level provides additional 10% compression ratio over the **Optimal** level at the cost of about 10x higher CPU usage.
- **Extreme** compression provides the smallest size of the file but reduces the performance. We recommend that you run backup proxies on computers with modern multi-core CPUs (6 cores recommended) if you intend to use the extreme compression level.
- [For backup copy job] **Auto** is the recommended compression level for backup copy jobs. Select this level to use compression settings of the copied backup files.

> **NOTE**
>
> If encryption is enabled for a job and the **Decompress backup data blocks before storing** check box is selected in the settings of the target backup repository, Veeam Backup & Replication does not compress VM data. Therefore, in the job statistics, you may observe a higher amount of transferred data (the **Transferred** counter) as compared to a job for which encryption is disabled. For details on job statistics, see Viewing Real-Time Statistics.

Changing Data Compression Settings

You can change data compression settings for existing jobs. After you change the settings, you will not need to create new full backups to use new settings. Veeam Backup & Replication will automatically apply the new compression level to the newly created backup files after you save the settings. Previously created backup files will not be affected.

However, if you use the reverse incremental backup method, the newly created backup files will contain a mixture of data blocks compressed at different levels. For example, you have a backup job that uses the reverse incremental backup method and the Optimal level of compression. After several job sessions, you change the compression level to High. In the reverse incremental backup chains, the full backup file is rebuilt with every job session to include new data blocks. As a result, the full backup file will contain a mixture of data blocks: data blocks compressed at the Optimal level and data blocks compressed at the High level. The same behavior applies to synthetic full backups: synthetic full backups created after the compression level change will contain a mixture of data blocks compressed at different levels.

If you want the newly created backup file to contain data blocks compressed at one level, you can create an active full backup. Veeam Backup & Replication will retrieve data for the whole VM image from the production infrastructure and compress it at the new compression level. All subsequent backup files in the backup chain will also use the new compression level.

Deduplication

Data deduplication decreases the size of files. With data deduplication enabled, Veeam Backup & Replication does not store to the resulting file identical data blocks and space that has been pre-allocated but not used.

We recommend you enable data deduplication if your backup jobs contain several VMs that have a lot of free space on their logical disks or VMs that have similar data blocks — for example, VMs that were created from the same template. However, note that data deduplication may decrease job performance.

Veeam Backup & Replication uses Veeam Data Movers to deduplicate VM data:

- Veeam Data Mover in the source side deduplicates VM data at the level of VM disks. Before the source-side Veeam Data Mover starts processing a VM disk, it obtains digests for the previous restore point in the backup chain from Veeam Data Mover in the target side. The source-side Veeam Data Mover consolidates this information with CBT information from the hypervisor and filters VM disk data based on it. If some data block exists in the previous restore point for this VM, the source-side Veeam Data Mover does not transport this data block to the target. In addition to it, in case of thin disks the source-side Veeam Data Mover skips unallocated space.
- Veeam Data Mover in the target side deduplicates VM data at the level of the backup file. It processes data for all VM disks of all VMs in the job. The target-side Veeam Data Mover uses digests to detect identical data blocks in transported data, and stores only unique data blocks to the resulting backup file.

You can change data deduplication settings for existing jobs. After you change the settings, you will not need to create new full backups to enable/disable the deduplication. Veeam Backup & Replication will automatically apply the change to the newly created backup files after you save the settings. Previously created backup files will not be affected.

Storage Optimization

To optimize job performance and storage usage, Veeam Backup & Replication allows you to choose the minimum data block size to process VMs. The optimal data block size depends on the type of storage you select as a target and the size of your files.

When selecting the data block size, consider the following aspects:

- When reading the VM image, Veeam Backup & Replication "splits" the VM image into blocks of the selected size. The more data blocks there are, the more time is required to process the VM image.
- [For replication and vCD replication] Veeam Backup & Replication writes information about every data block to the VM replica metadata stored in the backup repository. The more data blocks there are, the more metadata is written to the backup repository.
- [For changed block tracking enabled] During incremental job runs, Veeam Backup & Replication uses CBT to define changed data blocks in the VM. The larger the size of the found changed data block, the greater the amount of data that is transferred to the target site.

The incorrectly chosen data block size may decrease the performance. For example, when you deduplicate a large backup file to small data blocks, Veeam Backup & Replication produces a very large deduplication metadata table, which can potentially overgrow memory and CPU resources of your backup repository. For large backup files, it is better to use large data blocks.

Veeam Backup & Replication provides several storage optimization options with different block sizes used. The following table will help you to choose the optimal option in accordance with the size of your backup files and the storage type.

Storage optimization option	Block size	Description
Local target (large blocks)	4096 KB	Recommended for files that are larger than 16 TB. This option will provide the lowest deduplication ratio and the largest size of incremental files.
Local target	1024 KB	Recommended for backup and replication to SAN, DAS or local storage. This option provides the fastest job performance but reduces the deduplication ratio, because with larger data blocks it is less likely to find identical blocks.
LAN target	512 KB	Recommended for backup and replication to NAS, and on-site backup and replication. This option provides a better deduplication ratio and reduces the size of a file because of reduced data block sizes.
WAN target	256 KB	Recommended if you are planning to use WAN for off-site backup and replication. This option provides the maximum deduplication ratio and the smallest size of files that allows you to reduce the amount of traffic over WAN.

Changing Storage Optimization Settings

You can change storage optimization settings for existing jobs. New settings will not have any effect on previously created files in the chain. They will be applied to new files created after the settings were changed.

For Veeam Backup & Replication to apply the new settings, use the following instructions.

Backup Jobs

To apply new storage optimization settings in backup jobs, you must create an active full backup after you change storage optimization settings. Veeam Backup & Replication will use the new block size for the active full backup and subsequent backup files in the backup chain.

Backup Copy Jobs

To change data block size for backup copy jobs, you must perform the following actions:

1. Change data block size in settings of the initial backup job.
2. Create an active full backup with the initial backup job.
3. Create an active full backup with the backup copy job.

Data Exclusion

When you configure a backup or replication job, you can define what data you want to back up and replicate and exclude data that you do not need. Data exclusion helps reduce the size of the VM backup or replica and decrease the load on the network.

You can exclude data at the VM level and at the VM guest OS level.

At the VM level:

- VMs added as part of the container
- VM disks
- VM templates (only for backup)

At the VM guest OS level:

- Swap files on the VM guest OS
- Deleted file blocks on the VM guest OS (BitLooker)
- Files and folders on the VM guest OS

> **NOTE**
>
> To reduce the size of the backup file, Veeam Backup & Replication automatically excludes VM log files from processing.

VMs and VM Disks

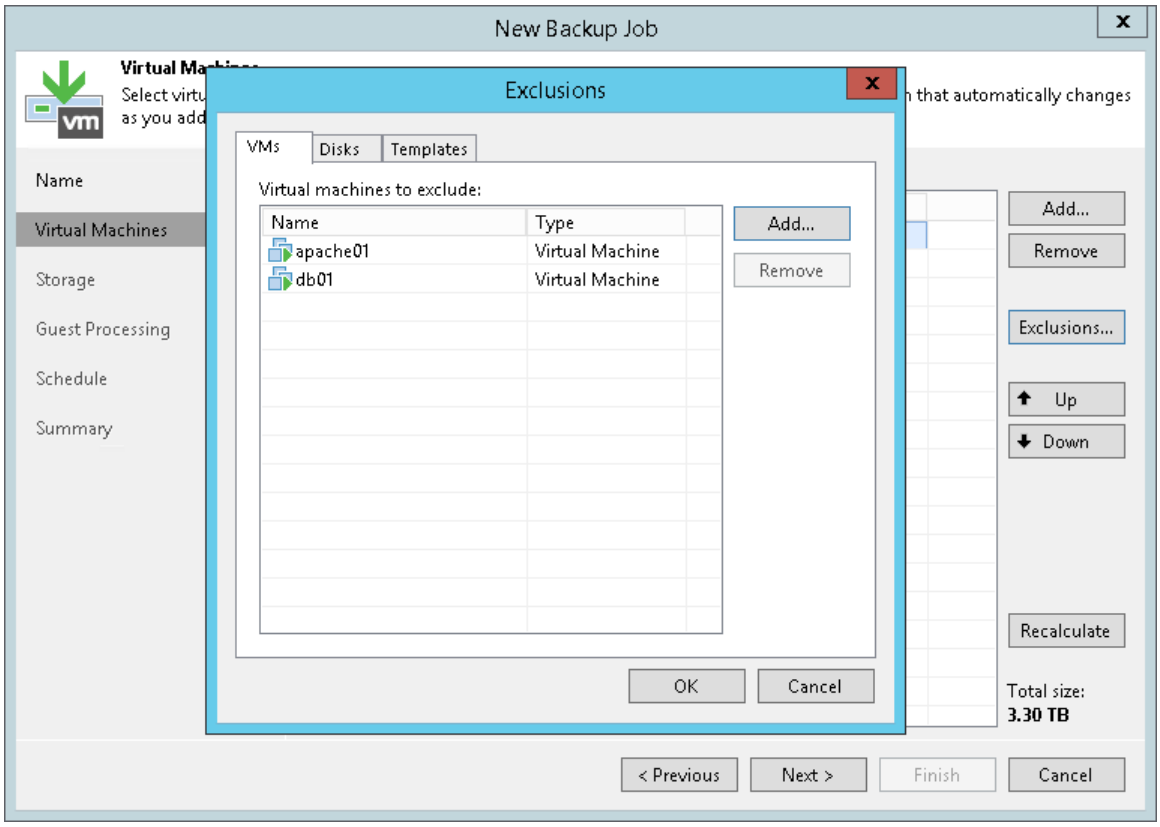
When you configure a backup or replication job, you can exclude the following objects from processing:

- VMs added as a part of a VM container
- Individual VM disks
- [For backup jobs] VM templates

VMs as Part of Container

If you want to back up or replicate a VM container that holds several VMs but want to skip some VMs, you can exclude specific VMs from the job processing. This option will help you reduce the size of the resulting backup or replica and increase the job performance.

You can define which VMs you want to skip at the **Virtual Machines** step of the backup or replication job wizard.



Individual VM Disks

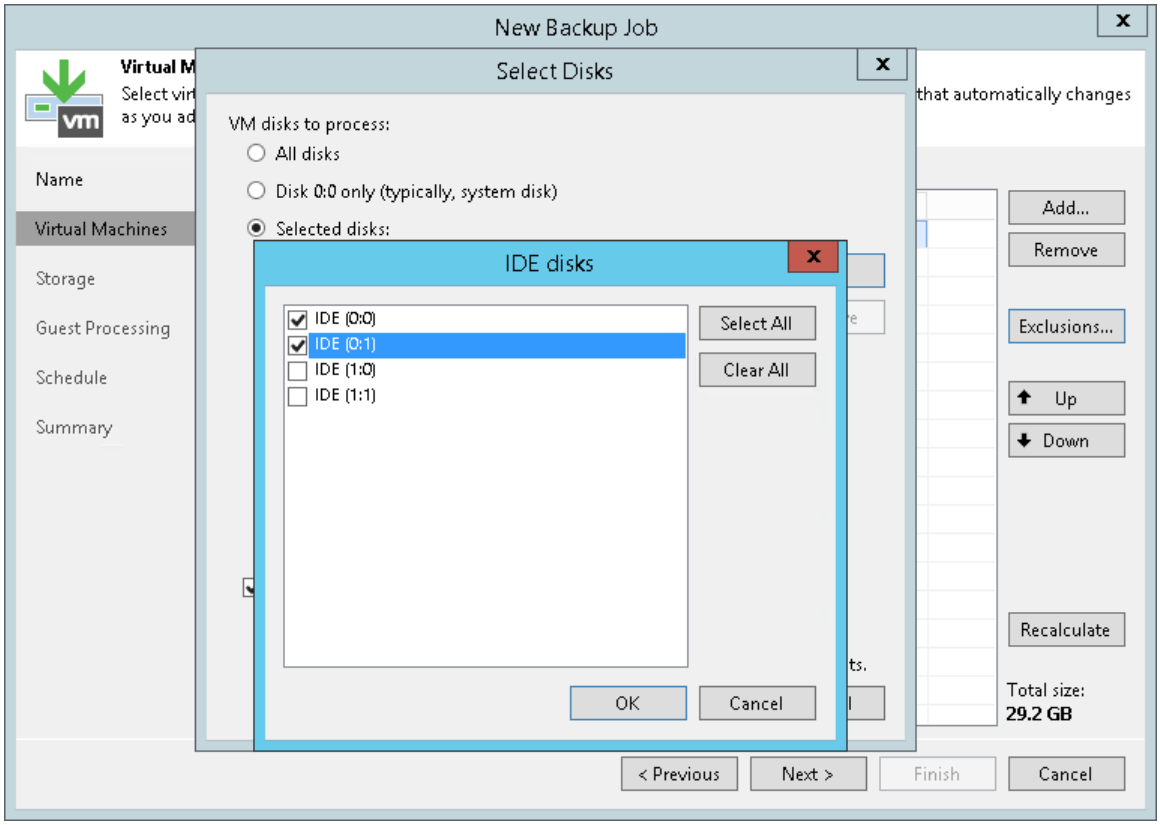
You can choose what VM disks you want to back up or replicate:

- All VM disks
- 0:0 disks (which are commonly the VM system disks)
- Specific IDE, SCSI or SATA disks

For example, you may want to back up or replicate only the system disk instead of creating a backup or replica of a full VM. VM disks exclusion reduces the size of the backup or replica.

You can define which VM disks you want to back up or replicate at the **Virtual Machines** step of the backup or replication job wizard. You can specify disk processing settings granularly for every VM in the job or for the whole VM container. In the latter case, Veeam Backup & Replication will apply the configured rule to all VMs in this container.

You can additionally instruct Veeam Backup & Replication to modify the configuration file of the VM. When you start a VM from the backup or fail over to the VM replica, you will be able to use such VM immediately. You will not have to edit its configuration file and remove excluded disks from it.

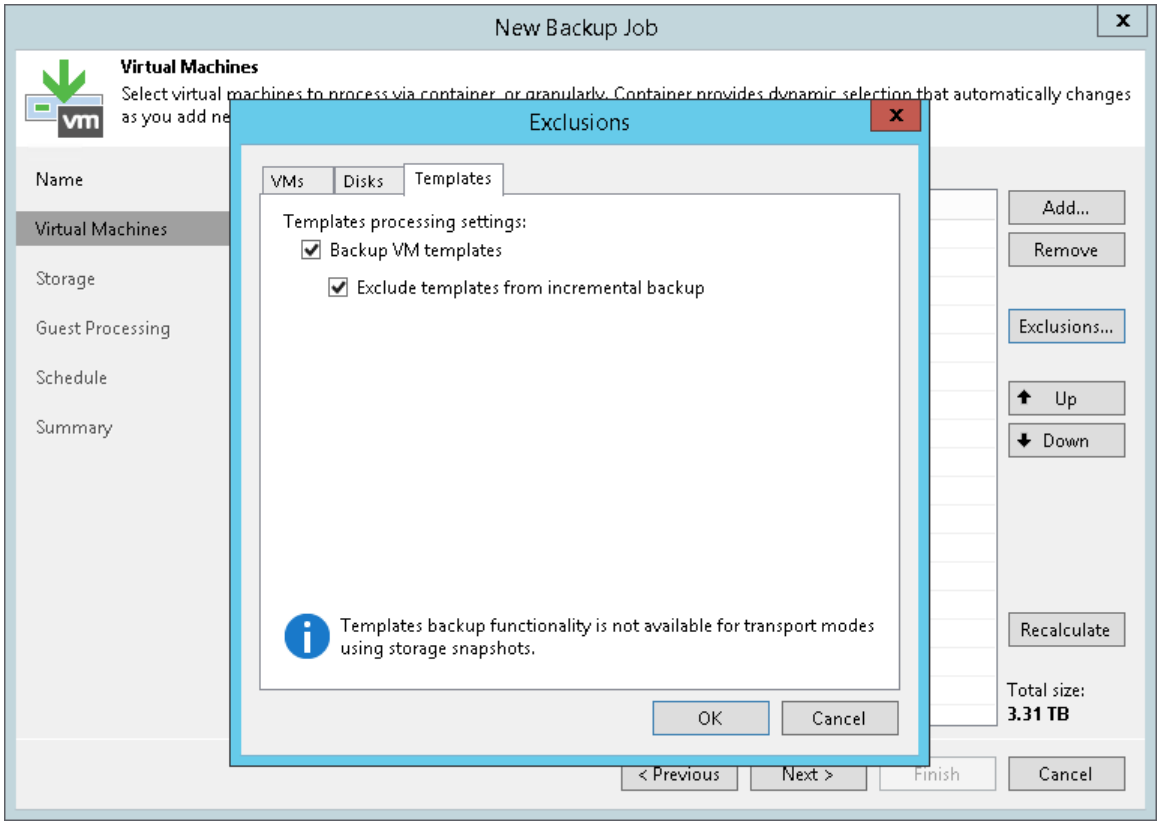


VM Templates

You can include VM templates in the backup. Backing up VM templates provides additional safety of your production environment but requires additional space in the backup repository.

Veeam Backup & Replication allows you to include a VM template only in the full backup and omit it in incremental backups. Note that Veeam Backup & Replication cannot use Direct SAN transport mode to back up VM templates data. If you have configured failing over to the Network mode, the Network mode will be used.

You can define how Veeam Backup & Replication must process VM templates at the **Virtual Machines** steps of the wizard.



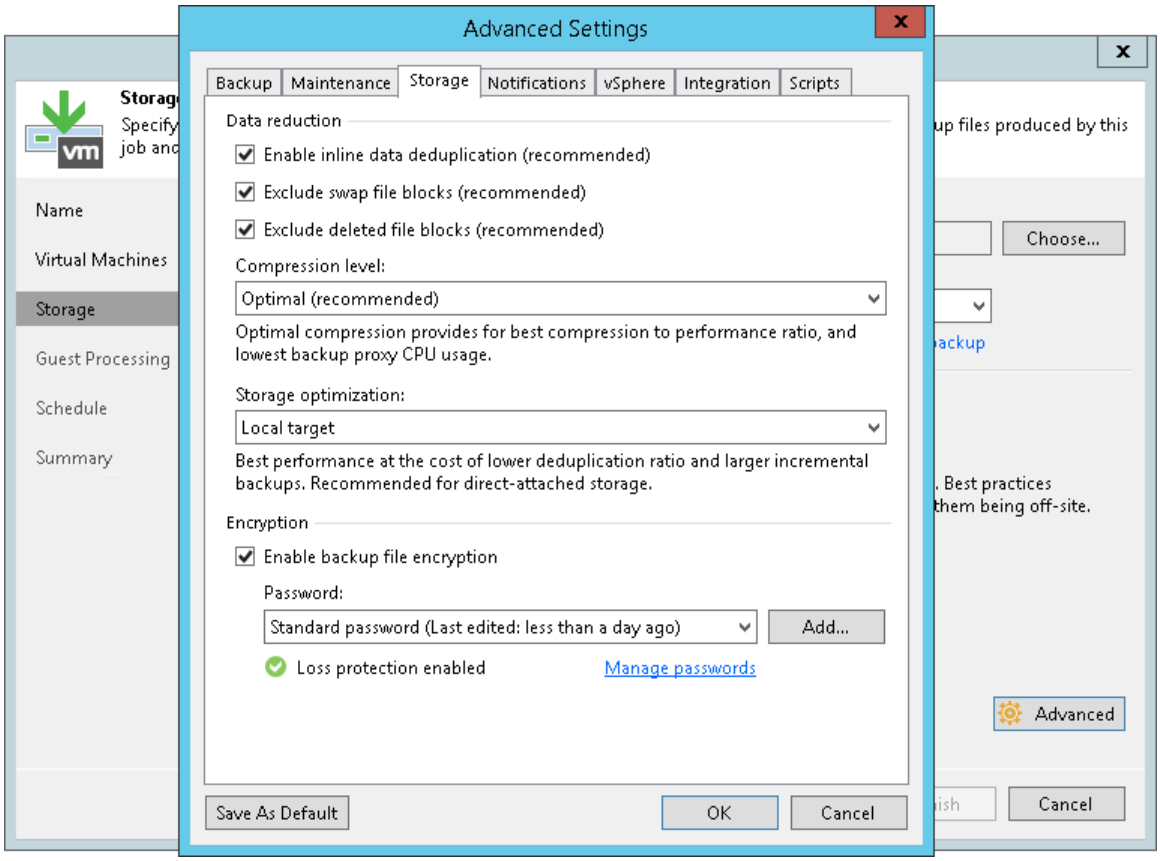
Deleted File Blocks (BitLooker)

By default, Veeam Backup & Replication does not copy "dirty" data blocks to the target location. Dirty data blocks are blocks marked as deleted after you remove them from the recycle bin. Omitting these blocks lets you reduce the size of the VM backup or replica and increase the job performance.

If you do not want to exclude deleted file blocks from backups or replicas, you can disable the **Exclude deleted file blocks** option in the backup or replication job settings.

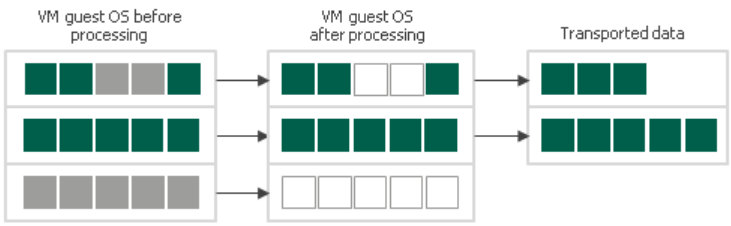
NOTE

If you enable or disable the **Exclude deleted file blocks** setting for the existing job, Veeam Backup & Replication will apply the new setting from the next job session.



With this option enabled, Veeam Backup & Replication performs the following operations during the job session:

1. Veeam Backup & Replication accesses the MFT file on the VM guest OS to identify deleted file blocks, and zeros out these blocks.
2. Veeam Backup & Replication processes and transports data blocks of the VM image in the following manner:
   - If a data block of the VM image contains only the deleted file blocks, Veeam Backup & Replication does not read this data block from the source datastore.
   - If a data block of the VM image contains zeroed out blocks and other data, Veeam Backup & Replication copies this block to the target. Due to data compression, data blocks that are marked as deleted are compressed, and the size of the resulting backup or replica file reduces.



Limitations for Deleted File Blocks Exclusion

Veeam Backup & Replication can exclude deleted file blocks only on the VM guest OS with Microsoft NTFS.

Swap Files

You can instruct Veeam Backup & Replication to exclude pagefile.sys and hiberfil.sys files from backups or replicas of Microsoft Windows VMs. As a result, the size of incremental backups and replicas will be smaller.

- hiberfil.sys is a system file created by the OS for correct work of the hibernate mode.
- pagefile.sys is a swap file. Swap files are dynamic in nature and can change intensively between job sessions, even if a VM itself does not change much.

How Swap File Exclusion Works

When you exclude pagefile.sys and hiberfil.sys files, Veeam Backup & Replication performs the following operations during the job session:

1. Veeam Backup & Replication accesses the MFT file on the VM guest OS to identify data blocks of pagefile.sys and hiberfil.sys files and zeros them out.
2. Veeam Backup & Replication processes and transports data blocks of the VM image in the following manner:
    - If a data block of the VM image contains only blocks of these files, Veeam Backup & Replication does not copy this data block to the target.
    - If a data block of the VM image contains blocks of these files and other data, Veeam Backup & Replication copies this block to the target.



Limitations for Swap File Exclusion

Veeam Backup & Replication can exclude blocks of pagefile.sys and hiberfil.sys files only on the VM guest OS with Microsoft Windows NTFS.

VM Guest OS Files

If you do not want to back up or replicate some files and folders on the VM guest OS, you can exclude them from the backup or replica. Files exclusion reduces the size of the backup or replica but may affect the job performance.

You can specify file exclusion settings granularly for every VM in the job or for the whole VM container. In the latter case, Veeam Backup & Replication will apply the configured rule to all VMs in this container.

To define which VM guest OS files must and must not be processed, you can use the following options:

- Disable file exclusion.

  Veeam Backup & Replication will back up or replicate the whole content of the VM guest file system.

- Exclude specific files and folders from the backup or replica.

  Veeam Backup & Replication will back up or replicate all files and folders except the specified ones.

- Include only specific files and folders in the backup or replica.

  Veeam Backup & Replication will back up or replicate only the specified files and folders.

Defining File Exclusions and Inclusions

To define a list of exclusions or inclusions, you can use the following methods:

- Specify a full path to a folder on the VM guest OS, for example, `C:\Documents\`.
- Specify a full path to a file on the VM guest OS, for example: `C:\Documents\MyReport.docx`.

  If a path is not full, Veeam Backup & Replication will expand it relatively the root directory on the computer volume and attempt to detect such files on all computer volumes. For example, you have `C`, `D` and `E` disks on the VM. In the list of exclusions, you specify `Document.docx`. Veeam Backup & Replication will scan the whole file system and exclude the following files (if any): `C:\Document.docx`, `D:\Document.docx`, `E:\Document.docx`. If there is a `C:\MyDocuments\Document.docx` file, it will not be excluded — this file is not located in the root directory.

- Use environmental variables, for example, `%TEMP%`, `%windir%`.

  Environment variables must be defined for the user account that you use to connect to the VM guest OS and under which the non-persistent runtime components or persistent agent components are started. For example, you connect to the VM guest OS under the *Administrator* account. If you want to use the `%windir%` variable in the list of exclusions or inclusions, you must make sure that the `%windir%` variable is added to the list of user variables for *Administrator* on the VM guest OS.

- Use file masks. You can use the following characters for masks:
  - (*) — a substitution for one or more characters in the file name or path. Can be used for any sequence of characters (including no characters). For example, `*.pdf`.
  - (?) — a substitution of one character in the file name or path. For example, `repor?.pdf`
  - (;) — mask separator, for example, `report.*;reports.*`.

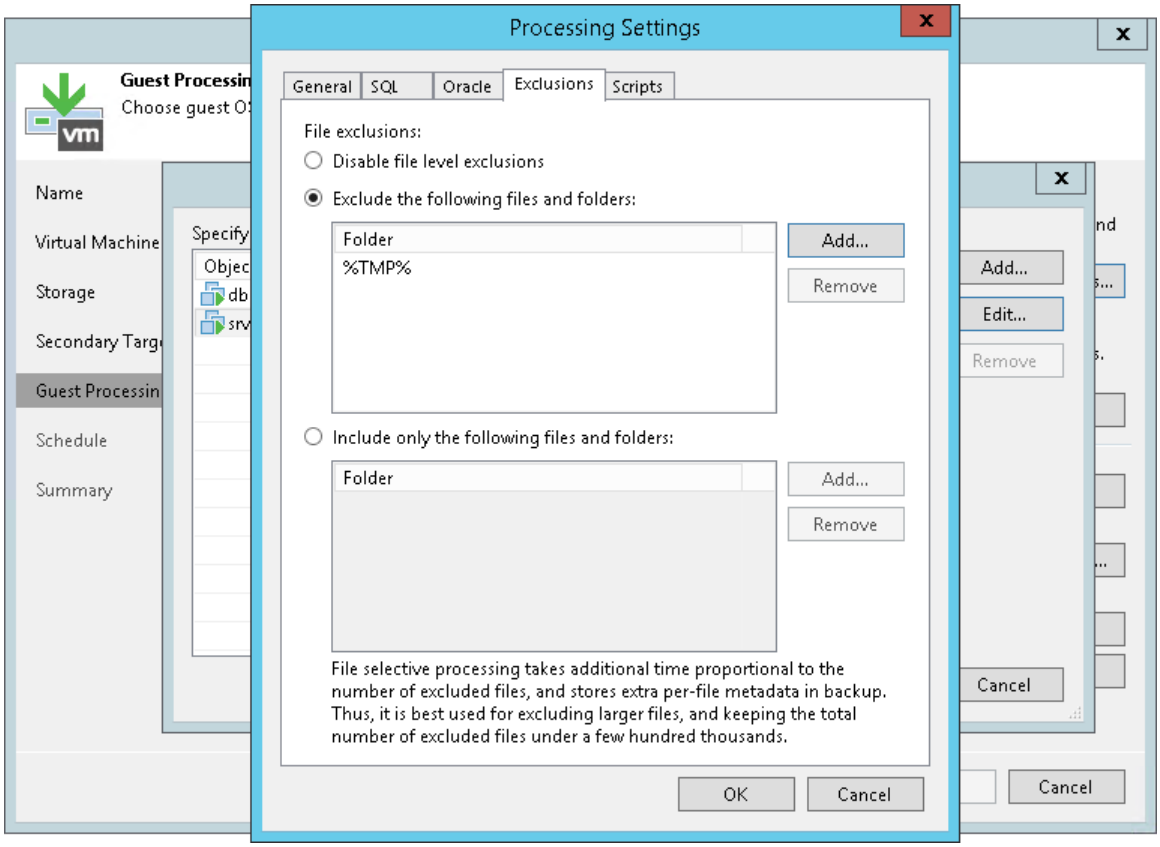
  In the table below, mask stands for any sequence of characters.

Mask format	Affects paths/files
`*mask*`	All paths that contain the given sequence.
`mask*`	If the asterisk character (*) is not specified at the beginning of the mask, the mask will be applied to all volumes on the VM guest OS, and Veeam Backup & Replication will include/exclude files and folders in the root folder on the volume: `A:\mask*`, `B:\mask*`, …, `Z:\mask*`.

Mask format	Affects paths/files
`<drive_letter>:\*mask*`	All paths on the specified volume that contain the given sequence.
`*mask1*;*mask2*;*mask3*`	All paths that contain at least one of the given character sequences: `*mask1*` or `*mask2*` or `*mask3*`.

IMPORTANT

Be careful when using masks with double wildcard characters. If you specify masks of such type, Veeam Backup & Replication will exclude all files and paths that contain the given mask. For example, if you specify the `*.doc*` mask, Veeam Backup & Replication will exclude files like `MyReport.docx`, `Report.doc.txt` and so on.



Requirements and Limitations for VM Guest OS File Exclusion

VM guest OS files exclusion has the following limitations:

- File exclusion works only on Microsoft Windows NTFS.
- File exclusion is included in the Veeam Universal License. When using a legacy socket-based license, Enterprise or higher edition is required.

- To exclude VM guest OS files, Veeam Backup & Replication must be able to deploy the non-persistent runtime components or use (in necessary, deploy) persistent agent components inside the VM. For this reason, the VM must be running and accessible by an IP address or through the Installer Service on VM, and credentials for application-aware processing must be valid.
- Veeam Backup & Replication supports both basic and dynamic disks. For the dynamic disks, simple type of volumes is supported. Spanned, mirrored and striped volumes are not supported.
- It is not recommended that you use VM guest files exclusion in Microsoft Windows for volumes with enabled Data Deduplication. If you decide to use VM guest files exclusion for such volumes and set up a list of inclusions, you must add the System Volume Information folder to the list of inclusions.

Consider the following:

- If you use file masks for file exclusion, Veeam Backup & Replication will scan the VM guest file system, and thus the time of VM disk processing will increase.
- The number of entries in the list of exclusions or inclusions must not exceed a few hundreds. The number of entries in the list influences the job performance — the more files are included or excluded from the backup or replica, the more time Veeam Backup & Replication requires to process these files.
- It is recommended that you do not exclude system files without the necessity.
  Veeam Backup & Replication does not perform any checks to verify the VM image integrity.
- Exclusion of small files (less than 2 KB in size) is ineffective and will not reduce the size of the backup or replica significantly.

Exclusion Rules for VMs with Several Volumes

The VM guest file exclusion and inclusion functionality works at the volume level. Consider the following situations:

Data exclusion

A VM has several volumes: `C:\`, `D:\` and `E:\`. You want to exclude from the backup the `Archive` folder that is present on all volumes of the VM. If you add the `C:\Archive` folder to the list of exclusions,
Veeam Backup & Replication will back up the following data:

- Whole content of the `C:\` volume except the `Archive` folder
- Whole content of `D:\` and `E:\` volumes

To exclude the `Archive` folder from all volumes of the VM, you must add a relative path to the `Archive` folder the list of exclusions: `..\Archive\`.

Data inclusion

A VM has several volumes: `C:\`, `D:\` and `E:\`. You want to include to the backup only the `D:\Documents` folder. If you add the `D:\Documents` folder to the list of inclusions, Veeam Backup & Replication will back up the following data:

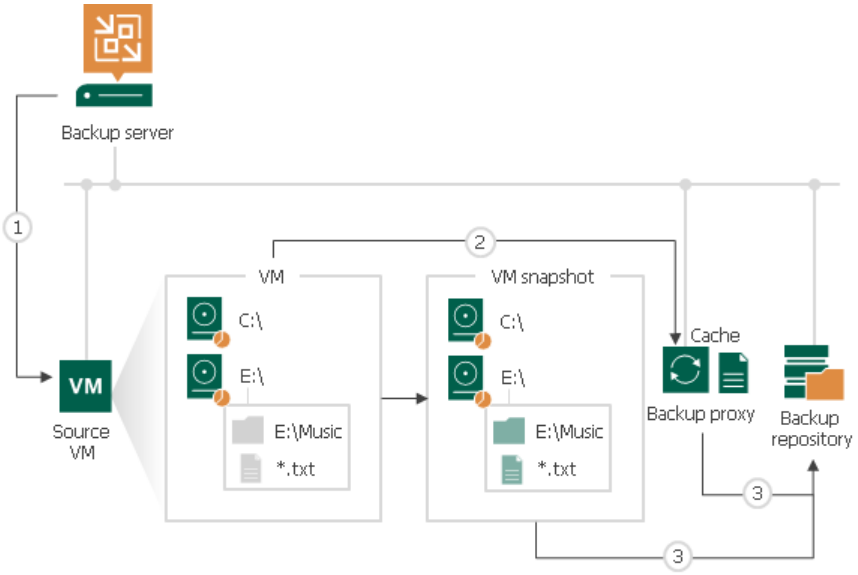
- `D:\Documents` folder
- Whole content of `C:\` and `E:\` volumes

To include only the `D:\Documents` folder to the backup, you must add the `D:\Documents` folder to the list of inclusions and, additionally, exclude unnecessary disks (that contain `C:\` and `E:\` volumes) at the **Virtual Machines** step of the wizard. For more information, see Exclude Objects from Backup Job .

How VM Guest OS File Exclusion Works

When you exclude VM guest OS files from the backup or replica, Veeam Backup & Replication performs the following operations:

1. Veeam Backup & Replication checks the job settings to identify what VM guest OS files must be excluded.
2. Veeam Backup & Replication opens the MFT file from the VM guest file system in the memory cache on the backup proxy, and marks data blocks of excluded files as deleted.
3. When Veeam Backup & Replication copies VM data to the target, it reads data both from the VM snapshot and memory cache on the backup proxy. On the target, Veeam Backup & Replication creates a "merged" version of VM disks that do not contain excluded VM guest OS files. Due to data compression, data blocks that are marked as deleted are compressed, and the size of the resulting backup or replica file reduces.



During the job session with file exclude, Veeam Backup & Replication makes changes to processed VM disks at the NTFS level using the cache on the backup proxy. However, these changes are not visible to the CBT mechanism. For this reason, Veeam Backup & Replication saves information about excluded data blocks in the backup file and replica metadata. During the next job session with use of CBT, Veeam Backup & Replication retrieves a list of data blocks that were excluded during the previous job session from the backup file or replica metadata and analyzes what data needs to be processed during the current job session. To do this, Veeam Backup & Replication regards the following data:

- Data blocks that are marked as new with CBT
- Data blocks that were excluded during the previous job session
- Data blocks that must be excluded during the current job session

VMware Tools Quiescence

When you back up or replicate a running VM, you need to quiesce or 'freeze' the VM to bring its file system and application data to a consistent state. Data quiescence is crucial for highly-transactional applications. It helps create transactionally consistent backups or replicas and guarantee safety of application data.

To create consistent backups and replicas for VMs that do not support Microsoft VSS (for example, Linux VMs), enable VMware Tools quiescence in job settings. In this case, Veeam Backup & Replication uses the VMware Tools to freeze the file system and application data on a VM before backup or replication.

VMware Tools also allow creating backups and replicas for Microsoft Windows-based VMs that support Microsoft VSS. For this, VMware Tools use VMware VSS component. For details on the supported OSes and quiescence features, see the VMware documentation.

However, to bring applications to a transactionally consistent state before freezing, you must use a pre-freeze script. To bring applications to the initial state after freezing, use a post-freeze script. Store these scripts on the backup server. When a backup or replica job starts, Veeam Backup & Replication uploads these scripts to the appropriate folders on the VM guest OS. For more information, see Pre-Freeze and Post-Thaw Scripts.

By default, VMware Tools quiescence is disabled. You can enable it at the job level. Note that VMware Tools quiescence will be enabled for all VMs added to the job. For more information on how to enable VMware Tools quiescence, see vSphere Settings.

Combined Approach

It is recommended to enable VMware Tools quiescence and application-aware processing when you add Microsoft Windows and Linux VMs to the same job. In this case, all VMs will be processed in a transactionally consistent manner — either with application-aware processing or VMware Tools quiescence.

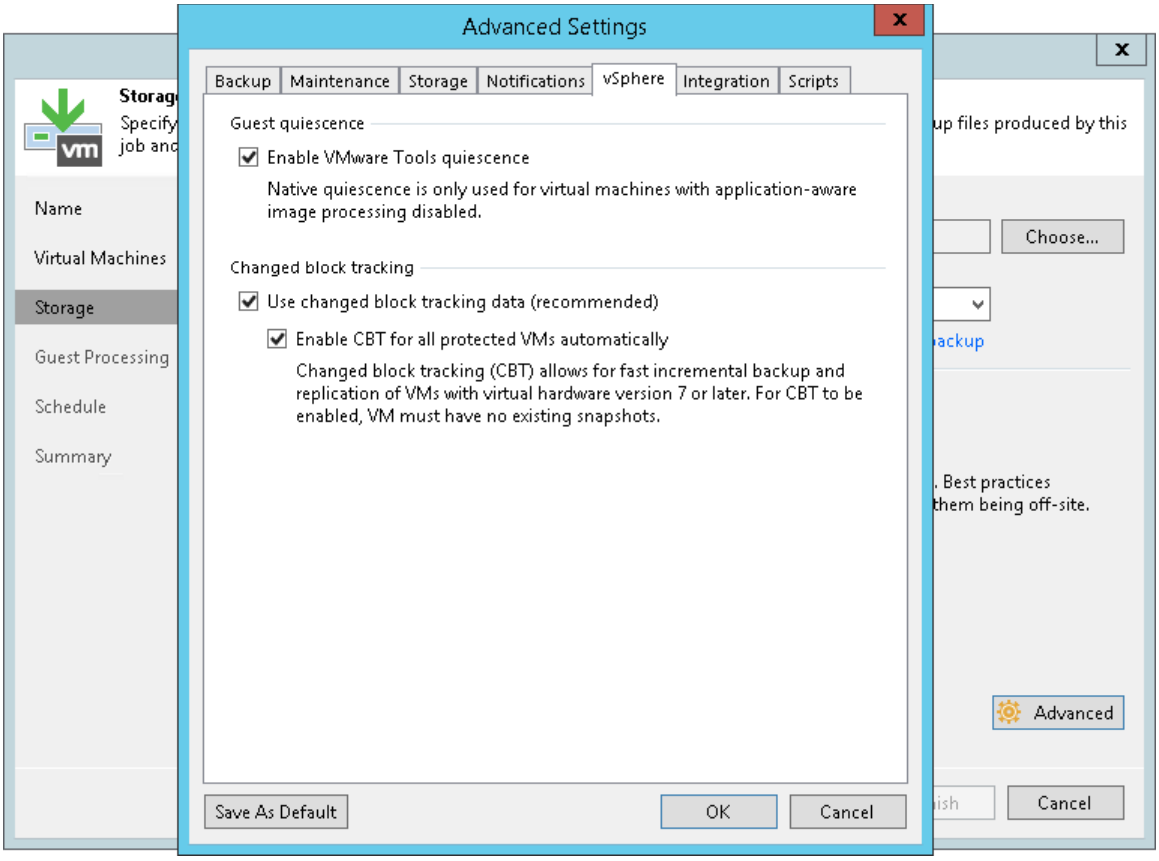
In such a scenario, Veeam Backup & Replication processes VMs in the job in the following way:

1. Veeam Backup & Replication first attempts to use application-aware processing to prepare VMs for backup or replication. If Veeam Backup & Replication manages to quiesce all VMs in the job with application-aware processing, it does not use VMware Tools quiescence.
2. If some VMs cannot be quiesced with application-aware processing or application-aware processing is disabled for some VMs in the job (the **Disable application processing** is set for VMs in the job settings), Veeam Backup & Replication uses VMware Tools quiescence to prepare these VMs for backup or replication.

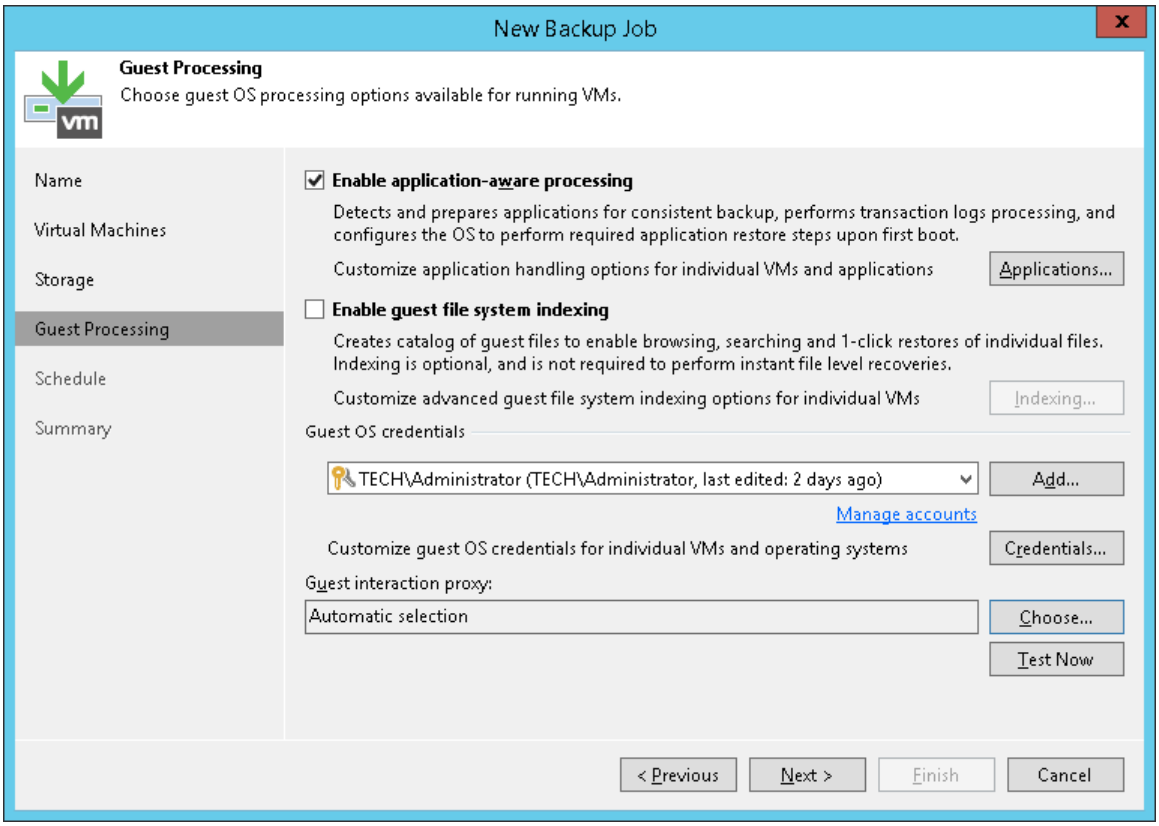
To enable both technologies:

1. At the **Storage** step of the wizard (for backup) or **Job Settings** step of the wizard (for replication), click **Advanced**.

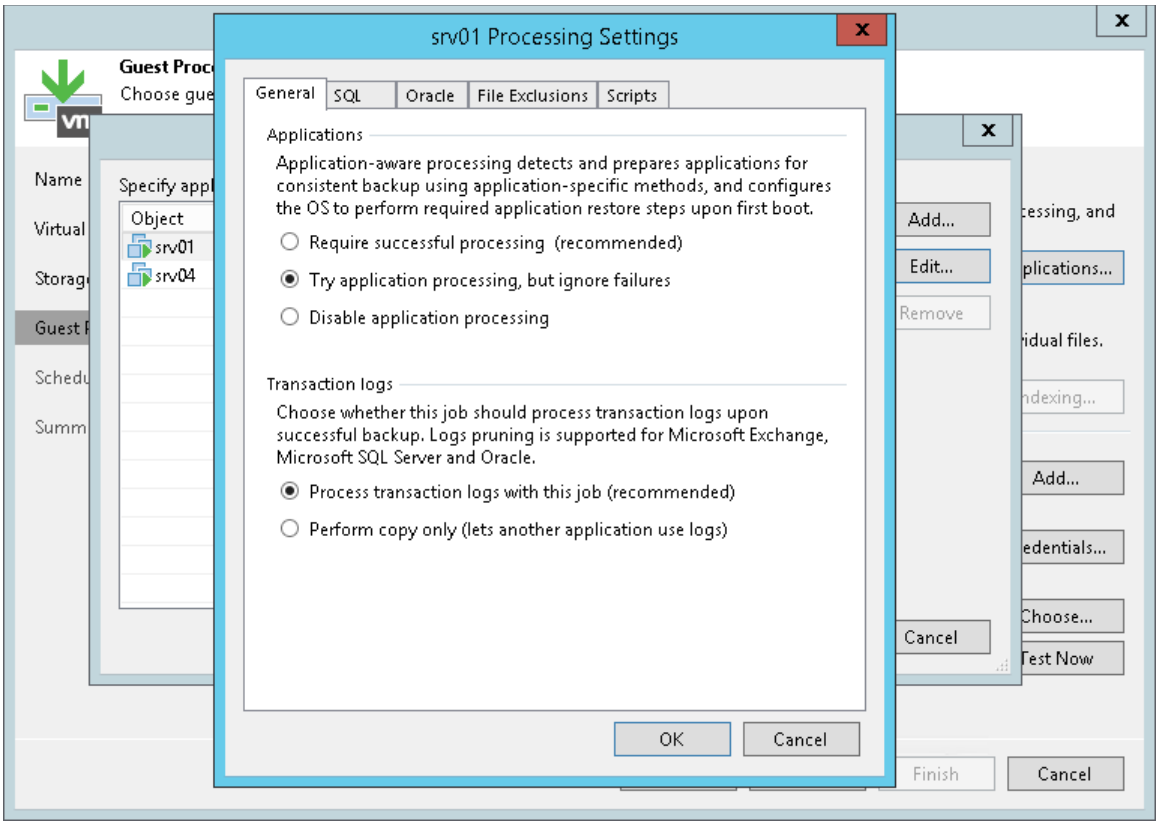
2. On the **vSphere** tab of the **Advanced Settings** window, select **Enable VMware Tools quiescence**.



3. At the **Guest Processing** step of the wizard, select the **Enable application-aware processing** check box.



4. When you configure advanced option for individual VMs, select **Try application processing, but ignore failures**. You can also select the **Disable application processing** option for VMs that you want to process with VMware Tools quiescence.



NOTE

If you enable application-aware processing and VMware Tools quiescence but do not select the **Ignore application processing failures** option, Veeam Backup & Replication uses only application-aware processing for the job.

Guest Processing

If you back up or replicate running VMs, you can enable guest processing options. Guest processing options are advanced tasks that require Veeam Backup & Replication to communicate with the VM guest OS.
Veeam Backup & Replication offers the following guest processing options:

- Application-aware processing. You can create transactionally consistent backups and replicas of VMs that run Microsoft Active Directory, Microsoft Exchange, Microsoft SharePoint, Microsoft SQL Server or Oracle Database. The transactionally consistent backup guarantees proper recovery of these applications without data loss. For information on system requirements for the applications, see System Requirements.
- Pre-freeze and post-thaw scripts. You can use pre-freeze and post-thaw scripts to quiesce VMs running applications that do not support Microsoft VSS.
- Transaction log truncation. You can set the backup or replication job to truncate transaction logs on the VM guest OS after the VM is successfully processed.
- Transaction logs backup for Microsoft SQL Server and Oracle Database. You can set up the backup job to back up transaction logs from Microsoft SQL Server and Oracle servers.
- VM guest file system indexing. You can set up the backup job to create a catalog of files and folders on the VM guest OS. The catalog lets you search for VM guest OS files and 1-click restore in Veeam Backup Enterprise Manager.

  VM guest file system indexing is optional. If you do not enable this option in the backup job settings, you will still be able to perform 1-click restore from the backup created with such backup job. For more information, see the Preparing for File Browsing and Searching section in the Enterprise Manager User Guide.
- VM guest OS files exclusion. You can exclude/include individual files and folders from/to backup or replica.

> **NOTE**
>
> If you plan to use Kerberos authentication for guest OS processing, mind the following:
>
> - Local accounts do not support Kerberos authentication. To authenticate with Microsoft Windows guest OS using Kerberos, specify an Active Directory account.
> - Networkless application-aware guest processing through VMware VIX/vSphere Web Services is not supported for VMs with guest OS where NTLM is restricted.
> - Backup infrastructure components (backup server, repositories, backup proxies, guest interaction proxies, and other components) must correctly resolve FQDNs of guest operating systems.
> - To process workloads where Kerberos is used, NTLM must be allowed in the backup infrastructure components. For details, see Kerberos Authentication for Guest OS Processing.

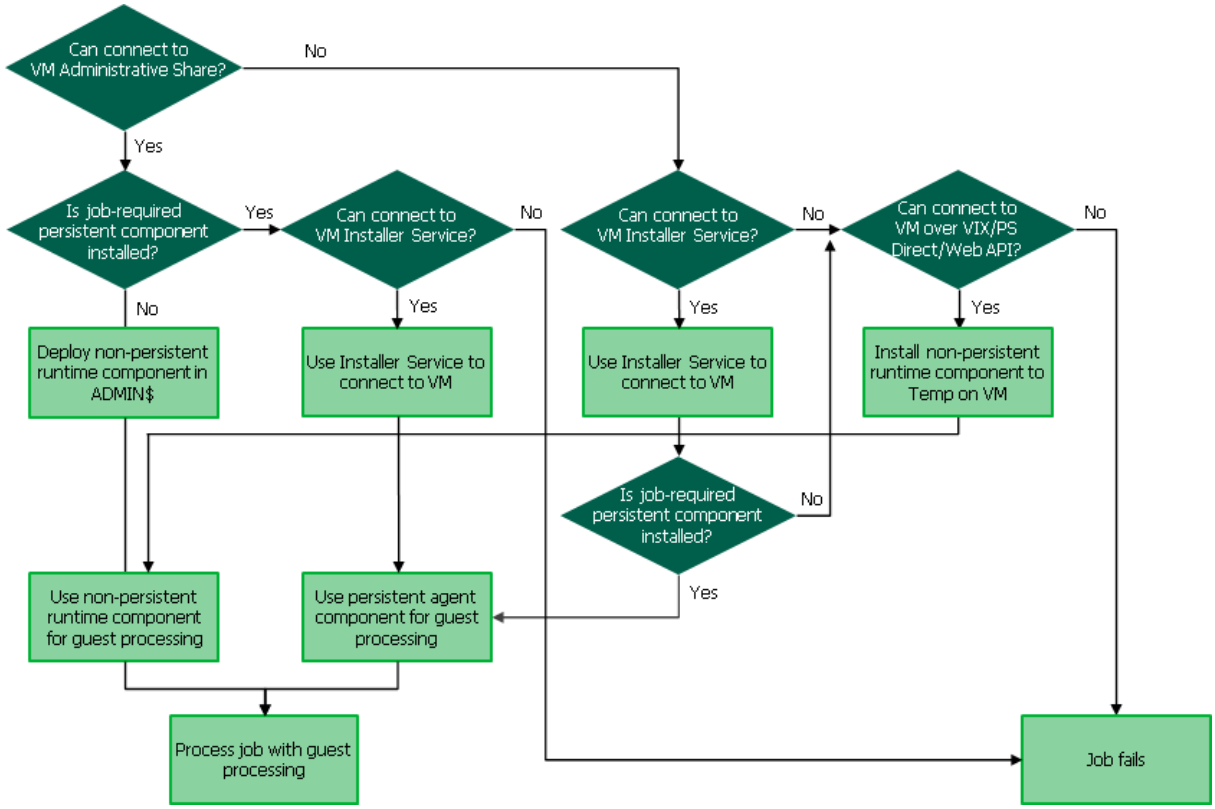
Non-Persistent Runtime Components and Persistent Agent Components

To perform guest processing tasks, Veeam Backup & Replication must use special components when protecting VMs. These components can be either non-persistent runtime components or persistent agent components.

Non-Persistent Runtime Components

The runtime components are non-persistent. They are deployed on every VM added to the job when the job starts. As soon as the job finishes, the components are removed. Use of the runtime components helps avoid agent-related drawbacks such as pre-installing, troubleshooting and updating. This method is used for guest processing by default. To enable it, make sure that you clear selection from the **Use persistent guest agent** check box when specifying application-aware processing settings, for example, as described in Application-Aware Processing for VM backup jobs.

When you do not select the **Use persistent guest agent** option, Veeam Backup & Replication performs processing of the backup jobs with enabled guest processing according to the following algorithm.



NOTE

If the **Use persistent guest agent** option is disabled, but persistent agent components were previously installed on the VM, Veeam Backup & Replication can use them for guest processing. In order to use non-persistent runtime components, uninstall the persistent agent components from protected VMs.

Persistent Agent Components

If you want to perform guest processing in a highly secure way, you can install persistent agent components (Guest Helper and Log Shipping Service) on protected VMs. Persistent agent components installed on protected VMs require very limited and clearly defined ports to communicate with Veeam Backup & Replication. That allows closing access to the VM administrative share (ADMIN$) and access to VM over VIX API/ vSphere Web Services. For more information on ports required to provide guest processing using persistent agent components, see Used Ports.

IMPORTANT

Mind that in order to use the persistent guest agent option, you must ensure that the Installer Service is installed on the VM, for which you schedule guest processing.

NOTE

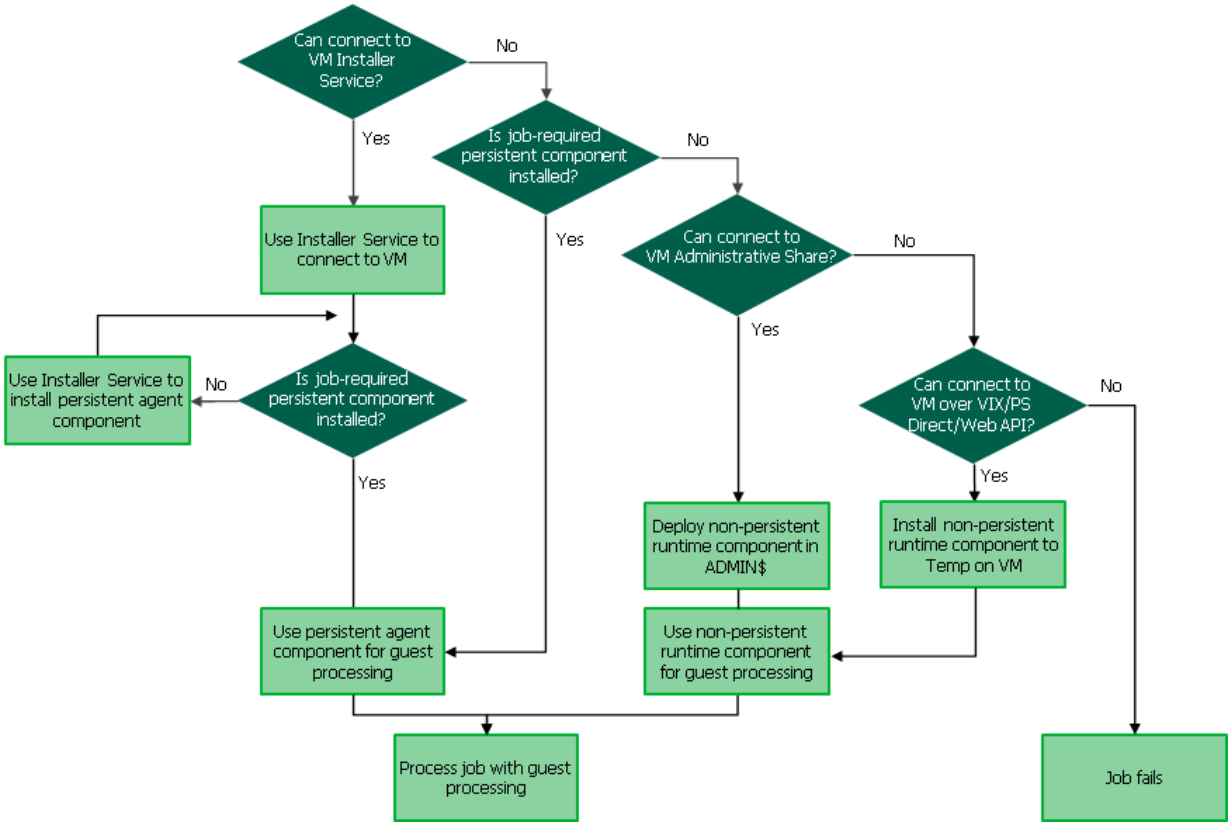
If the account used for guest processing is a local (non-domain) account, the connection may be blocked by the remote UAC. This is a default security policy in Windows that prevents local user accounts from being used to remotely connect to a server. In this case, either use the local administrator account or disable the remote UAC to connect to the persistent agent. For more information, see this Veeam KB article.

The Installer Service can be installed in one of the following ways:

- Manual installation by running the `VeeamInstallerSvc.msi` file located at `C:\Program Files\Veeam\Backup and Replication\Backup\Packages`. You can also manually preinstall the required components located in the same folder for future use. Or you can let the Installer Service install the components automatically during the first run of the backup job with enabled guest processing of this VM by persistent agents.
- Automatic installation by Veeam backup server when adding the VM to Managed Servers in the Veeam Backup & Replication infrastructure.
- Automatic installation by using a logon script.
- Automatic installation using domain group policies.
- Installation by using the standard third-party software distribution tool.

If the Installer Service is installed on the VM, for which you want to schedule guest processing, you can use persistent guest agents for guest processing. To enable this method, select the **Use persistent guest agent** check box when specifying application-aware processing settings, for example, as described in Application-Aware Processing for VM backup jobs. In this case, if persistent guest agents are missing on the VM, they will be installed by the Installer Service during the first run of the backup job.

When you select the **Use persistent guest agent** option, Veeam Backup & Replication performs processing of the backup jobs with enabled guest processing according to the following algorithm.



Deployment of Non-Persistent Runtime Components or Persistent Agent Components

Veeam Backup & Replication can deploy the non-persistent runtime components or persistent agent components on VMs in two ways:

- For VMs running Microsoft Windows, the non-persistent runtime components or persistent agent components are deployed using guest interaction proxies. For more information, see Guest Interaction Proxy.
- For VMs running OSes other than Microsoft Windows, for example, Linux, the non-persistent runtime components are deployed from the backup server.

> **NOTE**
>
> If there are no guest interaction proxies or guest interaction proxies fail for some reason, Veeam Backup & Replication will deploy the non-persistent runtime components or persistent agent components on Microsoft Windows VMs from the backup server.

When you start a job with guest processing tasks enabled, Veeam Backup & Replication performs the following operations:

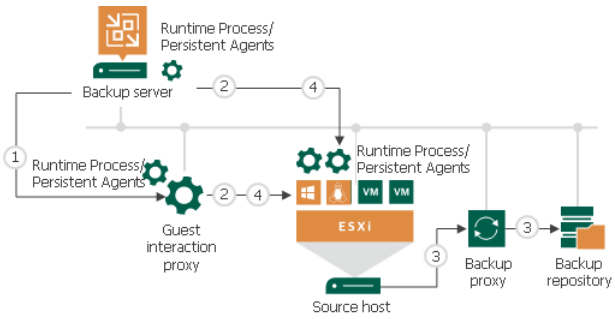
1. Veeam Backup & Replication defines the machines that will perform the guest interaction proxy role.

2. Veeam Backup & Replication detects what method it will use for guest processing: deploy the non-persistent runtime components or persistent agent components.

   Veeam Backup & Replication deploys the non-persistent runtime components or persistent agent components on VMs:

   - [For Microsoft Windows VMs] The guest interaction proxy connects to VMs and deploys the non-persistent runtime components or persistent agent components on them.
   - [For VMs running other OSes] The backup server connects to VMs and deploys the non-persistent runtime components on them.

3. The job session proceeds as usual.
4. [For the non-persistent runtime components] When the job session completes, Veeam Backup & Replication deletes the non-persistent runtime components on VMs.



If a network connection breaks during the job session, Veeam Backup & Replication makes attempts to re-establish the connection:

- If a network connection between the backup server/guest interaction proxy and VM guest OS breaks, Veeam Backup & Replication makes one attempt to reconnect.
- If a network connection between the backup server and guest interaction proxy breaks, Veeam Backup & Replication makes 10 attempts to reconnect.

If attempts are unsuccessful, guest processing tasks fail. The job proceeds with the scenario defined in the job settings. For example, if you have instructed a backup job to try application processing but ignore failures, Veeam Backup & Replication will not perform guest processing tasks but will proceed with the VM backup.

Application-Aware Processing

To create transactionally consistent backups or replicas of VMs that run Microsoft Active Directory, Microsoft SQL Server, Microsoft SharePoint, Microsoft Exchange or Oracle, you must enable application-aware processing in job settings.

Application-aware processing is the Veeam technology based on Microsoft VSS. Microsoft VSS is responsible for quiescing applications on the VM and creating a consistent view of application data on the VM guest OS. Use of Microsoft VSS ensures that there are no unfinished database transactions or incomplete application files when Veeam Backup & Replication triggers the VM snapshot and starts copying VM data to the target. For more information about Microsoft VSS, see Microsoft Docs.

Application-aware processing is supported for Microsoft Windows client versions starting from Windows Vista and for server versions starting from Windows Server 2008. To use application-aware processing, you must have VMware Tools and the latest updates installed on the VM guest OS.

IMPORTANT

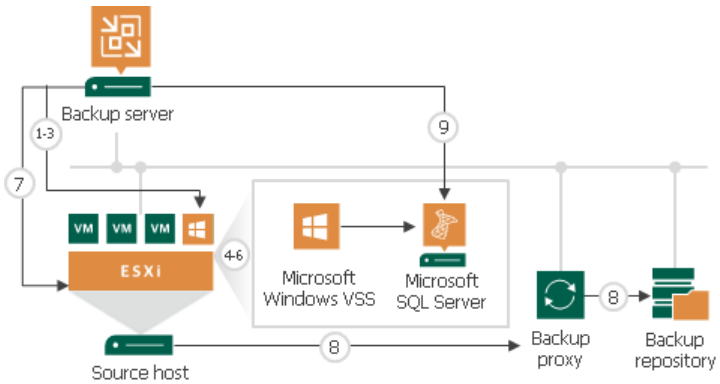
If a VM runs an application that does not support Microsoft VSS (there is no VSS writer for this particular type of application, for example, MySQL), Veeam Backup & Replication will not be able to utilize Microsoft VSS and application-aware processing for this VM. To process such VMs, you can use VMware Tools quiescence with pre-freeze and post-thaw scripts. For more information, see VMware Tools Quiescence and Pre-Freeze and Post-Thaw Scripts.

How Application-Aware Processing Works

If you enable application-aware processing in job settings, Veeam Backup & Replication performs the following operations as a part of the backup or replication process:

1. Veeam Backup & Replication deploys the non-persistent runtime components or, if necessary, persistent agent components on the VM and detects if the VM runs any of the supported applications.
2. Veeam Backup & Replication collects information about applications installed on VMs — this information is required for VSS-aware restore.

   VSS-aware restore is performed when the VM is started after you restore it from the backup or fail over to a VM replica.
3. Veeam Backup & Replication prepares applications for VSS-aware restore.
4. Microsoft VSS communicates with applications and quiesces I/O activities at a specific point in time.
5. Veeam Backup & Replication acts as a VSS requestor and triggers a VM VSS snapshot.
6. Veeam Backup & Replication triggers a VMware vSphere snapshot of the VM.
7. Microsoft VSS resumes quiesced I/O activities on the VM guest OS.
8. The job session proceeds as usual.
9. If you have instructed Veeam Backup & Replication to truncate transaction logs, Veeam Backup & Replication truncates transaction logs on the VM guest OS after the backup or replica are successfully created.



Pre-Freeze and Post-Thaw Scripts

If you back up or replicate VMs running applications that do not support Microsoft VSS, you can instruct Veeam Backup & Replication to run custom scripts for VMs. For example, the pre-freeze script may quiesce the file system and application data on the VM guest OS to bring the VM to a consistent state before Veeam Backup & Replication triggers a VM snapshot. After the VM snapshot is created, the post-thaw script may bring the VM and applications to their initial state.

You can use pre-freeze and post-thaw scripts for the following types of jobs:

- Backup job
- Replication job
- VM copy job

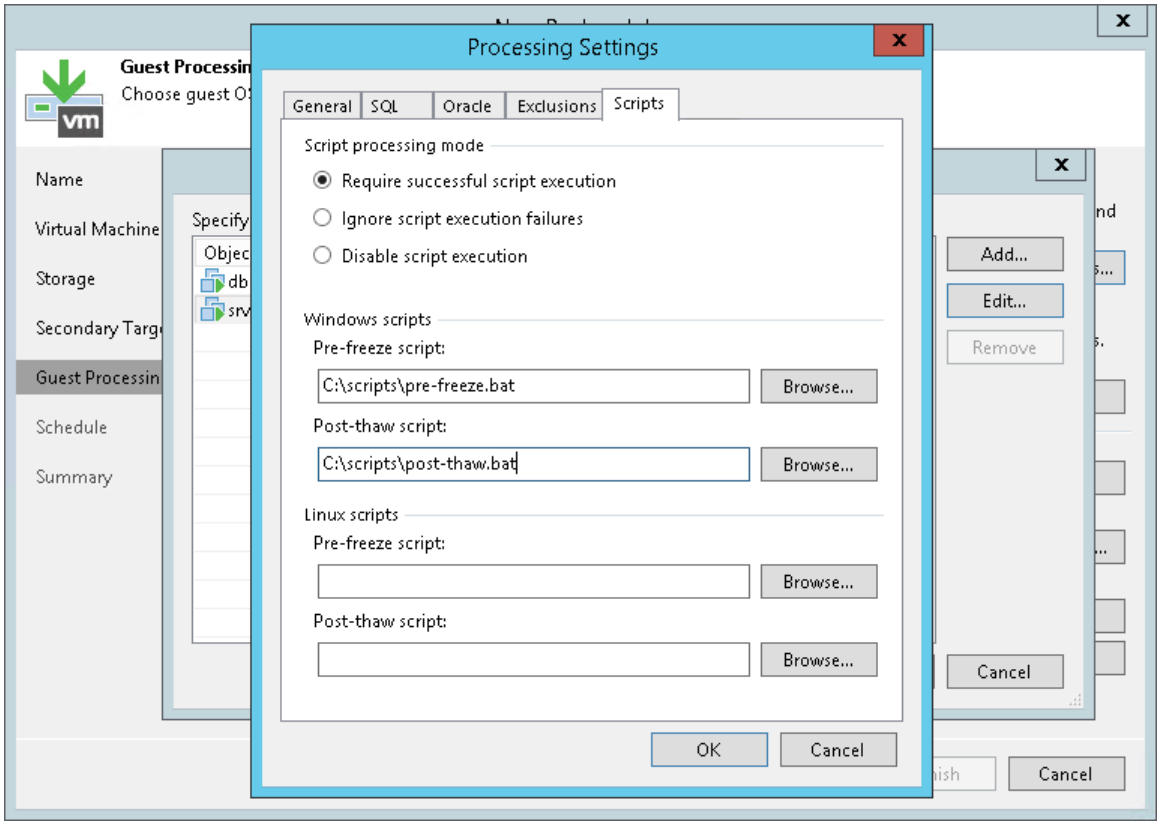
Scripts must be created beforehand. You must specify paths to them in the job settings. Script execution settings can be configured per VM or per container, depending on the objects included in the job.

When the job starts, Veeam Backup & Replication uploads scripts to the VM guest OS and executes them under the account specified in the **Guest OS credentials** section of the job settings.

- Scripts for Microsoft Windows VMs are uploaded to `\\<vmname>\admin$` over the network or VIX API/vSphere Web Services, if Veeam Backup & Replication fails to connect to the VM guest OS over the network. Scripts are executed from the `C:\Windows` directory.
- Scripts for Linux VMs are uploaded over SSH or VIX API/vSphere Web Services, if the SSH connection fails. Scripts are executed from the `/tmp` directory. If you use Renci or Rebex SSH library, and the Elevate account privileges automatically check box is selected for the user that you have specified in **Guest OS credentials**, the scripts will first be uploaded to the `/home/<username>` and then moved to `/tmp`.

The script is considered to be executed successfully if "0" is returned.

The default time period for script execution is 10 minutes. If the script fails to execute before the timeout expires, Veeam Backup & Replication displays an error message in the job session and error or warning messages issued during script execution.



Supported Script Formats

Pre-freeze and post-thaw scripts can be used for Microsoft Windows and Linux VMs.

- For Microsoft Windows VMs, Veeam Backup & Replication supports scripts in the EXE, BAT, CMD, WSF, JS, VBS and PS1 file format.
- For Linux VMs, Veeam Backup & Replication supports scripts in the SH file format.

Limitations for Pre-Freeze and Post-Thaw Scripts

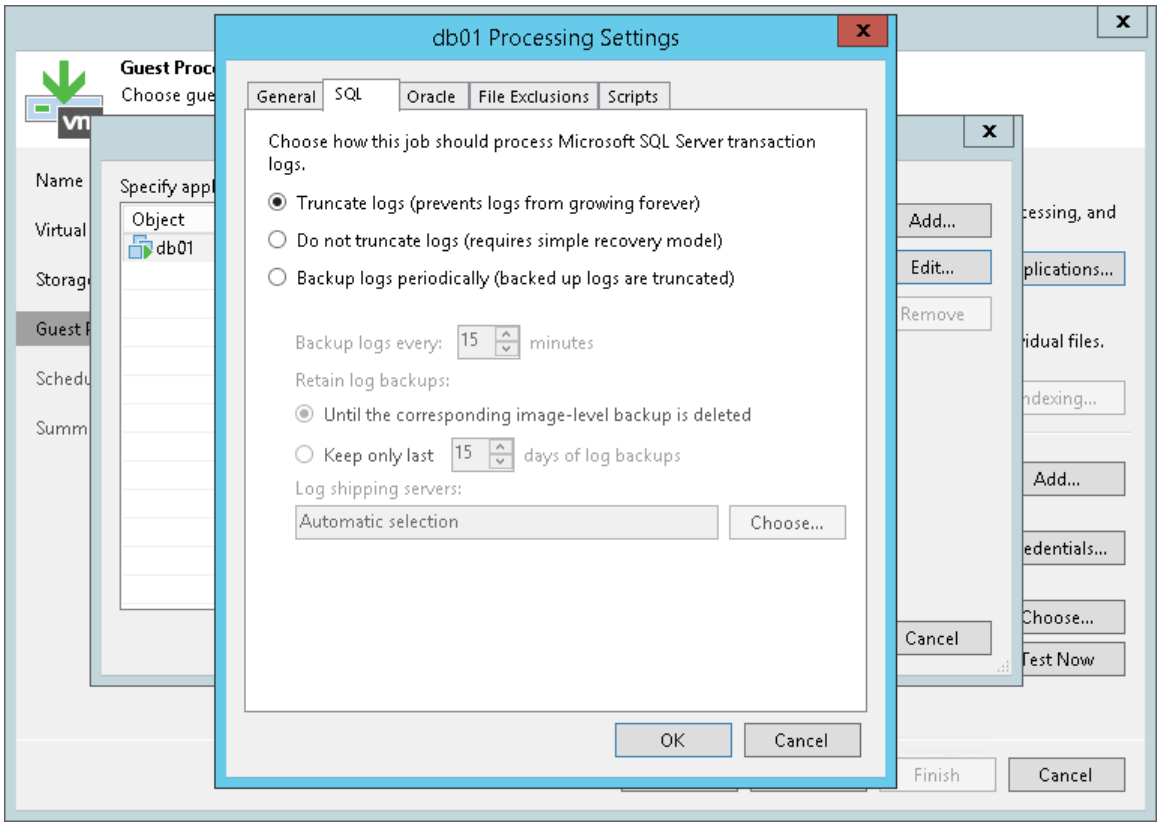
Veeam Backup & Replication has the following limitations for pre-freeze and post-thaw scripts:

- You cannot stop a job when the pre-freeze or post-thaw script is executed. If the script hangs up, Veeam Backup & Replication waits for 10 minutes and terminates the job.
- If you want to run several scripts that depend on each other, you must upload them to the VM guest OS manually. For example, you have *script1.bat* that sequentially starts *script2.bat*, *script3.bat* and *script4.bat*. In this case, you must specify a path to *script1.bat* in the job properties and upload *script2.bat*, *script3.bat* and *script4.bat* to the VM guest OS.
- You must not use the standard error (STDERR) stream for error output in Linux scripts. Scripts with STDERR cause failures in Veeam Backup & Replication.

Transaction Log Truncation

If you back up or replicate virtualized database systems that use transaction logs, for example, Microsoft Exchange or Microsoft SQL Server, you can instruct Veeam Backup & Replication to truncate transaction logs so that logs do not overflow the storage space on the VM. Veeam Backup & Replication provides the following options of transaction logs handling:

- Truncate logs
- Do not truncate logs
- Back up logs periodically



Truncate Logs

You can instruct Veeam Backup & Replication to truncate logs after a backup or VM replica is successfully created. With this option selected, Veeam Backup & Replication behaves in the following way:

- If the job completes successfully, Veeam Backup & Replication produces a backup file or VM replica and truncates transaction logs on the original VM. As a result, you have the backup file or replica that contains a VM image at a specific point in time.

  In this scenario, you can recover a database to the point in time when the backup file or replica was created. As transaction logs on the VM are truncated, you cannot use them to get the restored database to some point in time between job sessions.

- If the backup or replication job fails, Veeam Backup & Replication does not truncate transaction logs on the VM. In this scenario, you can restore a VM from the most recent backup or replica restore point and use database system tools to apply transaction logs and get the database system to the necessary point in time after the restore point.

Do not Truncate Logs

You can choose not to truncate transaction logs on the VM. This option is recommended if together with Veeam Backup & Replication you use another backup tool.

For example, you can use Veeam Backup & Replication to create a VM image backup and instruct the native Microsoft SQL Server log backup job to back up transaction logs. If you truncate transaction logs with Veeam Backup & Replication, the chain of transaction logs will be broken, and the Microsoft SQL Server log backup job will not be able to produce a consistent log backup.

With this option selected, Veeam Backup & Replication produces a backup file or VM replica and does not trigger transaction log truncation. As a result, you have a backup file or VM replica that contains a VM image captured at a specific point in time, and transaction logs on the VM. You can use transaction logs to restore the VM to any point in time between job sessions. To do this, you must recover the VM from the backup file or perform replica failover and use database system tools to apply transaction logs and get the database system to the necessary point in time.

Back Up Logs Periodically

This option can be used if you back up Microsoft SQL Server VMs and Oracle VMs.

You can choose to back up logs with Veeam Backup & Replication. For more information, see Microsoft SQL Server Logs Backup and Oracle Logs Backup.

Support for Database Availability Groups (DAG)

Veeam Backup & Replication supports any configuration of DAGs, in particular, with all databases active on one node, or with active databases on every node. Transaction logs will be truncated on all DAG members, no matter whether Veeam Backup & Replication backs up an active or passive database.

For more information and recommendations on Microsoft Exchange Server backup, you can also refer to this Veeam KB article.

Copy-Only Backup

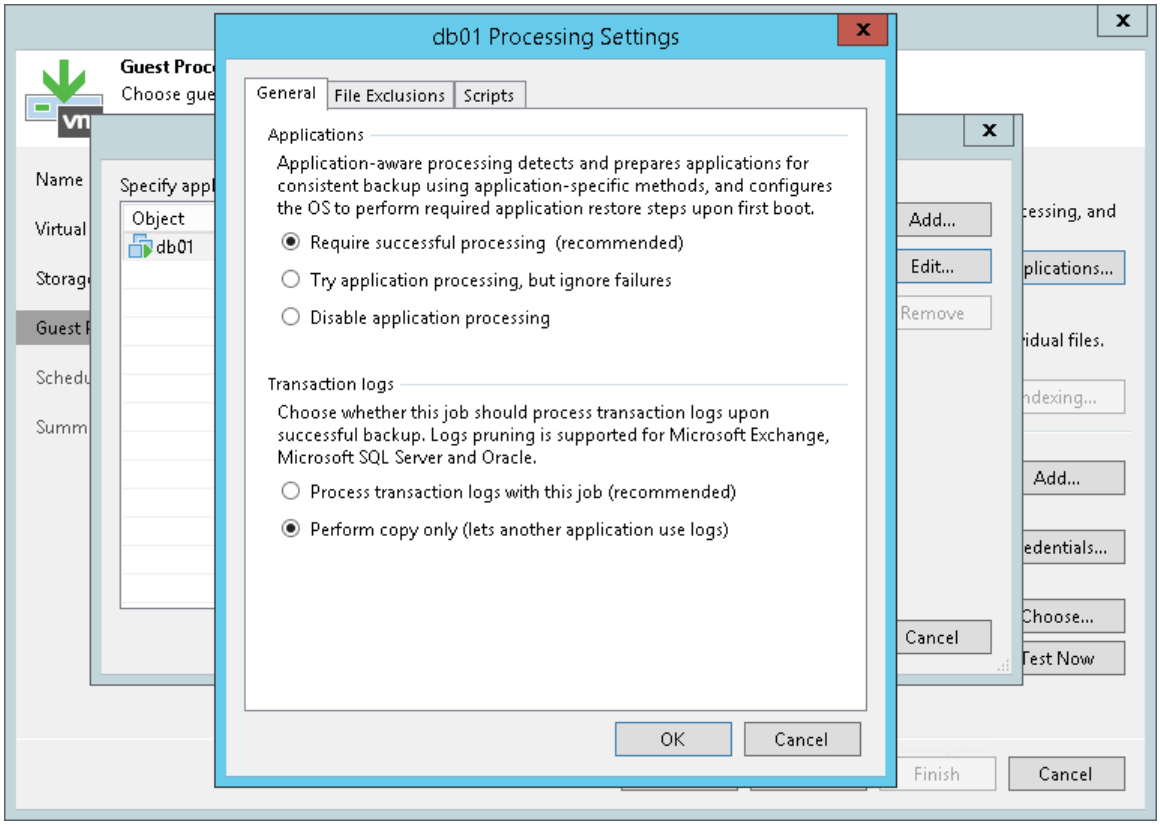
Some organizations prefer to back up Microsoft SQL Server databases and transaction logs with native Microsoft SQL Server tools or 3rd party backup tools. To restore database systems in a proper way, database administrators must be sure that they have database backups and a sequence of transaction log backups associated with these backups at hand.

If you use native Microsoft SQL Server tools or 3rd party backup tools and also want to back up Microsoft SQL Server VMs with Veeam Backup & Replication, you must enable the **Perform copy only** option in the job settings.

The **Perform copy only** option indicates that a chain of database backups is created with native Microsoft SQL Server means or by a 3rd party tool, and instructs Veeam to preserve this chain (backup history).
Veeam Backup & Replication backs up the Microsoft SQL Server VM using the *VSS_BS_COPY* method for snapshot creation. The *VSS_BT_COPY* method produces a copy-only backup — the backup that is independent of the existing chain of database backups and does not contain transaction logs data. As a result, the copy-only backup does not change the log sequence number and transaction log backup time.

IMPORTANT

Veeam Backup & Replication does not truncate transaction logs after copy-only backup. For this reason, if you instruct the backup job to perform copy-only backup, you cannot specify transaction log handing settings for this job.



VM Guest OS File Indexing

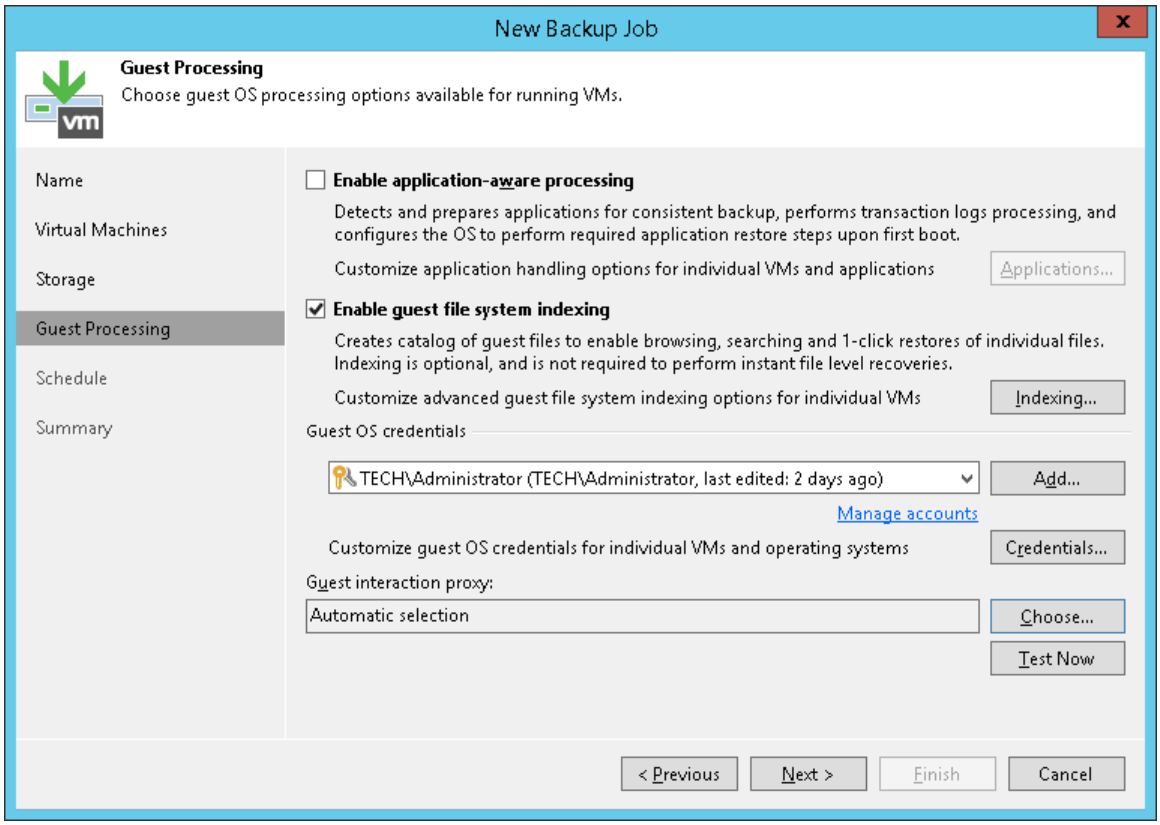
You can instruct Veeam Backup & Replication to create an index of files and folders on the VM guest OS during backup. VM guest OS file indexing allows you to search for VM guest OS files inside VM backups and perform 1-click restore in Veeam Backup Enterprise Manager.

VM guest OS file indexing is enabled at the job level. You can specify granular indexing settings for every VM in the job.

NOTE

VM guest OS file indexing is optional. If you do not enable this option in the backup job settings, you will still be able to perform 1-click restore from the backup created with such backup job. For more information, see the Preparing for File Browsing and Searching section in the Enterprise Manager User Guide.

Mind, however, that if you do not enable indexing in the backup job, during 1-click restore from Linux and other OS backups, Veeam Backup Enterprise Manager will not display symlinks to folders in the file system browser.



Requirements and Limitations for VM Guest OS File Indexing

Consider the following requirements and limitations for VM guest OS file indexing:

- Veeam Backup & Replication supports file indexing for VMs running Microsoft Windows and Linux OS.
- [For Linux VMs] To perform guest OS file indexing, Veeam Backup & Replication requires several utilities to be installed on the Linux VM: openssh, mlocate, gzip and tar. If these utilities are not found, Veeam Backup & Replication will prompt you to deploy them on the VM guest OS.

Veeam Backup Catalog

For VM guest OS file indexing, Veeam Backup & Replication uses the Veeam Guest Catalog Service. In the backup infrastructure, the Veeam Guest Catalog Service is installed on the Veeam backup server and Veeam Backup Enterprise Manager server.

- The Veeam Guest Catalog Service on the Veeam backup server works as a local catalog service. It collects indexing data for backup jobs and stores this data in the Veeam Backup Catalog folder.

  By default, the indexing data is stored in the `VBRCatalog` folder on the backup server.
  Veeam Backup & Replication creates the folder on a volume with the maximum amount of free space, for example, `C:\VBRCatalog`.

- The Veeam Guest Catalog Service on Veeam Backup Enterprise Manager works as a global, federal catalog service. It communicates with Veeam Guest Catalog Services on backup servers connected to Veeam Backup Enterprise Manager and performs the following tasks:

  - Replicates indexing data from backup servers to create a global catalog for the whole backup infrastructure.

    On the Veeam Backup Enterprise Manager server, the default folder for storing indexing data (the `VBRCatalog` folder) is located on a volume with the maximum amount of free space.

  - Maintains indexing data retention.

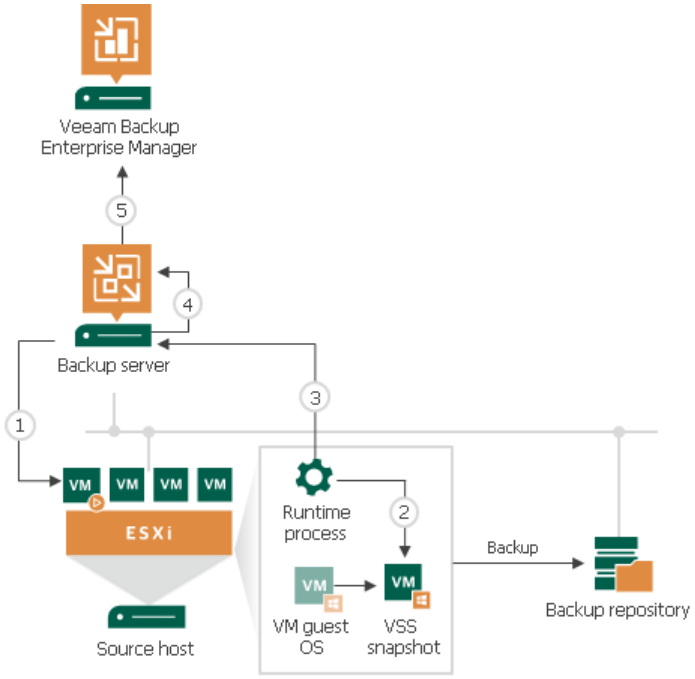
  - Allows you to search for VM guest OS files in current and archived backup files.

How VM Guest OS File Indexing Works

When you run a backup job with the file indexing option enabled, Veeam Backup & Replication performs the following operations:

1. When the backup job starts, Veeam Backup & Replication connects to the VM whose file system must be indexed and deploys non-persistent runtime components or connects to (if necessary, deploys) persistent agent components inside this VM. These components are responsible for coordinating indexing activities inside the VM.
2. The non-persistent runtime components or persistent agent components start indexing the VM file system. The indexing procedure is carried out in parallel with the backup procedure. If indexing takes long, Veeam Backup & Replication will not wait for the indexing procedure to complete. It will start copying VM data and continue file indexing inside the VM. If you have enabled application-aware processing for the VM, Veeam Backup & Replication performs indexing using the VSS snapshot, not the VM guest OS itself. As a result, the created file index exactly reflects the state of the backed-up VM.
3. When file indexing is complete, the non-persistent runtime components or persistent agent components collect indexing data and write it to the GuestIndexData.zip file. The GuestIndexData.zip file is stored to a temporary folder on the backup server.
4. When the backup job completes, Veeam Backup & Replication notifies the local Veeam Guest Catalog Service, and the service saves indexing data in the Veeam Catalog folder on the backup server.

5. During the next catalog replication session, the global Veeam Guest Catalog Service replicates data from the backup server to the Veeam Catalog folder on the Veeam Backup Enterprise Manager server.



Persistent VSS Snapshots

During application-aware processing, Veeam Backup & Replication uses a VSS writer for a corresponding application to freeze application data and bring it to a consistent state.

According to Microsoft limitations, applications cannot be kept frozen longer than for 60 seconds (20 seconds for Microsoft Exchange). If the Microsoft VSS writer keeps application data frozen longer than this period, a VSS processing timeout occurs, and Veeam Backup & Replication fails to create a transactionally consistent backup of the VM. The VSS processing timeout is a common problem for highly transactional applications such as Microsoft Exchange.

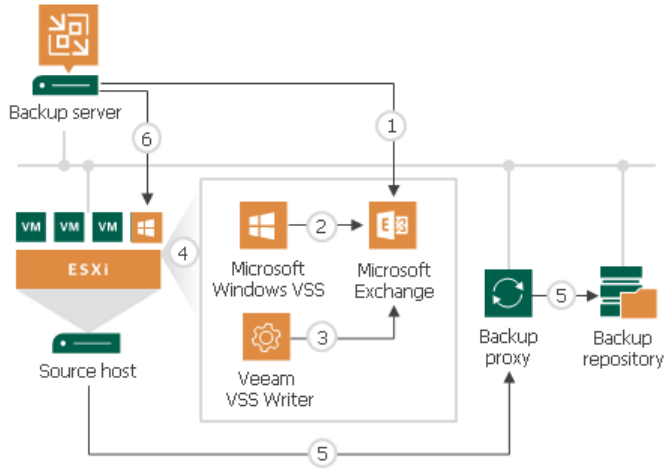
To overcome this limitation, Veeam Backup & Replication uses the persistent VSS snapshots technology for backup of Microsoft Exchange VMs. If Microsoft Exchange has to be kept frozen for a longer period of time than the allowed one, Veeam Backup & Replication automatically fails over to the persistent VSS snapshot mechanism.

Backup of Microsoft Exchange VMs is performed in the following way:

1. Veeam Backup & Replication triggers the Microsoft VSS framework to prepare Microsoft Exchange on the VM for backup.
2. The Microsoft VSS writer attempts to quiesce Microsoft Exchange.
3. If the Microsoft VSS writer fails to quiesce Microsoft Exchange within the allowed period of time, the control is passed to the native Veeam VSS writer. The Veeam VSS writer holds the freeze operation for the necessary amount of time.
4. After Microsoft Exchange data is brought to a consistent state, the control is passed to the Microsoft VSS provider. The Microsoft VSS framework creates a persistent VSS snapshot for VM disks except the system VM disk.
5. The job session proceeds as usual.

6. After the backup operation is complete, Veeam Backup & Replication triggers Microsoft VSS to remove the persistent VSS snapshot on the production VM. The persistent VSS snapshot holding consistent application data inside the created VM backup remains.

During entire VM restore, Veeam Backup & Replication recovers data from the backup and reverts VM disks to the persistent VSS snapshot inside the backup. As a result, the Microsoft Exchange VM is restored from the backup in a consistent state without data loss.



Limitations for Persistent VSS Snapshot Technology

Veeam Backup & Replication uses the persistent VSS snapshot technology if the VM meets the following requirements:

- The version of Microsoft Exchange installed on the VM is listed in Supported Applications.
- The VM does not perform the role of a domain controller.
- Microsoft Exchange databases and log files are located on a non-system disk of the VM. During backup, Veeam Backup & Replication does not trigger a persistent VSS snapshot for system VM disks. As a result, system disks are restored in a crash-consistent, not transactionally consistent state.

Microsoft SQL Server Log Backup

To protect Microsoft SQL Server VMs, you can instruct the backup job to create image-level VM backups and periodically back up database transaction logs. If Microsoft SQL Server fails, you can restore the Microsoft SQL Server VM from the necessary restore point of the image-level backup. After that, you can use Veeam Explorer for Microsoft SQL Server to apply transaction logs and get databases on the Microsoft SQL Server to the necessary state between backups.

Transaction Log Backup Jobs

To back up transaction logs, you must create a backup job, add Microsoft SQL Server VMs to it and specify advanced settings for transaction logs backup in the job settings. The resulting job will comprise two jobs:

- Parent backup job — the backup job that creates an image-level backup of the Microsoft SQL Server VM. The parent backup job is named *<job_name>*, for example, *DB Backup*. You can configure the parent job in the Veeam Backup & Replication console just like any other backup job.
- Child job — a transaction log backup job. To form a name of the child job, Veeam Backup & Replication adds a suffix to the name of the parent backup job: *<parent_job_name> + SQL Server Transaction Log Backup*, for example, *DB Backup SQL Server Transaction Log Backup*. Veeam Backup & Replication automatically creates the child job if it detects a backup job that is scheduled to back up at least one Microsoft SQL Server VM, and transaction log backup is enabled for this job. Session data of the transaction log backup job is stored in the configuration database and displayed in the Veeam Backup & Replication console.

The parent job runs in a regular manner — it starts by schedule or is started manually by the user. The transaction log backup job is triggered by the parent backup job. This sequence ensures that the VM (and the database) restore point is present when it comes to transaction log replay.

Sessions of Transaction Log Backup Jobs

The transaction log backup job runs permanently in the background, shipping transaction logs to the backup repository at a specific time interval (by default, every 15 minutes). A sequence of time intervals between sessions of the parent backup job makes up a session of the transaction log backup job.

The transaction log backup session starts and stops in the following way:

- The initial session starts when the parent backup job schedule is enabled. After that, the session starts with every new session of the parent backup job.
- The session ends before the next session of the parent backup job or when this parent backup job is disabled.
- When the session ends, Veeam Backup & Replication stops the non-persistent runtime components and uninstalls them from the VM guest OS. When a new session starts, the runtime components are deployed again.

How Microsoft SQL Server Log Backup Works

To perform transaction log backup, Veeam Backup & Replication installs the *Veeam Guest SQL Log Shipper* runtime component on the VM guest OS.

This component runs during the transaction log backup job session and checks the following details:

- Collects information about databases that require transaction log backup.

- Verifies that it is possible to ship logs directly to the backup repository. If it is not possible, Veeam Backup & Replication uses the log shipping server.

When the transaction log backup job session ends, the component is stopped and removed from the VM guest OS. When a new session starts, the component is installed on the VM guest OS again.

The transaction log backup is performed in the following way:

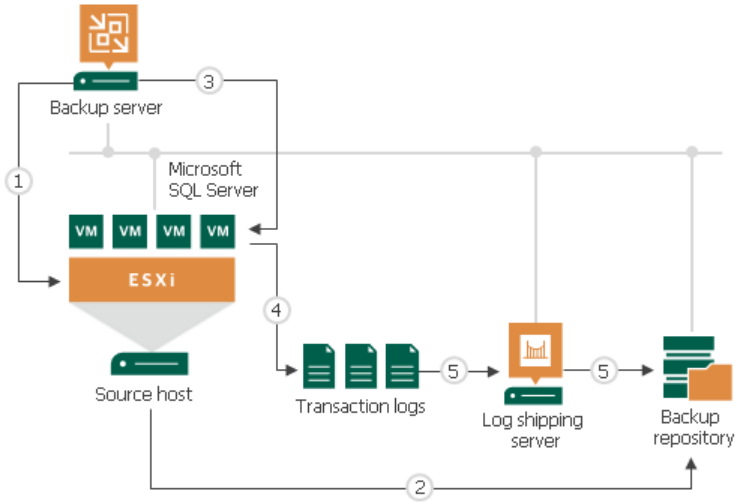
1. Veeam Backup & Replication launches the parent backup job by schedule.
2. The parent backup job creates an image-level backup of a Microsoft SQL Server VM and stores it in a backup repository.
3. A new session of the transaction log backup starts. Veeam Backup & Replication accesses the VM (directly or through the guest interaction proxy) and installs the runtime components for guest processing, database information collection and transaction log handing on the VM guest OS.
4. Veeam Backup & Replication detects what databases currently exist on the Microsoft SQL Server and maps this data with the information kept in the configuration database. This periodic mapping reveals the databases for which Veeam Backup & Replication must process transaction logs during this time interval.

   The runtime component backs up transaction log files and stores them as a *.bak file to a temporary folder on the VM guest file system.
5. Veeam Backup & Replication transports transaction log backup copies from the temporary folder on the Microsoft SQL Server VM to the backup repository, either directly or through the log shipping server, and saves them as VLB files. As soon as copies of transaction log backups are saved to the backup repository, transaction log backups in the temporary folder on the Microsoft SQL Server VM are removed.

The session of the transaction log backup job remains working until the next start of the parent backup job. When a new session of the parent job starts, the transaction log backup job stops the current session and then starts a new session, performing steps 1-5.

Transaction logs that for some reason were not processed during the log backup interval remain in the temporary folder and are processed during the next log backup interval. To detect these remaining logs, Veeam Backup & Replication enumerates log files in the temporary folder.

> **NOTE**
>
> If a new session of the transaction log backup starts and the parent backup job has not created a new restore point yet, the transaction log backup job will remain in the idle state, waiting for a new restore point to be created.



Retention for Transaction Log Backups

Transaction log backups are stored in files of the proprietary Veeam format — VLB. Veeam Backup & Replication keeps transaction log backups together with the VM image-level backup. The target location of VLB files depend on the type of the backup repository:

- If you store the VM image-level backup in a backup repository, Veeam Backup & Replication writes transaction log backups to the same folder where files of the image-level backup reside.
- If you store the VM image-level backup in a scale-out backup repository, Veeam Backup & Replication writes transaction log backups to the extent where the latest incremental backup file of the VM image-level backup is stored.

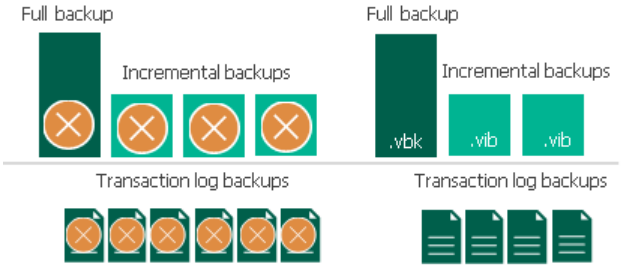
Veeam Backup & Replication removes transaction log backups by retention. You can choose one of the following retention methods:

- Retain logs according to the image-level backup
- Retain logs for the specified number of days

Retain Logs with Image-Level Backup

By default, Veeam Backup & Replication retains transaction log backups together with the corresponding image-level backup of the Microsoft SQL Server VM. Veeam Backup & Replication retains VM backup and log backups according to the short-term retention configured for VM backups. When Veeam Backup & Replication removes a restore point of the image-level backup from the backup chain, it also removes a chain of transaction logs relating to this image-level backup. Note that even if long-term retention is configured for the VM backup, Veeam Backup & Replication retains transaction log backups according to the short-term retention policy and deletes them after the short-term retention is exceeded.

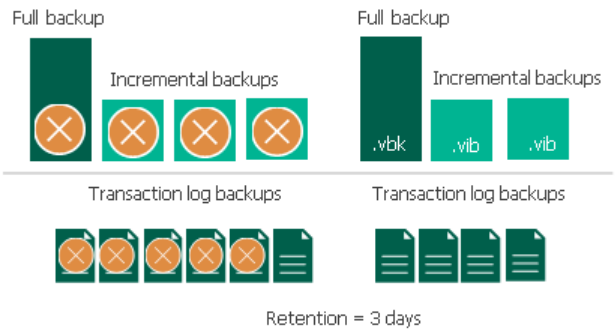
This method allows you to have both the image-level backup and necessary transaction log backups at hand. If you need to recover a database to some state, you can restore the Microsoft SQL Database from the necessary restore point and perform transaction log replay to bring the database to the desired state.



Retain Logs for a Number of Days

You can instruct Veeam Backup & Replication to keep transaction logs only for a specific period of time. This retention setting can be used, for example, if you want to save on storage space and plan to retain transaction log backups for the last few days. In this case, you will be able to restore the database only to one of the most recent states.

If you select this retention method, you must make sure that retention policies for the image-level backup and transaction log backup are consistent. The restore point of the image-level backup must always be preserved. If a backup of the database itself is missing, you will not be able to perform transaction log replay.



Related Topics

Microsoft SQL Server Transaction Log Settings

Log Shipping Servers

For every Microsoft SQL Server VM whose transaction logs you want to back up, Veeam Backup & Replication defines how to ship logs to the backup repository. Transaction logs can be shipped in the following ways:

- If it is possible to establish a direct connection between the VM guest OS and backup repository, log files will be shipped directly from the VM guest OS to the backup repository. This is the optimal method, as it does not involve additional resources and puts less load on the VM guest OS.
- Otherwise, files will be shipped through log shipping servers. You can instruct Veeam Backup & Replication to choose a log shipping server automatically from the list of available ones, or to use a specific server.

  Note that if direct connection is possible, files will be always transferred from VM guest to repository directly (regardless of the configured log shipping server, as this server will not be involved). This approach helps to optimize performance at file transfer.

A log shipping server is a Microsoft Windows server added to the backup infrastructure. You can explicitly define what servers you want to use for log shipping or instruct Veeam Backup & Replication to automatically choose an optimal log shipping server. Veeam Backup & Replication chooses the log shipping server based on two criteria: possible data transfer methods and location of the Microsoft SQL Server VMs and log shipping server.

Data Transfer Methods

Log shipping servers can transport data in two ways:

- Over the network. In this scenario, Veeam Backup & Replication obtains files from the VM guest OS and transfers them over the network.

  To offload the VM guest OS, logs are created one by one (not simultaneously). One log creation request is issued for every DB.
- Using VIX API/vSphere Web Services. In this scenario, Veeam Backup & Replication obtains transaction logs from the VM guest OS over the VIX API/vSphere Web Services, bypassing the network. For each Microsoft SQL Server instance one log creation request is created for all DBs (grouped by instance).

The default method is log shipping over the network.

Location of Log Shipping Server and VMs

When choosing a log shipping server for the job, Veeam Backup & Replication considers the location of the Microsoft SQL Server VM and log shipping server. Veeam Backup & Replication uses the following priority rules to select the log shipping server:

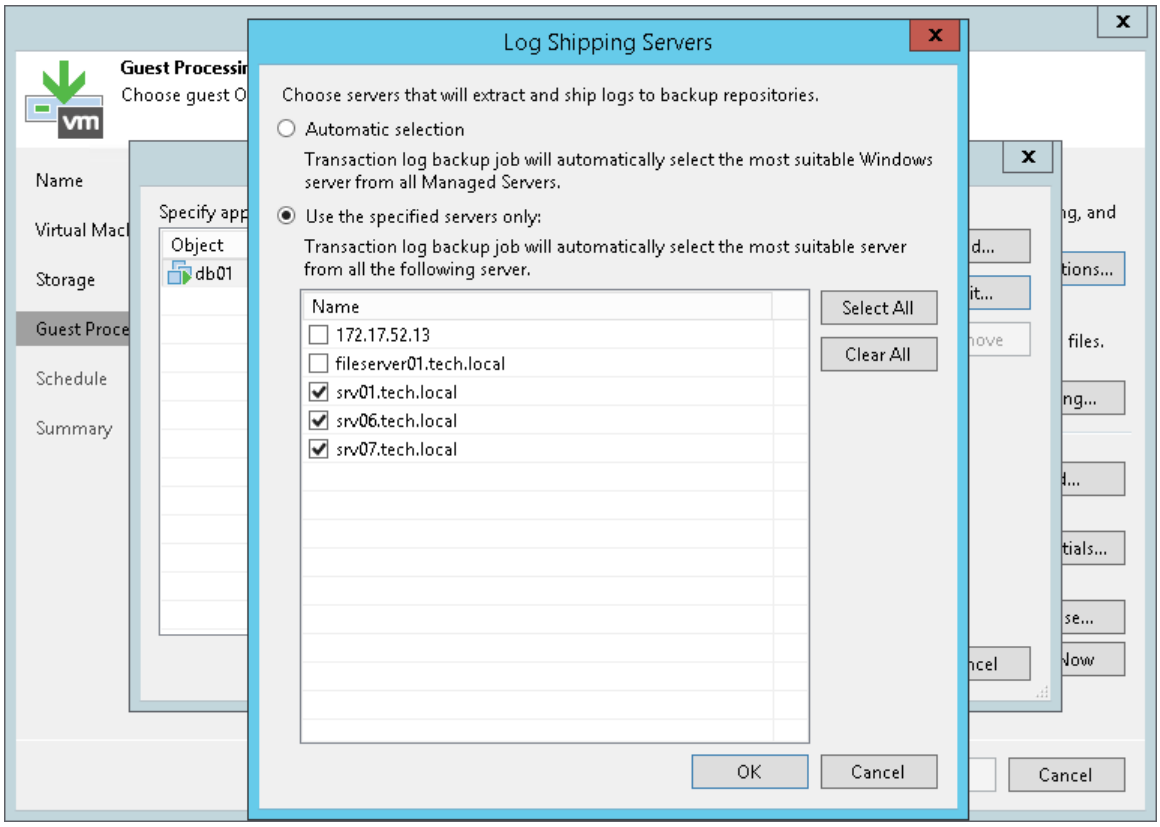
1. Log shipping server is located on the same ESXi host as the Microsoft SQL Server VM.
2. Log shipping server and Microsoft SQL Server VM are located in the same network.
3. Log shipping server and Microsoft SQL Server VM are located in different networks (the production infrastructure is isolated from the backup infrastructure).

That is, when choosing a log shipping server, Veeam Backup & Replication will give the top priority to a Microsoft Windows VM that is located on the same ESXi host as the Microsoft SQL Server VM and that has a network connection to the Microsoft SQL Server VM.

Log shipping servers are assigned per job session. When a new job session starts, Veeam Backup & Replication detects log shipping servers anew. Veeam Backup & Replication can also re-detect available servers during the job session. If a log shipping server becomes unavailable for some reason, Veeam Backup & Replication will fail over to another log shipping server.

IMPORTANT

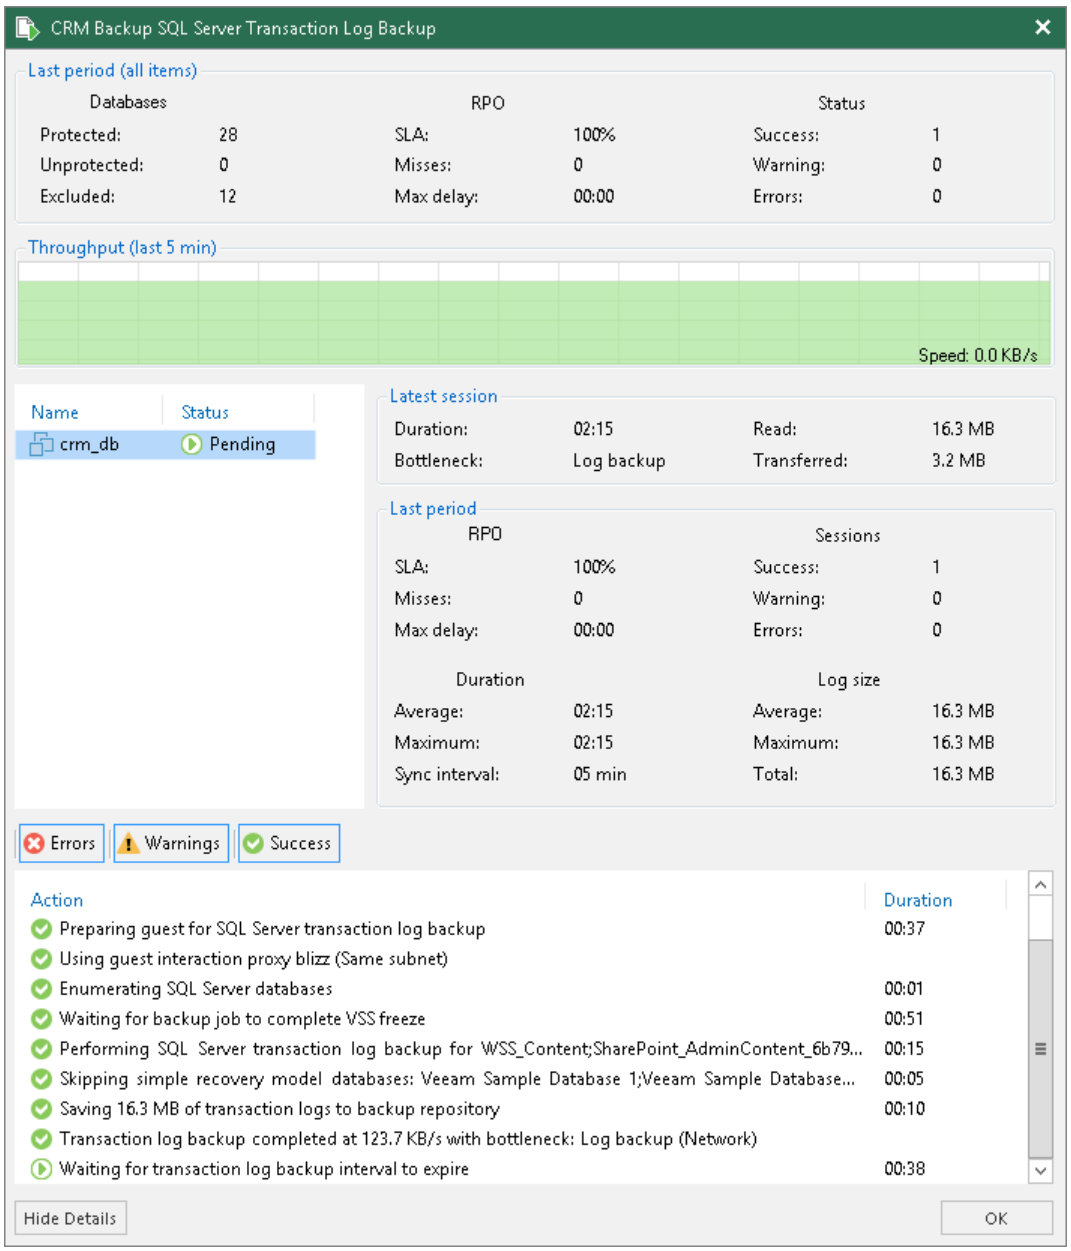
If you do not want to use some servers for transaction logs transport, you can manually define what server Veeam Backup & Replication must use as a log shipping server in the job settings. It is recommended that you assign the log shipping server role to a number of servers for availability purposes.



Transaction Log Backup Statistics

You can view the statistics of the transaction log backup job in the **History** view or in the **Home** view in Veeam Backup & Replication.

In the statistics window, you can examine the overall statistics for the transaction log backup job, as well as view per-VM information.



In the upper part of the statistics window, Veeam Backup & Replication displays information about the transaction log backup job for all VMs included in the parent backup job.

The **Last period (all VMs)** section contains statistics data for the selected session of the backup job.

In the **Databases** column, you can view the following information:

- *Protected* — number of databases that were backed up at least once during the last session
- *Unprotected* — number of databases that failed to be backed up during the last session

- *Excluded* — databases excluded from processing. Databases may be excluded for the following reasons:
  - The database status is *Offline.*
  - The database recovery model is set to *Simple.*
  - The database is read-only.
  - The database was deleted after the latest full backup.
  - The AutoClose property is enabled for the database.
  - The database was excluded from application-aware processing. For details, see this Veeam KB article.
  - The database was excluded from SQL log backup processing. For details, see this Veeam KB article.
  - The database belongs to vCenter Server. For details, see this Veeam KB article.

> **NOTE**
>
> Unprotected databases do not comprise *Excluded* databases, as they have different reasons for being non-processed.

In the **RPO** column, you can view the following information:

- *SLA value* — how many log backup intervals completed in time with successful log backup (calculated as percentage of total number of intervals).
- *Misses* — how many intervals were missed (number of intervals).
- *Max delay* — difference between the configured log backup interval and time actually required for log backup. If exceeded, a warning is issued.

In the **Status** column, the following information is displayed (per job): number of VMs processed successfully, with warnings or with errors.

The **Latest session** section displays the following information for the latest log processing interval for the selected VM:

- *Duration* — duration of log shipment from the VM guest OS to the backup repository since the current log processing interval has started
- *Bottleneck* — operation with the greatest duration in the last completed interval. The operation may have the following bottlenecks:

Display Name	Slowing-down Operation
Log backup	Saving BAK files to a temporary location on VM guest OS
Network	Uploading log files to the log shipping server
Target	Saving files to the target repository

- *Read* — amount of data read from the temporary folder on VM guest OS
- *Transferred* — amount of data transferred to the target repository

The **Last period** section displays the following statistics of log backups per VM for the latest session of the transaction log backup job:

- The **RPO** column displays statistics on log processing interval (calculated as described above)
- The **Sessions** column includes statistics of log backups per VM, calculated as follows:
  - *Success* — number of intervals when all database logs were backed up successfully
  - *Warning* — number of sequential intervals with failed log processing (if not more than 4 intervals in a sequence)
  - *Errors* — number of sequential intervals with failed log processing (more than 4 intervals in a sequence)
- The **Duration** column includes the following information:
  - *Average* — average duration of log data transfer (through all intervals in the session)
  - *Max* — maximal duration of log data transfer (through all intervals in the session)
  - *Sync interval* — duration of periodic intervals specified for log backup in the parent job settings (default is 15 min)
- The **Log size** column displays the following information:
  - *Average* — average amount of data read from the VM guest OS through all intervals
  - *Max* — maximal amount of data read from the VM guest OS over all 15-min intervals
  - *Total* — total amount of data written to the backup repository

NOTE

- Statistics on transaction log processing is updated periodically, simultaneously for the parent backup job and transaction log backup job.
- For Always On Availability groups, Veeam Backup & Replication collects logs only from one node. Thus, in reports, the status of database replicas will be the same for all nodes (Protected or Excluded).

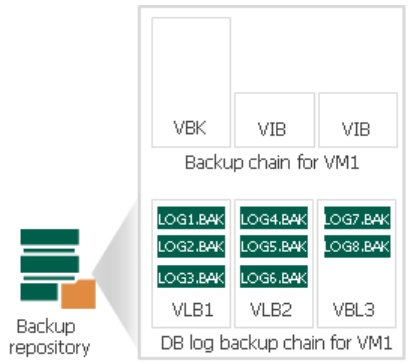
Log Files

At each start of the SQL Server backup job (parent), a new .VLB is created to store log backups in the repository:

- If the **Use per-machine backup files** option is selected for the repository, then Veeam Backup & Replication will create a separate .VLB for each server processed by the job.
- If this option is cleared, then a single .VLB will be created for all servers processed by the job.

For example, if a job processes only one SQL Server, the repository will contain a number of .VLB files for it (a so-called chain).

As described in the section above, during database log backup (child) job session, transaction log backup is performed by native means of the SQL Server and stored as .BAK file to a temporary folder in the SQL Server VM guest file system. Then Veeam Backup & Replication copies .BAK file to the current .VLB in the repository. When the new parent job session starts, another .VLB is created, and the .BAK files that appear after that will be stored there during the child job session. The resulting chain of .VLBs will look like shown below, depicted for a single SQL Server VM1:



Total number of all `LOG<N>.BAK` files stored at the moment in all VLBs is reported as **a number of restore points for the child job** that backs up database logs. So, in the example above, the log backup job for SQL Server VM1 has created 8 restore points by the moment.

In the Veeam Backup & Replication console this number of restore points for the log backup job can be seen in the **Restore Points** column of the preview pane.

Support for AlwaysOn Availability Groups

AlwaysOn Availability Groups allow you to increase fault tolerance between active and hot-standby databases without involving shared physical disks, which is quite important for virtualization of Microsoft SQL Servers. Veeam Backup & Replication supports AlwaysOn Availability Groups for virtualized Microsoft SQL Server 2012 or later.

Image-level Backup of Microsoft SQL Server VMs

During an image-level backup of a Microsoft SQL Server VM, Veeam Backup & Replication requests and analyzes information about databases that are included in the AlwaysOn Availability Groups. Depending on the retrieved information, Veeam Backup & Replication creates a VSS snapshot with or without *COPY_ONLY* flag. The *VSS_BS_COPY* flag for VSS snapshot is triggered if the VM represents a secondary node for at least one AlwaysOn Availability Group.

Veeam Backup & Replication also detects to what cluster the database belongs. If the backup job does not include all VMs from the cluster, an information message will be issued.

Retrieved information is saved for further log identification.

Transaction Log Backup

Transaction log backup can be performed only for those databases that were successfully backed up, either on the primary or on the secondary node of AlwaysOn Availability Group.

The transaction logs processing interval may be the same or may differ through VMs included in AlwaysOn Availability Group. If the interval is different, Veeam Backup & Replication will use minimal value (by default, 15 minutes).

At each log processing interval, Veeam Backup & Replication chooses the AlwaysOn Availability Group node for which transaction logs will be backed up.

Logs are backed up from one node of the AlwaysOn Availability Group. To become a subject for a log backup, the node must meet the following criteria:

- The necessary Veeam Backup & Replication components must be installed on this node and the VM included in AlwaysOn Availability Group must be running. For more information on the necessary components, see How Microsoft SQL Server Log Backup Works.
- The database backup preferences settings must allow a backup of the node that you want to process. For example, if you want to back up the primary node, you must not exclude this node from a backup, or select the **Secondary only** option in the database backup preferences settings.
- Databases in the AlwaysOn Availability Groups for this node were successfully backed up for the last two processing intervals.
- Veeam Backup & Replication can establish a network connection to the node or VIX connection, if a connection over the network cannot be established

NOTE

When you configure a backup job to process Distributed Availability Groups transaction logs, select either primary or secondary distributed availability group. Otherwise, the log chain of the distributed group databases might become inconsistent.

When you configure a backup job to back up transaction logs for other Distributed Availability Groups, use the Perform copy only mode. See Application-Aware Processing to learn more about the copy only mode. You can also use the exclude feature to prevent Guest-OS database from being processed. See Exclude Objects from Backup Job to learn more on excluding objects. To read about distributed availability group limitations, see Configure distributed availability group.

Oracle Log Backup

Veeam Backup & Replication supports backup of Oracle database archived logs and restore of Oracle databases.

Database archived logs are created by the Oracle system. The Oracle database can run in one of the following logging modes:

- ARCHIVELOG turned on — logs are saved and can be used for recovery purposes.
- ARCHIVELOG turned off — no logs are saved. This mode is not recommended as it does not provide proper disaster recovery.

With ARCHIVELOG turned on, the Oracle system stores database archived logs to a certain location on the VM guest OS, as specified by the database administrator. Veeam Backup & Replication allows you to set up the following ways of log handling:

- Instruct the backup job to collect log files from the Oracle VM and ship them to the backup repository where they are stored next to image-level backups of the Oracle VM.
- Skip log processing — log files remain untouched on the Oracle VM and are preserved within the image-level backup.

If you enable application-aware processing for an Oracle VM, during the job session Veeam Backup & Replication installs non-persistent runtime components or uses (if necessary, installs) persistent agent components on this VM to collect information about the database and process archived logs according to job settings. Application-specific settings are configured at the **Guest Processing** step of the backup job wizard — you can specify how logs should be backed up and\or deleted for Oracle databases.

Requirements for Archived Log Backup

- Veeam Backup & Replication supports archived log backup and restore for Oracle Database version 11 and later. Oracle Database may run on a Microsoft Windows VM or Linux VM.
- Automatic Storage Management (ASM) is supported for Oracle Database 11 and later.
- Oracle Database Express Edition (XE) are supported if running on Microsoft Windows machines only.

Archived Log Backup Jobs

To back up archived logs, you must create a backup job, add Oracle VMs to it and specify advanced settings for archived logs backup in the job settings. The resulting job will comprise two jobs:

- Parent backup job — the backup job that creates an image-level backup of the Oracle VM. The parent backup job is named *<job_name>*, for example, *Daily Job*. You can configure the parent job in the Veeam Backup & Replication console just like any other backup job.
- Child job — an archived log backup job. To form a name of the child job, Veeam Backup & Replication adds a suffix to the name of the parent backup job: *<parent_job_name>* + *Oracle Redo Log Backup*, for example, *Daily Job Oracle Redo Log Backup*. Veeam Backup & Replication automatically creates the child job if it detects a backup job that is scheduled to back up at least one Oracle VM, and archived log backup is enabled for this job. Session data of the archived log backup job is stored in the configuration database and displayed in the Veeam Backup & Replication console.

The parent job runs in a regular manner — it starts by schedule or is started manually by the user. The archived log backup job is triggered by the parent backup job. This sequence ensures that the VM (and the database) restore point is present when you need to use archived logs to restore the database.

Sessions of Archived Log Backup Jobs

The archived log backup job runs permanently in the background, shipping archived logs to the backup repository at a specific time interval (by default, every 15 minutes). A sequence of time intervals between sessions of the parent backup job makes up a session of the archived log backup job.

The archived log backup session starts and stops in the following way:

- The initial session starts when the parent backup job schedule is enabled. After that, the session starts with every new session of the parent backup job.
- The session ends before the next session of the parent backup job or when this parent backup job is disabled.
- When the session ends, Veeam Backup & Replication stops the non-persistent runtime components and uninstalls them from the VM guest OS. When a new session starts, the runtime components are deployed again.

How Oracle Archived Log Backup Works

The archived log backup for Oracle VMs is performed in the following way:

1. Veeam Backup & Replication launches the parent backup job by schedule.
2. The parent backup job creates an image-level backup of the Oracle VM and stores this backup to the backup repository.
3. A new session of the archived log backup starts. Veeam Backup & Replication accesses the VM guest OS to perform guest processing, collect database information and handle archived log.

   If Oracle runs on a Microsoft Windows server, Veeam Backup & Replication accesses the VM guest OS over a guest interaction proxy. You can instruct Veeam Backup & Replication to select the guest interaction proxy automatically or assign it explicitly.

   By default, Veeam Backup & Replication accesses the VM guest OS over the network:

   - For Linux VM guest OS — using SSH.
   - For Microsoft Windows VM guest OS — using RPC.

   If a network connection cannot be established, Veeam Backup & Replication accesses the VM guest OS over VIX API/vSphere Web Services.
4. Veeam Backup & Replication deploys the non-persistent runtime components or uses (if necessary, deploys) persistent agent components in the VM guest OS. The components scan the Oracle system and collect information about databases whose logs must be processed, including:

   - List of all databases
   - Database state — a database is on or off, in which logging mode it runs
   - Paths to all database files (configuration logs and so on) and other data required for backup

   Veeam Backup & Replication also detects whether it is possible to store logs to the backup repository through a direct access or a log shipping server is required.

   The non-persistent runtime components or persistent agent components copy archived log files from the log archive destination (set by the Oracle administrator) to a temporary folder on the VM guest file system.

5. Veeam Backup & Replication maps information about the Oracle system collected at step 4 with information kept in the configuration database. This periodic mapping helps reveal databases for which Veeam Backup & Replication must ship archived logs to the backup repository during this time interval.

6. Archived log backup files are transferred from the temporary location on the Oracle VM to the backup repository, either directly or through the log shipping server. to the backup repository and saves them as VLB files The source-side Veeam Data Mover compresses log data to be transferred according to its built-in settings. On the backup repository side, data is compressed according to the parent backup job settings.

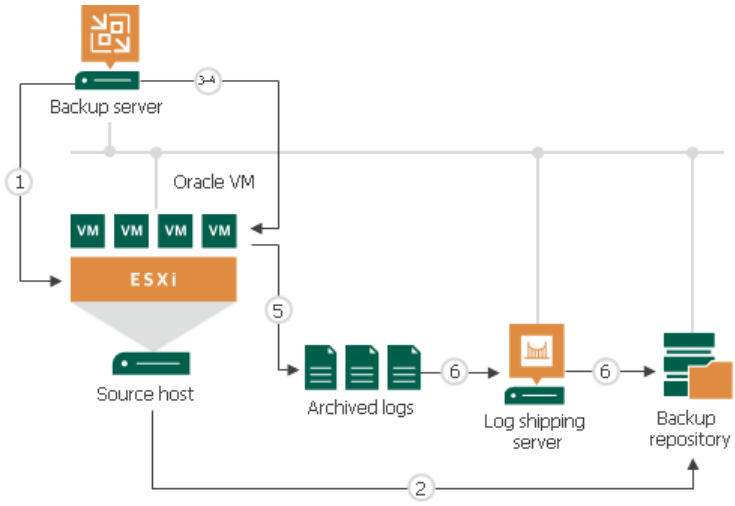
Archived logs that for some reason were not processed during the log backup interval remain in the temporary folder and are processed during the next log backup interval. To detect these remaining logs, Veeam Backup & Replication enumerates log files in the temporary folder.

NOTE

If a new session of the archived log backup starts and the parent backup job has not created a new restore point yet, the archived log backup job will remain in the idle state, waiting for a new restore point to be created.

IMPORTANT

Before backup, Veeam Backup & Replication shuts down databases in the NOARCHIVELOG mode. For details, see the Backing Up a Database in NOARCHIVELOG Mode section in the Oracle Database Backup and Recovery User Guide.



Retention for Archived Log Backup

Archived log backups are stored in files of the proprietary Veeam format — VLB. Veeam Backup & Replication keeps archived log backups together with the VM image-level backup. The target location of VLB files depend on the type of the backup repository:

- If you store the VM image-level backup in a backup repository, Veeam Backup & Replication writes archived log backups to the same folder where files of the image-level backup reside.
- If you store the VM image-level backup in a scale-out backup repository, Veeam Backup & Replication writes archived log backups to the extent where the latest incremental backup file of the VM image-level backup is stored.

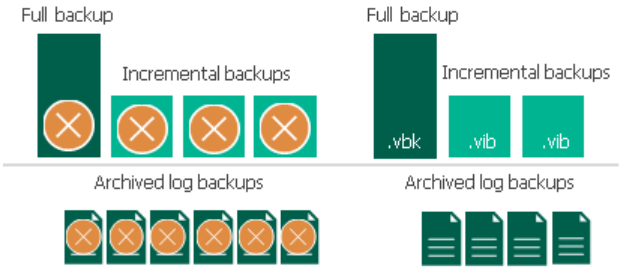
Veeam Backup & Replication removes archived log backups by retention. You can choose one of the following retention methods:

- Retain logs according to the image-level backup
- Retain logs for the specified number of days

Retain Logs with Image-Level Backup

By default, Veeam Backup & Replication retains archived log backups together with the corresponding image-level backup of the Oracle VM. Veeam Backup & Replication retains VM backup and log backups according to the short-term retention configured for VM backups. When Veeam Backup & Replication removes a restore point of the image-level backup from the backup chain, it also removes a chain of archived logs relating to this image-level backup. Note that even if long-term retention is configured for the VM backup, Veeam Backup & Replication retains archived log backups according to the short-term retention policy and deletes them after the short-term retention is exceeded.

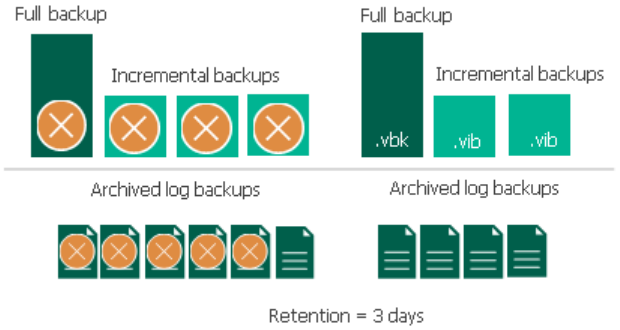
This method allows you to have both the image-level backup and necessary archived log backups at hand. If you need to recover a database to some state, you can restore the Oracle VM from the necessary restore point and use archived logs to bring the database to the desired state.



Retain Logs for a Number of Days

You can instruct Veeam Backup & Replication to keep archived logs only for a specific period of time. This retention setting can be used, for example, if you want to save on storage space and plan to retain archived log backups for the last few days. In this case, you will be able to restore the database only to one of the most recent states.

If you select this retention method, you must make sure that retention policies for the image-level backup and archived log backup are consistent. The restore point of the image-level backup must always be preserved. If a backup of the database itself is missing, you will not be able to use archived logs.



Log Shipping Servers

For every Oracle VM whose archived logs you want to back up, Veeam Backup & Replication defines how to ship logs to the backup repository. Archived logs can be transported in the following ways:

- Directly from the VM guest OS to the backup repository. This method is recommended — it does not involve additional resources and puts less load on the VM guest OS.
- Via log shipping servers. If it is not possible to establish a direct connection between the VM guest OS and backup repository, you can configure Veeam Backup & Replication to use a log shipping server.

A log shipping server is a Microsoft Windows or Linux server added to the backup infrastructure. You can explicitly define what servers you want to use for log shipping or instruct Veeam Backup & Replication to automatically choose an optimal log shipping server. Veeam Backup & Replication chooses the log shipping server based on two criteria: possible data transfer methods and location of the Oracle VM and log shipping server.

Data Transfer Methods

Log shipping servers can transport data in two ways:

- Over the network. In this scenario, Veeam Backup & Replication obtains files from the VM guest OS and transfers them over the network.
- Over VIX API/vSphere Web Services. In this scenario, Veeam Backup & Replication obtains archived logs from the VM guest OS over the VIX API/vSphere Web Services, bypassing the network.

The default method is log shipping over the network.

Location of Log Shipping Server and VMs

When choosing a log shipping server for the job, Veeam Backup & Replication considers the location of the Oracle VM and log shipping server. Veeam Backup & Replication uses the following priority rules to select the log shipping server:

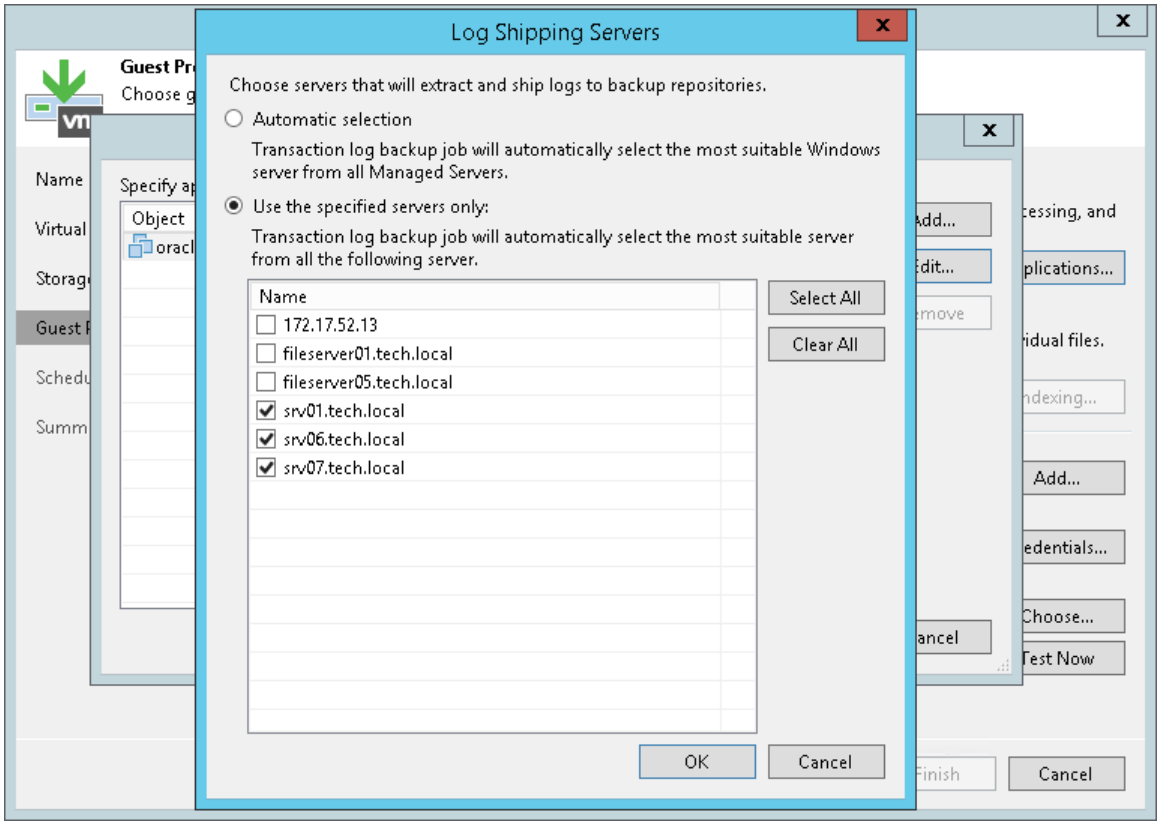
1. Log shipping server is located on the same ESXi host as the Oracle VM.
2. Log shipping server and Oracle VM are located in the same network.
3. Log shipping server and Oracle VM are located in different networks (the production infrastructure is isolated from the backup infrastructure).

That is, when choosing a log shipping server, Veeam Backup & Replication will give the top priority to a VM that is located on the same ESXi host as the Oracle VM and that has a network connection to the Oracle VM

Log shipping servers are assigned per job session. When a new job session starts, Veeam Backup & Replication detects log shipping servers anew. Veeam Backup & Replication can also re-detect available servers during the job session. If a log shipping server becomes unavailable for some reason, Veeam Backup & Replication will fail over to another log shipping server.

IMPORTANT

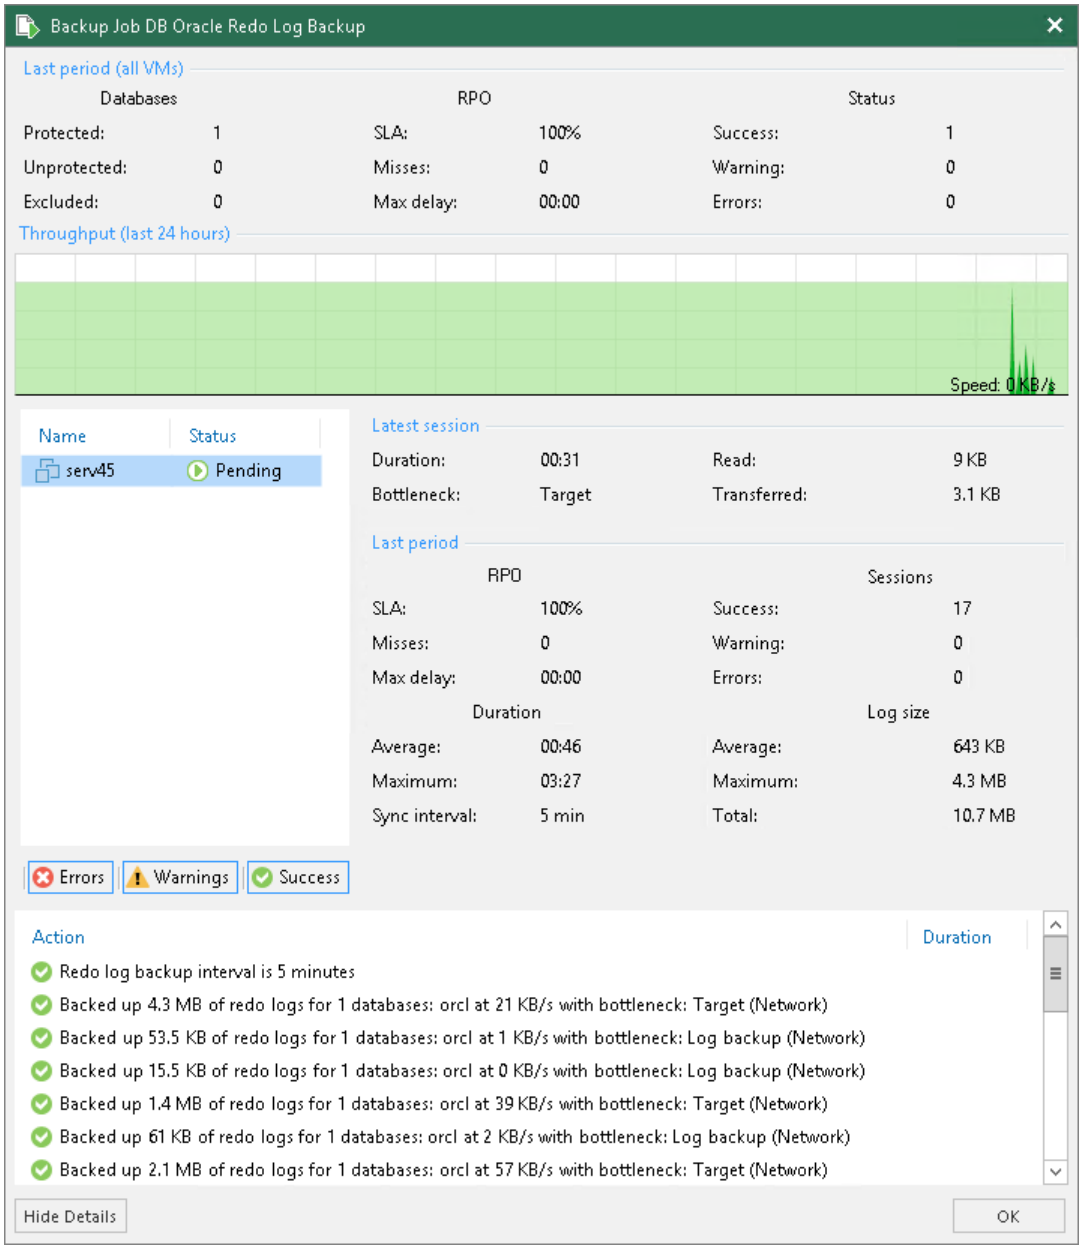
If you do not want to use some servers for archived logs transport, you can manually define what server Veeam Backup & Replication must use as a log shipping server in the job settings. It is recommended that you assign the log shipping server role to a number of servers for availability purposes.



Archived Log Backup Statistics

You can view the statistics of the archived log backup job in the **History** view or in the **Home** view in Veeam Backup & Replication.

In the statistics window, you can examine the overall statistics for the archived log backup job, as well as view per-VM information.



In the upper part of the statistics window, Veeam Backup & Replication displays information about the log backup job for all VMs included in the parent backup job.

The **Last period (all VMs)** section contains statistics data for the selected session of the backup job.

In the **Databases** column, you can view the following information:

- *Protected* — number of databases that were backed up at least once during the last session
- *Unprotected* — number of databases that failed to be backed up during the last session
- *Excluded*— databases excluded from processing. Databases may be excluded for the following reasons: ARCHIVELOG mode is turned off for the database (the database is in NOARCHIVELOG mode), database was deleted after the latest full backup, or database was added to the list of exclusions.

> **NOTE**
>
> Unprotected databases do not comprise **Excluded** databases, as they have different reasons for being non-processed. See also this Veeam KB article for information on processing databases in NOARCHIVE mode.

In the **RPO** column, you can view the following information:

- *SLA* — how many log backup intervals completed in time with successful log backup (calculated as percentage of total number of intervals).
- *Misses* — how many intervals failed to complete in time with successful log backup (number of intervals).
- *Max delay* — difference between the configured log backup interval and time actually required for log backup. If exceeded, a warning is issued.

In the **Status** column, the following information is displayed (per job): number of VMs processed successfully, with warnings or with errors.

The **Latest session** section displays the following information for the latest log processing interval for the selected VM:

- *Duration* — duration of log shipment from the VM guest OS to the backup repository since the current log processing interval has started
- *Bottleneck* — operation with the greatest duration in the last completed interval. The operation may have the following bottlenecks:

Display Name	Slowing-down Operation
Log backup	Saving archived log files to a temporary location on VM guest OS (to work around, see the Veeam KB article: this Veeam KB article)
Network	Uploading log files to the log shipping server
Target	Saving files to the target repository

- *Read* — amount of data read from the temporary folder on VM guest OS
- *Transferred* — amount of data transferred to the target repository

The **Last period** section displays the following statistics of log backups per VM for the latest session of the log backup job:

- The **RPO** column displays statistics on log processing interval (calculated as described above)
- The **Sessions** column includes statistics of log backups per VM, calculated as follows:
  - *Success* — number of intervals when all database logs were backed up successfully
  - *Warning* — number of sequential intervals with failed log processing (if not more than 4 intervals in a sequence)
  - *Errors* — number of sequential intervals with failed log processing (more than 4 intervals in a sequence)

- The **Duration** column includes the following information:
  - *Average* — average duration of log data transfer (through all intervals in the session)
  - *Maximum* — maximal duration of log data transfer (through all intervals in the session)
  - *Sync interval* — duration of periodic intervals specified for log backup in the parent job settings (default is 15 min)
- The Log size column displays the following information:
  - *Average* — average amount of data read from the VM guest OS through all intervals
  - *Maximum* — maximal amount of data read from the VM guest OS over all 15-min intervals
  - *Total* — total amount of data written to the backup repository

The pane below shows all actions performed during the job run. To filter out actions with the certain status, use the **Errors**, **Warnings** and **Success** buttons.

> **NOTE**
>
> Statistics on archived log processing is updated periodically, simultaneously for the VM backup job (parent) and archived log backup job (child job).

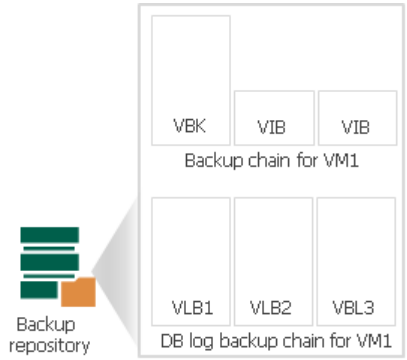
Log Files

At each start of the Oracle backup job ('parent'), a new .VLB is created to store log backups in the repository:

- If the **Use per-machine backup files** option is selected for the repository, then Veeam will create a separate .VLB for each server processed by the job.
- If this option is cleared, then a single .VLB will be created for all servers processed by the job.

For example, if a job processes only one Oracle server, the repository will contain a number of .VLB files for it (a so-called chain).

As described in the section above, during database log backup ('child') job session, log archiving is performed by native means of the Oracle server. Archived logs are stored to a temporary folder on the Oracle VM guest file system. Then Veeam copies archived log to the current .VLB in the repository. When the new 'parent' job session starts, another .VLB is created, and the archived log files that appear after that will be stored there during the 'child' job session. The resulting chain of .VLBs will look like shown below, depicted for a single Oracle VM1:



Total number of all archived logs files stored at the moment in all VLBs is reported as **a number of restore points for the 'child' job** that backs up database logs. So, in the example above, the log backup job for Oracle VM1 has created 8 restore points by the moment.

Backup Job Scheduling

You can start backup jobs manually or schedule them to start automatically at specific time. Veeam Backup & Replication lets you configure the following settings for the job:

- Scheduling settings
- Job retry settings
- Backup window settings

Automatic Startup Schedule

To run a job periodically without user intervention, you can schedule the job to start automatically. The Veeam Backup Service running on the backup server continuously checks configuration settings of all jobs configured on the backup server, and starts them according to their schedule.

Veeam Backup & Replication lets you configure the following scheduling settings for jobs:

- You can schedule jobs to run at specific time every day or on selected days
- You can schedule jobs to run periodically at specific time intervals
- You can schedule jobs to run continuously
- You can chain jobs

Jobs Started at Specific Time

You can schedule jobs to start at specific time daily, on specific week days or monthly on selected days.

This type of schedule requires that you define the exact time when the job must be started. For example, you can configure the job to start daily at 10:00 PM or every first Sunday of the month at 12:00 AM.



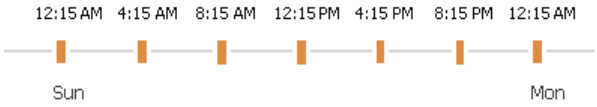
Jobs Started at Specific Time Intervals

You can schedule jobs to start periodically throughout a day at a specific time interval. The time interval between job sessions can be defined in minutes or hours. For example, you can configure a job to start every 30 minutes or every 2 hours.

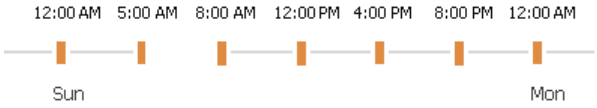
For periodically run jobs, reference time is midnight (12:00 AM). Veeam Backup & Replication always starts counting defined intervals from 12:00 AM, and the first job session will start at 12:00 AM. For example, if you configure a job to run with a 4-hour interval, the job will start at 12:00 AM, 4:00 AM, 8:00 AM, 12:00 PM, 4:00 PM and so on.



If necessary, you can specify an offset for periodically run jobs. The offset is an exact time within an hour when the job must start. For example, you can configure the job to start with a 4-hour interval and specify offset equal to 15 minutes. In this case, the job will start at 12:15 AM, 4:15 AM, 8:15 AM, 12:15 PM, 4:15 PM and so on.



If a session of a periodically run job does not fit into the specified time interval and overlaps the next planned job session, Veeam Backup & Replication starts the next backup job session at the nearest scheduled interval. For example, you set up a job to run with a 4-hour interval. The first job session starts at 12:00 AM, takes 5 hours and completes at 5:00 AM. In this case, Veeam Backup & Replication will start a new job session at 8:00 AM.



Jobs Run Continuously

You can schedule the job to run continuously — that is, in a non-stop manner. A new session of a continuously running job starts as soon as the previous job session completes. Continuously run jobs can help you implement near-continuous data protection (near-CDP) for the most critical applications installed on VMs.



Chained Jobs

In the common practice, data protection jobs configured in the virtual environment start one after another: when job *A* finishes, job *B* starts and so on. You can create a chain of jobs using scheduling settings. To do this, you must define the start time for the first job in the chain. For other jobs in the chain, you must select the **After this job** option and choose the preceding job from the list.

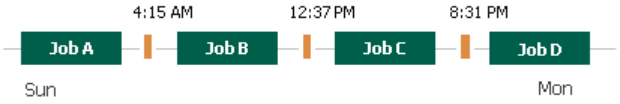
Job chaining is not limited to jobs of specific type only. You can create a chain of jobs of different types. For example, you can:

1. Set a backup job as the first job in the chain.
2. Configure a SureBackup job and chain with the backup job. In this case, Veeam Backup & Replication will automatically verify a backup file created with the backup job after the backup job completes.

NOTE

Mind the following:

- When you start the initial job manually, Veeam Backup & Replication does not start chained jobs in the *Disabled* state. If at least one job chained directly to the initial job is in the *Enabled* state, Veeam Backup & Replication offers you to start the chained jobs as well. Click **Yes** to start the whole job chain or **No** to start only the first job in the chain.
- If you start the initial job manually and chain another job while the initial job is running, the chained job will not start when the initial job completes.
- If a job in the chain fails or is canceled by a user, Veeam Backup & Replication still starts the next chained jobs. Note that if the failed or canceled job was started by a schedule, Veeam Backup & Replication will start the chained job only after all job retries.



Recommendations on Job Chaining

You should use job chaining wisely. Job chaining removes guesswork from job scheduling but has a number of drawbacks:

- You cannot predict precisely how much time the initial job will require and when jobs chained to it will start. Depending on the situation, the job schedule may shift, and some operations may even not be performed as planned.

  For example, you configure 2 jobs:

  - *Job 1* is scheduled to start at 10:00 PM daily and typically takes 1 hour.
  - *Job 2* is scheduled to start after *Job 1* daily. Synthetic full backup is scheduled on Saturday.

  Imagine that *Job 1* starts on Saturday and runs for 2.5 hours instead of 1 hour. *Job 2* will then start after midnight on Sunday, and the synthetic full backup planned on Saturday will not be created.

- Errors in job sessions may cause the job schedule to shift. For example, if the initial job in the chain fails, Veeam Backup & Replication will attempt to retry it, and the schedule for chained jobs will shift.
- Load on backup infrastructure resources may be not balanced. Some slots on backup proxies and backup repositories may be available but will not be used since jobs are queued to run one by one. And if you use a backup repository that supports multiple I/O streams, its resources will not be used efficiently.

Instead of job chaining, you can balance the load on backup infrastructure components. To do this, you must limit the number of concurrent tasks on backup proxies and backup repositories. For more information, see Limiting the Number of Concurrent Tasks.

Job Retry

You can instruct Veeam Backup & Replication to retry a job several times if the initial job pass fails. By default, Veeam Backup & Replication automatically retries a failed job for 3 times within one job session. If necessary, however, you can define a custom number of retries in the job settings.

Veeam Backup & Replication retries a job only if the previous job session has failed, and one or several VMs in the job have not been processed. Veeam Backup & Replication does not perform a retry if a job session has finished with the *Success* or *Warning* status. During the job retry, Veeam Backup & Replication processes only those VMs that have failed.

IMPORTANT

Veeam Backup & Replication does not perform automatic retry for jobs that were started or stopped manually.

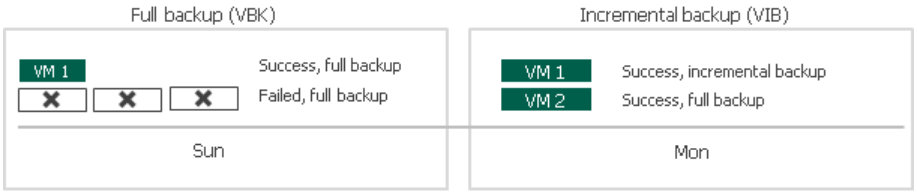
Veeam Backup & Replication always creates one backup file within one job session. If a job processes several VMs and some of them fail to be processed during the first job pass, Veeam Backup & Replication will create a backup file containing data for those VMs that have been successfully processed. During a job retry, Veeam Backup & Replication will attempt to process failed VMs. In case of success, Veeam Backup & Replication will write data of processed VMs to the backup file that was created at the initial job pass.

In some situations, Veeam Backup & Replication may fail to process VMs during all job retries. In this case, failed VMs will be processed during the next job session. Their data will be written to the backup file created within the current job session.

For example, you have configured a job for 2 VMs: *VM 1* and *VM 2*. The job uses the forward incremental method.

During the first job session, Veeam Backup & Replication successfully processed *VM 1* and created a full backup file for it. *VM 2* has failed to be processed during all 3 job retries. In this case, Veeam Backup & Replication will attempt to process the failed *VM 2* within the next job session. Data for *VM 2* will be written to the backup file created within this job session, which will be an incremental backup. As a result, at the end of the second backup job session, you will have 2 files:

- Full backup file containing a full restore point for *VM 1*
- Incremental backup file containing a full restore point for *VM 2* and an incremental restore point for *VM 1*



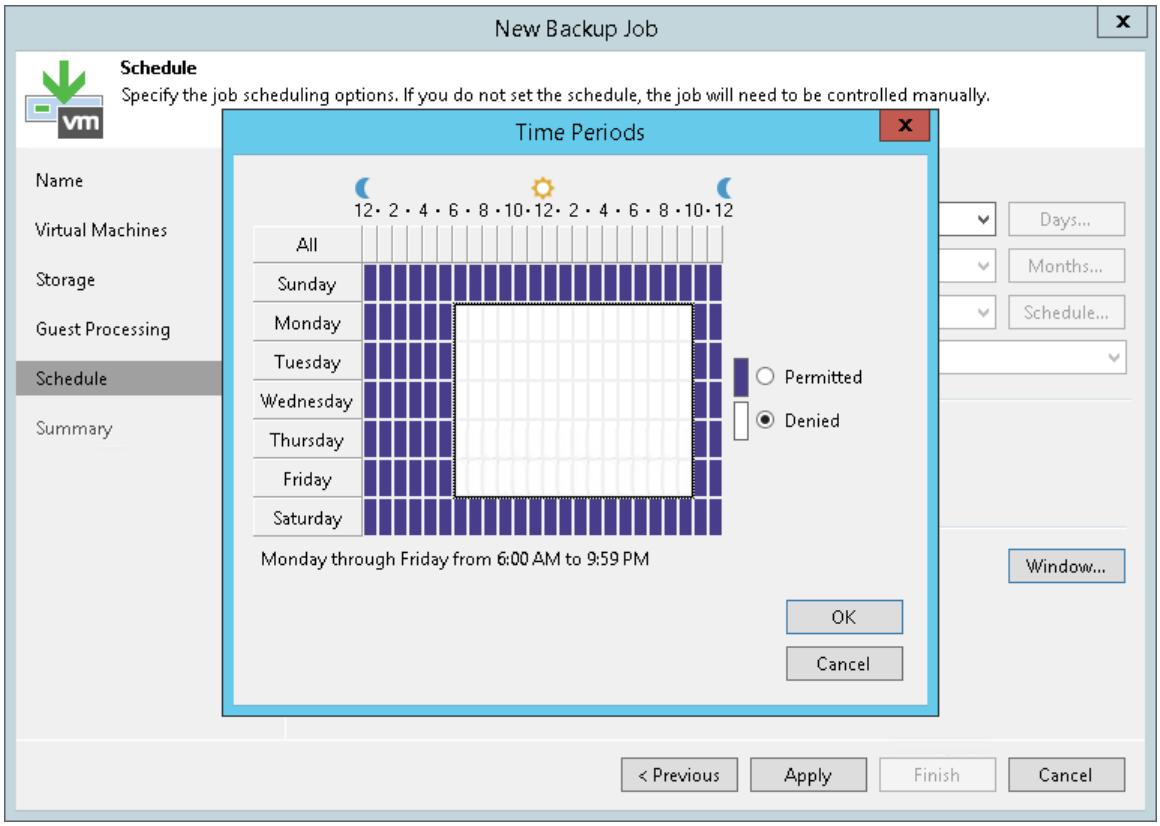
Backup Window

If necessary, you can specify a backup window for jobs. The backup window is a period of time on week days when jobs are permitted to run. If the job exceeds the allowed window, Veeam Backup & Replication will automatically stop this job.

The backup window can be helpful if you do not want data protection jobs to produce unwanted overhead for the production environment or do not want jobs to overlap production hours. In this case, you can define the time interval during which the job must not run.

IMPORTANT

The backup window affects only the data transport process and health check operations. Other transformation processes can be performed in the target repository outside the backup window. Linked jobs that process Microsoft SQL transaction logs and Oracle archived logs are not affected by the backup window settings.



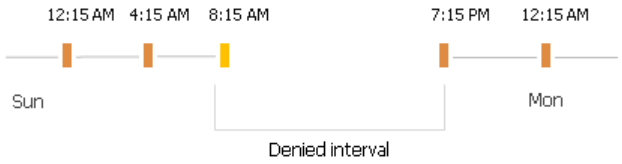
Backup Window for Periodically Run Jobs

If you define the backup window for a job that runs periodically at specific time intervals, Veeam Backup & Replication will immediately start the job after the denied window is over. All subsequent backup job sessions will be performed according to specified scheduling settings.

For example, you have configured a job to run with a 4-hour interval with an offset of 15 minutes. The allowed backup window for the job is 7:00 PM to 8:00 AM. Veeam Backup & Replication will run this job in the following way:

1. The first job session will start at 12:15 AM (since midnight is a reference time for periodically run jobs).
2. The next job session will start at 4:15 AM.
3. The job session at 8:15 AM will not be performed as it falls into the denied period of the backup window.
4. The next job session will start immediately after the denied period is over: at 7:15 PM.

5. After that, Veeam Backup & Replication will run the job by the defined schedule: at 8:15 PM, 12:15 AM and so on.

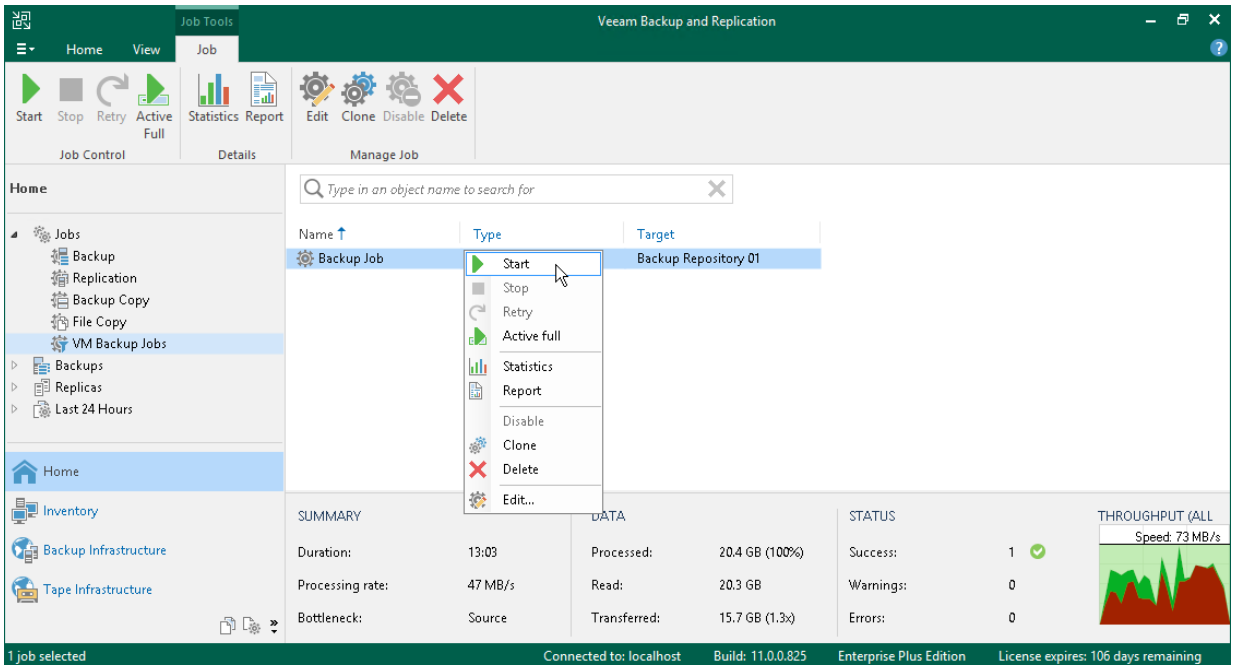


Manual Start of Backup Jobs

You can start jobs manually if you need to capture VM data at a specific point in time and do not want to re-configure job scheduling settings. For example, you can start a job to create a VM backup before you install new software on a VM or enable a new feature.

When you start the job manually, Veeam Backup & Replication runs a regular job session that produces a new restore point in the backup chain in the backup repository.

To start and stop jobs configured on the backup server, you can use the **Start** and **Stop** buttons on the ribbon or corresponding commands in the shortcut menu.



Manual Stop of Backup Jobs

You can stop job execution at any moment of time. For example, you can stop a job if the job processes several VMs but the workload appears to be greater than you expected. Or you can stop the job if there is not enough time to finish the job session.

You can stop a job in 2 ways:

- You can stop the job immediately. In this scenario, Veeam Backup & Replication terminates the job session and does not create a new restore point for VMs that are currently processed.
- You can stop the job gracefully. In this scenario, Veeam Backup & Replication creates a restore point for the VMs that are currently processed and then terminates the job session.

Immediate Stop of Jobs

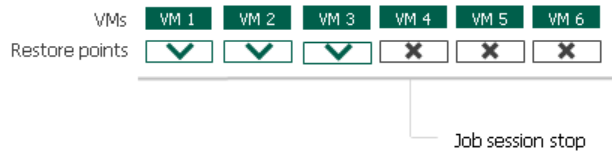
Immediate job stop terminates the job session instantly. The job finishes with the following results:

- VMs that Veeam Backup & Replication has succeeded to process by the time you stop the job will have new restore points.
- VMs that Veeam Backup & Replication is currently processing and VMs that Veeam Backup & Replication has not started to process will not have new restore points.

When you stop a job session immediately, Veeam Backup & Replication performs the following operations:

1. If a snapshot for a VM has already been created, Veeam Backup & Replication instructs VMware vSphere to remove the snapshot.
2. Veeam Backup & Replication terminates all job processes and tasks. The job is finished with the *Failed* error.

All restore points created with the previous job sessions remain untouched. You can use them for restore operations.



Graceful Stop of Jobs

Graceful job stop instructs Veeam Backup & Replication that it must create restore points for VMs that are currently being processed, and then terminate the job. The job finishes with the following results:

- VMs that Veeam Backup & Replication has succeeded to process and VMs that are being processed will have new restore points.
- VMs that Veeam Backup & Replication has not started to process will not have new restore points.

You can use graceful job stop for the following types of jobs:

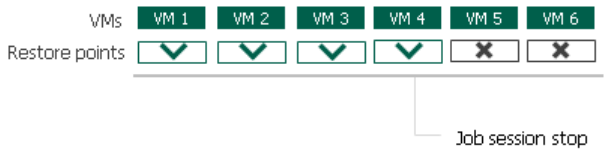
- Backup jobs
- VM copy jobs
- Replication jobs

You cannot use graceful job stop for the following types of jobs:

- File copy jobs
- Backup copy jobs
- Quick migration job (during quick migration, Veeam Backup & Replication processes all VMs in one task)
- Restore operations

VMs added to the job are processed in the order defined in job settings. Information about VMs that have already been processed and VMs that are being processed is displayed in job details.

If you stop the job gracefully before Veeam Backup & Replication starts processing the first VM in the job, the job will be finished with the *Failed* error. You will see the message *Operation was canceled by user* in job details.



Health Check for Backup Files

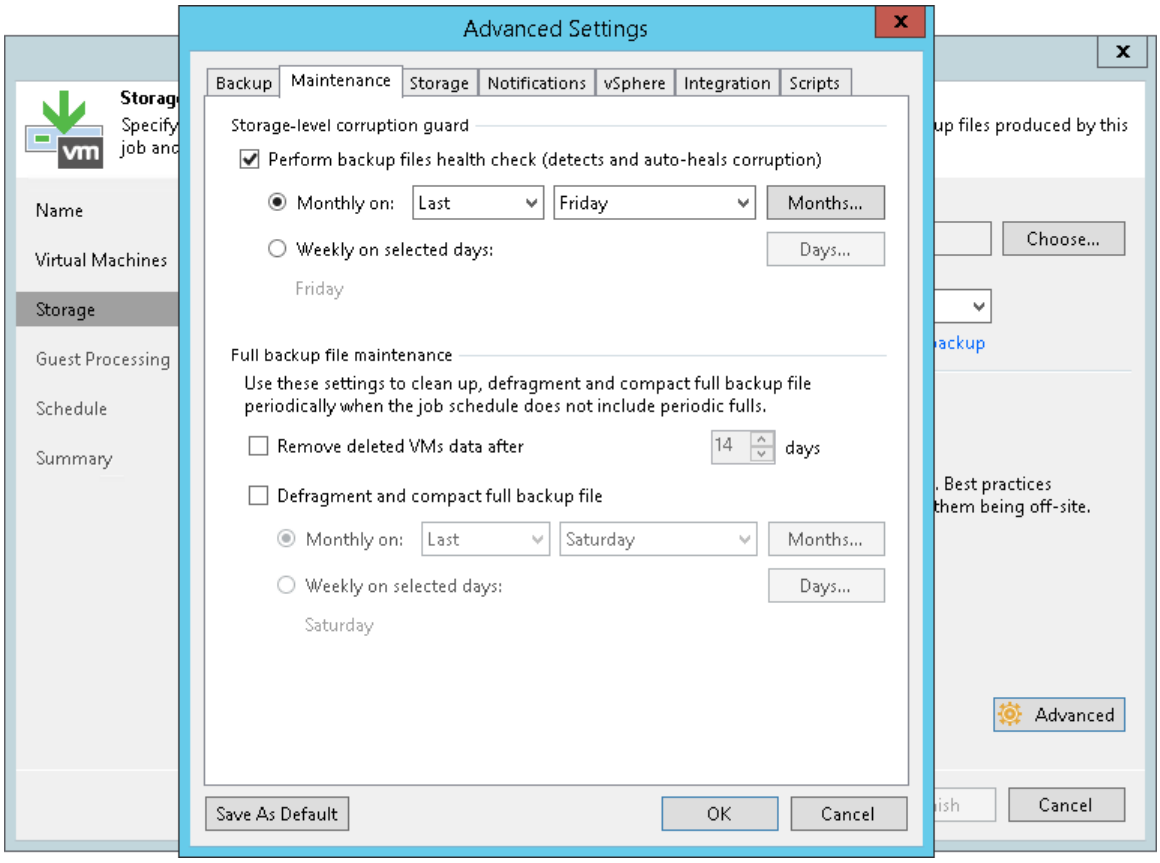
You can instruct Veeam Backup & Replication to periodically perform a health check for the latest restore point in the backup chain. During the health check, Veeam Backup & Replication performs a CRC check for metadata and a hash check for VM data blocks in the backup file to verify their integrity. The health check helps make sure that the restore point is consistent, and you will be able to restore data from this restore point.

The health check can be performed for all types of backup chains:

- Forever forward incremental
- Forward incremental
- Reverse incremental backup chains

To run the health check periodically, you must enable the **Perform backup files health check** option in the backup job settings and define the health check schedule. By default, the health check is performed on the last Friday of every month. You can change the schedule and run the health check weekly or monthly on specific days.

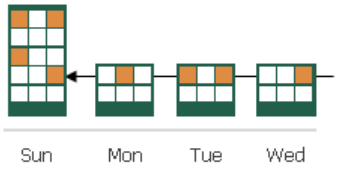
> **NOTE**
>
> Veeam Backup & Replication performs the health check during the first job session on the day when the health check is scheduled. If another job session runs on the same day, Veeam Backup & Replication will not perform the health check during this job session. For example, if the job is scheduled to run several times on Saturday, and the health check is scheduled on Saturday, the health check will only be performed during the first backup job session on Saturday.



Verification Content

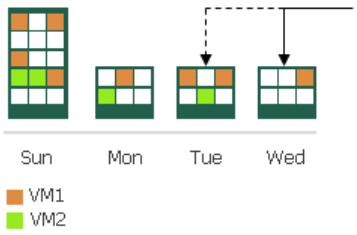
The health check always verifies only the latest restore point in the backup chain. In case of forever forward incremental and forward incremental backup chains, if the latest restore point is incomplete, the health check verifies the restore point preceding the latest one.

Bear in mind that the health check procedure verifies not the latest backup file in the backup chain, but the latest restore point for a VM. The latest restore point corresponds to the state of the VM at the date and time when the latest backup file for this VM was created. Data blocks that are required to "compose" the VM latest state are typically spread out across several backup files in the backup chain. Therefore, to verify the latest state of the VM, Veeam Backup & Replication must open several backup files in the backup chain and read data blocks from these backup files. For this reason, the health check procedure may take long.



The health check verifies only those virtual disks of a VM that are available in the latest restore point. For example, you added a VM with 3 virtual disks to a backup job. The VM was backed up Sunday through Tuesday. On Wednesday, you removed 1 virtual disk, and Veeam Backup & Replication run the health check for the VM. During the health check, Veeam Backup & Replication will verify only the 2 remaining virtual disks.

The health check verifies only those VMs that are available in the latest restore point. For example, you added 2 VMs to a backup job and run the job for some time. The health check verified 2 VMs. If you remove 1 VM from the backup job, the next scheduled health check run will verify the latest unverified restore point for the removed VM, and the latest restore point for the remaining VM. In future, the health check will verify only the restore point for the remaining VM in the job.



Limitations for Health Check

- The health check is not performed during an active full backup job session started manually or automatically by schedule.
- The health check is not performed for offloaded restore points. For more information, see Capacity Tier.
- [For per-machine backup chains] If you add a new VM to an existing backup job that has been run for some time, Veeam Backup & Replication will perform the health check for it during the next incremental backup job session for the added VM.

How Health Check Works

When Veeam Backup & Replication saves a new restore point to the backup repository, it calculates CRC values for backup metadata and hash values for data blocks of VM disk in the backup file, and saves these values in the metadata of the backup file, together with VM data. During the health check session,
Veeam Backup & Replication uses these values to make sure that a verified restore point is consistent.

> **NOTE**
>
> If you perform health check for encrypted backup files, Veeam Backup & Replication will pass encryption keys to the regular backup repository or cloud repository. For more information on encryption, see Data Encryption.

Veeam Backup & Replication uses different mechanisms of health check for different types of backup chains:

- Forever forward incremental and forward incremental backup chains
- Reverse incremental backup chains

Forever Forward Incremental and Forward Incremental Backup Chains

The health check for forward incremental backup chains is performed in the following way:

1. At the end of the backup job session, Veeam Backup & Replication performs the health check. It calculates CRC values for backup metadata and hash values for VM disks data blocks in the backup file and compares them with the CRC and hash values that are already stored in the backup file.

   During the health check, Veeam Backup & Replication verifies the latest restore point in the backup chain (restore point created with the current backup job session — the session during which the health check is performed). If the latest restore point in the backup chain is incomplete, Veeam Backup & Replication checks the restore point preceding the latest one.

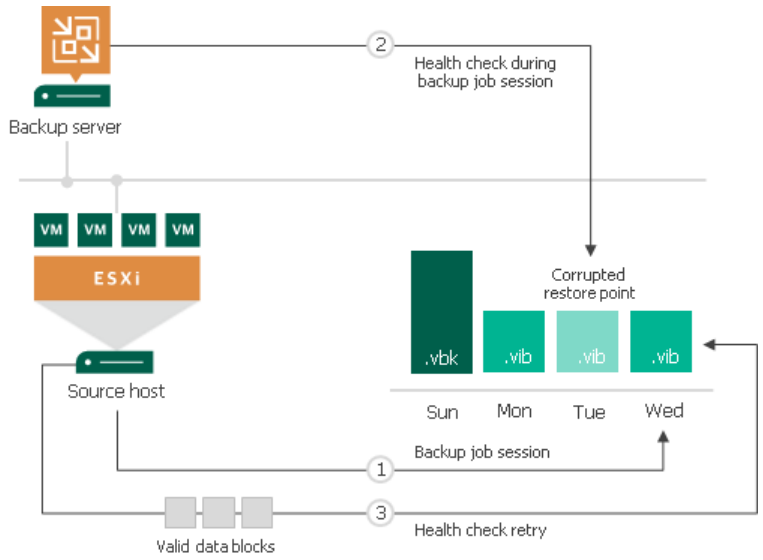
2. If the health check does not detect data corruption, the backup job session completes in a regular way.

   If the health check detects corrupted data, Veeam Backup & Replication completes the backup job with the *Error* status and starts the health check retry process. The health check retry starts as a separate backup job session.

   Depending on the revealed data corruption, Veeam Backup & Replication performs the following actions:

   - If the health check has detected corrupted backup metadata in the full backup file, Veeam Backup & Replication marks the backup chain starting from this full restore point as corrupted in the configuration database. During the health check retry, Veeam Backup & Replication transports data blocks of the whole VM image from the source datastore, creates a new full backup file in the backup repository and saves transported data blocks to it.
   - If the health check has detected corrupted backup metadata in the incremental backup file, Veeam Backup & Replication removes information about this incremental restore point and subsequent incremental restore points from the configuration database. During the health check retry, Veeam Backup & Replication transports incremental data relatively the latest valid restore point in the backup chain from the source datastore, creates a new incremental backup file in the backup repository and saves transported data blocks to it.

- If the health check has detected corrupted VM disk blocks in the full or incremental backup file, Veeam Backup & Replication marks the restore point that includes the corrupted data blocks and subsequent incremental restore points as corrupted in the configuration database. During the health check, Veeam Backup & Replication transports data blocks from the source datastore. In addition, Veeam Backup & Replication transports data blocks that have changed since the backup job session that has triggered the health check. Veeam Backup & Replication stores these data blocks to the latest restore point that has been created with the current backup job session (session that has triggered the health check retry).



Reverse Incremental Backup Chains

In case of reverse incremental backup chains, the health check always verifies only the latest restore point in the backup chain, which is always a full backup file.

The health check is performed in the following way:

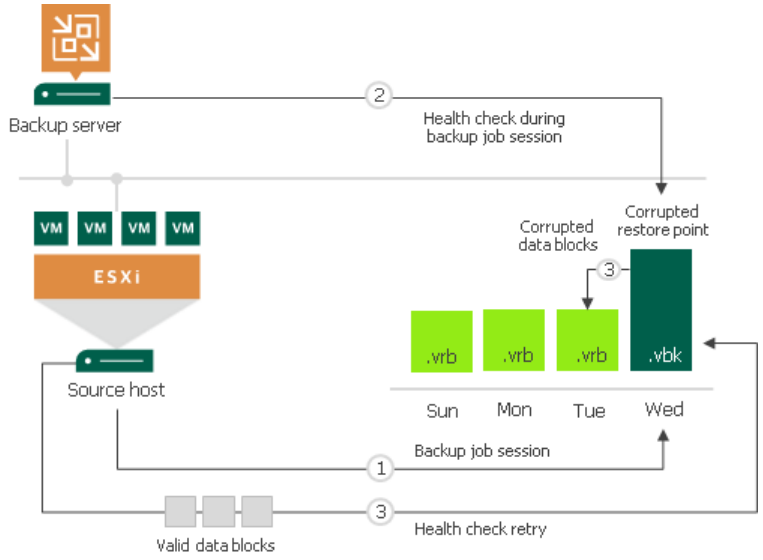
1. At the end of the backup job session, Veeam Backup & Replication verifies the full backup file. Veeam Backup & Replication calculates CRC values for backup metadata and hash values for VM disks data blocks in the full backup file, and compares them with the CRC and hash values that are already stored in the full backup file.
2. If the health check does not detect data corruption, the backup job session completes in a regular way.

   If the health check detects corrupted data, Veeam Backup & Replication completes the backup job with the *Error* status and starts the health check retry process. The health check retry starts as a separate backup job session.

   Depending on the revealed data corruption, Veeam Backup & Replication performs the following actions:

   - If the health check has detected corrupted backup metadata in the full backup file, Veeam Backup & Replication marks the whole backup chain (full backup file and preceding reverse incremental backup files) as corrupted in the configuration database. During the health check retry, Veeam Backup & Replication transports data blocks of the whole VM image from the source datastore, creates a new full backup file in the backup repository and saves transported data blocks to it.

- If the health check has detected corrupted VM disk blocks in the full backup file, Veeam Backup & Replication marks the full backup file and preceding reverse incremental backup files as corrupted in the configuration database. During the health check retry, Veeam Backup & Replication transports data blocks from the source datastore. In addition, Veeam Backup & Replication transports data blocks that have changed since the backup job session that has triggered the health check. Veeam Backup & Replication stores these data blocks to the existing full backup file in the backup repository. Corrupted data blocks that have been replaced with data blocks from the source datastore are stored to an existing reverse incremental backup file preceding the full backup file.



Health Check Retries

The health check itself is started during the backup job session or the job retry session if the backup job session has failed. If attempts are not successful, Veeam Backup & Replication performs the health check during the last job retry in any case.

If the health check detects corrupted data, Veeam Backup & Replication completes the backup job with the *Error* status and starts the health check retry process. The health check retry starts as a separate backup job session. During the health check retry, Veeam Backup & Replication attempts to transport data blocks for the corrupted restore point from the source datastore.

For scheduled jobs, the number of health check retries is equal to the number of job retries specified in the job settings. For jobs started manually, Veeam Backup & Replication performs 1 health check retry.

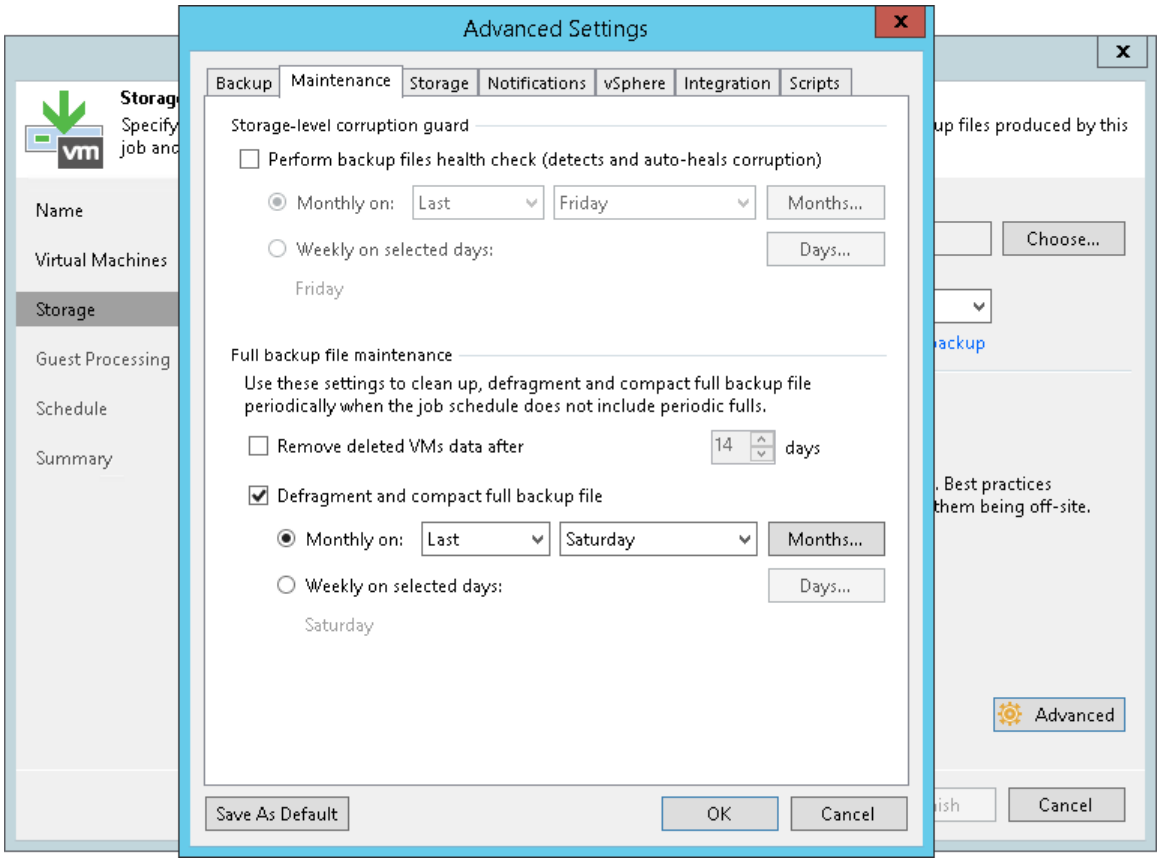
> **NOTE**
>
> If Veeam Backup & Replication fails to fix the corrupted data during all health check retries, you must retry the job manually. In this case, Veeam Backup & Replication will transport the required data blocks from the source datastore to fix the latest restore point. If the latest restore point in the backup chain is incomplete, Veeam Backup & Replication will attempt to fix the restore point preceding the latest one.

Compact of Full Backup File

If you use a forever forward incremental or reverse incremental backup method, the backup job constantly transforms the full backup file in the backup chain to meet retention policy settings. The transformation process, however, has a side effect. In the long run, the full backup file grows large and gets fragmented. The file data occurs to be written to non-contiguous clusters on disk, and operations of reading and writing data from and to the backup file slow down.

To resolve the fragmentation problem, you can instruct Veeam Backup & Replication to compact the full backup file periodically. During the file compact operation, Veeam Backup & Replication creates a new empty file and copies to it data blocks from the full backup file. As a result, the full backup file gets defragmented and the speed of reading and writing from and to the file increases.

To compact the full backup file periodically, you must enable the **Defragment and compact full backup file** option in the backup job settings and define the compact operation schedule. By default, the compact operation is performed on the last Saturday of every month. You can change the compact operation schedule and instruct Veeam Backup & Replication to perform it weekly or monthly on specific days.



Limitations for Full Backup File Compact

The full backup file compact has the following limitations:

- The **Defragment and compact full backup file** option works for forever forward incremental or reverse incremental backup chains. For this reason, you must not schedule active or synthetic full backups.

  Although you do not schedule active full backups for forever forward incremental or reverse incremental backup chains, full backups can be created. For example, you can create them manually or Veeam Backup & Replication can create them during the health check. On the day when active full backups are triggered, Veeam Backup & Replication does not create compact full backups. Veeam Backup & Replication will create them on another day during the backup job session.

- Veeam Backup & Replication does not compact full backup files that have been offloaded to cloud-based object storage. For more information, see Capacity Tier.
- The backup repository must have enough space to store a file of the full backup size. During the compact process, Veeam Backup & Replication creates auxiliary files that exist in the backup repository until the end of the compact operation.
- [For per-machine backup chains] If you add a new VM to an existing backup job that has been run for some time, Veeam Backup & Replication will perform the compact full operation for it during the next incremental backup job session for the added VM.
- If you change the block size in backup job settings, Veeam Backup & Replication does not change the block size in the compacted backup file till the next full backup. However, if you change compression settings in backup job settings, during the next compact file operation Veeam Backup & Replication changes the compression level for the compacted backup file.

Removal of Deleted VMs Data

During the compact operation, Veeam Backup & Replication does not copy all data blocks from the VBK file to the newly created file. It copies only data blocks of VMs whose information is stored in the configuration database. For example, if the VM is removed from the backup job, its data is not copied to the new full backup file. This approach helps reduce the size of the full backup file and remove unnecessary data from it.

VM Data Take Out

If the full backup file contains data for a VM that has only one restore point and this restore point is older than 7 days, during the compact operation Veeam Backup & Replication will extract data for this VM from the full backup file and write this data to a separate full backup file. Such backup will be displayed under the **Backups** > **Disk (Imported)** node in the **Home** view.

The mechanism works if the following conditions are met:

- The **Remove deleted VMs data** option is not enabled in the backup job settings.
- The **Use per-machine backup files** option is not enabled in backup repository settings.

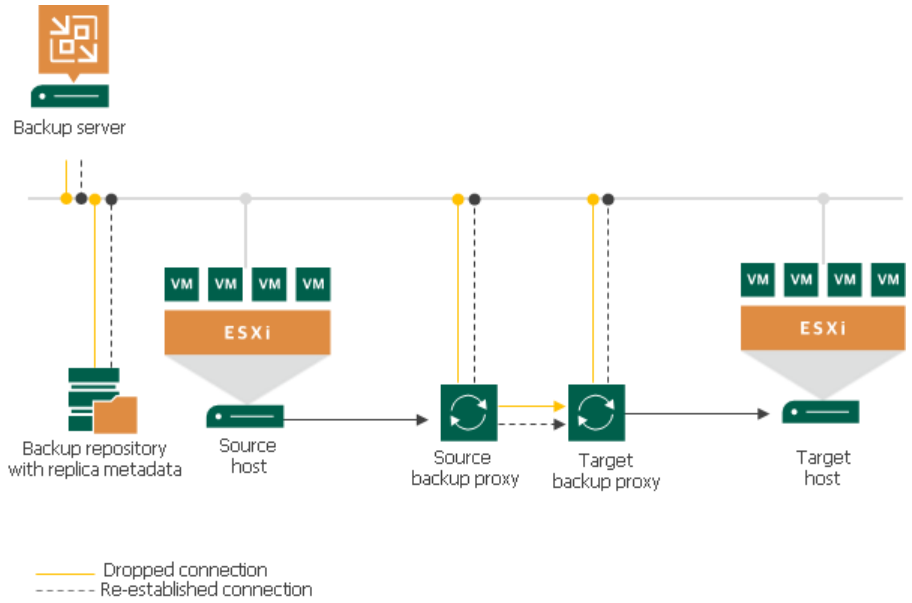
Resume on Disconnect

Veeam Backup & Replication can handle a situation of an unstable network during backup, backup copy and replication jobs. If a network connection drops for a short period of time during the data transport process, Veeam Backup & Replication automatically resumes the dropped network connection. The data transfer process starts from the point when the connection was lost. The resume on disconnect capability improves the reliability of remote data transfer, reduces the backup window and minimizes the network load.

Veeam Backup & Replication automatically re-establishes a connection between the following backup infrastructure components engaged in the data transfer process:

- Backup server
- Backup proxies
- Backup repository

Resume on disconnect works only for dropped network connections. Veeam Backup & Replication attempts to resume the connection with an interval of 15 seconds during 30 minutes. If the problem has any other nature, Veeam Backup & Replication retries the job in a regular manner.



Veeam Backup & Replication does not create a new restore point on resume: VM data is written to the same restore point that was created for the current job session. When resuming the data transfer process, Veeam Backup & Replication regards VM disks, not the whole VM.

For example, a VM has two disks: *disk A* and *disk B*. Before the connection dropped, Veeam Backup & Replication managed to transfer 20 GB of *disk A* and did not start transferring *disk B*. After the connection is re-established, Veeam Backup & Replication will start transferring the data for *disk A* from the 20 GB point; data of the whole *disk B* will be transferred anew.

Snapshot Hunter

The Snapshot Hunter is a Veeam technology used to detect and remove orphaned snapshots that may remain after backup or replication job sessions.

The Snapshot Hunter addresses the problem of "phantom" snapshots. Under some circumstances, VMware vSphere can report a successful removal of a snapshot but the snapshot actually remains on the datastore.

Phantom snapshots can take substantial space on the datastore or impact VM performance. They can even cause the production VMs to stop if the datastore runs out of free space.

To solve the problem of phantom snapshots, Veeam Backup & Replication starts the Snapshot Hunter during each backup or replication job session. The Snapshot Hunter looks for snapshot files not registered in vSphere. If there are no orphaned files, the Snapshot Hunter stops. If orphaned snapshot files are detected, the Snapshot Hunter removes them in the background mode.

The Snapshot Hunter runs in jobs that use VMware VM snapshots:

- Backup jobs: regular backup and backup from storage snapshots
- Replication jobs (the source VM snapshot): regular replication, replication from storage snapshots
- VeeamZIP

> **NOTE**
>
> During Snapshot Hunter analysis, Veeam Backup & Replication skips vCloud Director VMs.

How Snapshot Hunter Works

A temporary snapshot of the VM is taken and then removed during every backup or replication job session. To remove the snapshot, Veeam Backup & Replication triggers the VMware snapshot consolidation mechanism that includes two steps:

1. VMware vSphere removes the snapshot from the VM snapshots list.
2. VMware vSphere consolidates the data written to the delta file with the VM disks.

The problem occurs when the snapshot was removed successfully but the consolidation failed. This may happen, for example, if the files appear to be locked when VMware vSphere attempts to consolidate the snapshot files. In this case, the files remain on datastore.

The Snapshot Hunter is started as a separate process scheduled within every job session. The discovery of the phantom snapshots does not affect the job: if the phantom snapshots are discovered, the Veeam Backup Service schedules the snapshot consolidation, and the job runs in normal way.

Veeam Backup & Replication checks the datastore to discover orphaned snapshot files. To consolidate these files with the VM disks, Veeam Backup & Replication calls a consolidation algorithm. The algorithm consists of three steps, each representing a VMware method.

1. VMware Consolidate method

   As a first attempt, Veeam Backup & Replication calls the VMware Snapshot Consolidate method. This method is the same mechanism that VMware vSphere uses for VMs with the *Needs Consolidation* status.

2. Hard consolidation without quiesce

   If the first attempt fails, Veeam Backup & Replication creates a new snapshot and calls the VMware *Delete all snapshots* method. As a result, all VM snapshots and associated files are deleted. The snapshot is taken without quiescing the VM.

3. Hard consolidation with quiesce

   If the snapshot deletion still fails, Veeam Backup & Replication implies another VMware method that creates a quiesced snapshot and then removes all VM snapshots.

> **NOTE**
>
> Hard consolidation without quiesce and hard consolidation with quiesce are performed only if the VM does not have any user snapshots. In case there are one or more user snapshots, these steps will not be performed.

The 3-steps consolidation procedure is launched up to 4 times with a 4-hour interval.

In case all four attempts fail, Veeam Backup & Replication sends an e-mail notification informing that the user needs to manually troubleshoot the problem. Note that you need to have the global email notifications option enabled. For more information, see Specifying Email Notification Settings.

The Snapshot Hunter considers the backup window set for the job. If any of the attempts does not fit the backup window, Veeam Backup & Replication will not perform the consolidation and send the e-mail notification.

To view information on the Snapshot Hunter sessions, in the Veeam Backup & Replication console, open the **History** view and select **System** in the inventory pane.

In case no consolidation attempt could fit the backup window, the warning appears in the job statistics.

Creating Backup Jobs

To back up VMs, you must configure a backup job. The backup job defines how, where and when to back up VM data. One job can be used to process one or more VMs. Jobs can be started manually or scheduled to run automatically at a specific time.

Before you create a backup job, check prerequisites. Then use the **New Backup Job** wizard to configure the backup job.

Before You Begin

Before you create a backup job, check the following prerequisites:

- Backup infrastructure components that will take part in the backup process must be added to the backup infrastructure and properly configured. These include ESXi hosts on which VMs are registered, backup proxy and backup repository.
- The backup repository must have enough free space to store created backup files. To receive alerts about low space in the backup repository, configure global notification settings. For more information, see Specifying Other Notification Settings.
- For VM guest OS indexing on Linux-based VMs, a user account with root privileges on the VM is required. It is recommended that you create a separate user account for work with Veeam Backup & Replication on the Linux-based VM, grant root privileges to this account and specify settings of this account at the **Guest Processing** step of the **New Backup Job** wizard.
- To perform guest processing for Oracle databases on Linux servers, make sure that the `/tmp` directory is mounted with the `exec` option. Otherwise, you will get an error with the permission denial.
- If you plan to map a backup job to a backup file that already exists in the backup repository, you must perform the rescan operations for this backup repository. Otherwise, Veeam Backup & Replication will not be able to recognize backup files in the backup repository. For more information, see Rescanning Backup Repositories.
- If you plan to use an HPE StoreOnce repository as a target for backup jobs, check limitations and requirements for the repository.
- If you plan to configure a secondary destination for the backup job, you can create a backup copy job or backup to tape job beforehand. While creating the backup copy or backup to tape job, you can leave the source empty and link the source to the job later. For more information, see the Creating Backup Copy Jobs and Creating Backup to Tape Jobs sections in the Veeam Backup & Replication User Guide.
- If you plan to use pre-job and post-job scripts and/or pre-freeze and post-thaw scripts, you must create scripts before you configure the backup job.
- To back up Microsoft SQL transaction logs with Veeam Backup & Replication, you must make sure that the recovery model is set to *Full* or *Bulk-logged* recovery model for required databases on Microsoft SQL Server VMs. If the recovery model is set to *Simple*, Veeam Backup & Replication will not detect and process transaction logs on Microsoft SQL Server VMs.
- Veeam Backup & Replication excludes from application-aware processing Microsoft SQL databases that are mounted to the Microsoft SQL Server using a remote UNC path. If at least one file of the database is located on a network shared folder, this database will be backed up in the crash-consistent state. Other databases on this server will be backed up in the transactionally consistent state. For more information, see this Veeam KB article.
- By default, system databases (master, model, msdb) are skipped from transaction log processing and are not a part of the Veeam Explorer for Microsoft SQL Server restore workflow. To recover these databases, you can use file-level restore.

  If you want to exclude other databases from the transaction log processing workflow, see this Veeam KB article. Consider that exclusion configured this way will be treated as a global setting.
- Transaction log backups cannot be copied to capacity tier.
- To back up Oracle transaction logs with Veeam Backup & Replication, you must make sure that ARCHIVELOG is turned on for required databases on Oracle VMs. If ARCHIVELOG is turned off, Veeam Backup & Replication will not detect and process transaction logs on Oracle VMs.

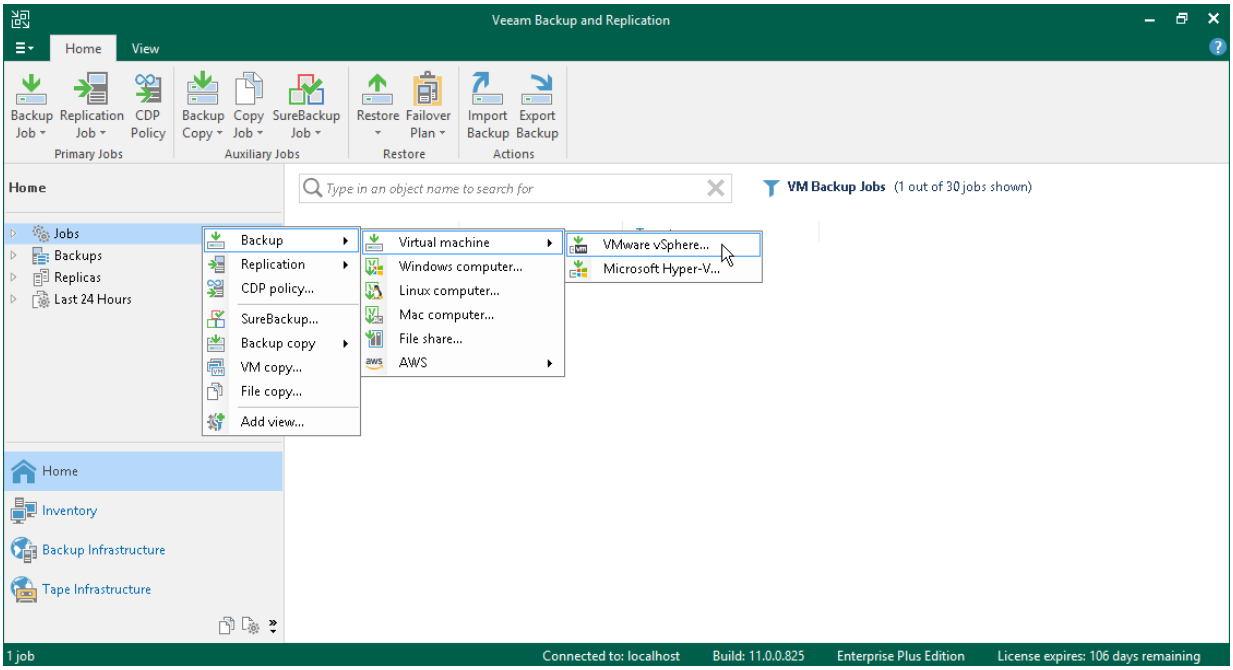
Mind the following:

- If you plan to periodically perform maintenance operations with backup files, mind the following limitations: Health Check for Backup Files, Retention Policy for Deleted VMs, Compact of Full Backup File.
- Due to Microsoft limitations, you cannot use Microsoft Azure Active Directory credentials to perform guest processing on VMs running Microsoft Windows 10 (or later).
- [For Dell EMC Data Domain backup repository] The length of forward incremental and forever forward incremental backup chains that contain one full backup and a set of subsequent incremental backups cannot be greater than 60 restore points. To overcome this limitation, schedule full backups (active or synthetic) to split the backup chain into shorter series. For example, to perform backups at 30-minute intervals, 24 hours a day, you must schedule synthetic fulls every day. In this scenario, intervals immediately after midnight may be skipped due to the duration of synthetic processing. For more information, see How Synthetic Full Backup Works.
- If you assign the role of a backup proxy to a VM, you should not add this VM to the list of processed VMs in a job that uses this backup proxy. Such configuration may result in degraded job performance. Veeam Backup & Replication will assign this backup proxy to process other VMs in the job first, and processing of this VM itself will be put on hold. Veeam Backup & Replication will report the following message in the job statistics: *VM is a backup proxy, waiting for it to stop processing tasks*. The job will start processing this VM only after the backup proxy deployed on the VM finishes its tasks.
- If you use tags to categorize virtual infrastructure objects, check limitations for VM tags. For more information, see VM Tags.
- Veeam Backup & Replication supports backup of Microsoft Exchange, SharePoint and SQL Server databases existing in mount point volumes.
- If a job is unable to complete within 21 days period, it will be stopped with the *Failed* status.

Step 1. Launch New Backup Job Wizard

To launch the **New Backup Job** wizard, do one of the following:

- On the **Home** tab, click **Backup Job** > **Virtual machine** > **VMware vSphere**.
- Open the **Home** view. In the inventory pane, right-click **Jobs** and select **Backup** > **Virtual machine** > **VMware vSphere**.
- Open the **Inventory** view. In the working area select the VMs, click **Add to Backup** on the ribbon and select **New job** or right-click the VMs and select **Add to backup job** > **New job**. Veeam Backup & Replication will start the **New Backup Job** wizard and add the VMs to this job. You can add other VMs to the job later on, when you pass through the wizard steps.
- You can quickly add the VMs to an already existing job. To do this, open the **Inventory** view, in the working area select the VMs and click **Add to Backup** > *name of the job* on the ribbon or right-click the VMs and select **Add to backup job** > *name of the job*.

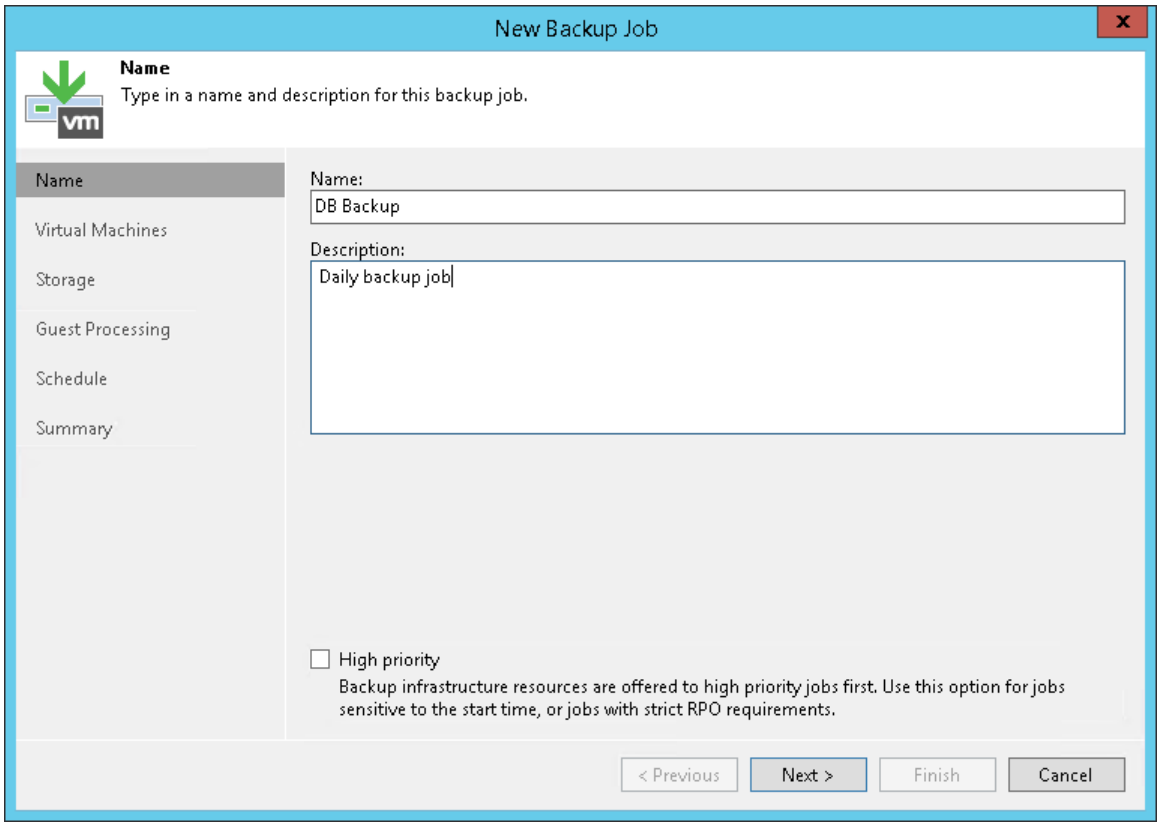


Step 2. Specify Job Name and Description

At the **Name** step of the wizard, specify a name and description for the backup job.

1. In the **Name** field, enter a name for the backup job.
2. In the **Description** field, provide a description for future reference. The default description contains information about the user who created the job, date and time when the job was created.
3. Select the **High priority** check box if you want the resource scheduler of Veeam Backup & Replication to prioritize this job higher than other similar jobs and to allocate resources to it in the first place. For more information on job priorities, see Job Priorities.

> **TIP**
>
> In the UI, jobs with the **High priority** option enabled will be marked with a special flag ( ).



Step 3. Select VMs to Back Up

At the **Virtual Machines** step of the wizard, select VMs and VM containers (hosts, clusters, folders, resource pools, VirtualApps, datastores or tags) that you want to back up:

1. Click **Add**.
2. In the **Add Object** window, select the necessary VMs or VM containers and click **Add**. If you select VM containers and add new VMs to this container in future, Veeam Backup & Replication will update backup job settings automatically to include these VMs.

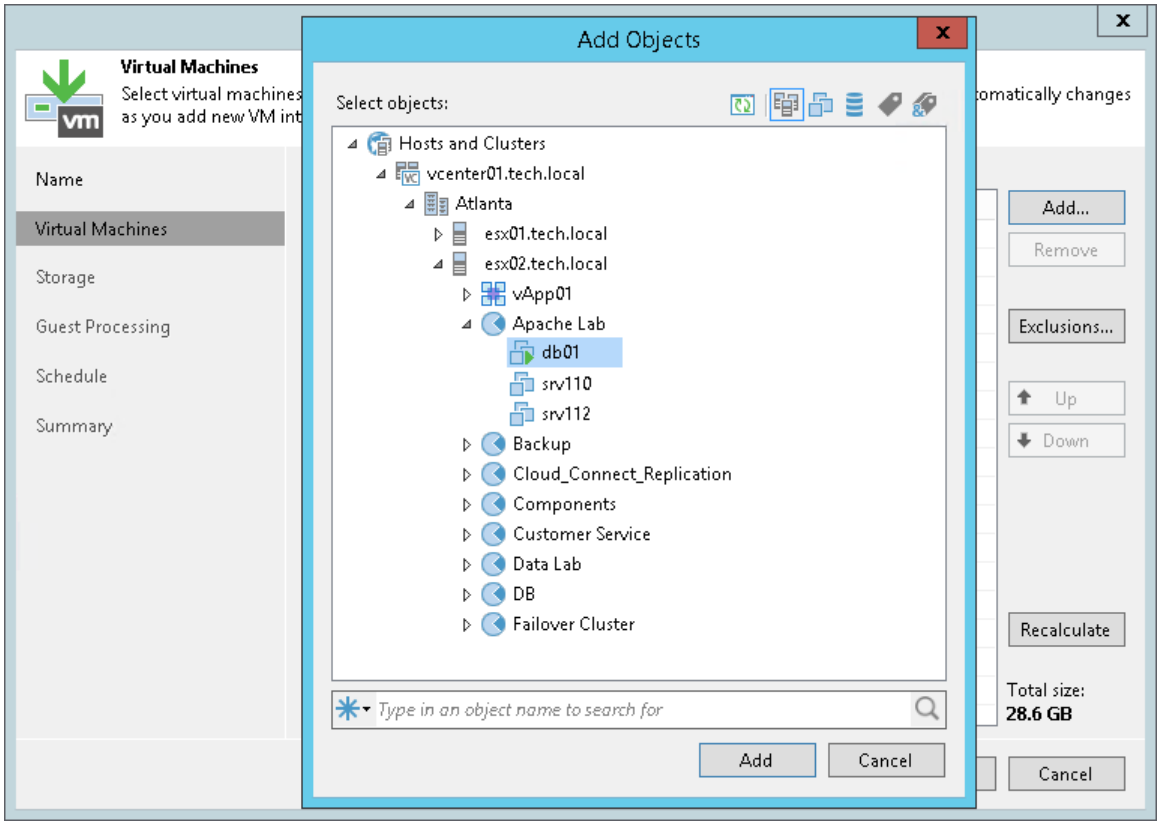
   You can use the toolbar at the top right corner of the window to switch between views. Depending on the view you select, some objects may not be available. For example, if you select **Tags combination** view, no resource pools, hosts or clusters will be displayed in the tree. In the **Tags combination** view, you can select multiple tags and only those VMs that have all the selected tags will be processed by the job.

   To quickly find the necessary VMs, you can use the search field at the bottom of the **Add Object** window. If you want to switch between types of VMs you want to search through, use the button to the left of the search field.

NOTE

You can use a regular backup job to process VMs that are part of vApps created in vCenter Server. To back up vCloud Director vApps, you must use specifically developed vCD backup jobs. For more information, see Backup and Restore of vApps.

The total size of objects added to the job is displayed in the **Total size** field. Use the **Recalculate** button to refresh the total size value after you add a new object to the job.



Step 4. Exclude Objects from Backup Job

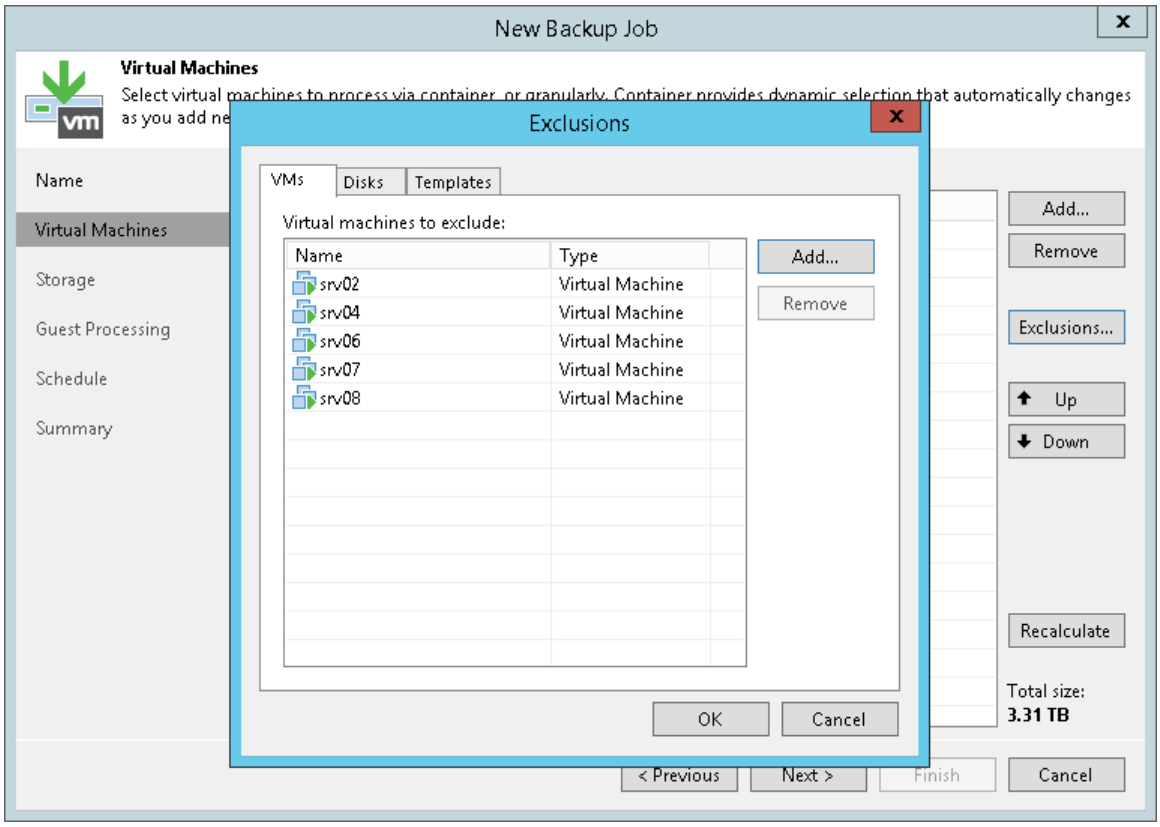
After you have added VMs and VM containers to the job, you can specify which objects you want to exclude from the backup. You can exclude the following types of objects:

- VMs from VM containers
- Specific VM disks
- VM templates

> **NOTE**
>
> Veeam Backup & Replication automatically excludes VM log files from backup to make the backup process faster and reduce the size of the backup file.

To exclude VMs from a VM container:

1. At the **Virtual Machines** step of the wizard, click **Exclusions**.
2. Click the **VMs** tab.
3. Click **Add**.
4. Use the toolbar at the top right corner of the window to switch between views. Depending on the view you select, some objects may not be available. For example, if you select **Tags combination** view, no resource pools, hosts or clusters will be displayed in the tree. In the **Tags combination** view, you can select multiple tags and only those VMs that have all the selected tags will be excluded from the job.
5. In the displayed tree, select the necessary object and click **Add**. Use the **Show full hierarchy** check box to display the hierarchy of all VMware Servers added to the backup infrastructure.
6. Click **OK**.

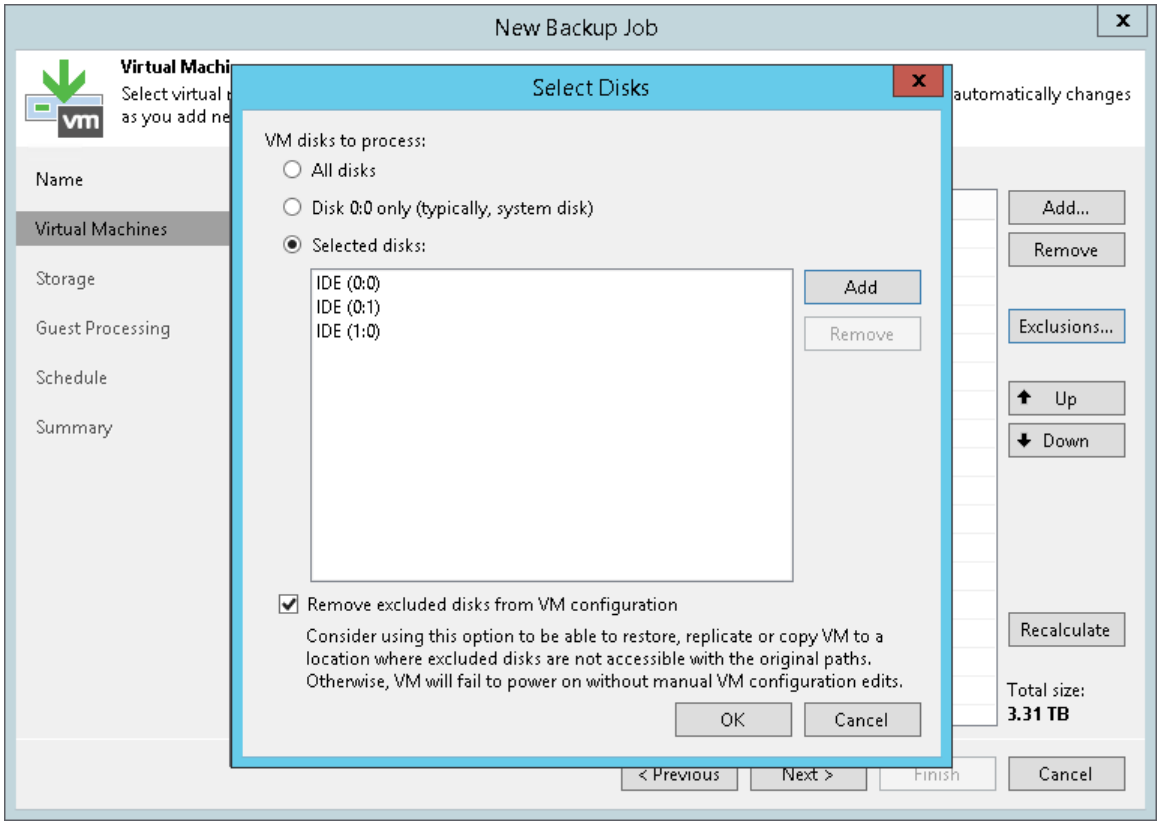


To exclude VM disks:

1. At the **Virtual Machines** step of the wizard, click **Exclusions**.
2. Click the **Disks** tab.
3. Select the VM in the list and click **Edit**. If you want to exclude disks of a VM added as part of the container, click **Add** to include the VM in the list as a standalone object.
4. Choose disks that you want to back up. You can choose to process all disks, 0:0 disks (typically, system disks) or add to the list custom IDE, SCSI or SATA disks.
5. Select the **Remove excluded disks from VM configuration** check box. Veeam Backup & Replication will modify the VMX file of a backed-up VM to remove excluded disks from the VM configuration. If you restore this VM from the backup file to a location where excluded disks are not accessible with the original paths, you will not have to manually edit the VM configuration file to be able to power on the VM.

NOTE

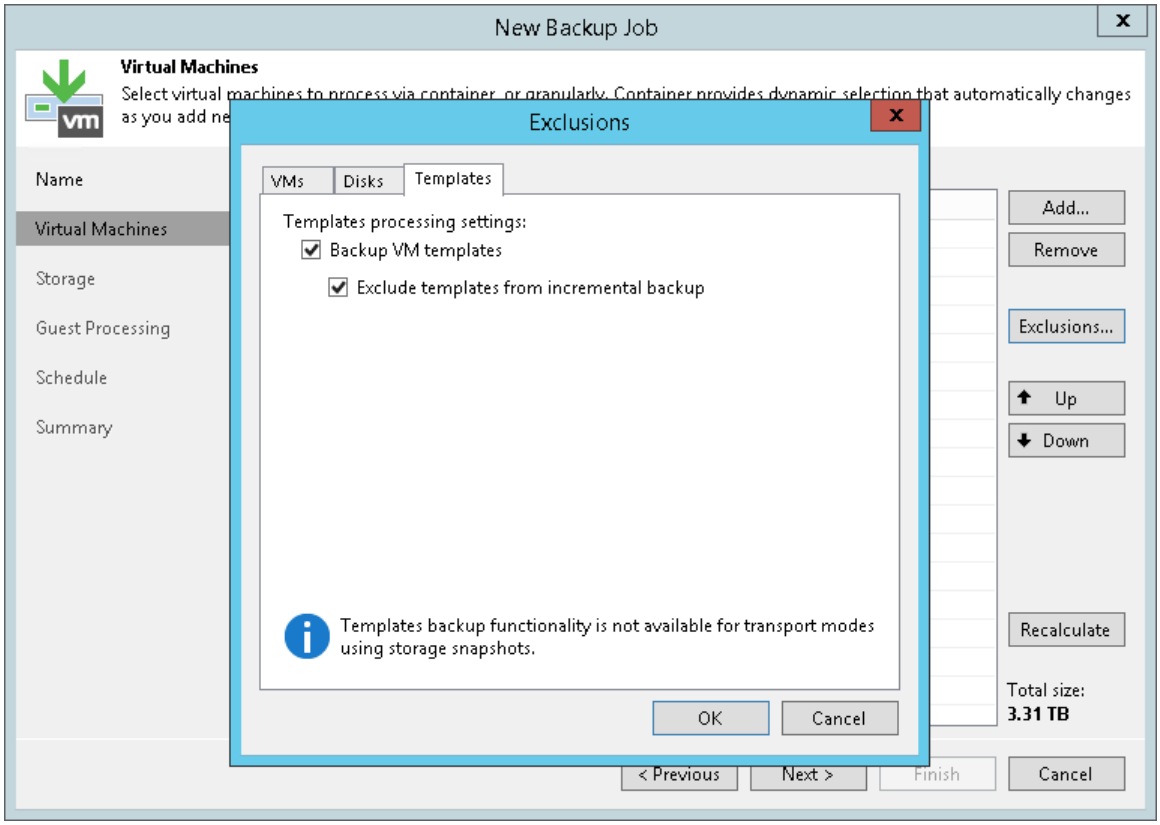
If you exclude disks from a backup and enable application-aware processing, Veeam Backup & Replication will still perform application-aware processing for the excluded disks. This means that VSS will process disk data.



To exclude VM templates:

1. At the **Virtual Machines** step of the wizard, select a VM container and click **Exclusions**.
2. Click the **Templates** tab.
3. Clear the **Backup VM templates** check box.

4. If you want to include VM templates into the full backup only, leave the **Backup VM templates** check box selected and select the **Exclude templates from incremental backup** check box.



Step 5. Define VM Backup Order

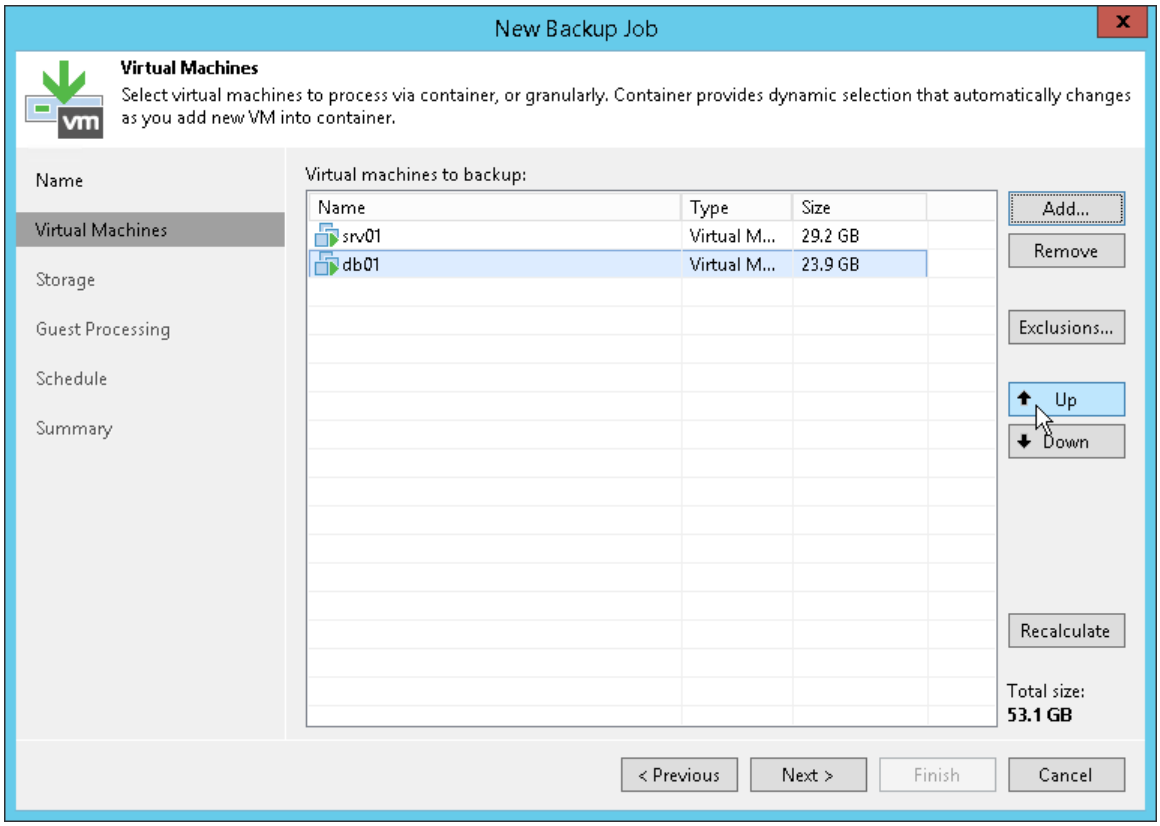
You can define the order in which the backup job must process VMs. Setting VM order can be helpful, for example, if you add some mission-critical VMs to the job and want the job to process them first. You can set these VMs first in list to ensure that their processing fits the backup window.

VMs inside a VM container are processed at random. To ensure that VMs are processed in the defined order, you must add them as standalone VMs, not as a part of the VM container.

To define the VM backup order:

1. At the **Virtual Machines** step of the wizard, select a VM or VM container.
2. Use the **Up** and **Down** buttons on the right to move the VM or VM container up or down in the list.

> **NOTE**
>
> VMs may be processed in a different order. For example, if backup infrastructure resources for a VM that is higher on the priority list are not available, and resources for a VM that is lower on the list are available, Veeam Backup & Replication will start processing the VM that is lower on the list first.



Step 6. Specify Backup Storage Settings

At the **Storage** step of the wizard, select backup infrastructure components for the job — backup proxy and backup repository, and specify backup storage settings.

1. Click **Choose** next to the **Backup proxy** field to select a backup proxy.
   - If you choose **Automatic selection**, Veeam Backup & Replication will detect backup proxies that have access to the source datastore and automatically assign an optimal backup proxy to process VMs in the job.

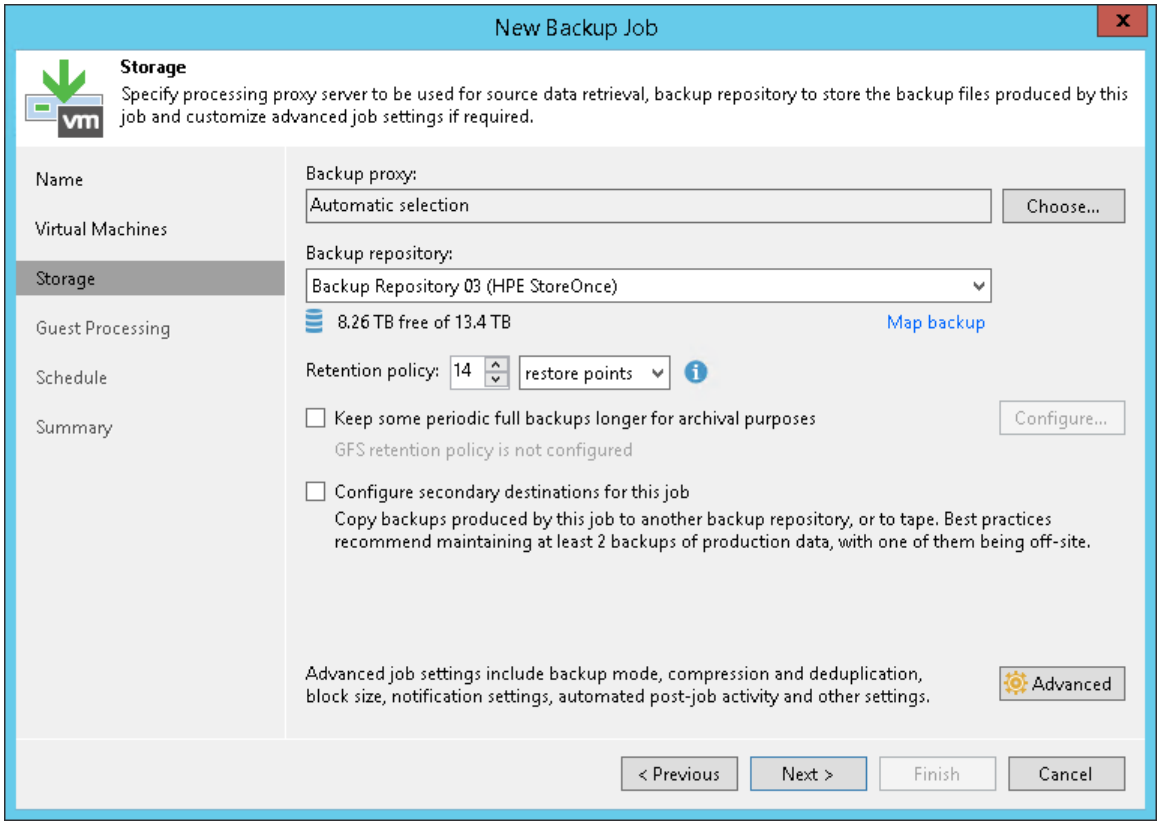
     Veeam Backup & Replication assigns backup proxies to VMs included in the backup job one by one. Before processing a new VM in the VM list, Veeam Backup & Replication checks available backup proxies. If more than one backup proxy is available, Veeam Backup & Replication analyzes transport modes that the backup proxies can use to retrieve VM data and the current workload on the backup proxies to select the most appropriate one for VM processing.
   - If you choose **Use the selected backup proxy servers specified below**, you can explicitly select backup proxies that the job must use. It is recommended that you select at least two backup proxies to ensure that the backup job starts if one of the proxies fails or loses its connectivity to the source datastore.
2. From the **Backup repository** list, select a backup repository where the created backup files must be stored. When you select a backup repository, Veeam Backup & Replication automatically checks how much free space is available in the backup repository.
3. You can map the job to a specific backup stored in the backup repository. Backup job mapping can be helpful if you have moved backup files to a new backup repository and want to point the job to existing backups in this new backup repository. You can also use backup job mapping if the configuration database got corrupted and you need to reconfigure backup jobs.

   To map the job to a backup, click the **Map backup** link and select the backup in the backup repository. Backups can be easily identified by job names. To find the backup, you can also use the search field at the bottom of the window.
4. In the **Retention Policy** field, specify retention policy settings for restore points:
   - If you want to keep the last N restore points, select *restore points* from the drop-down list. Then specify the number of restore points.
   - If you want to keep all restore points created during the last N days, select *days* from the drop-down list. Then specify the number of days.

   When the specified number is exceeded, the earliest restore point is removed from the backup chain or merged with the next closest restore point. For more information, see Short-Term Retention Policy.

5. If you want to archive backup files created with the backup job to a secondary destination (backup repository or tape), select the **Configure secondary destination for this job** check box. With this option enabled, the **New Backup Job** wizard will include an additional step — **Secondary Target**. At the **Secondary Target** step of the wizard, you can link the backup job to the backup copy job or backup to tape backup job.

   You can enable this option only if a backup copy job or backup to tape job is already configured on the backup server.



Step 7. Configure Long-Term Retention

If you want to ignore the short-term retention policy for some full backups and store them for long-term archiving, you can configure long-term retention policy (or GFS retention policy) for the backup job. For more information on GFS and its limitations, see Long-Term Retention Policy (GFS).

1. At the storage step of the wizard, select the **Keep some periodic full backups longer for archival purposes** check box.
2. Click **Configure**.
3. In the **Configure GFS** window, do the following:
   - If you want to create weekly restore points, select the **Keep weekly full backups for** check box. Then specify the number of weeks during which you want to prevent restore points from being modified and deleted.

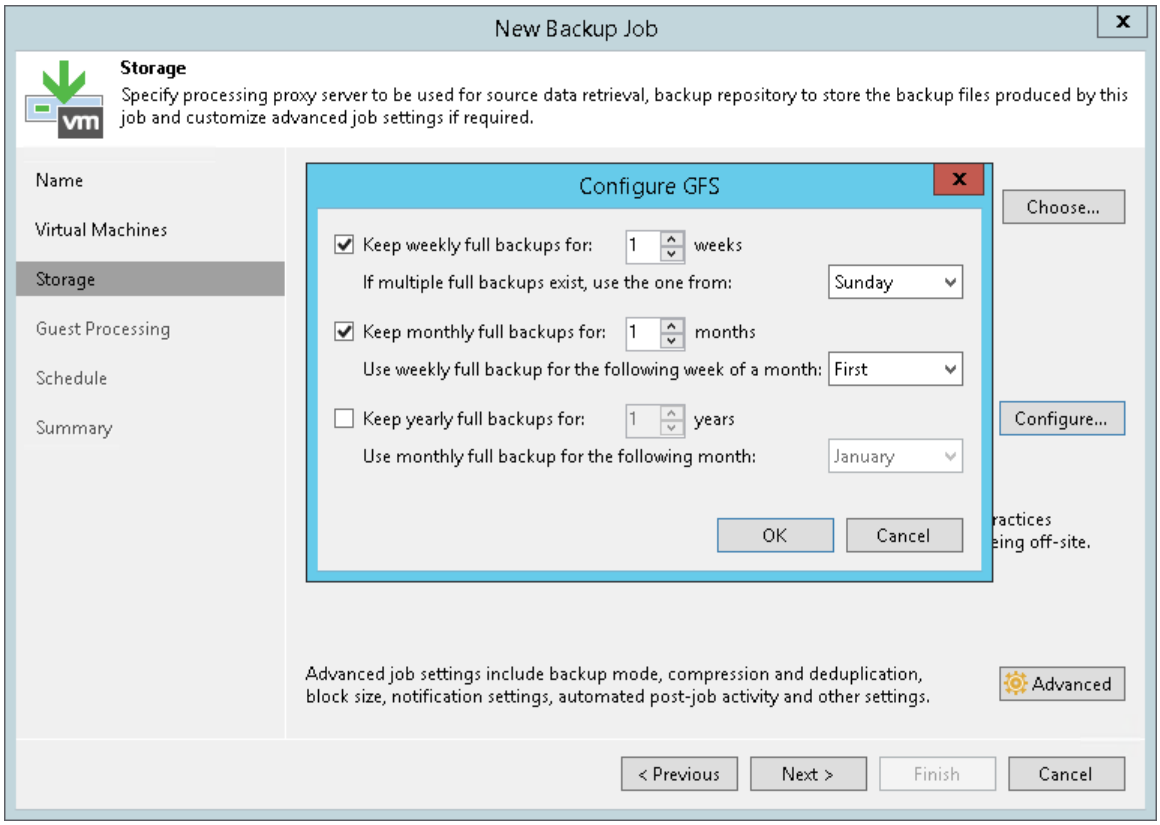
     From the **If multiple full backups exist, use the one from** drop-down list, select a week day when Veeam Backup & Replication must assign the weekly GFS flag to a full restore point.
   - If you want to create monthly restore points, select the **Keep monthly full backups for** check box. Then specify the number of months during which you want to prevent restore points from being modified and deleted.

     From the **Use weekly full backup for the following week of a month** drop-down list, select a week when Veeam Backup & Replication must assign the monthly GFS flag to a full restore point. A week equals 7 calendar days; for example, the first week of May is days 1–7, and the last week of May is days 25–31.
   - If you want to create yearly restore points, select the **Keep yearly full backups for** check box. Then specify the number of years during which you want to prevent restore points from being modified and deleted.

     From the **Use monthly full backup for the following month** drop-down list, select a month when Veeam Backup & Replication must assign the yearly GFS flag to a full restore point.

NOTE

If you select to assign multiple types of GFS flags, the flags begin to depend on each other. For more information on this dependency, see Assignment of Multiple GFS Flags.



Step 8. Specify Advanced Backup Settings

At the **Storage** step of the wizard, specify advanced settings for the backup job:

- Backup settings
- Maintenance settings
- Storage settings
- Notification settings
- vSphere settings
- Integration settings
- Script settings

> **TIP**
>
> After you specify necessary settings for the backup job, you can save them as default settings. To do this, click **Save as Default** at the bottom left corner of the **Advanced Settings** window. When you create a new backup job, Veeam Backup & Replication will automatically apply the default settings to the new job.

Backup Settings

To specify settings for a backup chain created with the backup job:

1. At the **Storage** step of the wizard, click **Advanced**.
2. On the **Backup** tab, select the backup method that you want to use to create the backup chain in the backup repository:
   - To create a reverse incremental backup chain, select **Reverse Incremental**.

     Dell EMC Data Domain and HPE StoreOnce do not support the reverse incremental backup method. Do not select this option for backup jobs targeted at these types of backup repositories.
   - To create an incremental backup chain, select **Incremental** and enable synthetic full and/or active full backups (see items 3-4).
   - To create a forever forward incremental backup chain, select **Incremental** and do not enable synthetic full and/or active full backups (see items 3-4).

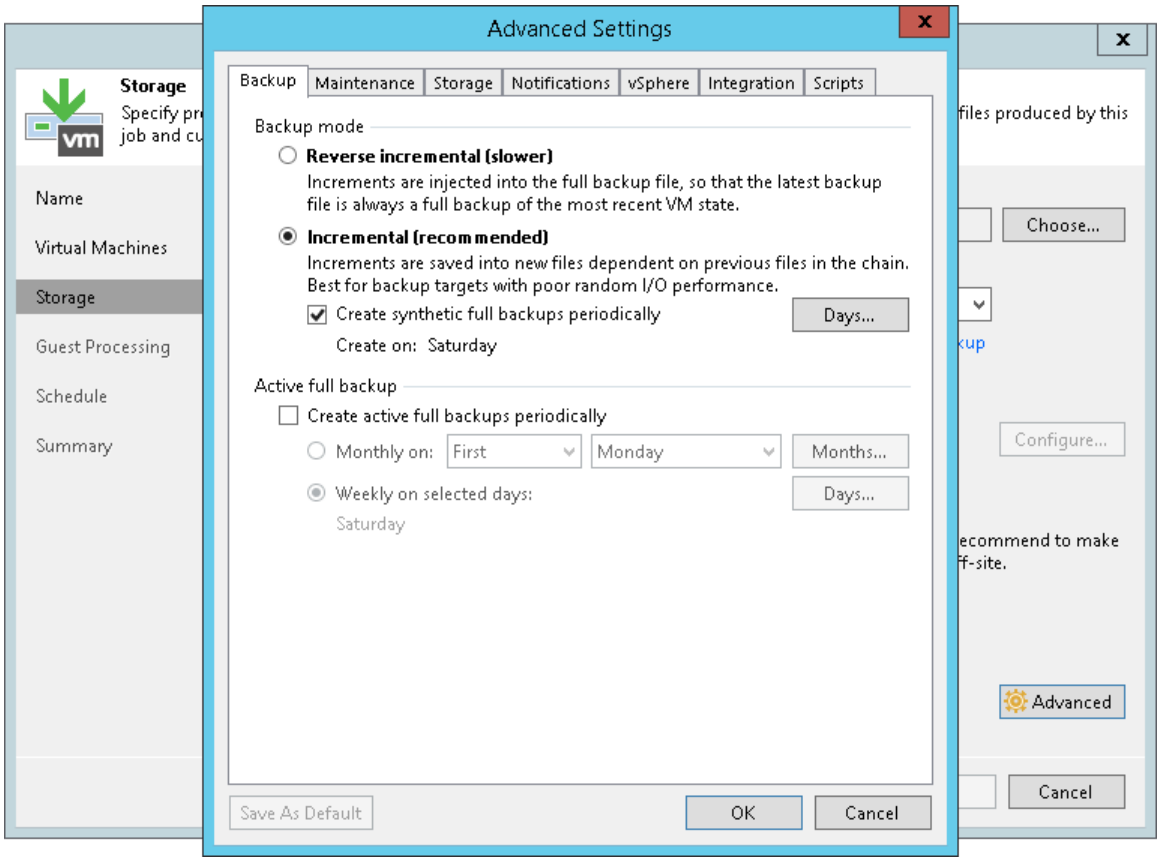
   For more information, see Backup Methods.
3. If you choose the incremental backup method, you can select to periodically create synthetic full backups and/or active full backups.
   - To create a synthetic full backup, select the **Create synthetic full backups periodically** check box and click **Days** to schedule synthetic full backups on the necessary week days.

> **NOTE**
>
> For backup jobs created in Veeam Backup & Replication version prior to 11, you could enable the **Transform previous backup chains into rollbacks** option. Starting from Veeam Backup & Replication version 11, this option is deprecated. However, the transformation functionality is still working for backup jobs where the **Transform previous backup chains into rollbacks** option was already enabled. For more information on the transformation, see the Backup Chain Transformation section in Veeam Backup & Replication User Guide for version 10.

- To create full backups regularly, select the **Create active full backups periodically** check box. Use the **Monthly on** or **Weekly on selected days** options to define scheduling settings.

Before you schedule periodic full backups, you must make sure that you have enough free space in the backup repository. As an alternative, you can create active full backups manually when needed. For more information, see Active Full Backup.

> **IMPORTANT**
>
> Mind the following:
>
> - If you schedule the active full backup and synthetic full backup on the same day, Veeam Backup & Replication will perform only active full backup. Synthetic full backup will be skipped.
> - If you set the job schedule to start after another job, and this job has not run yet, Veeam Backup & Replication will not perform active full backup according to the schedule. For more information on the job schedule options, see Define Job Schedule.



Maintenance Settings

You can instruct Veeam Backup & Replication to periodically perform maintenance operations — service actions that will help make sure that the backup chain remains valid and consistent.

To specify maintenance settings for the backup job:

1. At the **Storage** step of the wizard, click **Advanced**.
2. Click the **Maintenance** tab.
3. To periodically perform a health check for the latest restore point in the backup chain, in the **Storage-level corruption guard** section select the **Perform backup files health check** check box and specify the time schedule for the health check.

   An automatic health check can help you avoid a situation where a restore point gets corrupted, making all dependent restore points corrupted, too. If during the health check Veeam Backup & Replication detects corrupted data blocks in the latest restore point in the backup chain (or, in case of forever forward incremental and forward incremental chains, the restore point before the latest one if the latest restore point is incomplete), it will start the health check retry and transport valid data blocks from the source datastore to the backup repository. The transported data blocks are stored to a new backup file or the latest backup file in the backup chain, depending on the data corruption scenario. For more information, see Health Check for Backup Files.
4. Select the **Remove deleted items data after** check box and specify the number of days for which you want to keep backup data for deleted VMs. If a VM is no longer available (for example, it was deleted or excluded from the job), Veeam Backup & Replication will keep its data in the backup repository for the period that you have specified. When this period is over, data of the deleted VM will be removed from the backup repository.

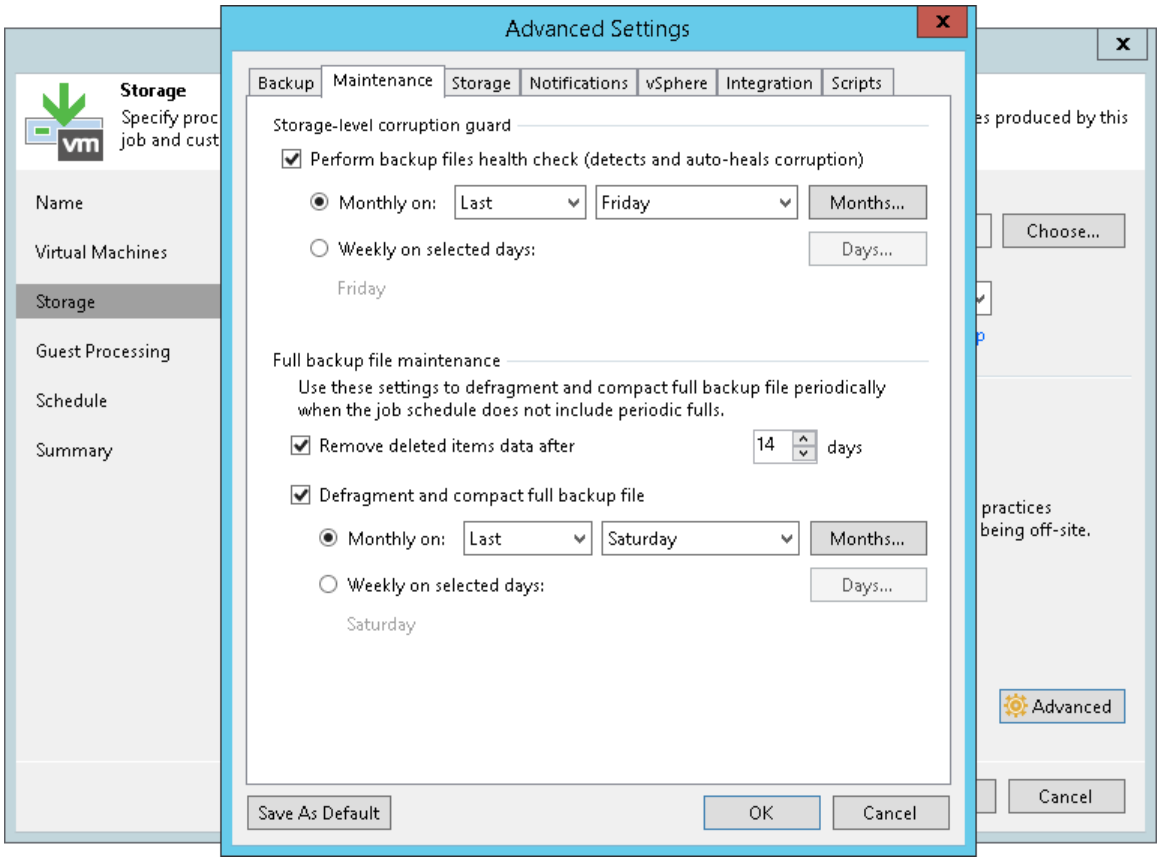
   By default, the retention period for deleted VM data is 14 days. It is strongly recommended that you set the retention period to 3 days or more to prevent unwanted data loss. For more information, see Retention Policy for Deleted VMs.
5. To periodically compact a full backup, select the **Defragment and compact full backup file** check box and specify the schedule for the compact operation.

   During the compact operation, Veeam Backup & Replication creates a new empty file and copies to it data blocks from the full backup file. As a result, the full backup file gets defragmented and the speed of reading and writing from/to the backup file increases.

   If the full backup file contains data blocks for deleted VMs, Veeam Backup & Replication will remove these data blocks. If the full backup file contains data for a VM that has only one restore point, and this restore point is older than 7 days, Veeam Backup & Replication will perform the take out operation. For more information, see Compact of Full Backup File.

IMPORTANT

If you schedule periodic full backups, the **Remove deleted items data after** and **Defragment and compact full backup file** check boxes do not apply.



Storage Settings

To specify storage settings for the backup job:

1. At the **Storage** step of the wizard, click **Advanced**.
2. Click the **Storage** tab.
3. By default, Veeam Backup & Replication deduplicates VM data before storing it in the backup repository. Data deduplication provides a smaller size of the backup file but may reduce the backup job performance. For more information on data deduplication, see Deduplication.

   To disable data deduplication, clear the **Enable inline data deduplication** check box. When you disable data deduplication, you also change the workflow for incremental backups. If Changed Block Tracking (CBT) is enabled for the job, Veeam Backup & Replication will save all data blocks marked by CBT as new to the destination storage without performing an additional check or using Veeam filtering mechanism. This will result in faster creation of incremental backups. However, backup files may increase in size.

4. By default, Veeam Backup & Replication checks the NTFS MFT file on VMs with Microsoft Windows OS to identify data blocks of the `hiberfil.sys` file (file used for the hibernate mode) and `pagefile.sys` file (swap file), and excludes these data blocks from processing. The swap file is dynamic in nature and changes intensively between backup job sessions, even if the VM itself does not change much. Processing of service files reduces the job performance and increases the size of incremental backup files.

   If you want to include data blocks of the `hiberfil.sys` file and `pagefile.sys` file to the backup, clear the **Exclude swap file blocks** check box. For more information, see Swap Files.

5. By default, Veeam Backup & Replication does not copy deleted file blocks ("dirty" blocks on the VM guest OS) to the target location. This option lets you reduce the size of backup files and increase the job performance. If you want to include dirty data blocks to the backup, clear the **Exclude deleted file blocks** check box. For more information, see Deleted File Blocks.

6. From the **Compression level list**, select a compression level for the backup: *None*, *Dedupe-friendly*, *Optimal*, *High* or *Extreme*. For more information on data compression, see Compression and Deduplication.

7. In the **Storage optimization** section, select what type of backup target you plan to use. Depending on the selected type, Veeam Backup & Replication will use different data block sizes to process VMs. For more information on the available storage optimization options, the used data blocks and how they affect performance, see Storage Optimization.

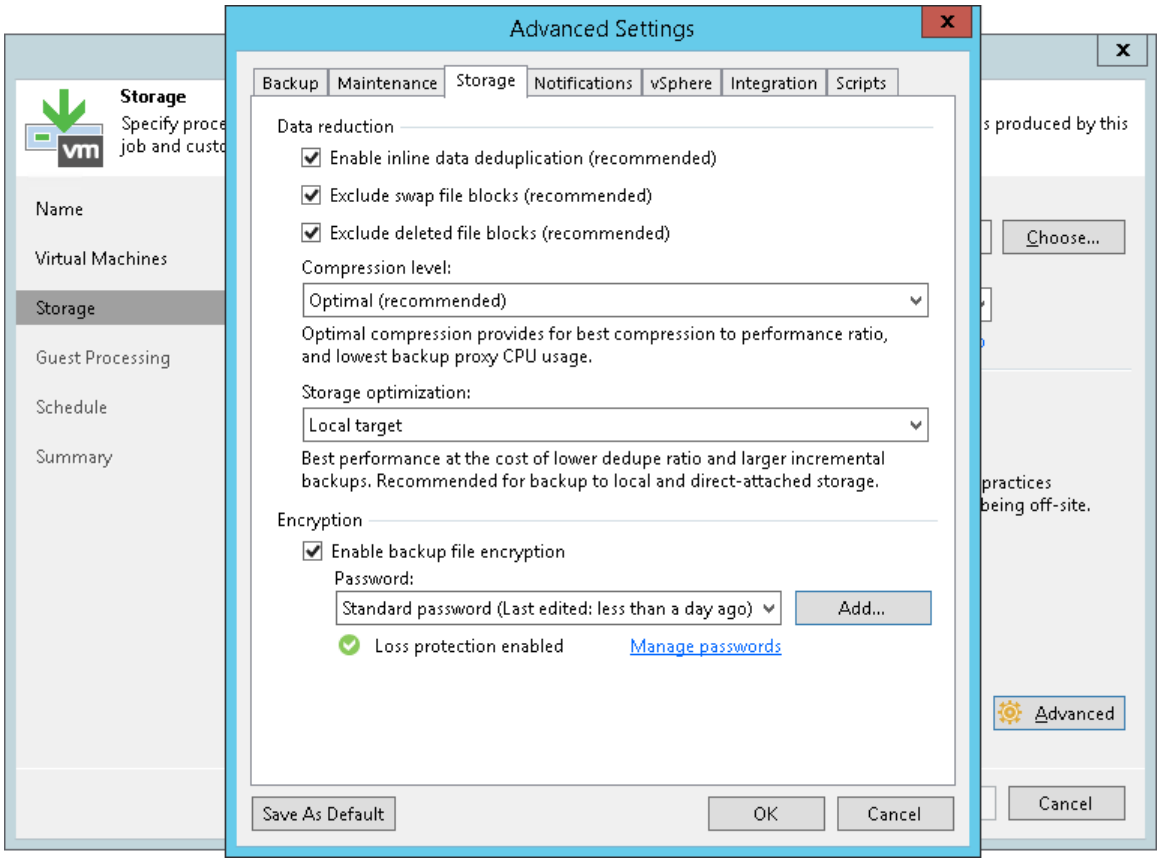
8. To encrypt the content of backup files, select the **Enable backup file encryption** check box. In the **Password** field, select a password that you want to use for encryption. If you have not created the password beforehand, click **Add** or use the **Manage passwords** link to specify a new password. For more information, see Managing Passwords for Data Encryption.

   If the backup server is not connected to Veeam Backup Enterprise Manager, you will not be able to restore data from encrypted backups in case you lose the password. Veeam Backup & Replication will display a warning about it. For more information, see Decrypting Data Without Password.

NOTE

If you enable encryption for an existing backup job, during the next job session Veeam Backup & Replication will create a full backup file. The created full backup file and subsequent incremental backup files in the backup chain will be encrypted with the specified password.

Encryption is not retroactive. If you enable encryption for an existing job, Veeam Backup & Replication does not encrypt the previous backup chain created with this job. If you want to start a new chain so that the unencrypted previous chain can be separated from the encrypted new chain, see this Veeam KB article.



Notification Settings

To specify notification settings for the backup job:

1. At the **Storage** step of the wizard, click **Advanced**.
2. Click the **Notifications** tab.
3. Select the **Send SNMP notifications for this job** check box if you want to receive SNMP traps when the job completes successfully.

   SNMP traps will be sent if you specify global SNMP settings in Veeam Backup & Replication and configure software on recipient's machine to receive SNMP traps. For more information, see Specifying SNMP Settings.

4. Select the **Send email notifications to the following recipients** check box if you want to receive notifications about the job completion status by email. In the field below, specify recipient's email address. You can enter several addresses separated by a semicolon.

   Email notifications will be sent if you configure global email notification settings in Veeam Backup & Replication. For more information, see Configuring Global Email Notification Settings.

5. You can choose to use global notification settings or specify custom notification settings.

   - To receive a typical notification for the job, select **Use global notification settings**. In this case, Veeam Backup & Replication will apply to the job global email notification settings specified for the backup server. For more information, see Configuring Global Email Notification Settings.
   - To configure a custom notification for the job, select **Use custom notification settings specified below** check box. You can specify the following notification settings:

     i. In the **Subject** field, specify a notification subject. You can use the following variables in the subject: *%Time%* (completion time), *%JobName%*, *%JobResult%*, *%ObjectCount%* (number of VMs in the job) and *%Issues%* (number of VMs in the job that have finished with the *Warning* or *Failed* status).

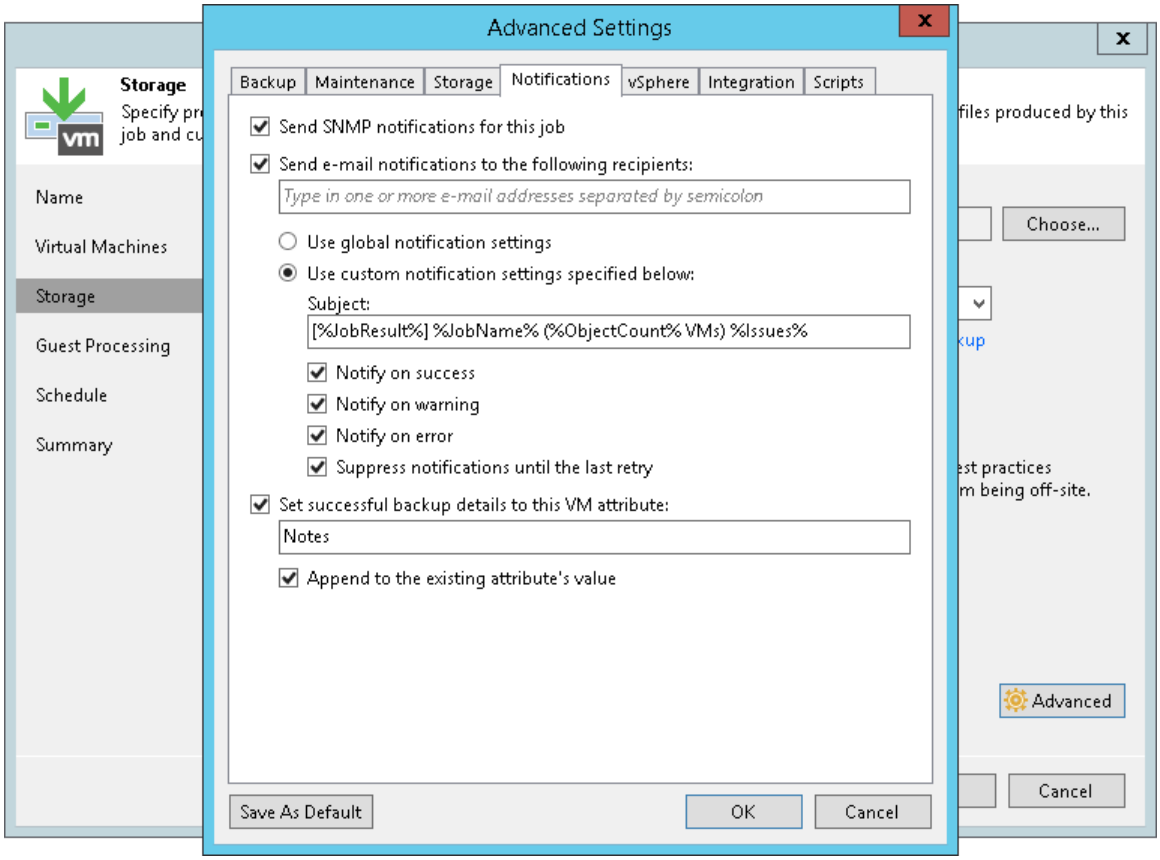
     ii. Select the **Notify on success**, **Notify on warning** and/or **Notify on error** check boxes to receive email notification if the job completes successfully, fails or completes with a warning.

     iii. Select the **Suppress notifications until the last retry** check box to receive a notification about the final job status. If you do not enable this option, Veeam Backup & Replication will send one notification per every job retry.

   > **NOTE**
   >
   > If you specify the same email recipient in both job notification and global notification settings, Veeam Backup & Replication will send two separate notifications only if a subject for the email message specified in job notification is different from the subject specified in global notification settings. Otherwise, Veeam Backup & Replication will suppress global notification settings and will send job notifications only.

6. Select the **Set successful backup details to this VM attribute** check box to write information about successfully performed backup and backup results (backup date and time, backup server name and path to the backup file) to a VM attribute. In the field below, enter a name of the attribute. If the specified attribute does not exist, Veeam Backup & Replication will create it.

7. Select the **Append to the existing attribute's value** check box to append information about successfully performed backup to an existing value of the attribute. In this case, Veeam Backup & Replication will keep values added by the user in the attribute, and will overwrite only the value added by the backup job. If you do not select this option, Veeam Backup & Replication will overwrite the existing attribute values (made both by the user and backup job).



vSphere Settings

To specify VMware vSphere settings for the backup job:

1. At the **Storage** step of the wizard, click **Advanced**.
2. Click the **vSphere** tab.
3. Select the **Enable VMware tools quiescence** check box to freeze the file system of processed VMs during backup.

   Depending on the VM version, Veeam Backup & Replication will use the VMware FileSystem Sync Driver (vmsync) driver or VMware VSS component in VMware Tools for VM snapshot creation. These tools are responsible for quiescing the VM file system and bringing the VM to a consistent state suitable for backup. For more information, see VMware Tools Quiescence.
4. In the **Changed block tracking** section, configure VMware vSphere CBT:
   a. Make sure that the **Use changed block tracking data** check box is selected if you want to enable CBT.
   b. Make sure that the **Enable CBT for all processed VMs automatically** check box is selected if you want to force using CBT even if CBT is disabled in VM configuration.

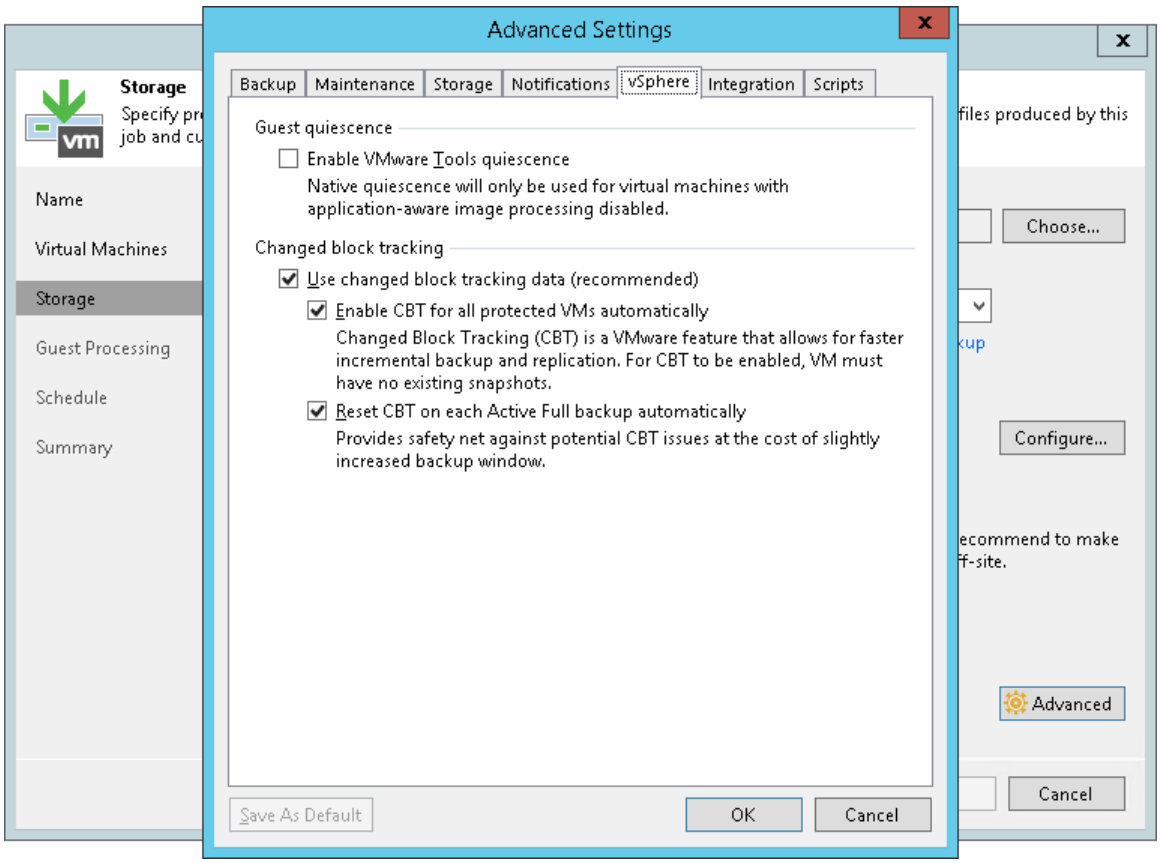
c. Make sure that the **Reset CBT on each Active Full backup automatically** check box is selected if you want to reset CBT before Veeam Backup & Replication creates active full backups.

CBT reset helps avoid issues, for example, when CBT returns incorrect changed data.

For more information on CBT, see Changed Block Tracking.

IMPORTANT

You can use CBT for VMs with virtual hardware version 7 or later. These VMs must not have existing snapshots.



Integration Settings

On the **Integration** tab, you can define whether you want to use the Backup from Storage Snapshots technology to create the backup. Backup from Storage Snapshots lets you leverage storage snapshots for VM data processing. The technology improves RPOs and reduces the impact of backup activities on the production environment.

Before you start using the Backup from Storage Snapshots technology, check prerequisites. For more information, see the Requirements and Limitations section in the Veeam Backup & Replication User Guide.

To specify storage integration settings for the backup job:

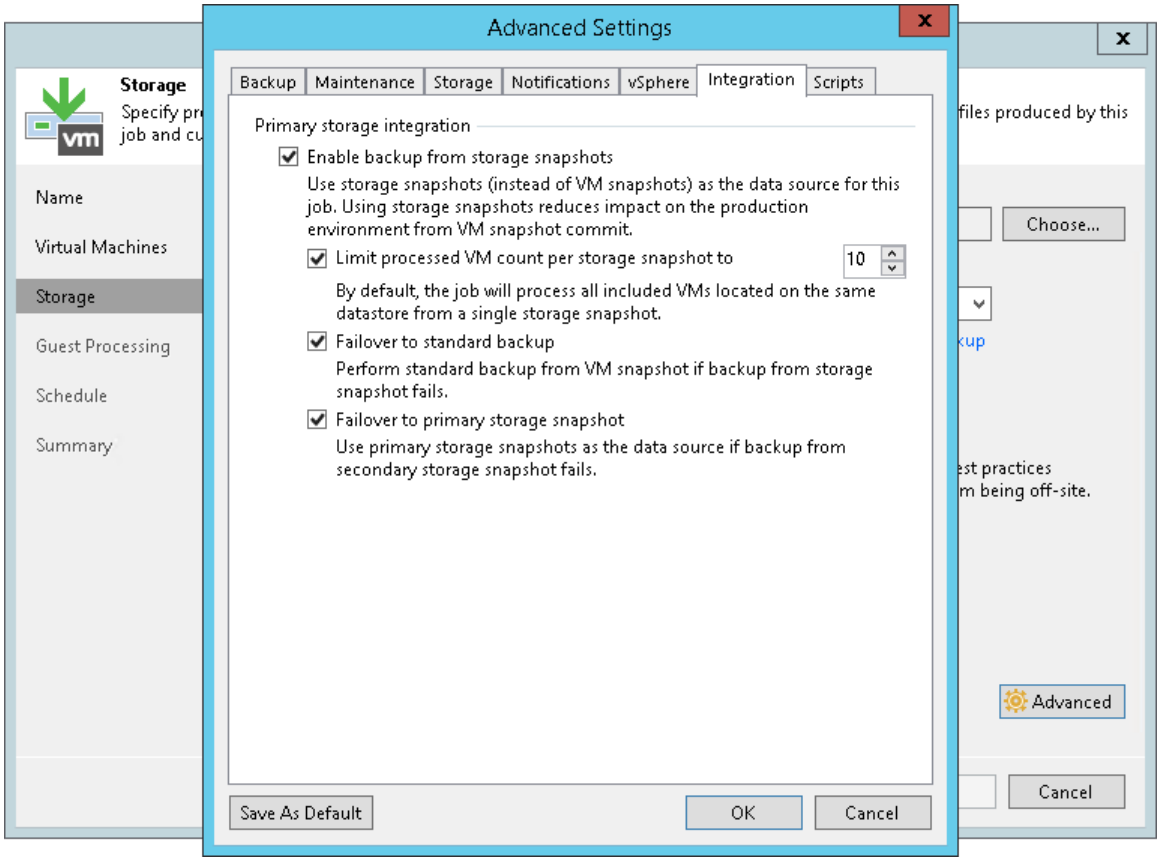
1. At the **Storage** step of the wizard, click **Advanced**.
2. Click the **Integration** tab.
3. By default, the **Enable backup from storage snapshots** option is enabled. If you do not want to use Backup from Storage Snapshots, clear this check box. For more information, see the Configuring Backup from Storage Snapshots section in the Veeam Backup & Replication User Guide.

4. If you add to the job many VMs whose disks are located on the same volume or LUN, select the **Limit processed VM count per storage snapshot to <N>** check box and specify the number of VMs for which one storage snapshot must be created. Veeam Backup & Replication will divide VMs into several groups and trigger a separate storage snapshot for every VM group. As a result, the job performance will increase.

   For more information, see the Limitation on Number of VMs per Snapshot section in the Veeam Backup & Replication User Guide.

5. If Veeam Backup & Replication fails to create a storage snapshot, VMs whose disks are located to the storage system will not be processed by the job. To fail over to the regular VM processing mode and back up or replicate such VMs in the regular processing mode, select the **Failover to standard backup** check box.

6. [For secondary NetApp, HPE Nimble and HPE 3PAR storage systems] If Veeam Backup & Replication cannot create a storage snapshot on secondary storage arrays, the job will not back up VMs whose disks are located on the storage system. To fail over to Backup from Storage Snapshots on the production storage, select the **Failover to primary storage snapshot** check box. If Veeam Backup & Replication fails to create a storage snapshot on secondary storage arrays, it will trigger the storage snapshot on primary storage arrays and use it as a source for backup. Note, however, that Backup from Storage Snapshots on primary storage arrays will put additional load on the production environment.



Script Settings

To specify script settings for the backup job:

1. At the **Storage** step of the wizard, click **Advanced**.
2. Click the **Scripts** tab.

3. If you want to execute custom scripts before and/or after the backup job, select the **Run the following script before the job** and **Run the following script after the job** check boxes and click **Browse** to choose executable files from a local folder on the backup server. The scripts are executed on the backup server.

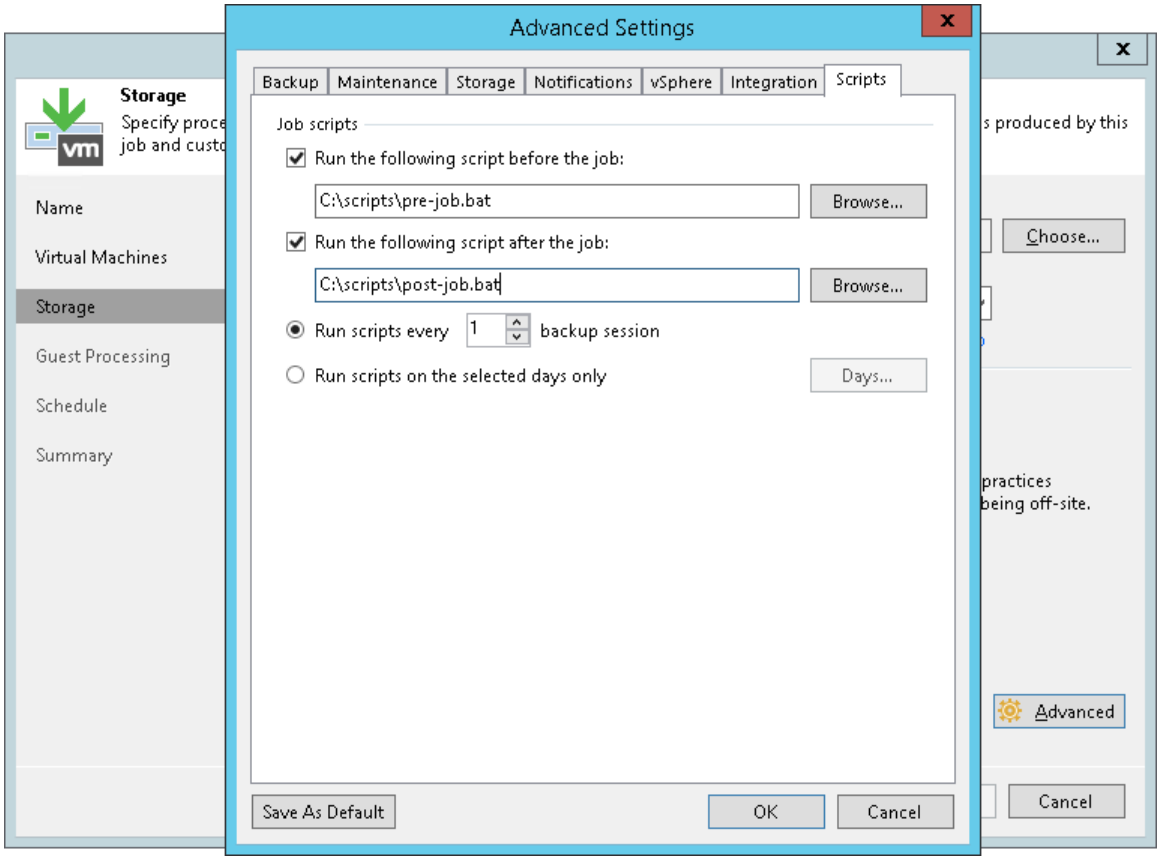
   You can select to execute pre- and post-backup actions after a number of backup sessions or on specific week days.

   - If you select the **Run scripts every <N> backup session** option, specify the number of the backup job sessions after which the scripts must be executed.
   - If you select the **Run scripts on the selected days only** option, click **Days** and specify week days on which the scripts must be executed.

NOTE

Mind the following:

- Custom scripts that you define in the advanced job settings relate to the backup job itself, not the VM quiescence process. To add pre-freeze and post-thaw scripts for VM image quiescence, use the **Guest Processing** step of the wizard.
- If you select the **Run scripts on the selected days only** option, Veeam Backup & Replication executes scripts only once on each selected day — when the job runs for the first time. During subsequent job runs, scripts are not executed.
- To run the script, Veeam Backup & Replication uses the Service Account under which the Veeam Backup Service is running.



Step 9. Specify Secondary Target

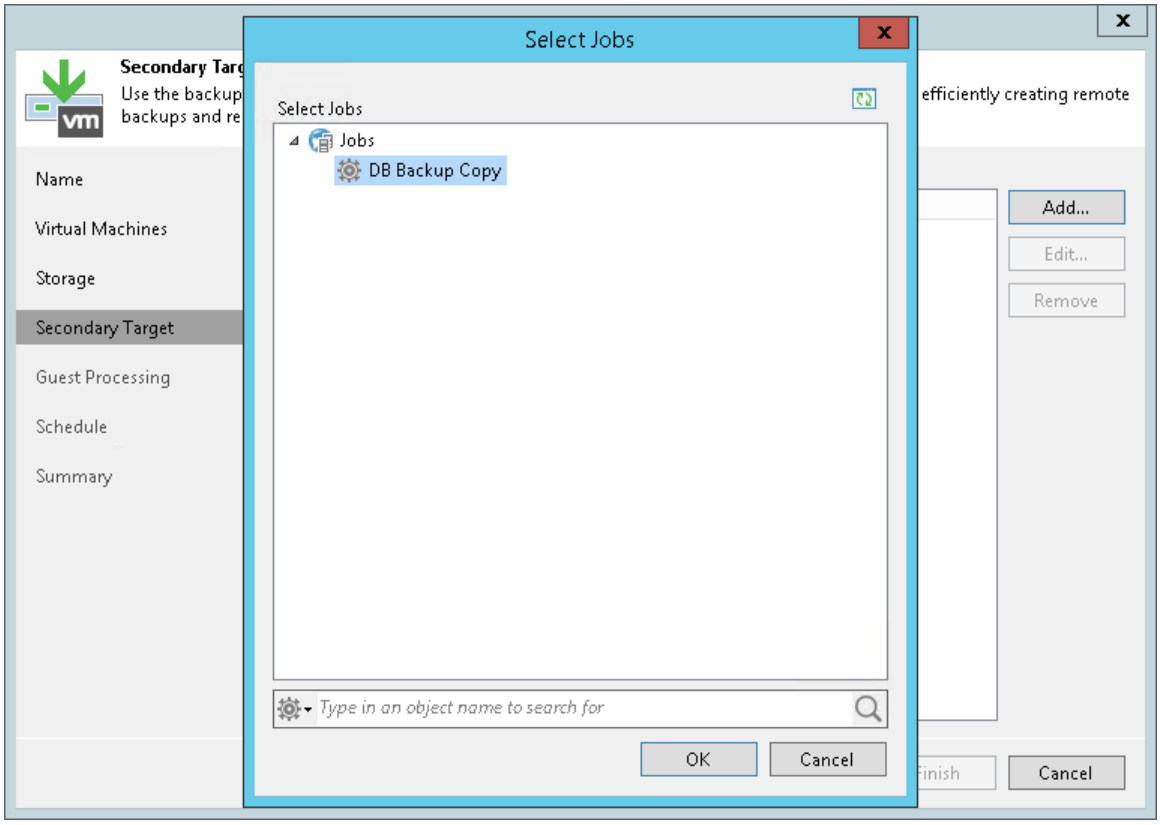
The **Secondary Target** step of the wizard is available if you have enabled the **Configure secondary destination for this job** option at the **Storage** step of the wizard.

At the **Secondary Target** step of the wizard, you can link the backup job to a backup to tape or backup copy job. As a result, the backup job will be added as a source to the backup to tape or backup copy job. Backup files created with the backup job will be archived to tape or copied to the secondary backup repository according to the secondary jobs schedule. For more information, see the Linking Backup Jobs to Backup Copy Jobs and Linking Backup Jobs to Backup to Tape Jobs sections in the Veeam Backup & Replication User Guide.

The backup to tape job or backup copy job must be configured beforehand. You can create these jobs with an empty source. When you link the backup job to these jobs, Veeam Backup & Replication will automatically update the linked jobs to define the backup job as a source for these jobs.

To link jobs:

1. Click **Add**.
2. From the jobs list, select a backup to tape or backup copy job that must be linked to the backup job. You can link several jobs to the backup job, for example, one backup to tape job and one backup copy job. To quickly find the job, use the search field at the bottom of the wizard.



Step 10. Specify Guest Processing Settings

At the **Guest Processing** step of the wizard, you can enable the following settings for VM guest OS processing:

- Application-aware processing
- Microsoft SQL server transaction log settings
- Oracle archived log settings
- VM guest OS file exclusion
- Use of pre-freeze and post-thaw scripts
- VM guest OS file indexing

To coordinate guest processing activities, Veeam Backup & Replication deploys non-persistent runtime components or uses (if necessary, deploys) persistent agent components on the VM guest OS. For more information, see the Non-Persistent Runtime Components and Persistent Agent Components section.

The non-persistent runtime components run only during guest processing and are stopped immediately after the processing is finished (depending on the selected option, during the backup job session or after the backup job completes).

You must specify a user account that will be used to connect to the VM guest OS and deploy the non-persistent runtime components or use persistent agent components:

1. From the **Guest OS credentials** list, select a user account that has enough permissions. For more information on the permissions and requirements for the user account, see Required Permissions for Guest Processing.

   If you have not set up credentials beforehand, click the **Manage accounts** link or click **Add** on the right to add credentials. For more information on adding credentials, see the Credentials Manager section.

   By default, Veeam Backup & Replication uses the **Log on as a batch job** policy to connect to guest OS. If the connection fails, Veeam Backup & Replication switches to **Interactive Logon**.

   > **NOTE**
   >
   > [For Kerberos authentication] Mind the following:
   >
   > - Local accounts do not support Kerberos authentication. To authenticate with Microsoft Windows guest OS using Kerberos, specify an Active Directory account.
   > - Networkless application-aware guest processing through VMware VIX/vSphere Web Services is not supported for VMs with guest OS where NTLM is restricted.
   > - Veeam backup infrastructure servers must correctly resolve FQDNs of guest operating systems.
   > - To back up VMs where Kerberos is used, NTLM must be allowed in the Veeam backup infrastructure machines. For details, see Kerberos Authentication for Guest OS Processing.

2. By default, Veeam Backup & Replication uses the same credentials for all VMs in the job. If some VM requires a different user account, click **Credentials** and enter custom credentials for the VM.

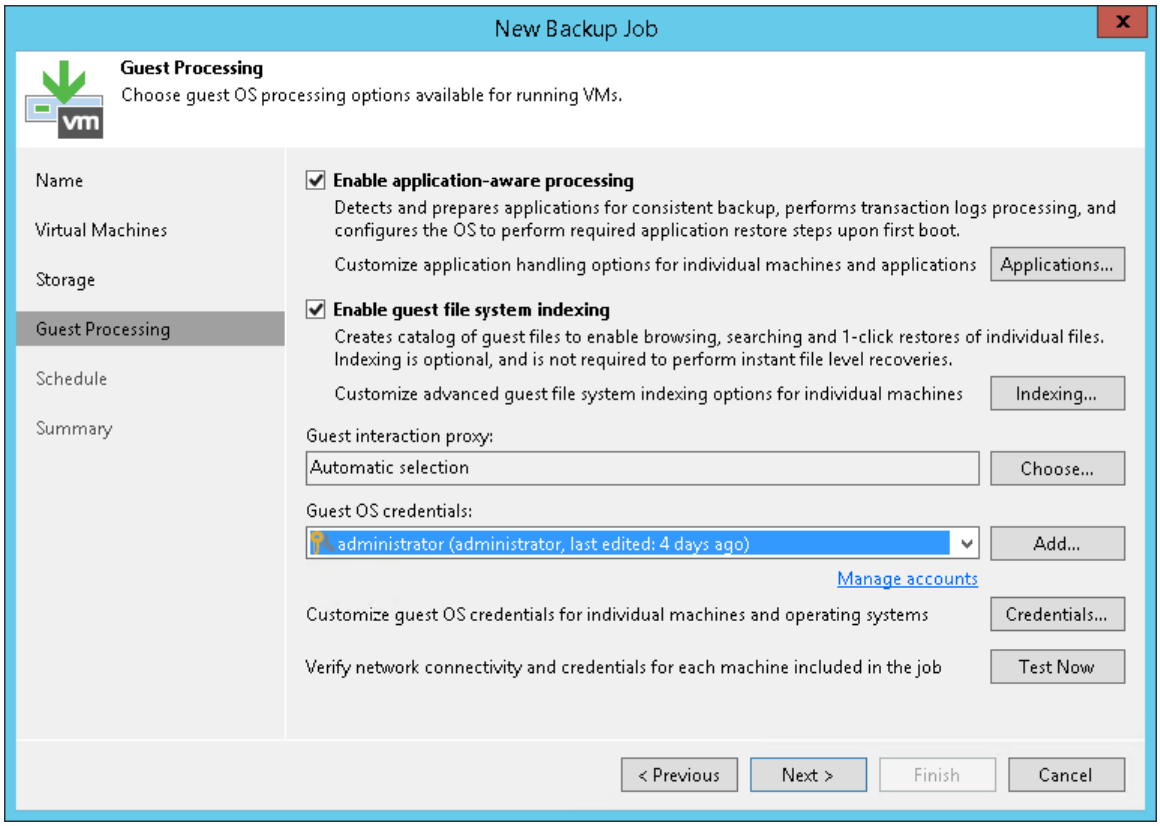
3. If you have added Microsoft Windows VMs to the job, specify which guest interaction proxy Veeam Backup & Replication can use to deploy the non-persistent runtime components or connect to (if necessary, deploy) persistent agent components on the VM guest OS. On the right of the **Guest interaction proxy** field, click **Choose**.

   - Leave **Automatic selection** to let Veeam Backup & Replication automatically select the guest interaction proxy.

- Select **Use the selected guest interaction proxy servers only** to explicitly define which servers will perform the guest interaction proxy role. The list of servers contains Microsoft Windows servers added to the backup infrastructure.

To check if Veeam Backup & Replication can communicate with VMs added to the job and deploy the non-persistent runtime components or use (if necessary, deploy) persistent agent components on their guest OSes, click **Test Now**. Veeam Backup & Replication will use the specified credentials to connect to all VMs in the list.

> **NOTE**
>
> The guest interaction proxy functionality is included in the Veeam Universal License. When using a legacy socket-based license, Enterprise or higher edition is required.



Application-Aware Processing

If you add to the backup job VMs running VSS-aware applications, you can enable application-aware processing to create a transactionally consistent backup. The transactionally consistent backup guarantees proper recovery of applications on VMs without data loss.

To enable application-aware processing:

1. At the **Guest Processing** step of the wizard, select the **Enable application-aware processing** check box.
2. Click **Applications**.
3. In the displayed list, select the VM and click **Edit**.

   To define custom settings for a VM added as a part of a VM container, you must include the VM to the list as a standalone object. To do this, click **Add** and choose the VM whose settings you want to customize. Then select the VM in the list and define the necessary settings.

4. On the **General** tab, in the **Applications** section specify the behavior scenario for application-aware processing:
   - Select **Require successful processing** if you want Veeam Backup & Replication to stop the backup process if any error occurs during application-aware processing.
   - Select **Try application processing, but ignore failures** if you want to continue the backup process even if an error occurs during application-aware processing. This option guarantees completion of the backup job. However, the resulting backup will not be transactionally consistent but crash consistent.

     > **IMPORTANT**
     >
     > If application-aware processing fails, the backup job will not process SQL transaction logs until Veeam Backup & Replication creates a new image-level backup of the Microsoft SQL Server VM.

   - Select **Disable application processing** if you do not want to enable application-aware processing for the VM.
5. [For Microsoft Exchange, Microsoft SQL and Oracle VMs] In the **Transaction logs** section, specify if Veeam Backup & Replication must process transaction logs or copy-only backups must be created.
   a. Select **Process transaction logs with this job** if you want Veeam Backup & Replication to process transaction logs.

      [For Microsoft Exchange VMs] With this option selected, the non-persistent runtime components or persistent components running on the VM guest OS will wait for backup to complete successfully and then trigger truncation of transaction logs. If the backup job fails, the logs will remain untouched on the VM guest OS until the next start of the non-persistent runtime components or persistent components.

      [For Microsoft SQL Server VMs and Oracle VMs] You will have to specify settings for transaction log handling on the **SQL** and **Oracle** tabs of the **VM Processing Settings** window. For more information, see Microsoft SQL Server Transaction Log Settings and Oracle Archived Log Settings.

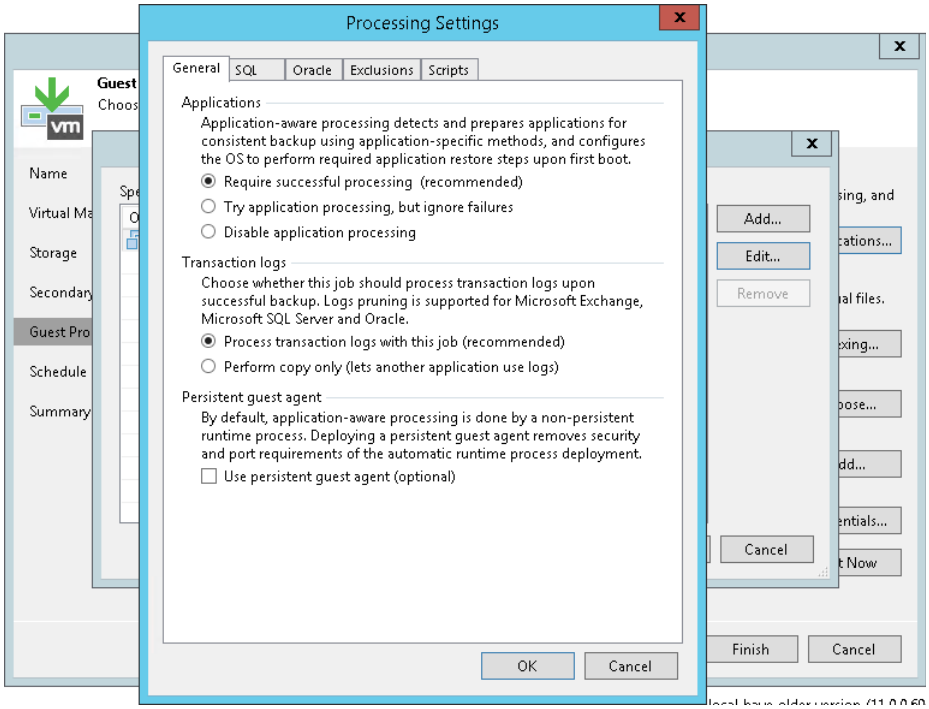
   b. Select **Perform copy only** if you use another backup tool to perform VM guest level backup, and this tool maintains consistency of the database state. Veeam Backup & Replication will create a copy-only backup for the selected VM. The copy only backup preserves the chain of full/differential backup files and transaction logs on the VM. For more information, see Microsoft Docs.
6. In the **Persistent guest agent** section, specify if Veeam Backup & Replication must use persistent guest agents on each protected VM for application-aware processing.

   By default, Veeam Backup & Replication uses non-persistent runtime components.
   Veeam Backup & Replication deploys runtime components on each protected VM when the backup job starts, and removes the runtime components as soon as the backup job finishes.

   Select the **Use persistent guest agent check** box to enable persistent agent components for guest processing. For more information, see Non-Persistent Runtime Components and Persistent Agent Components.

IMPORTANT

If both Microsoft SQL Server and Oracle Server are installed on one VM, and this VM is processed by a job with log backup enabled for both applications, Veeam Backup & Replication will back up only Oracle transaction logs. Microsoft SQL Server transaction logs will not be processed.



Microsoft SQL Server Transaction Log Settings

If you back up a Microsoft SQL VM, you can specify how Veeam Backup & Replication must process transaction logs on this VM:

1. At the **Guest Processing** step of the wizard, select the **Enable application-aware processing** check box.
2. Click **Applications**.
3. In the displayed list, select the Microsoft SQL Server VM and click **Edit**.
4. In the **Transaction logs** section, select **Process transaction logs with this job**.
5. In the **VM Processing Settings** window, click the **SQL** tab.
6. Specify how transaction logs must be processed. You can select one of the following options:
   - Select **Truncate logs** to truncate transaction logs after successful backup. The non-persistent runtime components or persistent components running on the VM guest OS will wait for the backup to complete successfully and then truncate transaction logs. If the job does not manage to back up the Microsoft SQL Server VM, the logs will remain untouched on the VM guest OS until the next start of the non-persistent runtime components or persistent components.

> **NOTE**
>
> If the account specified at the Guest Processing step does not have enough rights, Veeam Backup & Replication tries to truncate logs using the local SYSTEM account for Microsoft SQL Server 2008 and 2008 R2. For other Microsoft SQL Server versions, Veeam Backup & Replication uses NT AUTHORITY\SYSTEM account.
>
> Make sure that these accounts have permissions listed in Required Permissions for Guest Processing.

- Select **Do not truncate logs** to preserve transaction logs. When the backup job completes, Veeam Backup & Replication will not truncate transaction logs on the Microsoft SQL Server VM.

  It is recommended that you enable this option for databases that use the *Simple* recovery model. If you enable this option for databases that use the *Full* or *Bulk-logged* recovery model, transaction logs on the VM guest OS may grow large and consume all disk space. In this case, the database administrators must take care of transaction logs themselves.

- Select **Backup logs periodically** to back up transaction logs with Veeam Backup & Replication. Veeam Backup & Replication will periodically copy transaction logs to the backup repository and store them together with the image-level backup of the Microsoft SQL Server VM. During the backup job session, transaction logs on the VM guest OS will be truncated.

  For more information, see Microsoft SQL Server Log Backup.

If you have selected to back up transaction logs with Veeam Backup & Replication, you must specify settings for transaction log backup:

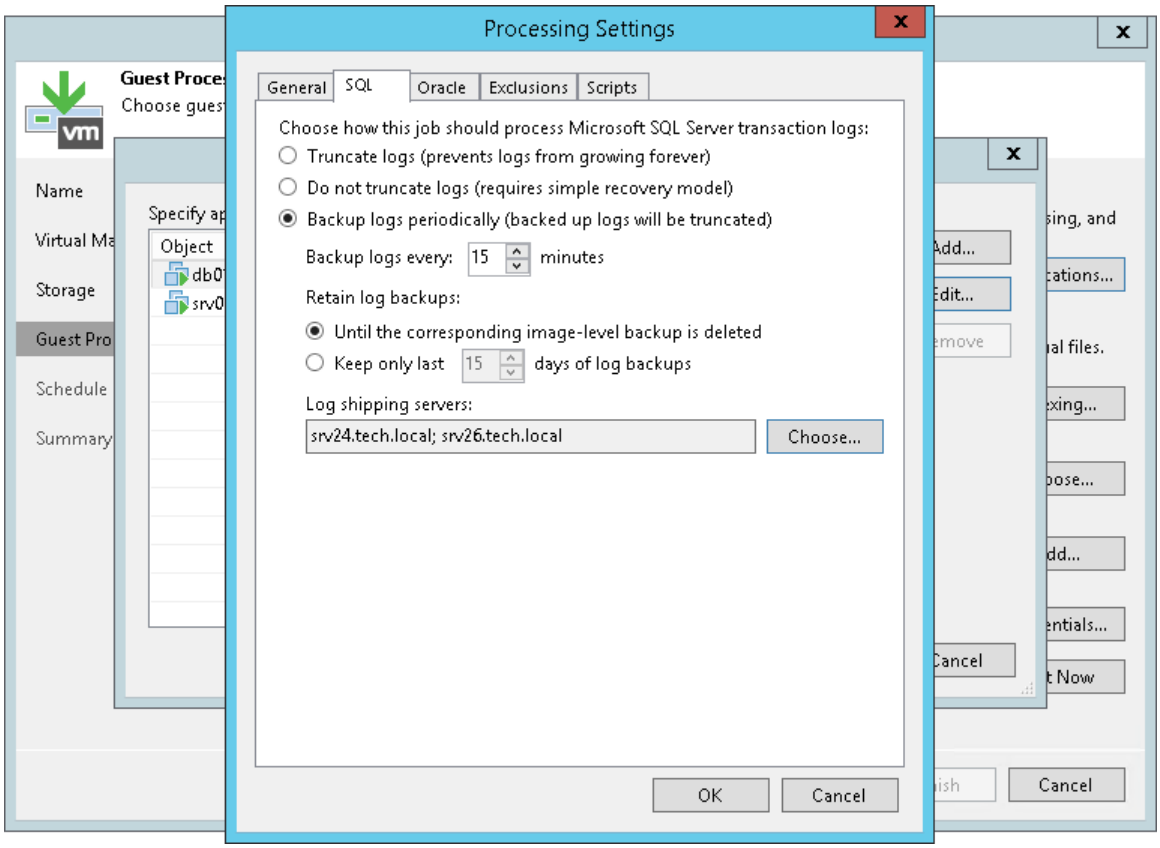
1. In the **Backup logs every <N> minutes** field, specify the frequency for transaction log backup. By default, transaction logs are backed up every 15 minutes. The maximum log backup interval is 480 minutes.
2. In the **Retain log backups** section, specify retention policy for transaction logs stored in the backup repository.
   - Select **Until the corresponding image-level backup is deleted** to apply the same retention policy for image-level backups and transaction log backups.
   - Select **Keep only last <N> days of log backups** to keep transaction logs for a specific number of days. By default, transaction logs are kept for 15 days. If you select this option, you must make sure that retention for transaction logs is not greater than retention for the image-level backups. For more information, see Retention for Transaction Log Backups.
3. In the **Log shipping servers** section, click **Choose** to select what log shipping server you want to use to transport transaction logs:
   - Select **Automatic selection** if you want Veeam Backup & Replication to choose an optimal log shipping server automatically. If the optimal shipping server is busy, Veeam Backup & Replication will direct the data flow to another shipping server not to lose data and to comply with RPO. The process of transaction logs shipment does not require a dedicated server — Veeam Backup & Replication can use any Microsoft Windows server added to the backup infrastructure.
   - To define a log shipping server explicitly, select **Use the specified servers only** and select check boxes next to servers that you want to use as log shipping servers. The server list includes all Microsoft Windows servers added to the backup infrastructure.

     Make sure that you select a server that is not engaged in other resource-consuming tasks. For example, you may want not to use a server that performs the WAN accelerator role as a log shipping server. For load balance and high availability purposes, it is recommended that you select at least 2 log shipping servers.

IMPORTANT

Veeam Backup & Replication automatically excludes its configuration database from application-aware processing during backup if the database is hosted without using SQL Server AlwaysOn Availability Group. Transaction logs for the configuration database are not backed up.

If the Veeam Backup & Replication configuration database is hosted using SQL Server AlwaysOn Availability Group, you should manually exclude this database from application-aware processing during backup as described in this Veeam KB article. Otherwise, job processing will fail with the following error: Failed to freeze guest over network, wait timeout.



Oracle Archived Log Settings

If you back up a VM where Oracle Database is deployed, you can specify how Veeam Backup & Replication must process archived redo logs on this VM:

1. At the **Guest Processing** step of the wizard, select the **Enable application-aware processing** check box.
2. Click **Applications**.
3. In the displayed list, select an Oracle VM and click **Edit**.
4. In the **Transaction logs** section, select **Process transaction logs with this job**.
5. In the **VM Processing Settings** window, click the **Oracle** tab.

6. In the **Specify Oracle account with SYSDBA privileges** section, specify a user account that Veeam Backup & Replication will use to connect to the Oracle database. The account must have privileges for the Oracle database listed in Required Permissions for Guest Processing.

   You can select **Use guest credentials** in the list of user accounts. In this case, Veeam Backup & Replication will use the account specified at the **Guest Processing** step of the wizard to access the VM guest OS and connect to the Oracle database.

7. In the **Archived logs** section, specify how Veeam Backup & Replication must process archived redo logs on the Oracle VM:

   - Select **Do not delete archived logs** if you want Veeam Backup & Replication to preserve archived logs on the VM guest OS. When the backup job completes, the non-persistent runtime components or persistent components will not delete archived logs.

     It is recommended that you select this option for databases in the NOARCHIVELOG mode. If the database is in the ARCHIVELOG mode, archived logs on the VM guest OS may grow large and consume all disk space. In this case, database administrators must take care of archived logs themselves.

   - Select **Delete logs older than <N> hours** or **Delete logs over <N> GB** if you want Veeam Backup & Replication to delete archived logs that are older than <N> hours or larger than <N> GB. The log size threshold refers not to the total size of all logs for all databases, but to the log size of each database on the selected Oracle VM.

     When the parent backup job (job creating an image-level backup) runs, Veeam Backup & Replication will wait for the backup to complete successfully, and then trigger archived logs deletion on the Oracle VM over Oracle Call Interface (OCI). If the primary job does not manage to back up the Oracle VM, the logs will remain untouched on the VM guest OS until the next start of the non-persistent runtime components or persistent components.

   > **TIP**
   >
   > Veeam Backup & Replication removes redo logs only after the parent backup job session. To remove redo logs more often, you can schedule the parent job to run more often.

8. To back up Oracle archived logs with Veeam Backup & Replication, select the **Backup log every <N> minutes** check box and specify the frequency for archived log backup. By default, archived logs are backed up every 15 minutes. The maximum log backup interval is 480 minutes.

9. In the **Retain log backups** section, specify retention policy for archived logs stored in the backup repository:

   - Select **Until the corresponding image-level backup is deleted** to apply the same retention policy for image-level backups and archived log backups.

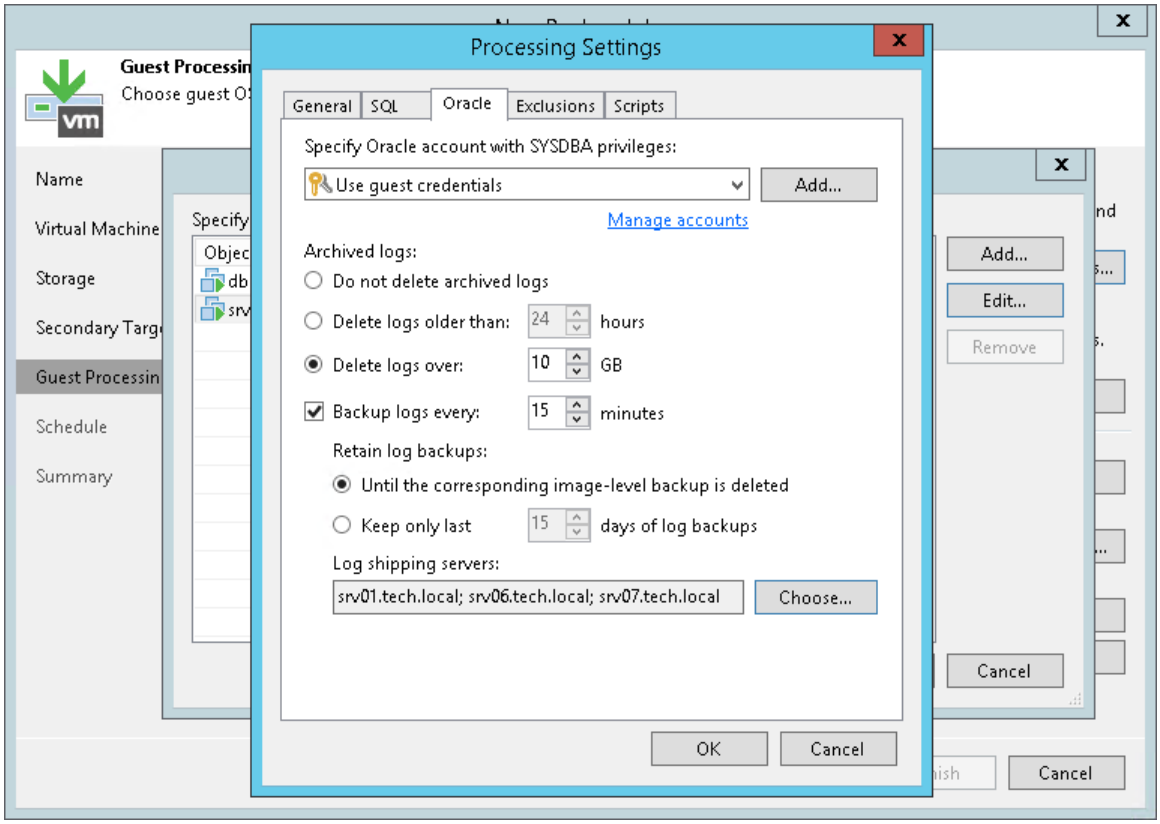
   - Select **Keep only last <n> days** to keep archived logs for a specific number of days. By default, archived logs are kept for 15 days. If you select this option, you must make sure that retention for archived logs is not greater than retention for the image-level backups. For more information, see Retention for Archived Log Backups.

10. In the **Log shipping servers** section, click **Choose** to select what log shipping server you want to use to transport archived logs:

    - Select **Automatic selection** if you want Veeam Backup & Replication to select an optimal log shipping server automatically. The process of archived logs shipment does not require a dedicated server — Veeam Backup & Replication can use any Microsoft Windows or Linux server added to the backup infrastructure.

- Select **Use the specified servers only** to define a log shipping server explicitly. In the **Log Shipping Servers** window, select check boxes next to servers that you want to use as log shipping servers. The server list includes all Microsoft Windows servers added to the backup infrastructure.

  Make sure that you select a server that is not engaged in other resource-consuming tasks. For example, you may want not to use a server that performs the WAN accelerator role as a log shipping server. For load balance and high availability purposes, it is recommended that you select at least 2 log shipping servers.



VM Guest OS File Exclusion

If you do not want to back up specific files and folders on the VM guest OS, you can exclude them from the backup.

> **NOTE**
>
> VM guest OS file exclusion functionality is included in the Veeam Universal License. When using a legacy socket-based license, Enterprise or higher edition is required.

To define what files and folders must be excluded:

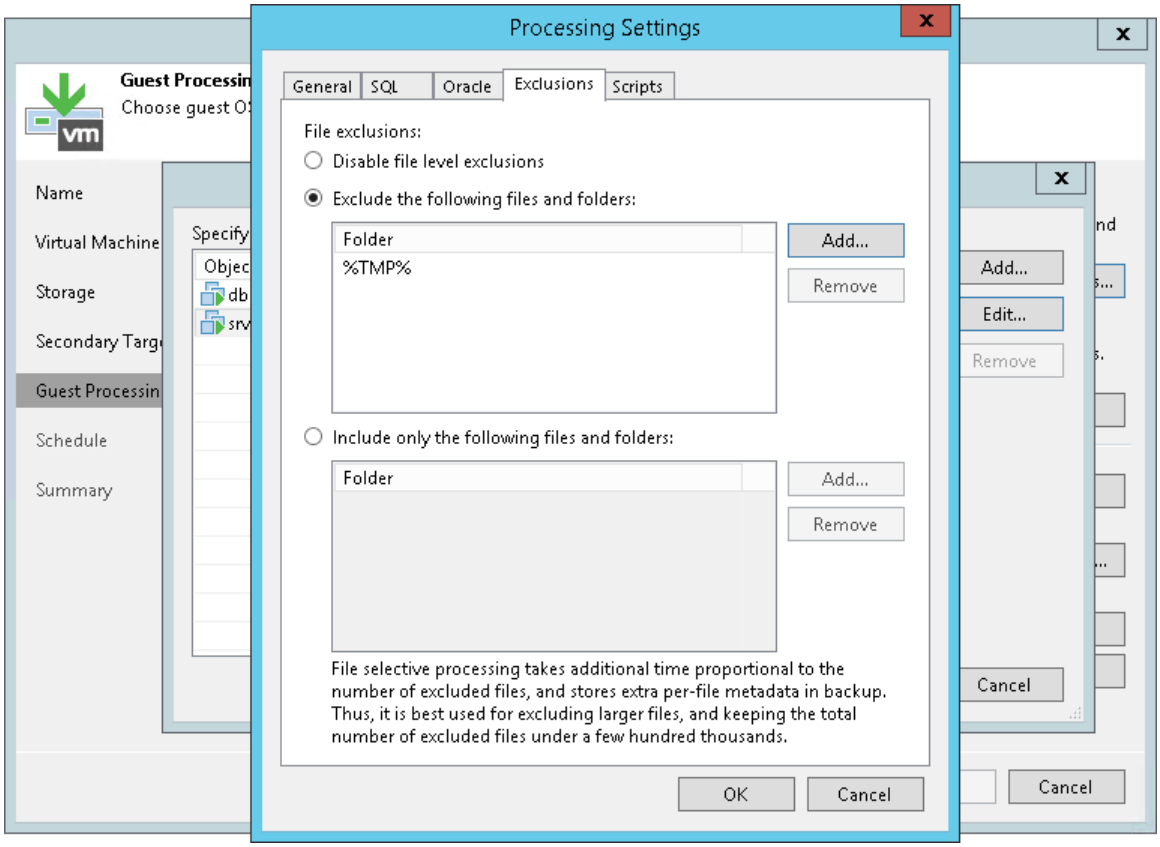
1. At the **Guest Processing** step of the wizard, select the **Enable application-aware processing** check box.
2. Click **Applications**.
3. In the displayed list, select the VM and click **Edit**.

   To define custom settings for a VM added as part of a VM container, you must include the VM in the list as a standalone object. To do this, click **Add** and choose a VM whose settings you want to customize. Then select the VM in the list and define the necessary settings.

4. Click the **Exclusions** tab and specify what files must be excluded from the backup:
   - Select **Exclude the following files and folders** to remove the individual files and folders from the backup.
   - Select **Include only the following files and folders** to leave only the specified files and folders in the backup.
5. Click **Add** and specify what files and folders you want to include or exclude. To form the list of exclusions or inclusions, you can use full paths to files and folders, environmental variables and file masks with the asterisk (*) and question mark (?) characters. For more information, see VM Guest OS Files.

> **NOTE**
>
> When you choose files to be included or excluded, mind requirements and limitations that are listed in the Requirements and Limitations for VM Guest OS File Exclusion section.

6. Click **OK**.
7. Repeat steps 5-6 for every file or folder that you want to exclude or include.



Pre-Freeze and Post-Thaw Scripts

If you plan to back up VMs running applications that do not support VSS, you can specify what scripts Veeam Backup & Replication must use to quiesce the VM. The pre-freeze script quiesces the VM file system and application data to bring the VM to a consistent state before Veeam Backup & Replication triggers a VM snapshot. After the VM snapshot is created, the post-thaw script brings the VM and applications to their initial state.

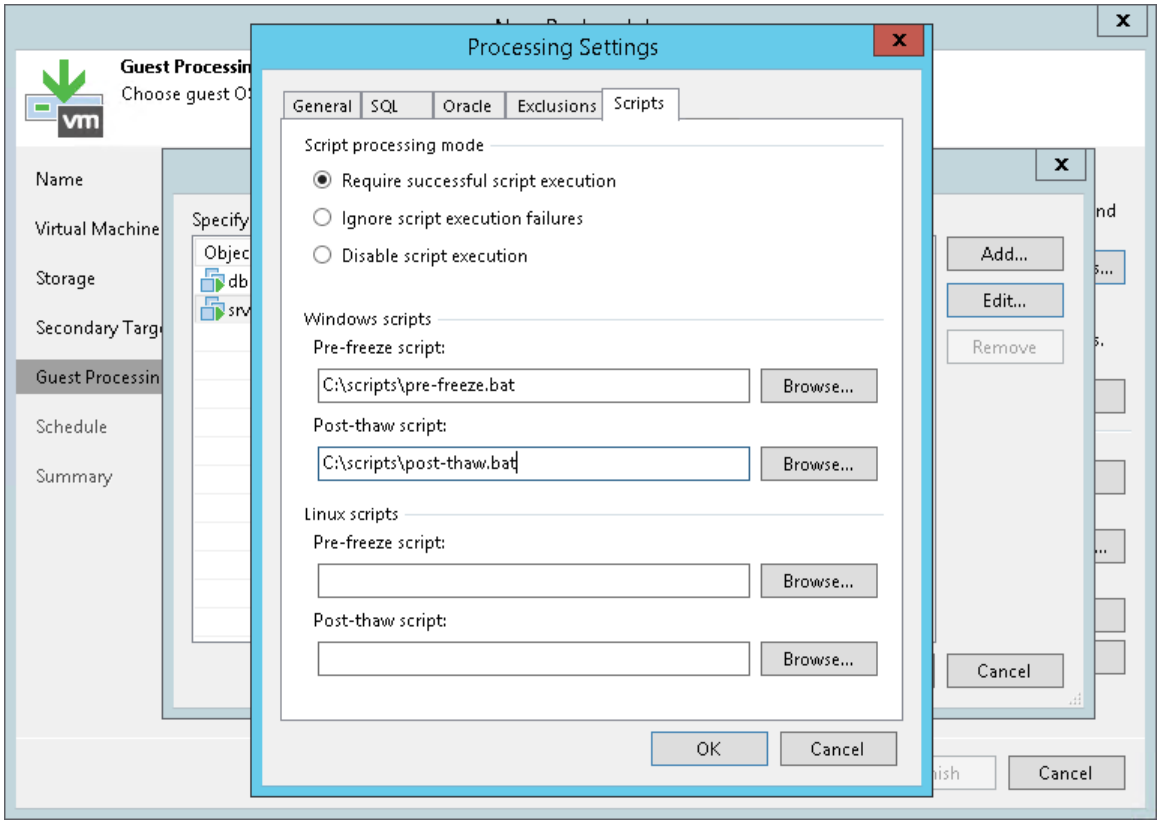
To specify pre-freeze and post-thaw scripts for the job:

1. At the **Guest Processing** step, click **Applications**.
2. In the displayed list, select the VM and click **Edit**.
3. Click the **Scripts** tab.
4. In the **Script processing mode** section, specify the scenario for scripts execution:
   - Select **Require successful script execution** if you want Veeam Backup & Replication to stop the backup process if the script fails.
   - Select **Ignore script execution failures** if you want to continue the backup process even if script errors occur.
   - Select **Disable script execution** if you do not want to run scripts for the VM.
5. In the **Windows scripts** section, specify paths to pre-freeze and post-thaw scripts for Microsoft Windows VMs. For the list of supported script formats, see Pre-Freeze and Post-Thaw Scripts.
6. In the **Linux scripts** section, specify paths to pre-freeze and/or post-thaw scripts for Linux VMs. For the list of supported script formats, see Pre-Freeze and Post-Thaw Scripts.

   If you have added to the job a VM container with Microsoft Windows and Linux VMs, you can select to execute both Microsoft Windows and Linux scripts for the VM container. When the job starts, Veeam Backup & Replication will automatically determine what OS type is installed on the VM and use corresponding scripts to quiesce this VM.

TIP

Beside pre-freeze and post-thaw scripts for VM quiescence, you can instruct Veeam Backup & Replication to run custom scripts before the job starts and after the job completes. For more information, see Script Settings.



VM Guest OS File Indexing

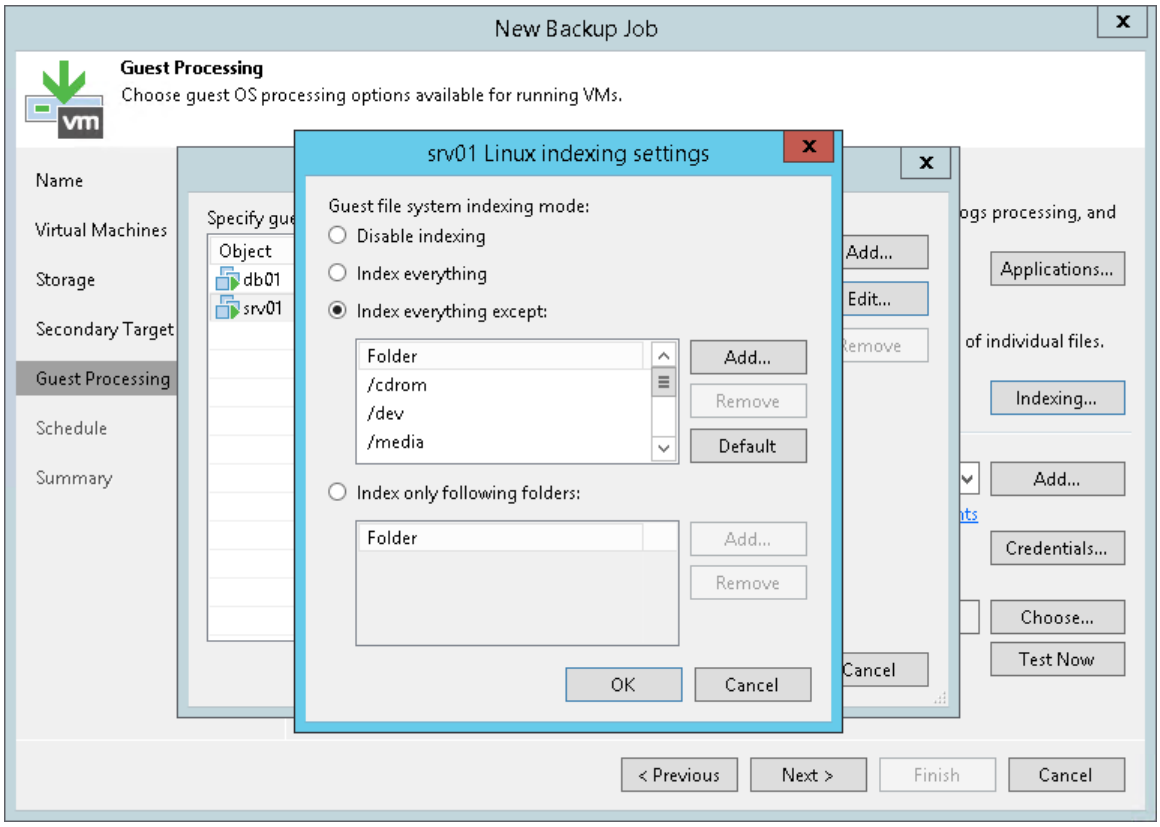
To specify VM guest OS indexing options for a VM:

1. At the **Guest Processing** step of the wizard, click **Indexing**.
2. Select a VM in the list and click **Edit** > **Windows indexing** or **Linux indexing**.
3. Specify the indexing scope:
   - Select **Disable indexing** if you do not want to index guest OS files of the VM.
   - Select **Index everything** if you want to index all VM guest OS files.
   - Select **Index everything except** if you want to index all VM guest OS files except those defined in the list. By default, system folders are excluded from indexing. You can add or delete folders using the **Add** and **Remove** buttons on the right. You can also use system environment variables to form the list, for example: *%windir%, %ProgramFiles%* and *%Temp%*.

     To reset the list of folders to its initial state, click **Default**.
   - Select **Index only following folders** to define folders that you want to index. You can add or delete folders to index using the **Add** and **Remove** buttons on the right. You can also use system environment variables to form the list, for example: *%windir%, %ProgramFiles%* and *%Temp%*.

NOTE

To perform guest OS file indexing on Linux VMs, Veeam Backup & Replication requires several utilities to be installed on the Linux VM: openssh, mlocate, gzip and tar. If these utilities are not found, Veeam Backup & Replication will prompt you to deploy them on the VM guest OS.



Step 11. Define Job Schedule

At the **Schedule** step of the wizard, select to run the backup job manually or schedule the job to run on a regular basis.

To specify the job schedule:

1. Select the **Run the job automatically** check box. If this check box is not selected, you will have to start the job manually to create the VM backup.
2. Define scheduling settings for the job:
   - To run the job at specific time daily, on defined week days or with specific periodicity, select **Daily at this time**. Use the fields on the right to configure the necessary schedule.
   - To run the job once a month on specific days, select **Monthly at this time**. Use the fields on the right to configure the necessary schedule.

     > **NOTE**
     >
     > When you configure the job schedule, keep in mind possible date and time changes (for example, related to daylight saving time transition).

   - To run the job repeatedly throughout a day with a specific time interval, select **Periodically every**. In the field on the right, select the necessary time unit: *Hours* or *Minutes*. Click **Schedule** and use the time table to define the permitted time window for the job. In the **Start time within an hour** field, specify the exact time when the job must start.

     A repeatedly run job is started by the following rules:

     - Veeam Backup & Replication always starts counting defined intervals from 12:00 AM. For example, if you configure to run a job with a 4-hour interval, the job will start at 12:00 AM, 4:00 AM, 8:00 AM, 12:00 PM, 4:00 PM and so on.
     - If you define permitted hours for the job, after the denied interval is over, Veeam Backup & Replication will immediately start the job and then run the job by the defined schedule.

     For example, you have configured a job to run with a 2-hour interval and defined permitted hours from 9:00 AM to 5:00 PM. According to the rules above, the job will first run at 9:00 AM, when the denied period is over. After that, the job will run at 10:00 AM, 12:00 PM, 2:00 PM and 4:00 PM.

   - To run the job continuously, select the **Periodically every** option and choose **Continuously** from the list on the right. A new backup job session will start as soon as the previous backup job session finishes.
   - To chain jobs, use the **After this job** field. In the common practice, jobs start one after another: when job *A* finishes, job *B* starts and so on. If you want to create a chain of jobs, you must define the time schedule for the first job in the chain. For the rest of the jobs in the chain, select the **After this job option** and choose the preceding job from the list.
3. In the **Automatic retry** section, define whether Veeam Backup & Replication must attempt to run the backup job again if the job fails for some reason. During a job retry, Veeam Backup & Replication processes failed VMs only. Enter the number of attempts to run the job and define time intervals between them. If you select continuous backup, Veeam Backup & Replication will retry the job for the defined number of times without any time intervals between the job runs.

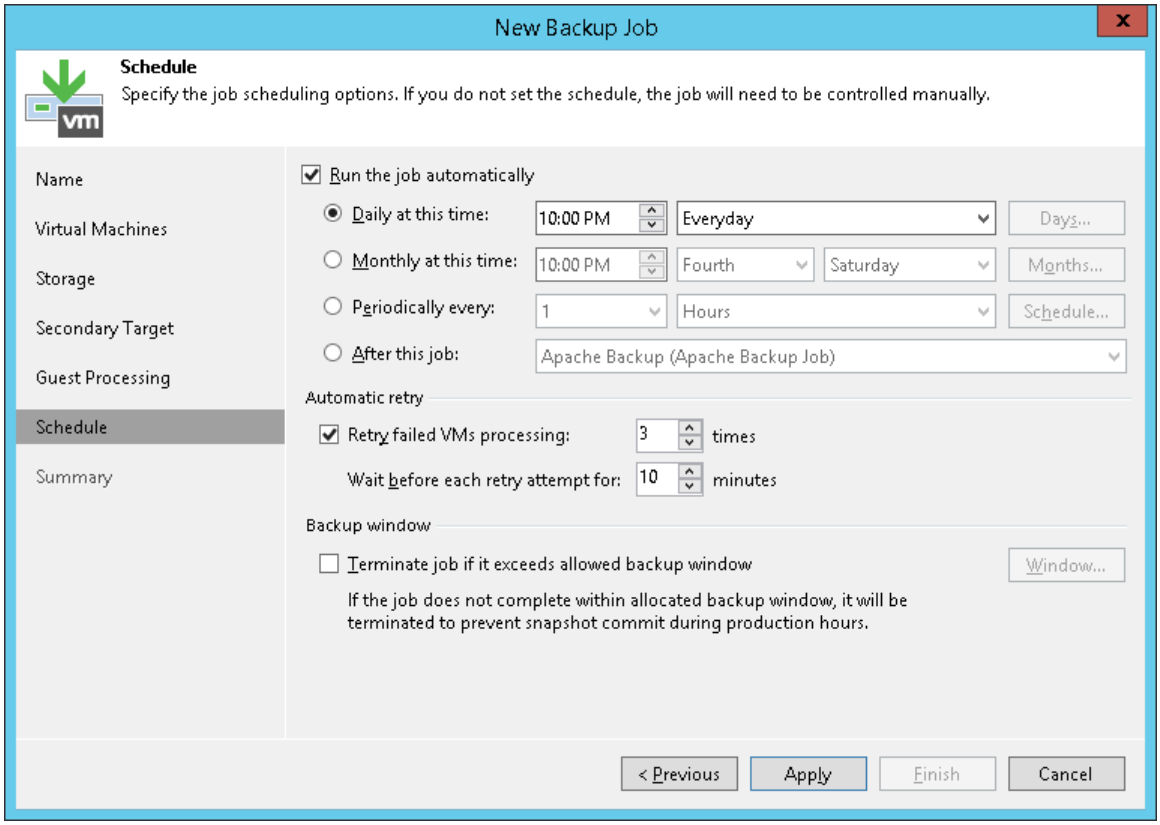
4. In the **Backup window** section, define the time interval within which the backup job must complete. The backup window prevents the job from overlapping with production hours and ensures that the job does not provide unwanted overhead on the production environment. To set up a backup window for the job:

   a. Select the **Terminate job if it exceeds allowed backup window** check box and click **Window**.

   b. In the **Time Periods** window, define the allowed hours and prohibited hours for backup. If the job exceeds the allowed window, it will be automatically terminated.

NOTE

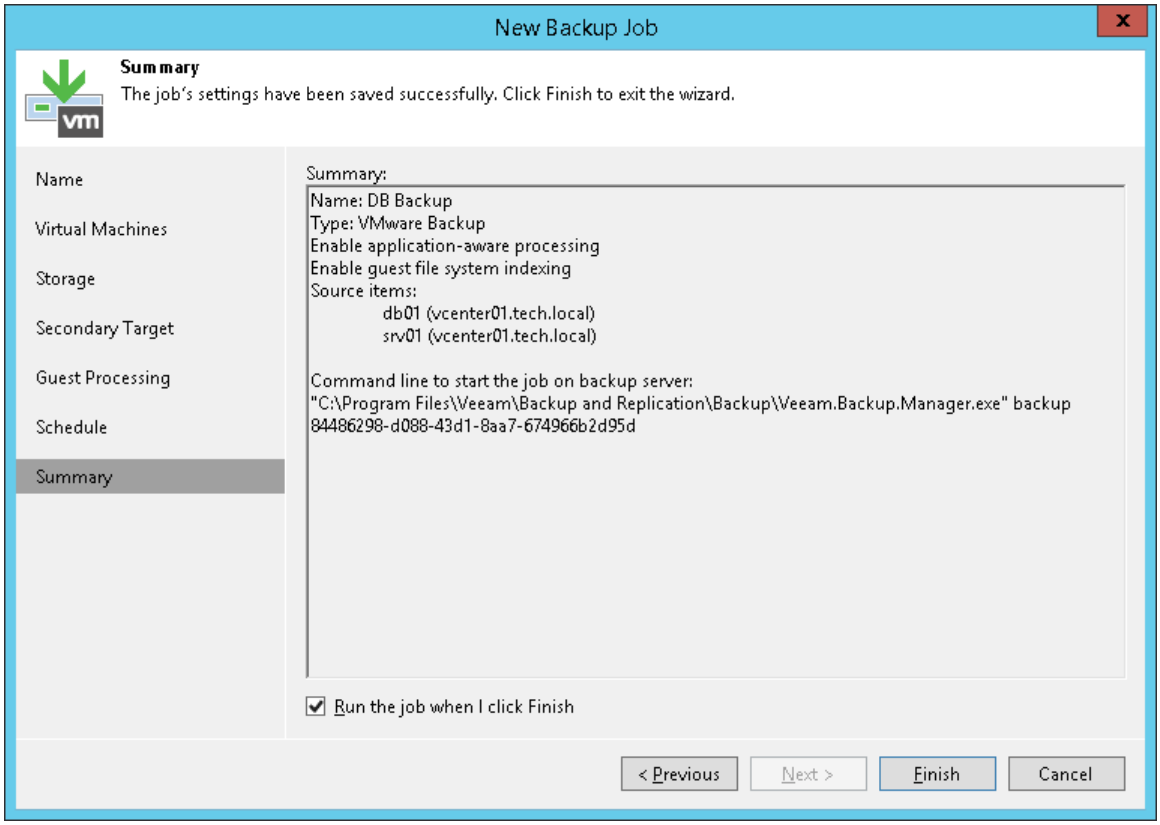
The **After this job** function will automatically start a job if the first job in the chain is started automatically by schedule. If you start the first job manually, Veeam Backup & Replication will display a notification. You will be able to choose whether Veeam Backup & Replication must start the chained job as well.



Step 12. Finish Working with Wizard

At the **Summary** step of the wizard, complete the procedure of backup job configuration.

1. Review details of the backup job.
2. Select the **Run the job when I click Finish** check box if you want to start the job right after you finish working with the wizard.
3. Click **Finish** to close the wizard.

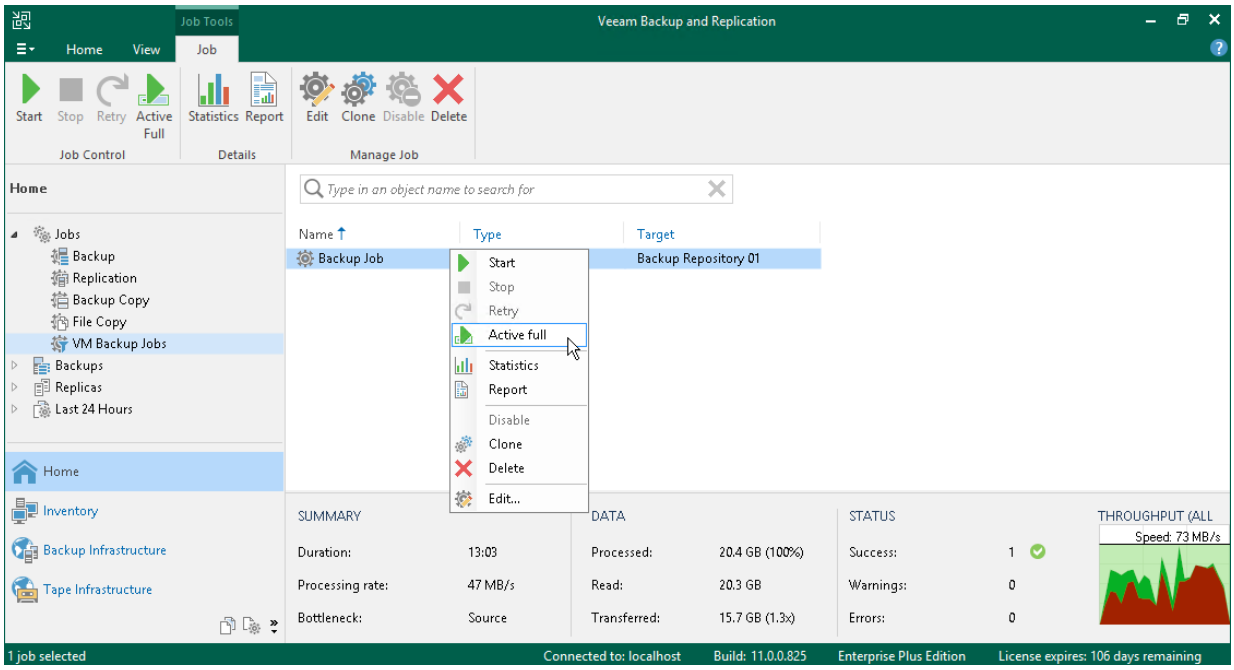


Performing Active Full Backup

You can create an ad-hoc full backup — active full backup, and add it to the backup chain in the backup repository. The active full backup resets the backup chain. All subsequent incremental backups use the active full backup as a starting point. The previously used full backup will remain in the backup repository until it is removed from the backup chain according to the retention policy.

To perform active full backup:

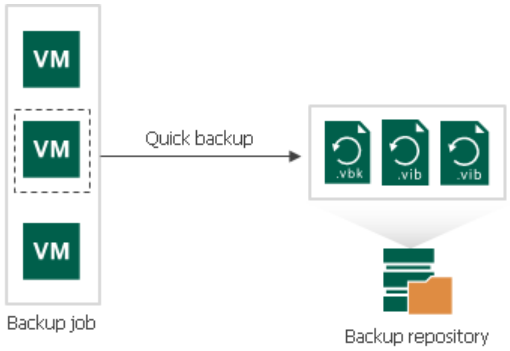
1. Open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, select the job and click **Active Full** on the ribbon or right-click the job and select **Active Full**.



Quick Backup

Quick backup lets you perform on-demand incremental backup for VMs. You can use quick backup if you want to produce an additional restore point for one or more VMs in a backup job and do not want to configure a new job or modify the existing one. Quick backup can be run for both incremental and reverse incremental backup chains.

Quick backup is an incremental backup task: Veeam Backup & Replication copies only changed data for selected VMs and saves this data to a new restore point in the backup chain. Similar to incremental backup, quick backup can only be run for VMs that have been successfully backed up at least once and has a full restore point. If there is no full restore point for a VM, quick backup cannot be performed.



To perform quick backup, Veeam Backup & Replication uses an existing backup job. When you start a quick backup task for a VM, Veeam Backup & Replication verifies that a backup job processing this VM exists on the backup server. If such job is detected, Veeam Backup & Replication triggers a job and creates an incremental restore point for the VM. If a backup job for the VM does not exist, quick backup is terminated.

You can run quick backup for one VM or more VMs at once. If you start quick backup for several VMs and these VMs are processed by different backup jobs, Veeam Backup & Replication triggers a set of backup jobs. Each triggered job creates a separate restore point and stores this restore point in a corresponding backup chain.

In some cases, a VM may be processed by several backup jobs on the backup server. In this case, Veeam Backup & Replication starts the job that has created the most recent restore point for the VM.

For example, *VM01* is processed by 2 jobs:

- *Backup job 1* created the most recent restore point on Monday
- *Backup job 2* created the most recent restore point on Tuesday

When you start quick backup for *VM01*, Veeam Backup & Replication will trigger *Backup job 2* to create a new incremental restore point.

> **NOTE**
>
> If the quick backup task overlaps the scheduled backup job, the backup job waits for the quick backup task to complete.

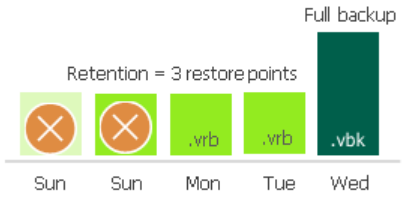
Limitations for Quick Backup

You cannot perform quick backup for vCloud Director VMs processed with vCloud Director jobs. However, if you process a vCloud Director VM with a regular backup job, you can switch to the **Computer** view and start the quick backup operation for this VM.

Retention Policy for Quick Backups

When you perform quick backup, Veeam Backup & Replication creates a single VM incremental restore point. Unlike a regular incremental restore point that contains data for all VMs in a job, single VM incremental restore point contains data only for a specific VM.

A single VM restore point is not regarded as full-fledged restore point in the backup chain. From the retention policy perspective, single VM restore point is grouped with a regular restore point following it. When Veeam Backup & Replication needs to delete a single VM restore point by retention, it waits for the next regular restore point to expire — that is, Veeam Backup & Replication increases the retention by one restore point for some time. After the next regular restore point expires, Veeam Backup & Replication deletes two restore points at once.



If the backup chain is stored in a repository with the Use per-machine backup files option enabled, the retention increases in a different way. As a rule, the retention increases by the number of VMs from this chain for which quick backup was performed. This applies to the reverse incremental and forward incremental backup chains.

Performing Quick Backup

You can create an ad-hoc incremental backup for one or more VMs — quick backup, and add it to the backup chain in the backup repository. Quick backup can be helpful if you want to produce an additional restore point for one or more VMs in the backup job and do not want to configure a new job or modify the existing one.

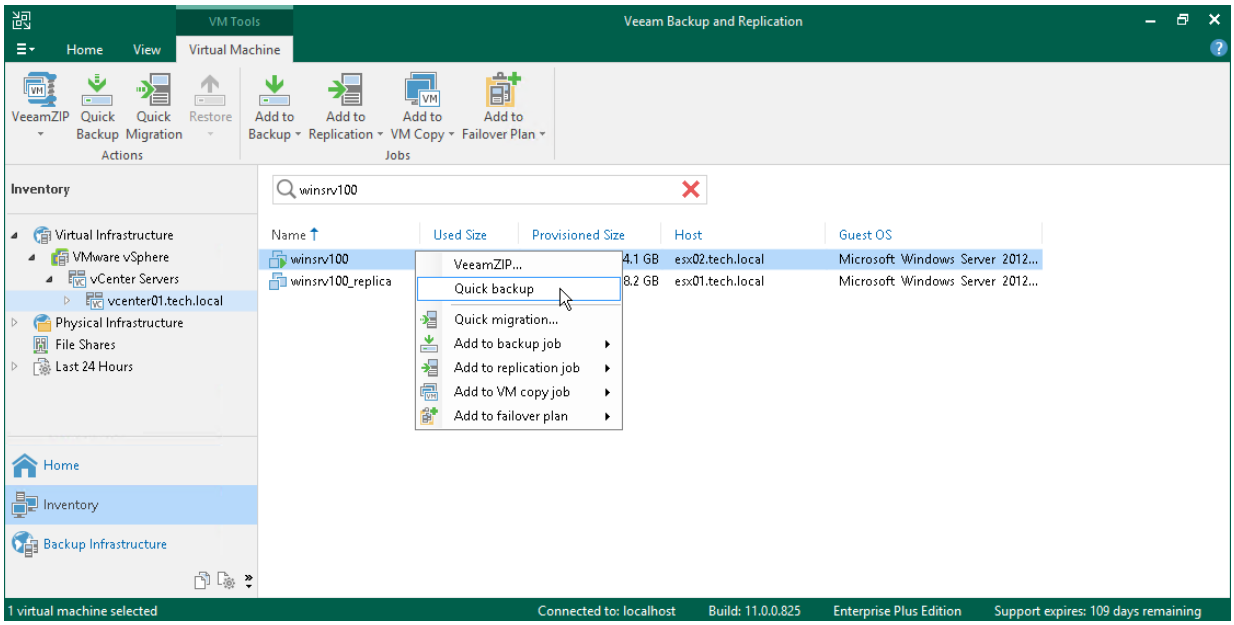
Quick backup can be performed for VMs that meet the following requirements:

1. A backup job processing the VM exists on the backup server.
2. A full backup file for the VM exists in the backup repository configured in the backup infrastructure.

To perform quick backup:

1. Open the **Inventory** view.
2. In the infrastructure tree, select a host or VM container (host, cluster, folder, resource pool, VirtualApp, datastore or tag) in which the VMs that you want to back up reside.
3. In the working area, select the VMs and click **Quick Backup** on the ribbon. You can also right-click the VMs and select **Quick Backup**.

   Veeam Backup & Replication will trigger a backup job to create a new incremental restore point for selected VMs. Details of a running quick backup task are displayed in the job session window.



Importing Backups Manually

You may need to import backups to Veeam Backup & Replication in the following situations:

- The backup server has failed and you have restored it in a new location. You want to restore VM data from backups created by the backup server that has failed.
- You want to restore VM data from backups created on another backup server.
- You want to restore VM data from backups in the backup repository that is not added to the backup infrastructure (for example, if you removed it earlier).
- You want to restore VM data from VeeamZIP files created on your backup server or another backup server.

The imported backup becomes available in the Veeam Backup & Replication console. You can use any restore operation to recover VM data from this backup.

> **IMPORTANT**
>
> You can import backups created by a newer version of Veeam Backup & Replication, however, restore operations for this backup will not be available.

Before you import a backup, check the following prerequisites:

- The server from which you plan to import backups must be added to the backup infrastructure. Otherwise you will not be able to access backup files.
- To be able to restore VM data from previous backup restore points, make sure that you have all required incremental backup files (forward or reverse) in the same folder where the full backup file resides.

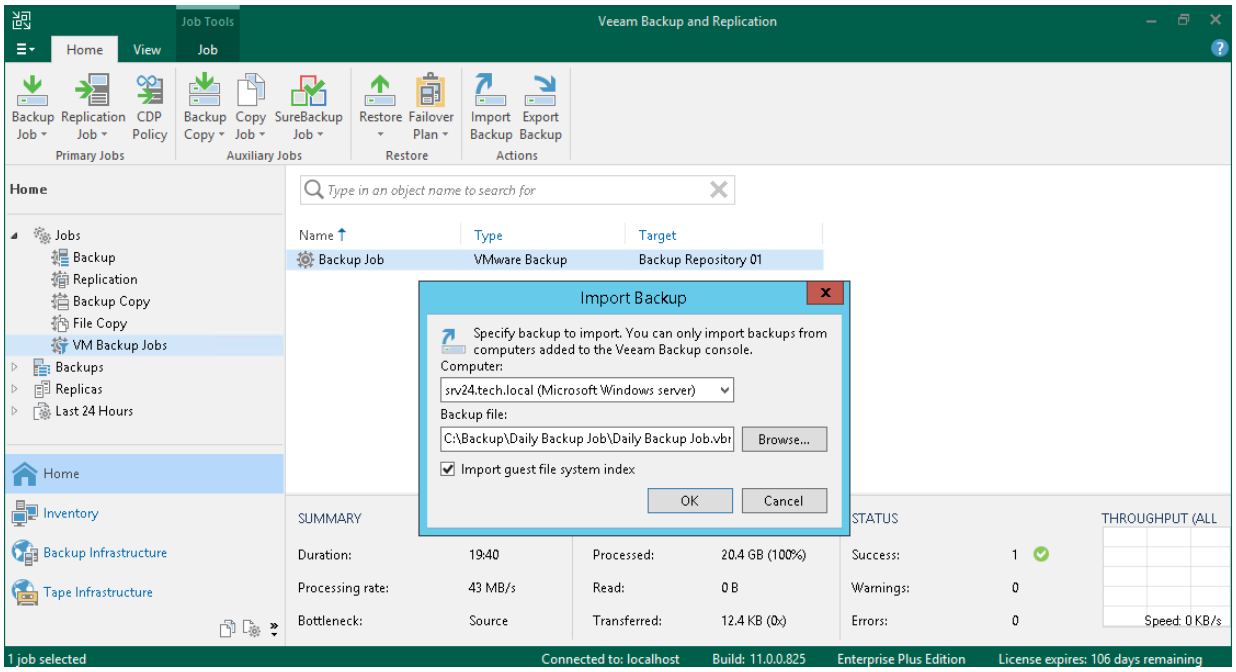
To import a backup to the Veeam Backup & Replication console:

1. On the **Home** tab, click **Import Backup**.
2. From the **Computer** list, select the server on which the backup you want to import is stored.
3. Click **Browse** and select the necessary VBM or VBK file. If you select the VBM file, the import process will be notably faster. It is recommended that you select the VBK file only if the VBM file is not available.
4. By default, index data of the VM guest OS file system is not imported with the backup to speed up the import process. If you want to import index data, select the **Import guest file system index** check box.
5. Click **OK** to import the backup. The imported backup will be displayed in the **Home** view, under the **Backups** > **Imported** node in the inventory pane. Backups are imported using the original name of the backup job with the `_imported` suffix appended.

TIP

If you need to import all backups stored on a server, assign a backup repository role to it and enable import at the Review step of the wizard. If the repository is already added to the backup infrastructure, you can rescan it. Veeam Backup & Replication will automatically import backups. For more information, see Rescanning Backup Repositories.

Note that Veeam Backup & Replication will not be able to import backups if VBM files are not available. In this case you will have to import backups manually using the VBK files.



Importing Encrypted Backups

You can import backups that were encrypted on this backup server or on another backup server.

To import an encrypted backup file:

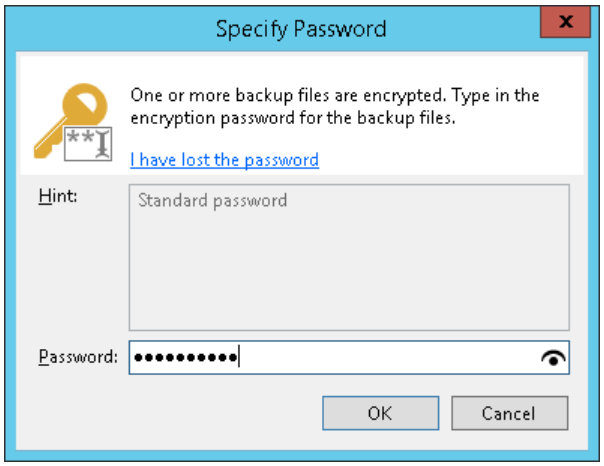
1. On the **Home** tab, click **Import Backup**.
2. From the **Computer** list, select the host on which the backup you want to import is stored.
3. Click **Browse** and select the VBM or VBK file.
4. Click **OK**. The encrypted backup will appear under the **Backups** > **Disk (encrypted)** node in the inventory pane.
5. In the working area, select the imported backup and click **Specify Password** on the ribbon or right-click the backup and select **Specify password**.
6. In the **Password** field, enter the password for the backup file.

   If you changed the password one or several times while the backup chain was created, you must enter passwords in the following manner:

   - If you select a VBM file for import, you must specify the latest password that was used to encrypt files in the backup chain.
   - If you select a VBK file for import, you must specify the whole set of passwords that were used to encrypt files in the backup chain.

If you enter correct passwords, Veeam Backup & Replication will decrypt the backup file. The backup will be moved under the **Backups** > **Disk (Imported)** node in the inventory pane.

> **NOTE**
>
> You can recover data from encrypted backups even if the password is lost. Restoring data without a password is included in the Veeam Universal License. When using a legacy socket-based license, Enterprise or higher edition is required. Also, your backup server must be connected to Veeam Backup Enterprise Manager. For more information, see Decrypting Data Without Password.



Importing Transaction Logs

You cannot import transaction log backups without VM backups (as there will be no restore point to which the transaction logs can be applied).

To import a VM backup with transaction log backups, do either of the following:

- Import a backup metadata file (VBM). In this case, Veeam Backup & Replication will automatically import the database backup and log backups.
- Import a full backup file (VBK). In this case, Veeam Backup & Replication will browse to corresponding log backups and import them, too.

Importing Backup Files from Scale-Out Backup Repositories

You cannot import a backup directly from the scale-out backup repository. When you perform backup import, you cannot browse through all extent of the scale-out backup repository. Veeam Backup & Replication lets you browse only through individual extents.

To import a backup from the scale-out backup repository, you must place backup files from all extents to one staging folder. The staging folder can reside on any server added to the backup infrastructure. After that, you can import the backup as usual.

> **TIP**
>
> If you need to import all backups stored in a scale-out backup repository, rescan the repository. In this case, you do not need to place files in one folder, Veeam Backup & Replication will import backups automatically. For more information, see Rescanning Backup Repositories.

Exporting Backups

Exporting backups allows you to synthesize a complete and independent full backup file out of selected restore points that are located in your backup repositories. That is, you can transform any incremental or reverse-incremental backup chain (i.e all dependent .VBK, .VIB or .VRB files) into a standalone .VBK file.

Export applies to *Full, Incremental* and *Reverse-incremental* restore points located in:

- Backup repositories.
- Object storage repositories.
- Backup repositories of cloud service providers and their tenants.

Consider the following:

- The restore point that is being exported as a new full backup file is saved to the same repository, wherein the source selected restore points reside.
- Once export is complete, exported backup files will be displayed under the **Backups** > **Disk (VeeamZIP)** node.
- [For scale-out backup repositories] Backup files exported to the capacity tier will be displayed under the **Backups** > **Object Storage (VeeamZIP)** node, and backup files exported to the archive tier will be displayed under the **Backups** > **Archive (VeeamZIP)** node.
- If a restore point that is being exported resides on the tenant side, a new full backup file will also be exported to the same repository (on the tenant side) from which the source restore point is being taken.
- If a tenant initiates export of a restore point that resides in the *subtenant* directory, a new full backup file will be exported to the *tenant* directory.
- If you select a backup job consisting of multiple virtual machines, Veeam will synthesize a separate full backup file per each machine.
- When exporting VMs from vCloud Director (vCD) backups, all the VMs will be exported without vApps, that is, a new full backup file will be exported as a simple VMWare backup, not vCD backup. For more information about vCD backups, see Backup of vCloud Director VMs.
- Export session results are saved to the configuration database and available for viewing, as described in Viewing Session Statistics.

Performing Export

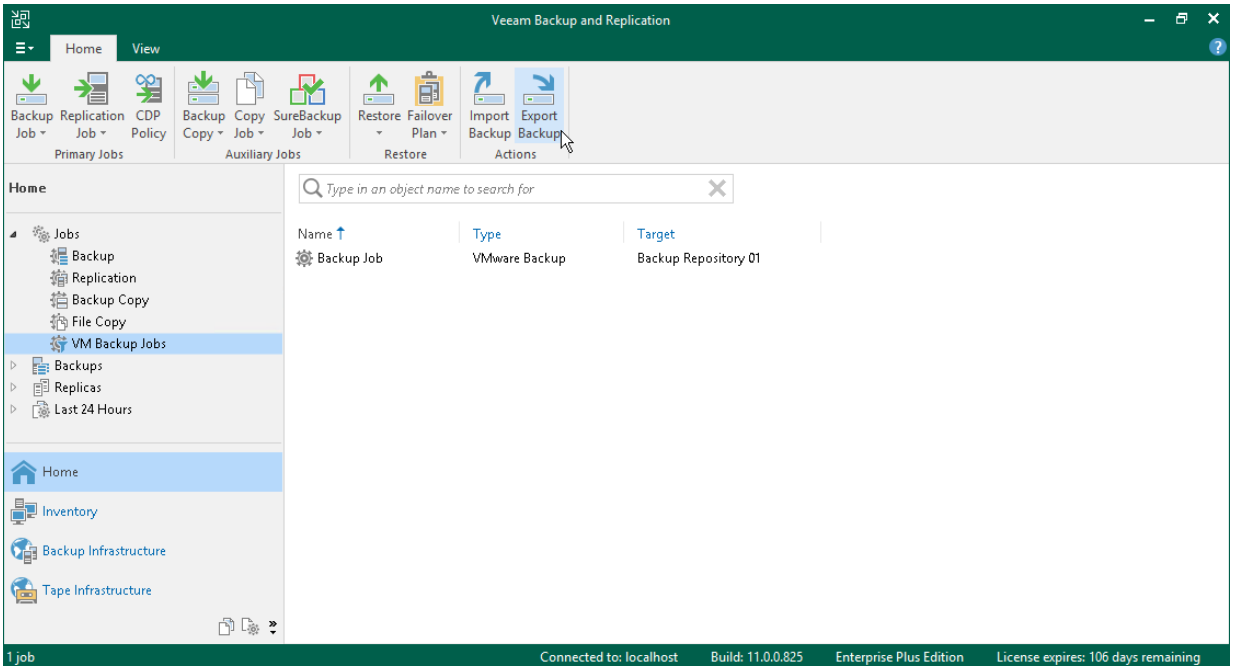
To export data, do the following:

1. Launch New Export Wizard
2. Select Restore Points to Export
3. Specify Export Reason
4. Finish Working with Wizard

Step 1. Launch New Export Wizard

To launch the **New Export** wizard, do either of the following:

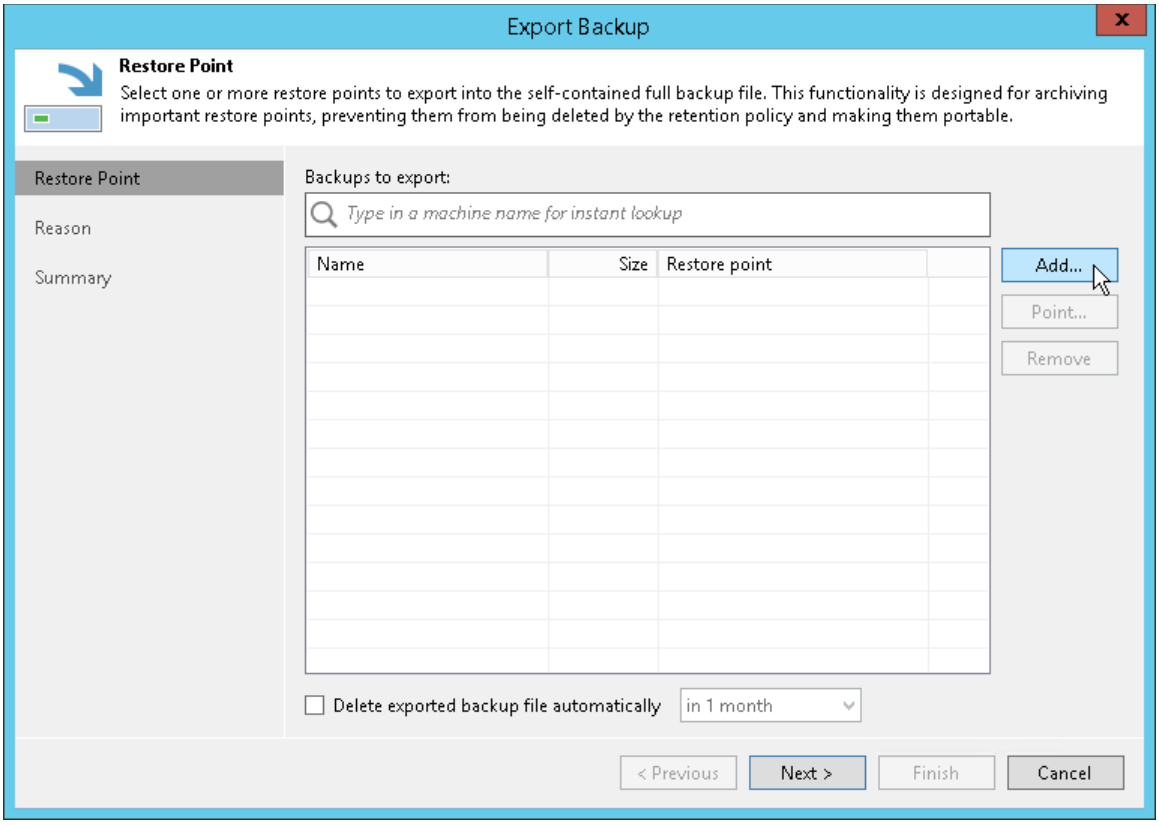
- On the **Home** tab, click **Export Backup**.
- In the **Home** view under the **Backups > Disks** node, select a VM you want to transform into a full backup file and click **Export backup**.



Step 2. Select Restore Points to Export

At the **Restore Point** step of the wizard, do the following:

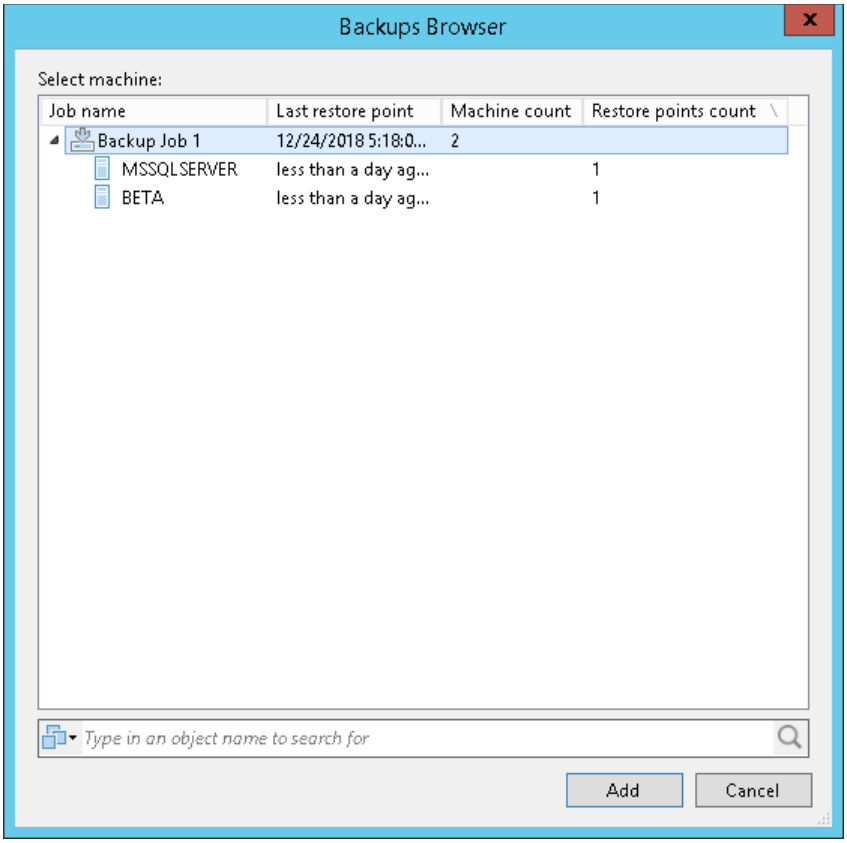
1. Click **Add** to select a VM, the restore points of which you want to transform into full backup files.



2. In the **Backups Browser** dialog, select a backup job or virtual machine.

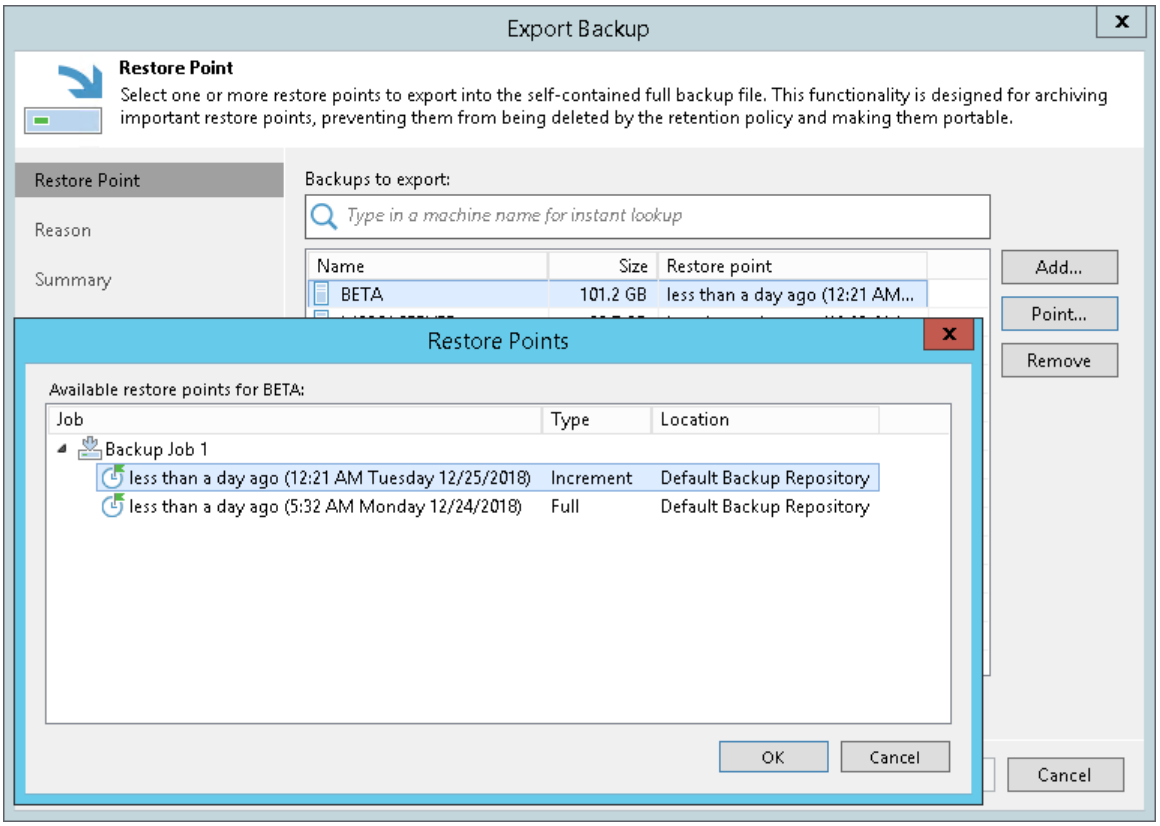
   When selecting a backup job consisting of multiple machines, then each machine will be exported as an independent full backup file.

Use the search field at the bottom of the dialog to find particular VMs in the list.



3. Select a VM from the table and click **Point** to select a restore point that you want to transform into a full backup file.

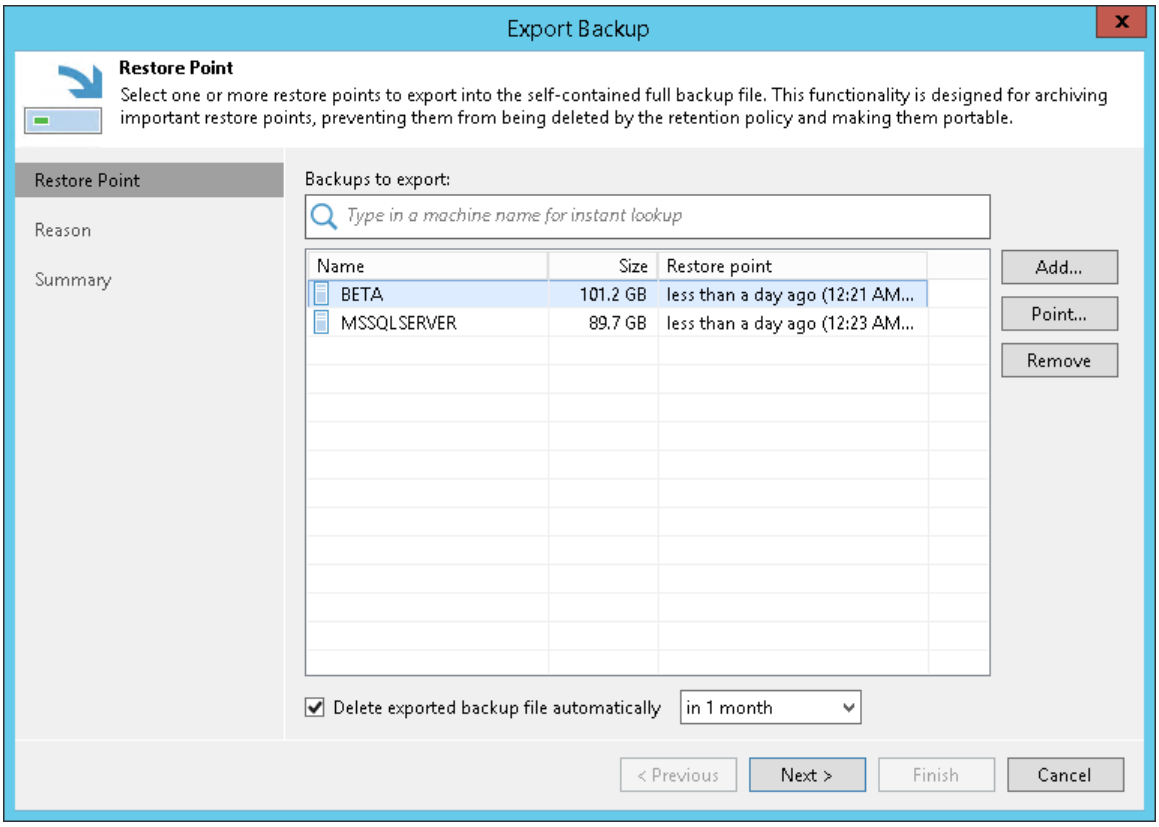
   By default, the latest available restore point is selected.



4. Select the **Delete exported backup file automatically** check box and set up the value in the drop-down list to apply a retention policy for files that are being exported.

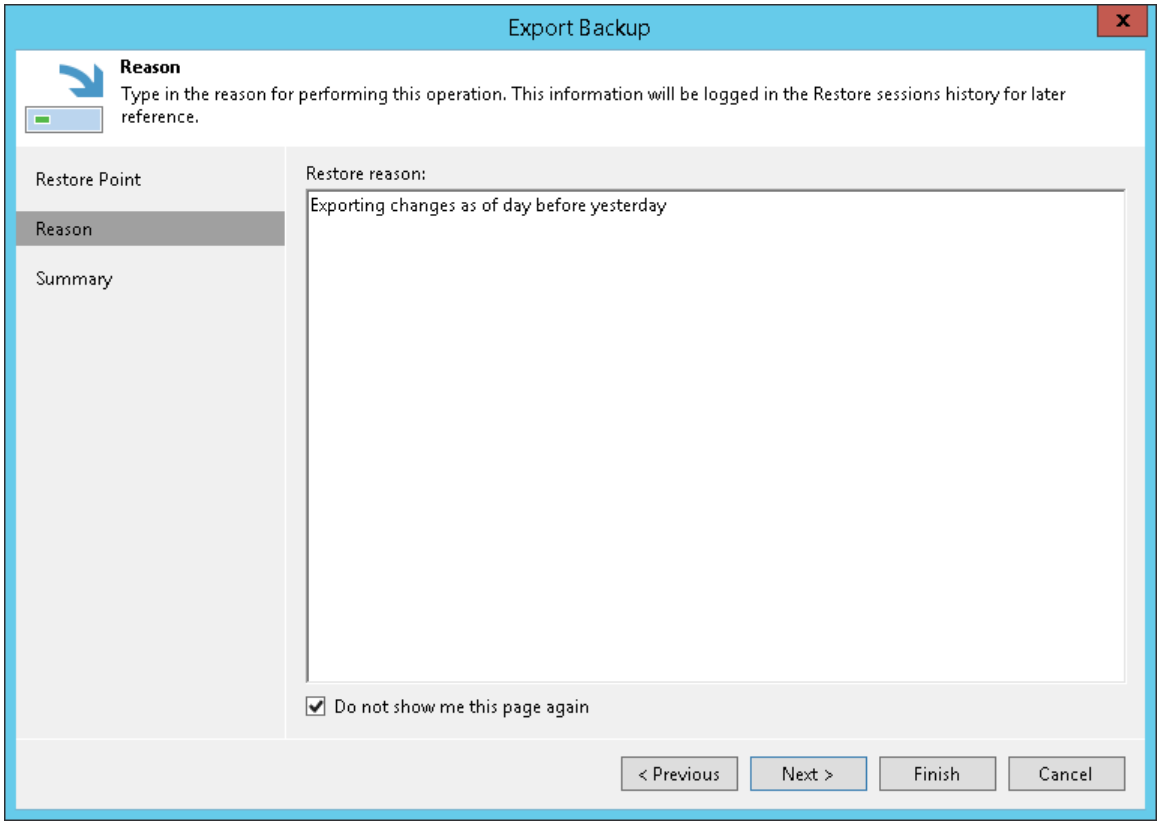
   All the exported backup files that fall under the retention policy will be removed automatically.

   To remove a restore point from the list, select a machine and click **Remove**.



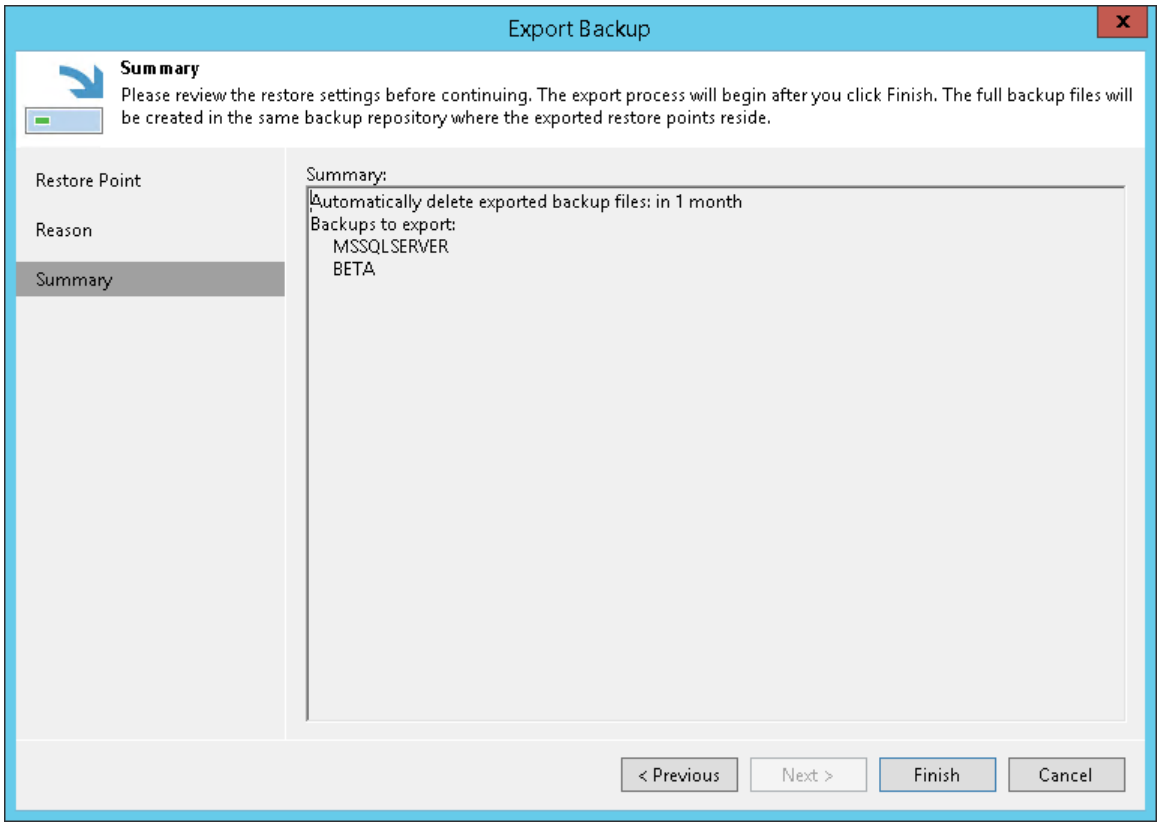
Step 3. Specify Export Reason

At the **Reason** step of the wizard, provide the reason of restore.



Step 4. Finish Working with Wizard

At the **Summary** step of the wizard, review summary information, click **Finish** and wait until the restore session, which is described in Viewing Session Statistics, is complete.



Viewing Session Statistics

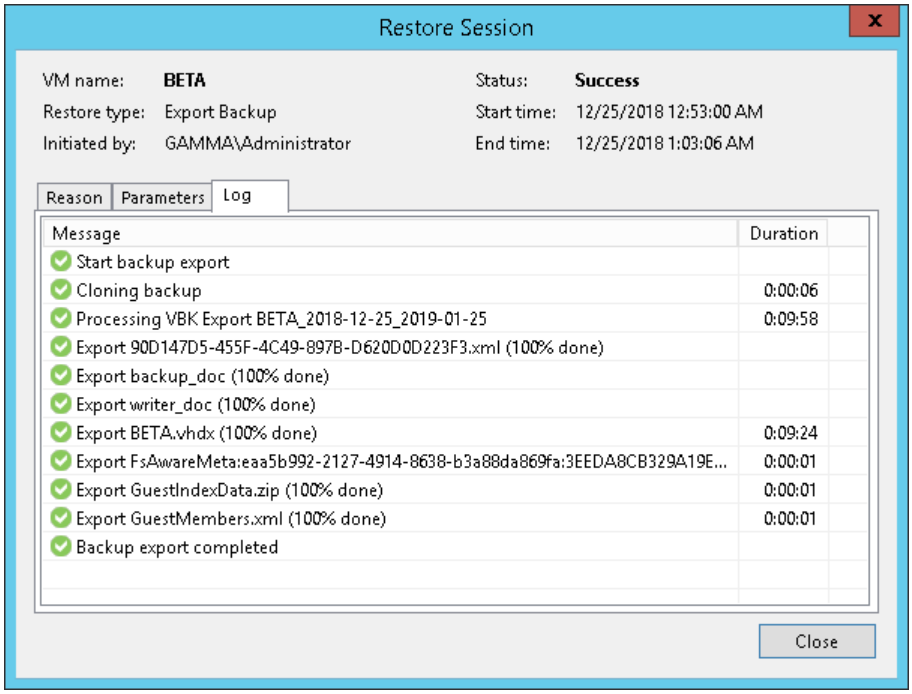
Once you invoke the export procedure, Veeam shows the **Restore Session** progress dialog that informs you of the current export status.

You can close the dialog by clicking the **Close** button in the lower-right corner and let Veeam perform export in the background.

As each export session saves its results to the configuration database, you can review them at any time.

To review the export session results, do the following:

1. In the inventory pane, go to the **History** view and select the **Restore** > **Export** node.
2. In the working area, double-click a machine for which you want to review the session results or right-click a machine and select **Statistics**.

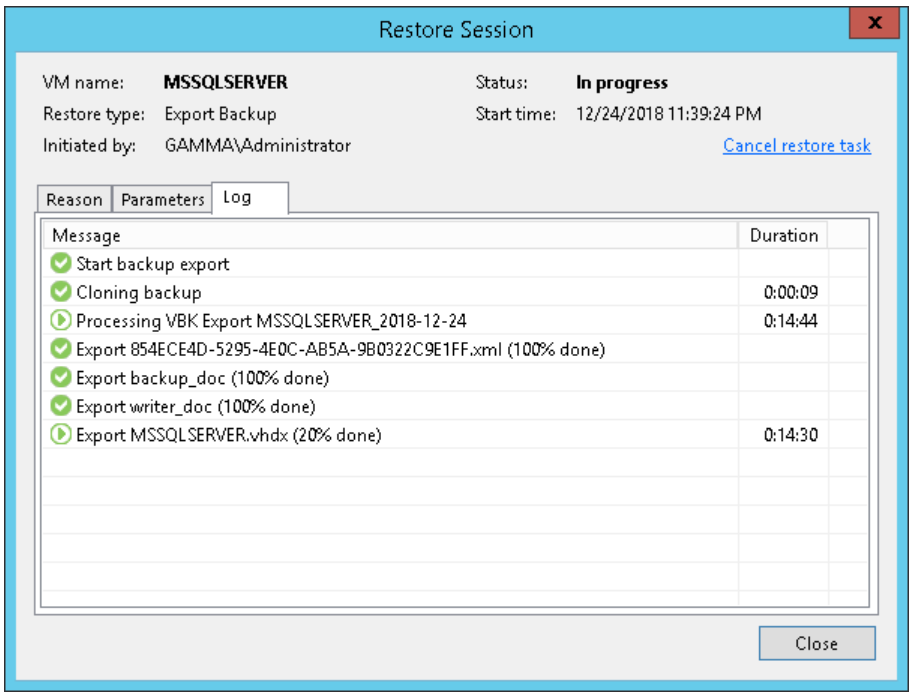


The **Restore Session** dialog contains the following tabs:

- The **Reason** tab — shows you the reason of export you may have provided at the Specify Export Reason step of the wizard.
- The **Parameters** tab — shows you the date when the exported backup files will be removed due to the retention policy you may have configured at the Select Restore Points to Export step of the wizard. In this tab you can also find a backup name and Date/time of a restore point that was synthesized into a full backup file.
- The **Log** tab — shows you the actual export progress.

Canceling Session

To cancel a session, open the **Restore Session** dialog, as described above, and click **Cancel restore task** in the upper-right corner of the dialog.



Managing Backups

You can perform the following operations with backups:

- View backup properties
- Remove a backup from configuration
- Delete a backup from disks
- Delete a backup from object storage repository
- Remove missing restore points

Viewing Backup Properties

You can view summary information about created backups. The summary information provides the following data:

- Available restore points
- Date of restore points creation
- Compression and deduplication ratios
- Data size and backup size
- GFS retention policy applied to restore points (W — weekly; M — monthly; Y — yearly)

In the **Backup Properties** window, you can see the following icons:

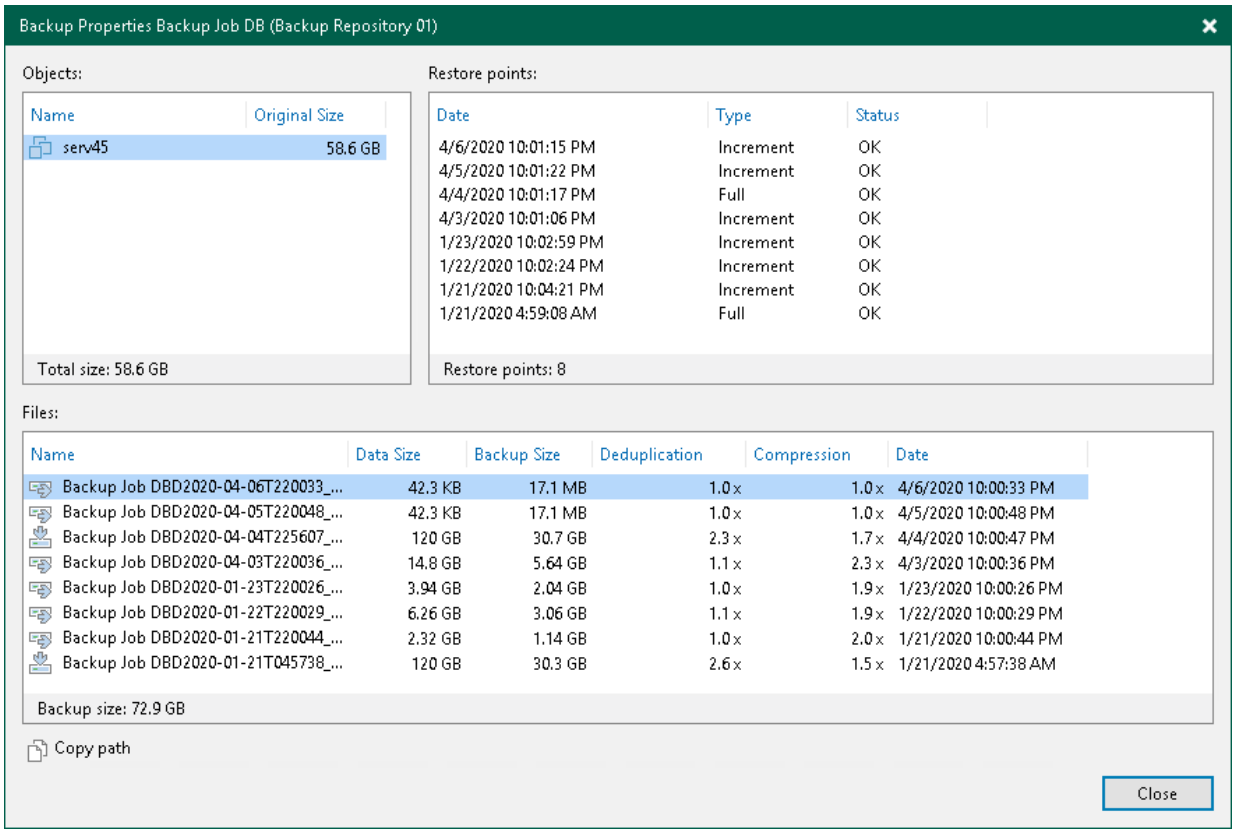
Icon	State
	Full restore point
	Incremental restore point
	Reverse incremental restore point
	Missing full restore point
	Missing incremental restore point
	Missing reverse incremental restore point

For backups stored in the scale-out backup repository icons differ. For more information, see Backup State Indicators.

To view summary information for backups:

1. Open the **Home** view.
2. In the inventory pane, select **Backups**.
3. In the working area, right-click the backup and select **Properties**.

4. To see the list of available restore points, select the required object from the **Objects** list.



Removing Backups from Configuration

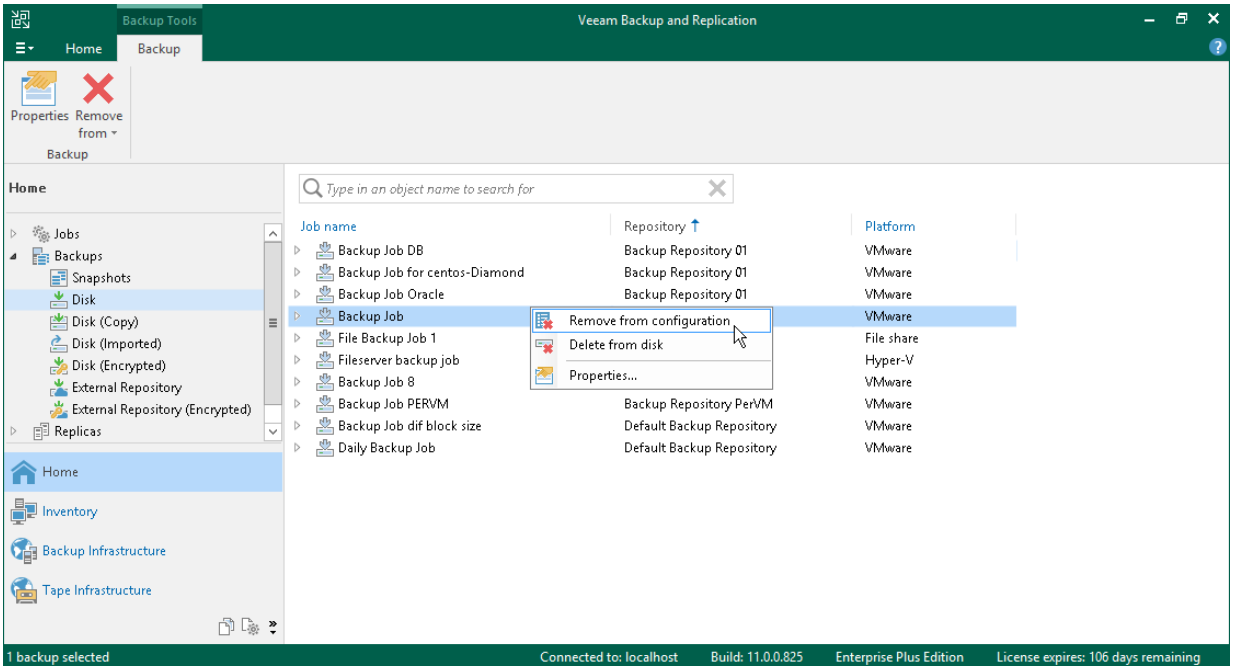
If you want to remove records about backups from the Veeam Backup & Replication console and configuration database, you can use the **Remove from configuration** operation.

When you remove a backup from the configuration, backup files (VBK, VIB, VRB, VBM, VBLOB) remain in the backup repository. You can import the backup later and restore VM data/files and folders from it.

When you remove an encrypted backup from configuration, Veeam Backup & Replication removes encryption keys from the configuration database. If you import such backup on the same backup server or another backup server, you will have to specify the password or unlock the backup with Veeam Backup Enterprise Manager. For more information, see Importing Encrypted Backups.

To remove a backup from the configuration:

1. Open the **Home** view.
2. In the inventory pane, select **Backups** or **Replicas**.
3. In the working area, select the backup and click **Remove from** > **Configuration** on the ribbon. You can also right-click the backup and select **Remove from configuration**.



Deleting Backups from Disk

The **Delete from disk** operation is needed if you want to remove records about backups from the Veeam Backup & Replication console and configuration database. This option allows you to delete the following type of data:

- Backup files from the backup repository
- Separate VMs from backups

When you delete backup files from a disk, Veeam Backup & Replication deletes the whole chain from the backup repository. Thus, on the next run of the backup job, Veeam Backup & Replication will create full backups for VMs included in the job.

Mind the following:

- Do not delete backup files from the backup repository manually. Use the **Delete from disk** option instead. If you delete backup files manually, subsequent backup or replication job sessions will fail.
- If the per-machine functionality is enabled, you can perform the **Delete from disk** operation for separate VMs in the backup. If you delete backup files for one VM, on the next run of the backup job Veeam Backup & Replication will create a full backup for VMs whose backup files are deleted. For all other VMs, Veeam Backup & Replication will create increments.

  To learn more about per-machine backup files, see Per-Machine Backup Files
- When you delete a separate VMs from a backup and if the per-machine functionality is enabled for this VM, Veeam Backup & Replication behaves differently:
  - [If per-machine is enabled] Veeam Backup & Replication deletes backup files of the selected VMs.

    On the next run of the job, Veeam Backup & Replication will create full backups for VMs whose backup files were deleted. For all other VMs, Veeam Backup & Replication will create increments.
  - [If per-machine is disabled] Veeam Backup & Replication only marks data block that belong to the deleted VMs as empty — the size of backup files does not change. However, Veeam Backup & Replication will use these data blocks during such operations as merging backup files. To reduce the size of full backup files in forever forward incremental and reverse incremental backup chains, you can compact full backup files. In forward incremental backup chains, files with block marked as empty will be deleted by retention.

    On the next run of the job, Veeam Backup & Replication will create full backups for VMs whose backup files were deleted. Note that full backups of these VMs will be stored in an incremental file. For all other VMs, Veeam Backup & Replication will create incremental backups.
- If you use the scale-out backup repository, keep in mind that the **Delete from disk** operation will remove the backups not only from the performance tier, but also from the capacity and archive tier. If you want to remove backups from the performance tier only, you should move those backups to the capacity tier instead. For details, see Moving to Capacity Tier.

Deleting Backups

To delete backup files from the backup repository, do the following:

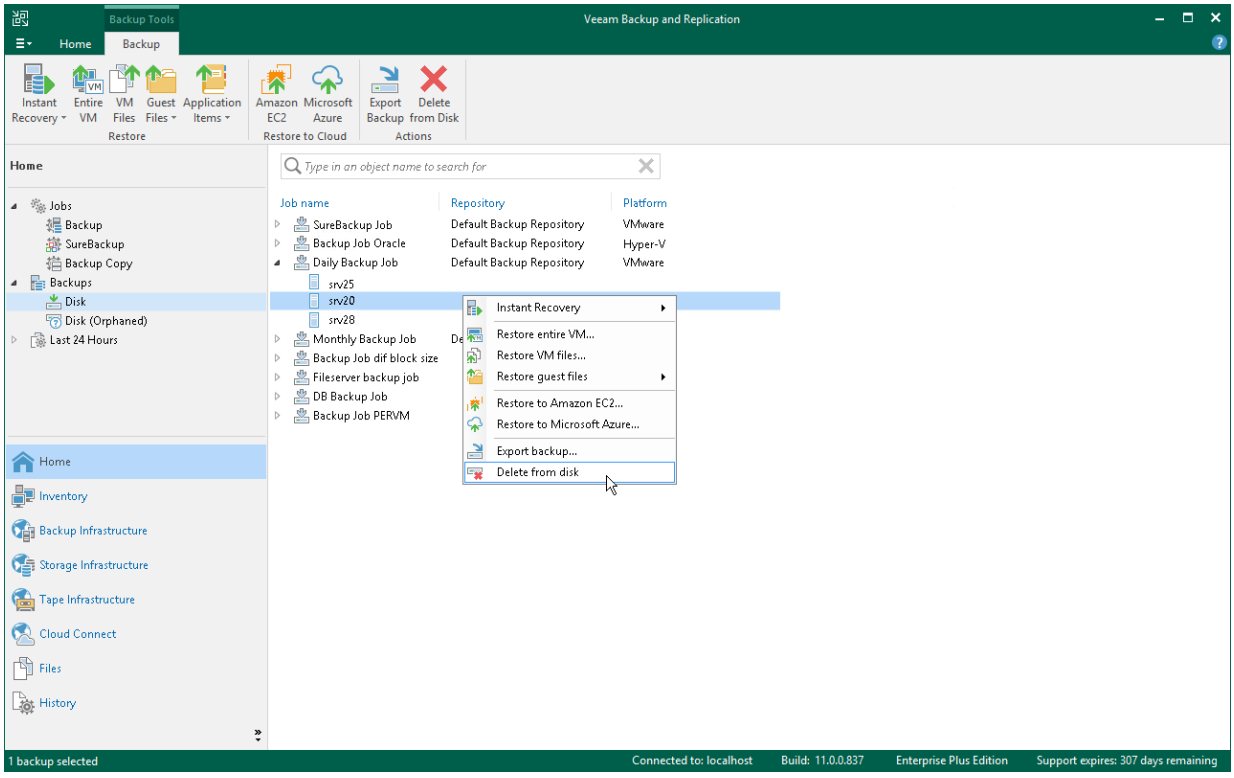
1. Open the **Home** view.
2. In the inventory pane, select **Backups** or **Replicas**.
3. In the working area, select the backup and click **Remove from** > **Disk** on the ribbon. You can also right-click the backup and select **Delete from disk**.

4. To remove backups with GFS flags (weekly, monthly and yearly), select the **Remove GFS full backups** check box and click **Yes**.

Deleting VM from Backups

To delete a VM from a backup, do the following:

1. Open the **Home** view.
2. In the inventory pane, select **Backups** or **Replicas**.
3. In the working area, expand the necessary backup, select the VM you want to delete and click **Remove from** > **Disk** on the ribbon. You can also right-click the backup and select **Delete from disk**.
4. To remove backups with GFS flags (weekly, monthly and yearly), select the **Remove GFS full backups** check box and click **Yes**.

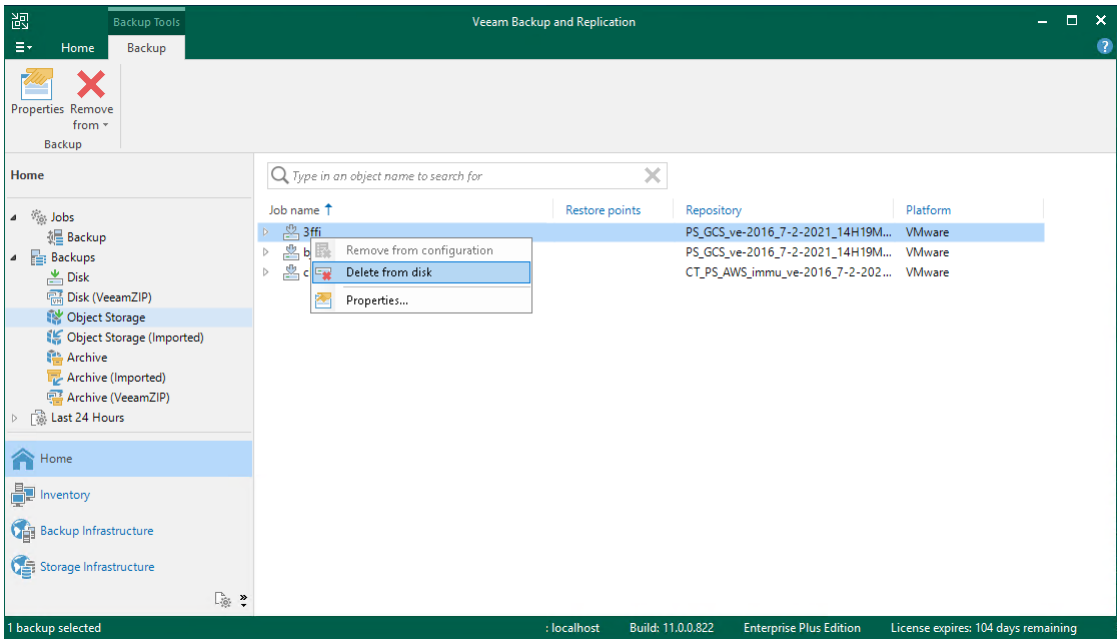


Deleting Backups from Object Storage

This section explains how to delete offloaded backups from object storage repositories. For more information, see Removing Backups from Capacity or Archive Tier.

To delete an offloaded backup from an object storage repository, do the following:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Backups** > **Object Storage** node if you want to delete a backup from the capacity tier, to the **Backups** > **Archive** node if you want to delete backups from the archive tier.
3. In the working area, select a backup or VM and click **Delete from Disk** on the ribbon. Alternatively, you can right-click a backup and select **Delete from disk**.



Removing Missing Restore Points

In some cases, one or more restore points in the backup chain may be inaccessible. This can happen, for example, if the backup repository is put to the maintenance mode (for scale-out backup repositories), the backup repository is not available or some backup file is missing in the backup chain. Backup chains that contain missing restore points get corrupted — you cannot perform backup or restore VM data from the missing restore point, and restore points that depend on the missing restore point.

You can perform two operations with missing restore points:

- Forget — you can remove records about missing restore points from the configuration database. Veeam Backup & Replication will "forget" about missing restore points and will not display them in the console. The actual backup files will remain on disk (if backup files are still available).
- Remove — you can remove records about missing restore points from the configuration database and delete backup files from disk (if backup files are still available).

NOTE

Consider the following:

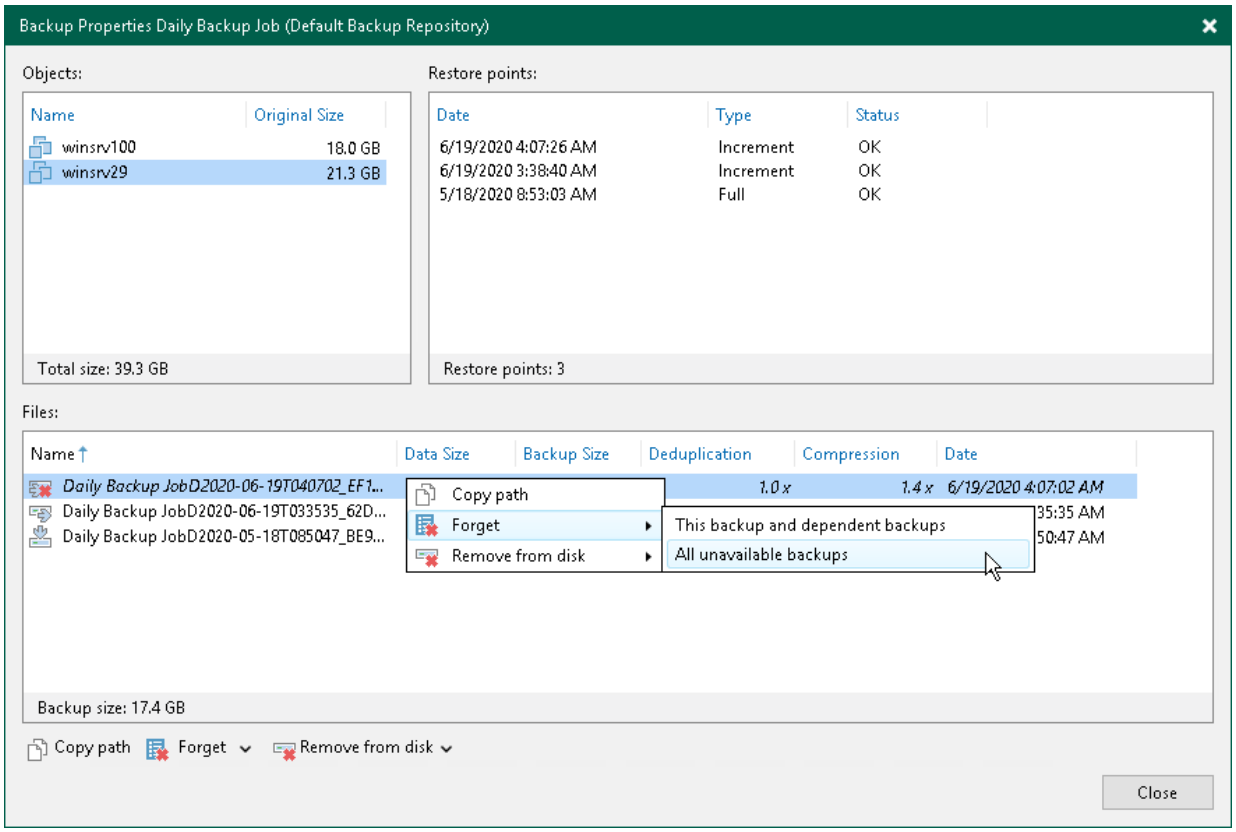
- The **Forget** and **Remove from disk** options are available only for restore points that are missing from the backup chain or points that depend on missing ones. If the restore point is available in the backup chain and does not depend on a missing restore point, you will not be able to use the **Forget** and **Remove from disk** options for it.
- You can manually update information about missing restore points. For this, disable a backup job and rescan the backup repository that is the target for the backup copy job. For more information, see Disabling and Removing Jobs and Rescanning Backup Repositories.

  Manual update can be required because Veeam Backup & Replication requires some time to update information in the configuration database for restore points that were removed from a backup chain or became inaccessible. That is why such restore points may not be displayed in the console as missing restore points.
- Veeam Backup & Replication does not track missing restore points in backups that reside in the cloud repository.
- If you apply the **Forget** or **Remove from disk** options to a missing restore point in a scale-out backup repository, the backup file associated with the missing restore point will be deleted from capacity tier and archive tier on the next offload and archiving job run.

To remove records about missing restore points from the configuration database:

1. Open the **Home** view.
2. In the inventory pane, select **Disk** under **Backups**.
3. In the working area, select the backup and click **Properties** on the ribbon or right-click the backup and select **Properties**.
4. In the **Backup Properties** window, right-click the missing restore point and select **Forget**.
   - To remove only the selected restore point and restore points that depend on it (that is, a part of the backup chain starting from this restore point), select **This and dependent backups**.

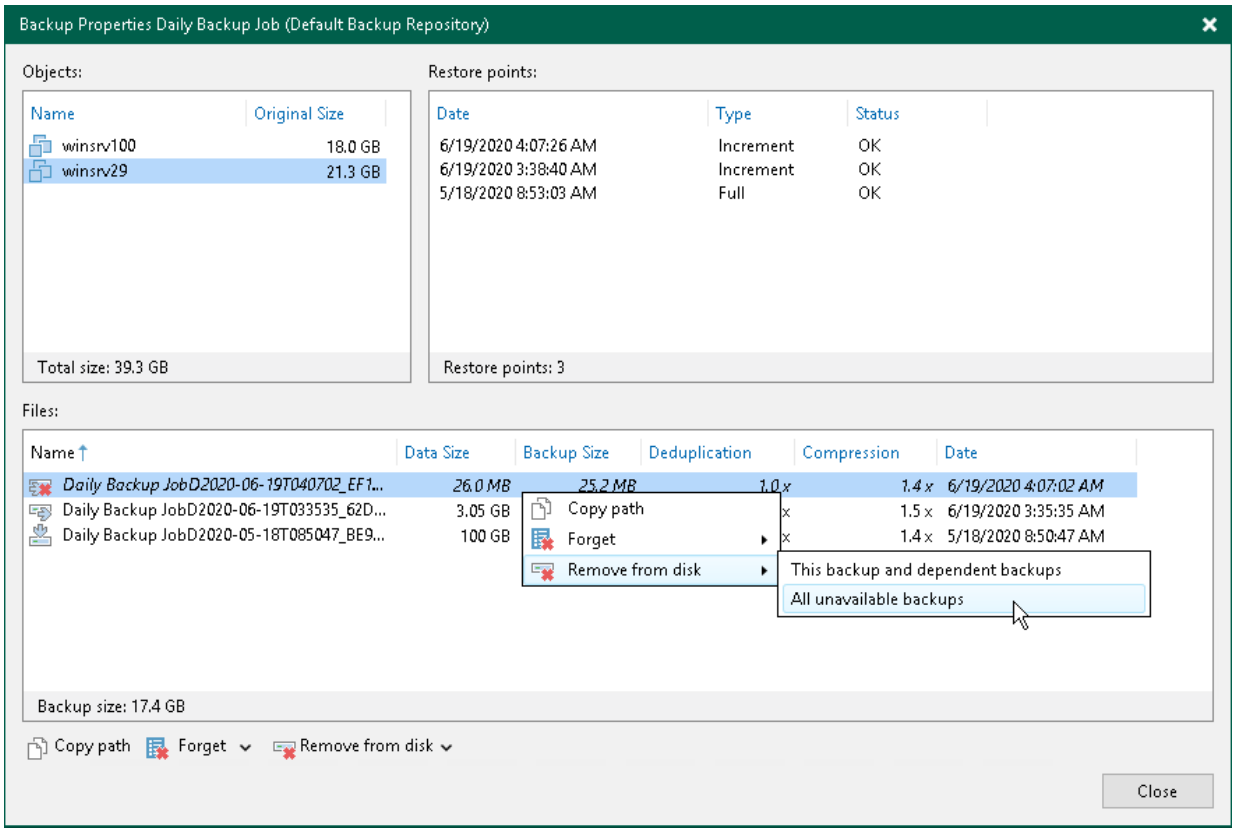
- To remove all missing restore points, select **All unavailable backups**.



To remove missing restore points from the configuration database and disk:

1. Open the **Home** view.
2. In the inventory pane, click **Disk** under **Backups**.
3. In the working area, select the backup and click **Properties** on the ribbon or right-click the backup and select **Properties**.
4. In the **Backup Properties** window, right-click the missing restore point and select **Remove from disk**.
   - To remove only the selected restore point and restore points that depend on it (that is, a part of the backup chain starting from this restore point), select **This and dependent backups**.

- To remove all missing restore points, select **All unavailable backups**.



Managing Jobs

To view all jobs configured on the backup server, open the **Home** view and select the **Jobs** node in the inventory pane. The list of available jobs is displayed in the working area. You can edit job properties, start and stop jobs, restart failed jobs, clone jobs, view job statistics and delete unnecessary jobs.

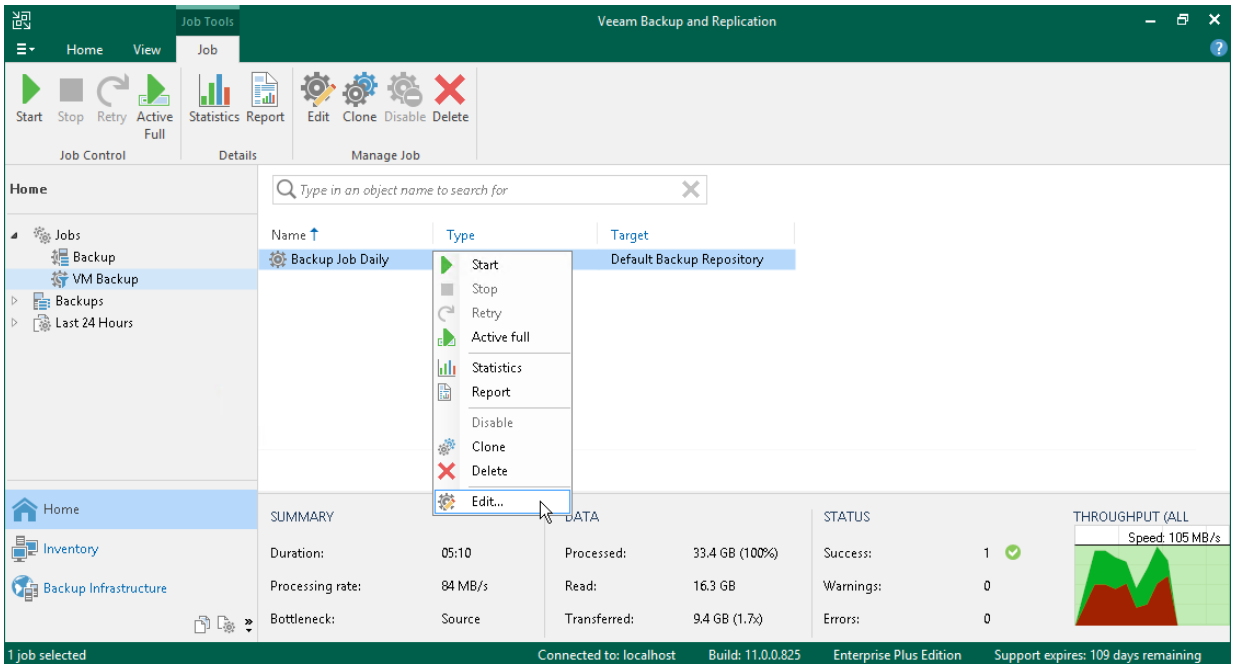
Editing Job Settings

You can edit configured jobs at any moment. For example, you may want to change scheduling settings for the job or add some VMs to the job.

To edit job settings:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, select the job and click **Edit** on the ribbon or right-click the job and select **Edit**.

You will follow the same steps as you have followed when creating the job and can change job settings as required.



Cloning Jobs

You can create new jobs by means of job cloning. Job cloning allows you to create an exact copy of any job with the same job settings. Configuration information of the created job copy are written to the configuration database that stores information of the original job.

To create multiple jobs with similar settings, you can configure a set of jobs that will be used as 'job templates'. You can then clone these 'job templates' and edit settings of cloned jobs as required.

The name of the cloned job is formed by the following rule: *<job_name_clone1>*, where *job_name* is the name of the original job and *clone1* is a suffix added to the original job name. If you clone the same job again, the number in the name will be incremented, for example, *job_name_clone2*, *job_name_clone3* and so on.

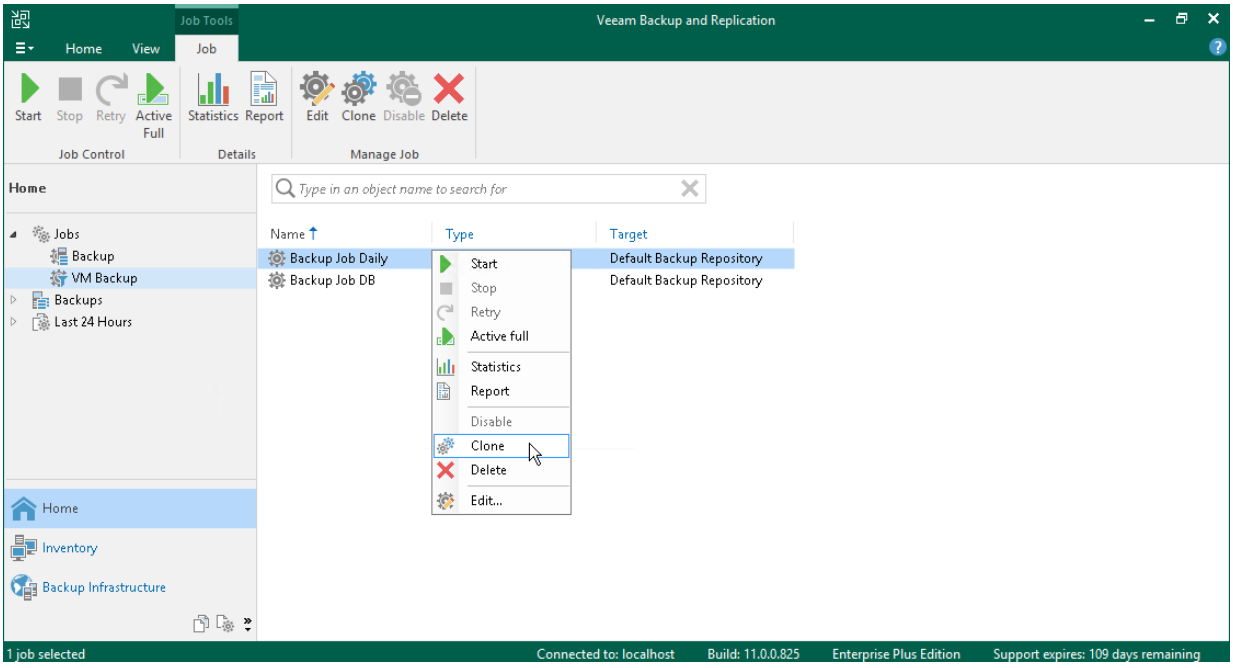
When cloning job, Veeam Backup & Replication can change some job settings so that cloned jobs do not hinder original jobs.

- If the original job is scheduled to run automatically, Veeam Backup & Replication disables the cloned job. To enable the cloned job, select it in the job list and click **Disable** on the ribbon or right-click the job and select **Disable**.
- If the original job is configured to use a secondary target, the cloned job is created without the secondary target settings.

To clone a job:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, select the job and click **Clone** on the ribbon or right-click the job and select **Clone**.S
4. After a job is cloned, you can edit all its settings, including the job name.

> **NOTE**
>
> The job cloning functionality is included in the Veeam Universal License. When using a legacy socket-based license, Enterprise or higher edition is required.



Disabling and Deleting Jobs

You can temporary disable scheduled jobs. The disabled job is paused for some period of time and is not run by the specified schedule. You can enable a disabled job at any time. You can also permanently delete a job from Veeam Backup & Replication and from the configuration database.

Disabling Job

To disable a job:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, select the job and click **Disable** on the ribbon or right-click the job and select **Disable**.

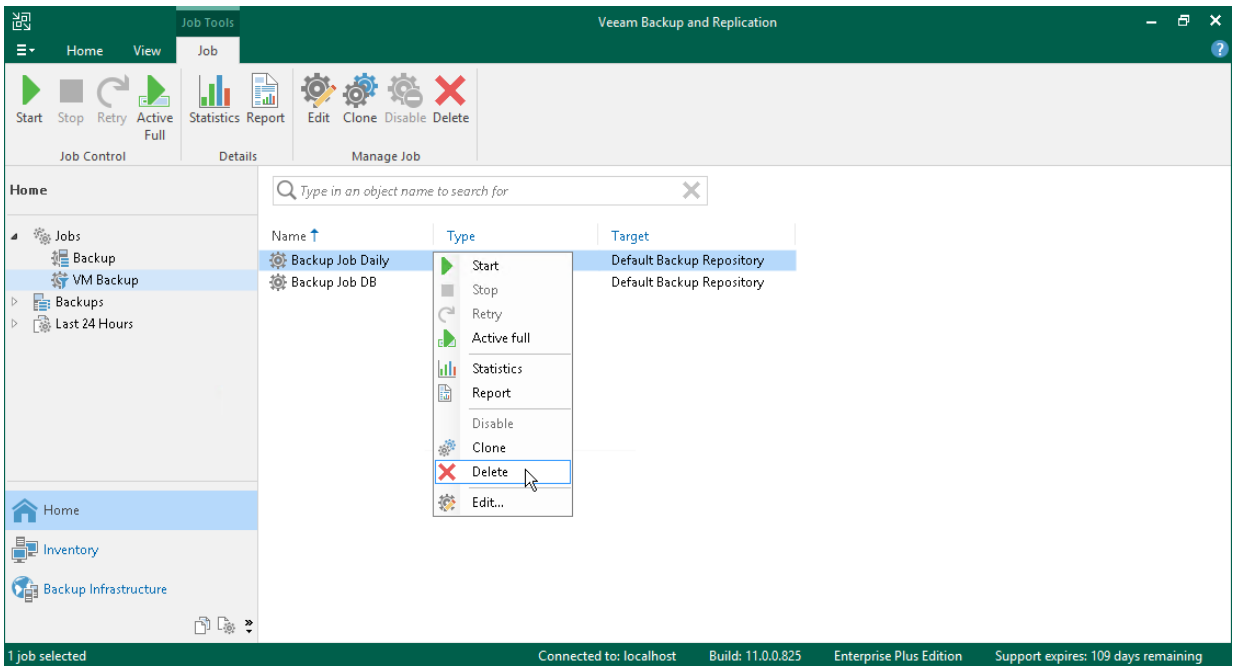
To enable a disabled job, select it in the list and click **Disable** on the ribbon once again.

Deleting Job

To remove a job:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, select the job and click **Delete** on the ribbon or right-click the job and select **Delete**.

After the job is deleted, the backups created by this job are displayed under the **Backups** > **Disk (Orphaned)** node. If the backup files created by this job were also stored in capacity tier or archive tier, they will also be displayed under the **Backups** > **Object Storage (Orphaned)** or **Backups** > **Archive (Orphaned)** nodes.



Starting and Stopping Jobs

You can start job manually, for example, if you want to create an additional restore point for a VM backup or replica and do not want to change the job schedule. You can also stop a job, for example, if VM processing is about to take long, and you do not want the job to produce workload on the production environment during business hours.

Starting Jobs

To start a job:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, select the backup job and click **Start** on the ribbon or right-click the job and select **Start**.

Stopping Jobs

You can stop a job in one of the following ways:

- Stop job immediately. In this case, Veeam Backup & Replication will produce a new restore point only for those VMs that have already been processed by the time you stop the job.
- Stop job after current VM. In this case, Veeam Backup & Replication will produce a new restore point only for those VMs that have already been processed and for VMs that are being processed at the moment.

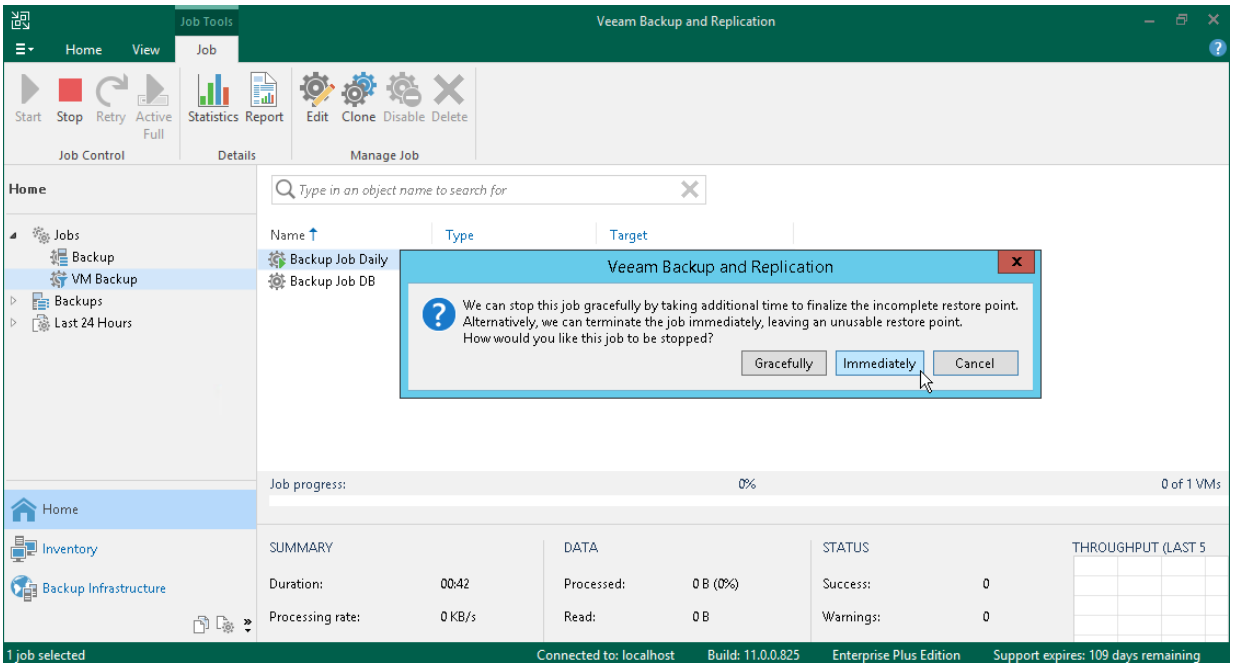
To stop a job:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, select the backup job and click **Stop** on the ribbon or right-click the job and select **Stop**. In the displayed window, click **Immediately**.

To stop the job after the current VM:

1. Open the **Home** view.
2. In the inventory pane, click **Jobs**.

3. In the working area, right-click the job and select **Stop**. In the displayed window, click **Gracefully**.



Starting and Stopping Transaction Log Backup Jobs

If you create a backup job and instruct it to ship transaction logs, the backup job comprises 2 jobs:

1. A parent backup job creating an image-level backup of the VM on which the database runs. This job is named as a regular backup job, for example: *Daily Job*.
2. A transaction log backup job responsible for shipping transaction logs to the backup repository. This job is named by the following pattern:
   - For MS SQL: *<job_name> SQL Server Transaction Log Backup*. For example, *Daily Job SQL Server Transaction Log Backup*.
   - For Oracle: *<job_name> Oracle Redo Log Backup*. For example, *Daily Job Oracle Redo Log Backup*.

The transaction log backup job is created automatically by Veeam Backup & Replication if it detects that you have added to the backup job at least one Microsoft SQL Server or Oracle VM, enabled application-aware processing and instructed Veeam Backup & Replication to back up transaction logs periodically.

Starting Transaction Log Backup Jobs

A parent backup job is started manually when you click **Start** on the toolbar, or automatically by schedule. The transaction log backup job is initially started when you enable schedule for the parent backup job. The transaction log backup works continuously in the background. A new session of the transaction log backup job starts every time the parent backup job is launched.

Stopping Transaction Log Backup Jobs

You can stop transaction log processing in one of the following ways:

- Disable transaction log shipping
- Disable the parent backup job

If you want the backup job to create image-level backups of the VM but do not want it to ship transaction logs anymore, you can disable transaction log backup in the backup job settings.

To disable transaction log shipping:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, select the backup job and click **Edit** on the ribbon or right-click the backup job and select **Edit**.
4. Pass to the **Guest Processing** step of the wizard and click **Applications**.
5. In the **Application-Aware Processing Options** window, select the VM and click **Edit**.
6. On the **SQL** or **Oracle** tab of the **VM Processing Settings** window, disable transaction log backup.
7. Click **Finish** to save the job settings.

If you do not want to create image-level backups of the VM and back up database transaction logs, you can disable scheduling for the parent backup job. Veeam Backup & Replication will instruct the transaction log backup job to complete log processing for all VMs added to the parent backup job, and will switch the parent backup job to the non-scheduled mode. The parent backup job will no longer be started automatically by schedule — you will have to run it manually.

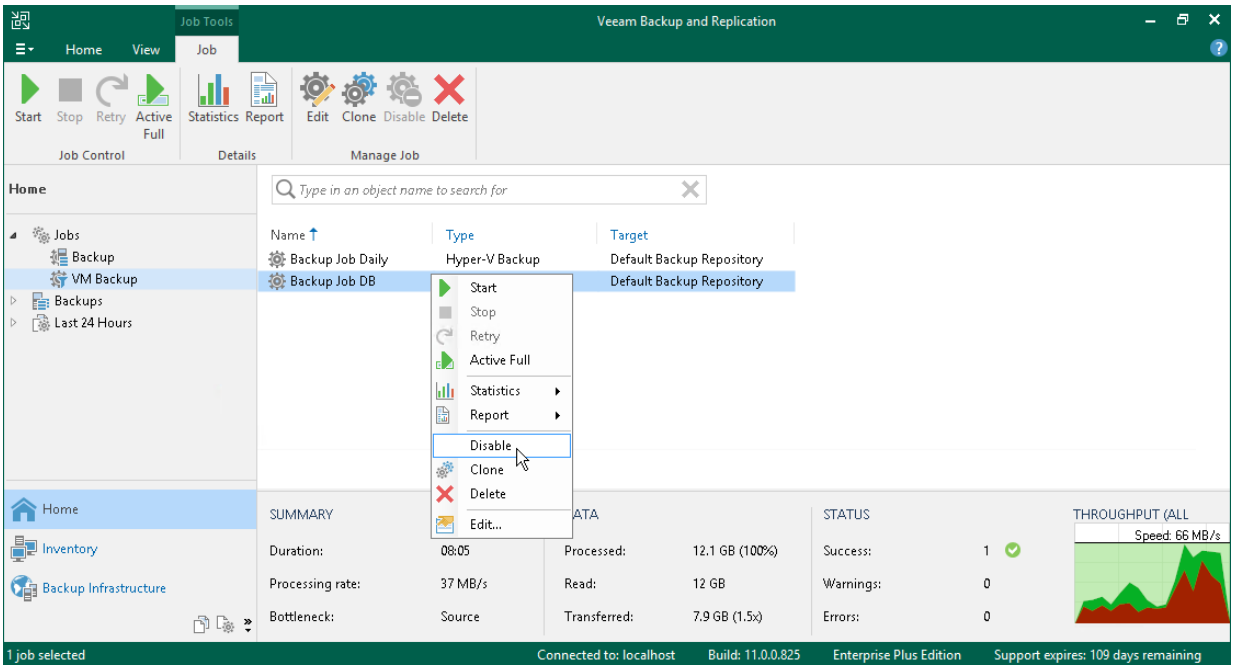
To disable scheduling for the parent backup job:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, select the backup job and click **Edit** on the ribbon. Alternatively, you can right-click the backup job and select **Edit**.
4. Pass to the **Schedule** step of the wizard and clear the **Run the job automatically** check box.
5. Click **Finish** to save the job settings.

Alternatively, you can disable the parent backup job:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, select the backup job and click **Disable** on the ribbon or right-click the job and select **Disable**.

To re-activate transaction log processing for all VMs in the parent backup job, select the job in the list and click **Disable** on the ribbon once again.



Reconfiguring Jobs with Microsoft SQL Server VMs

In some situations, you may need to reconfigure a backup job that processes a Microsoft SQL Server VMs and ships transaction logs. For example, you may want to create a separate backup job to process the virtualized database, and delete the VM running the database from the previously created job.

When you configure a new job, mind the restriction on transaction logs shipping. By default, the new backup job that processes the VM will not ship transaction logs if transaction logs for this VM have been shipped for the last 7 days by another backup job on the same backup server.

You can overcome this restriction with registry keys. For more information, contact Veeam Customer Support.

Reporting

When you run a job, Veeam Backup & Replication saves the job statistics and operation data to the configuration database. You can view real-time statistics for any performed job and generate reports with statistics data for any job or separate job session.

Viewing Real-Time Statistics

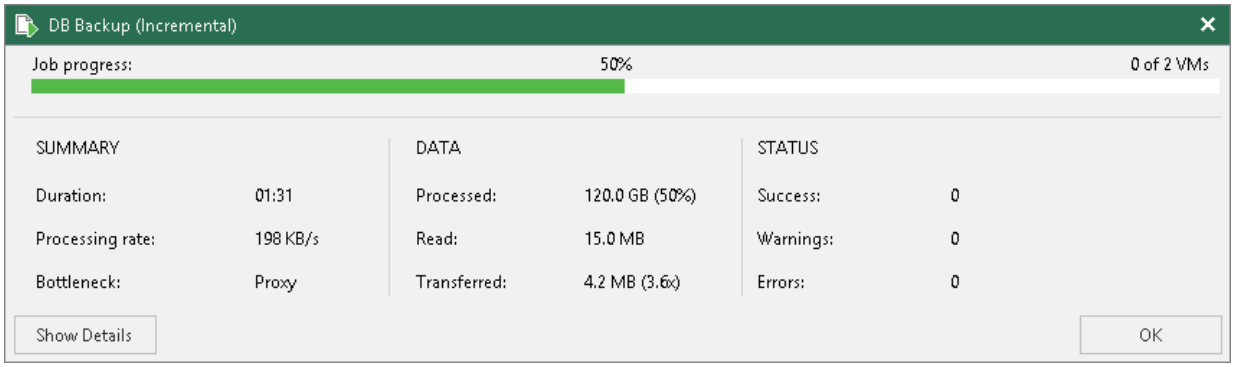
To view real-time statistics for a job, do one of the following:

- Open the **Home** view, in the inventory pane select one of the following nodes: **Jobs**, **Last 24 hours** or **Running**. In the working area, right-click the job and select Statistics.
- Open the **Home** view, in the inventory pane select one of the following nodes: **Jobs**, **Last 24 hours** or **Running**. In the working area, double-click the running job.

The real-time statistics provides detailed data on job sessions: job progress, duration, processing rate, performance bottlenecks, amount of processed data, read and transferred data and details of the session performance, for example, warnings and errors that have occurred in the process of operation.

In addition to overall job statistics, the real-time statistics provides information on each object processed with the job. To view the processing progress for a specific object, select it in the list on the left.

> **TIP**
>
> Mind the following:
>
> - To collapse and expand the real-time statistics window, use **Hide Details** and **Show Details** buttons at the bottom left corner of the window.
> - To switch between the job sessions backward and forward, use left and right arrow keys on the keyboard.



Statistics Counters

Veeam Backup & Replication displays jobs statistics for the following counters:

- The **Job progress** bar shows percentage of the job completion.
- The **Summary** box shows general information about the job:
  - **Duration** – time from the job start till the current moment or job end.
  - **Processing rate** – average speed of VM data processing. This counter is a ratio between the amount of data that has actually been read and job duration.
  - **Bottleneck** – bottleneck in the data transmission process. To learn about job bottlenecks, see Detecting Performance Bottlenecks.
- The **Data** box shows information about processed VM data:
  - **Processed** – total size of all VM disks processed by the job.

- **Read** — amount of data read from the datastore by the source-side Data Mover prior to applying compression and deduplication. For incremental job runs, the value of this counter is typically lower than the value of the **Processed** counter. Veeam Backup & Replication reads only data blocks that have changed since the last job session, processes and copies these data blocks to the target.
- **Transferred** — amount of data transferred from the source-side Veeam Data Mover to the target-side Veeam Data Mover after applying compression and deduplication. This counter does not directly indicate the size of the resulting files. Depending on the backup infrastructure and job settings, Veeam Backup & Replication can perform additional activities with data: deduplicate data, decompress data prior to writing the file to disk and so on. The activities can impact the size of the resulting file.

- The **Status** box shows information about the job results. This box informs how many tasks have completed with the *Success*, *Warning* and *Error* statuses (1 task per 1 VM).
- The pane at the lower left corner shows a list of objects included in the job.
- The pane at the lower right corner shows a list of operations performed during the job. To see a list of operations for a specific object included in the job, click the object in the pane on the left. To see a list of operations for the whole job, click anywhere on the blank area in the left pane.

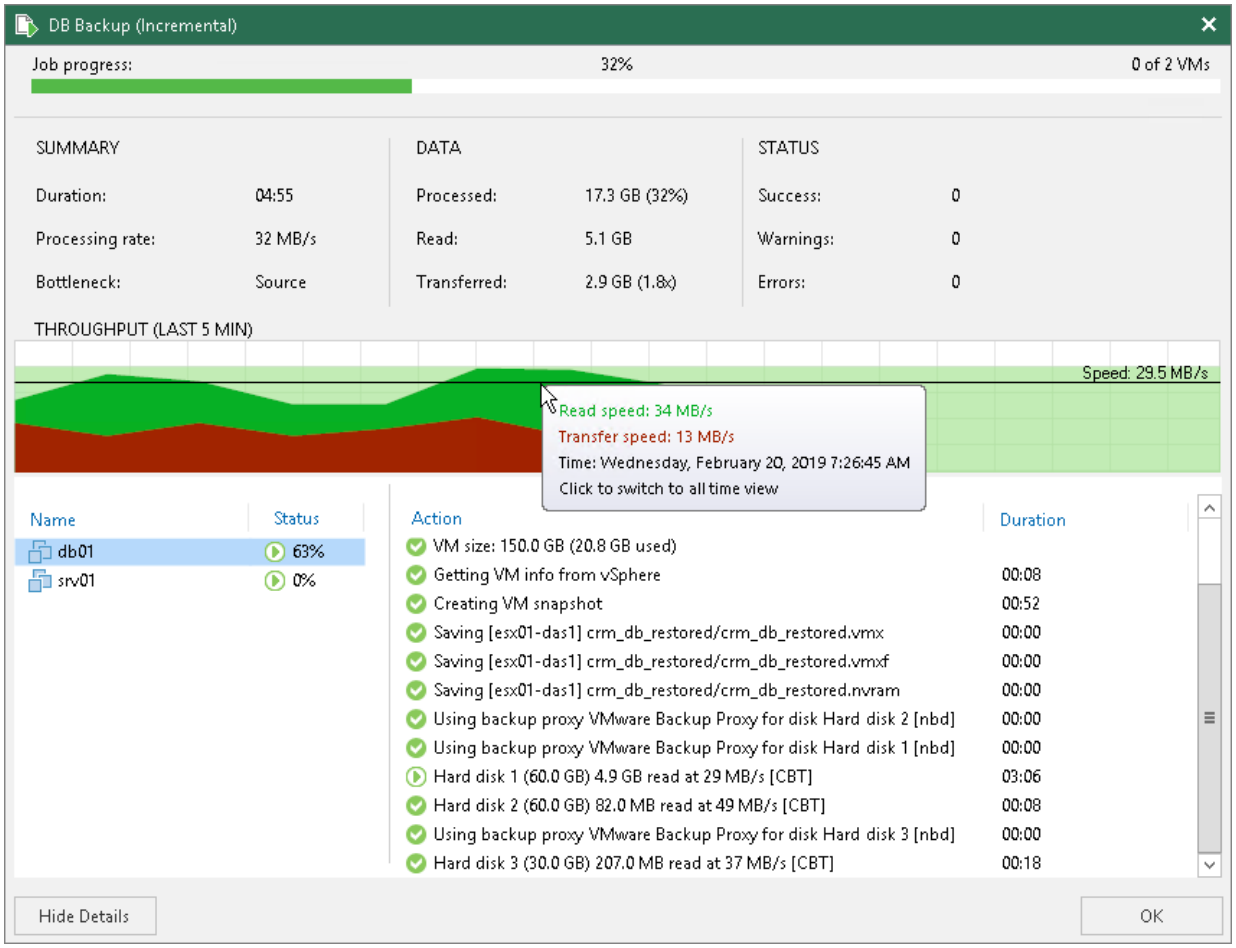
Colored Graph

To visualize the data transfer process, Veeam Backup & Replication displays a colored graph in the real-time statistics window:

- The green area defines the amount of data read from source.
- The brown area defines the amount of data transported to target.
- The horizontal line defines the current data processing speed.

If the job session is still being performed, you can click the graph to view data rate for the last 5 minutes or the whole processing period. If the job session has already ended, the graph will display information for the whole processing period only.

The colored graph is displayed only for the currently running job session or the latest job session. If you open real-time statistics for past sessions other than the latest one, the colored graph will not be displayed.



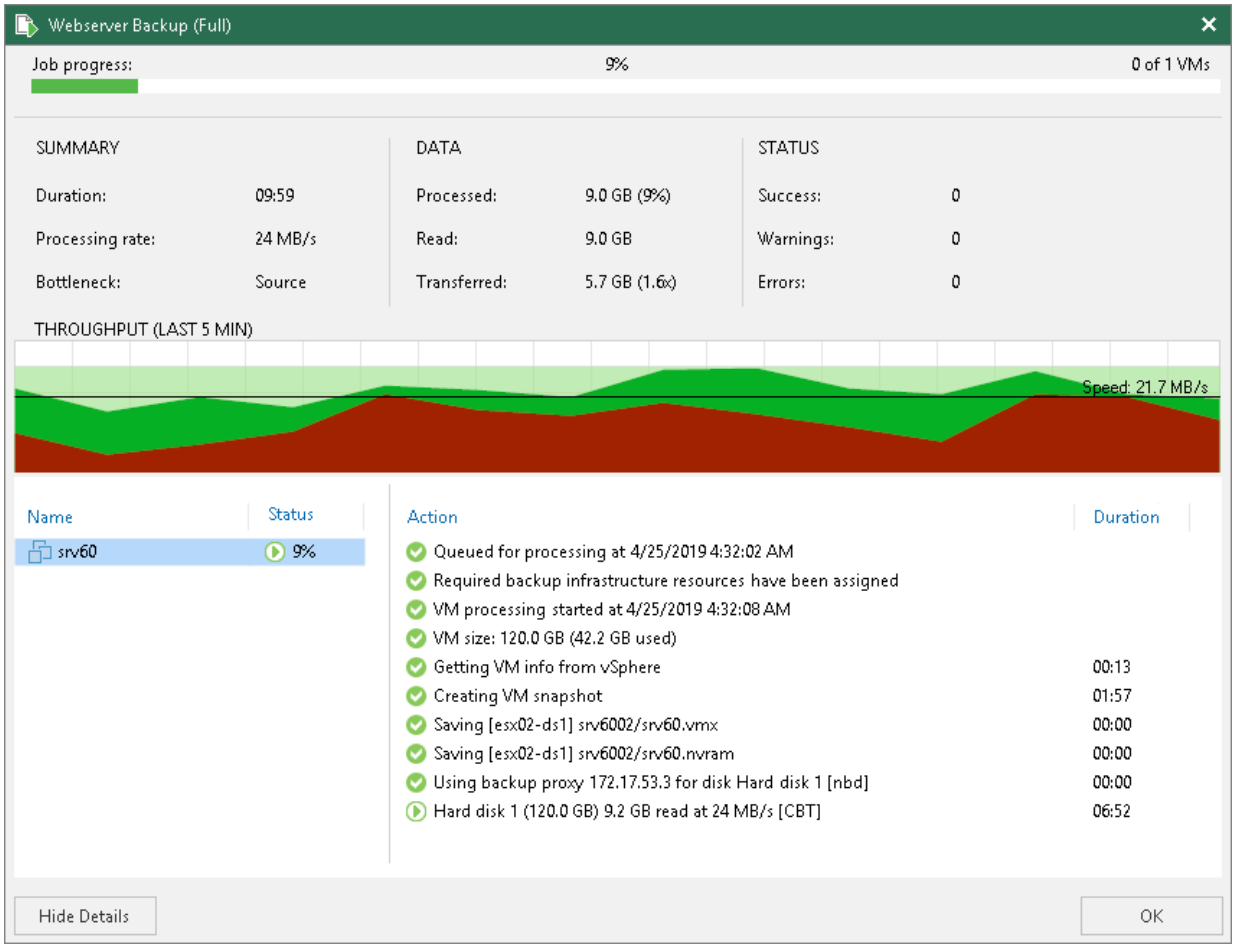
Viewing Job Session Results

You can view detailed statistics on every job session.

To view statistics for a selected job session, do either of the following:

- Open the **History** view. In the inventory pane select **Jobs**. In the working area, double-click the necessary job session.
- Open the **History** view. In the inventory pane select **Jobs**. In the working area, right-click the necessary job session and select **Statistics**.

> **TIP**
>
> To switch between past job sessions, use left and right arrow keys on the keyboard.



Viewing Job and Job Session Reports

You can generate reports with details about all sessions of a job or a single session only.

Job Report

The job report contains data on all sessions initiated for a specific job, that is, job history. The report shows data for all sessions stored in the configuration database.

To generate a job report:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, select the necessary job and click **Report** on the ribbon. You can also right-click the job and select **Report**.

For more information on counters in the report, see Report Counters.

> **TIP**
>
> Generated reports are stored in the `C:\Users\<username>\AppData\Local\Temp` folder.

Session Report

To generate a report for a single session:

1. Open the **History** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, select the necessary session and click **Report** on the ribbon. You can also right-click the necessary session and select **Report**.

Backup job: DB Backup — Success
2 of 2 VMs processed

Friday, April 26, 2019 9:26:08 AM

Success	2	**Start time**	9:26:08 AM	**Total size**	177.0 GB	**Backup size**	42.3 MB
Warning	0	**End time**	9:28:24 AM	**Data read**	456.0 MB	**Dedupe**	1.0x
Error	0	**Duration**	0:02:15	**Transferred**	21.0 MB	**Compression**	2.7x

Details

Name	Status	Start time	End time	Size	Read	Transferred	Duration	Details
appsrv02	Success	9:26:38 AM	9:27:19 AM	50.0 GB	100.0 MB	4.7 MB	0:00:41	
db01	Success	9:26:38 AM	9:27:33 AM	127.0 GB	356.0 MB	16.3 MB	0:00:55	

Report Counters

Veeam Backup & Replication displays the following counters in reports:

- The **Success**, **Warning** and **Error** counters show how many workloads were processed with the *Success*, *Warning* and *Error* statuses.
- The **Start time** and **End time** counters show when the job started and completed.

- The **Duration** counter shows time from the job start till the current moment or job end.
- The **Total size** counter shows provisioned size (the maximum configured size) of all workload disks in the job.
- The **Data read** counter shows the amount of data read from the datastore before compression and deduplication. The value of this counter is typically lower than the value of the **Total size** counter. Veeam Backup & Replication reads only data blocks that have changed since the last job session, processes and copies these data blocks to the target.
- The **Transferred** counter shows the amount of data transferred from the source side to the target side after applying compression and source-side deduplication. This counter does not directly indicate the size of the resulting files. Depending on the backup infrastructure and job settings, Veeam Backup & Replication can perform additional activities with data: perform target-side deduplication, decompress data prior to writing the file to disk and so on. The activities can impact the size of the resulting file.
- The **Backup size** counter shows the resulting backup file size.
- The **Dedupe** counter shows the deduplication level.
- The **Compression** counter shows the compression level.

In the **Details** section, you can see similar counters for each workload processed by the job.

Replication

Replication is a technology that helps you protect mission-critical VMware virtual machines. When you replicate a VM, Veeam Backup & Replication creates an exact copy of the VM in the native VMware vSphere format on the target host. Veeam Backup & Replication maintains this copy in sync with the original VM. Replication provides minimum recovery time objective (RTO) in case a disaster strikes because VM replicas are in a ready-to-start state.

We recommend you to replicate VMs for which recovery point objective (RPO) of hours is required. If you need RPO of seconds, consider continuous data protection (CDP).

Data Replication

To replicate VMs, Veeam Backup & Replication leverages VMware vSphere snapshot capabilities. During replication, Veeam Backup & Replication requests VMware vSphere to create a VM snapshot. The VM snapshot can be thought of as a point-in-time copy of a VM that includes virtual disks, system state, configuration and so on. Veeam Backup & Replication uses the snapshot as a source of data for replication.

During the first replication cycle, Veeam Backup & Replication copies data of the original VM running on the source host, and creates its full replica on the target host. Unlike backup files, replica virtual disks are stored decompressed in their native format. All subsequent replication cycles are incremental.
Veeam Backup & Replication copies only those data blocks that have changed since the last replication job session. To keep track of changed data blocks, Veeam Backup & Replication uses different approaches. For more information, see Changed Block Tracking.

Veeam Backup & Replication lets you perform on-site replication for high availability scenarios and remote (off-site) replication for disaster recovery scenarios. To facilitate replication over the WAN or slow connections, Veeam Backup & Replication optimizes traffic transmission. It filters out unnecessary data blocks such as duplicate data blocks, zero data blocks, blocks of swap files and blocks of excluded VM guest OS files, and compresses replica traffic. Veeam Backup & Replication also allows you to use WAN accelerators and apply network throttling rules to prevent replication jobs from consuming the entire network bandwidth.

To replicate a VM, you need to configure required backup infrastructure components and create a replication job.

Recovery

If a disaster strikes and the production VM stops working properly, you can fail over to its replica.

When you fail over to a replica, the replica takes over the role of the original VM. After your original VM is repaired, you can fail back to it and transfer all changes that occurred to replica to the original VM. If your original VM cannot be repaired, you can perform permanent failover, that is, permanently switch from the original VM to the VM replica and use this replica as the original VM. For more information, see Failover and Failback for Replication.

Requirements and Limitations

Replication has the following requirements and limitations:

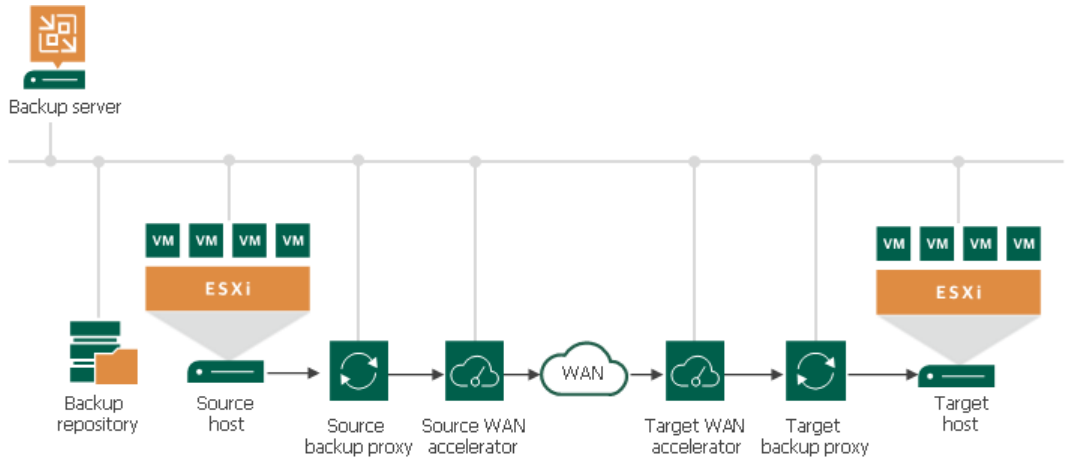
- Due to VMware vSphere limitations, if you change the size of VM disks on the source VM, Veeam Backup & Replication deletes all available restore points (represented as VM snapshots) on the VM replica during the next replication job session. For more information, see this VMware KB article.
- If you assign the role of a backup proxy to a VM, you should not add this VM to the list of processed VMs in a job that uses this backup proxy. Such configuration may result in degraded job performance. Veeam Backup & Replication will assign this backup proxy to process other VMs in the job first, and processing of this VM itself will be put on hold. Veeam Backup & Replication will report the following message in the job statistics: *VM is a backup proxy, waiting for it to stop processing tasks*. The job will start processing this VM only after the backup proxy deployed on the VM finishes its tasks.
- Replication of VM templates is not supported.
- If you use tags to categorize virtual infrastructure objects, check limitations for VM tags. For more information, see VM Tags.
- Due to Microsoft limitations, you cannot use Microsoft Azure Active Directory credentials to perform application-aware processing on VMs running Microsoft Windows 10 (or later).
- If a job is unable to complete within 21 days period, it will be stopped with the *Failed* status.

Backup Infrastructure for Replication

Veeam Backup & Replication uses the following components for the replication process:

- Backup server
- Source and target hosts
- Backup proxies
- Backup repository
- [Optional] WAN accelerators

The amount and placement of these components depend on a replication scenario you use. For more information, see Replication Scenarios.



Backup Server

During the replication process, the backup server coordinates replication tasks, controls resource allocation and replica job scheduling. The backup server runs the Veeam Backup Service and Veeam Broker Service that coordinate and interact with the virtual infrastructure.

For more information on the backup server, see Backup Server.

Source and Target Hosts

The source host is the host where VMs that you plan to replicate are located. The target host is the host where VM replicas will be created and maintained in the ready-to-start state.

The role of a target host can be assigned to a standalone ESXi host or ESXi cluster. If you assign a cluster or vCenter Server as a target, the replication process becomes more sustainable — the replication process will not fail if there is at least one available host in the cluster.

To replicate data from and to hosts, they must be first added to the Veeam Backup & Replication infrastructure. For more information on how to add hosts, see the Adding VMware vSphere Servers section.

Backup Proxies

A backup proxy collects, transforms and transports VM data during the replication process. For more information on backup proxies, requirements and limitations for them, see Backup Proxy.

For replication, you can deploy backup proxies on the following machines:

- **Physical machines**. In this case, Veeam Backup & Replication uses the Network transport mode to populate replica disk files. For more information, see Network Mode.
- **VMs**. The virtual backup proxy must be registered on an ESXi host that has a direct connection to the target datastore. In this case, the backup proxy will be able to use the Virtual appliance transport mode to populate replica disk files. This results in increased writing speed and fail-safe replication to ESXi targets. For more information, see Virtual Appliance (HotAdd). Note that if the Virtual appliance transport mode cannot be used, the backup proxy can fail over to the network mode if you configure it while adding a backup proxy.

We recommend you to use at least two backup proxies to ensure that the job will be performed if one of backup proxies fails or loses its connectivity to the source datastore. For more information on how assign the role of a backup proxy, see Adding VMware Backup Proxies. For more information on how to assign proxies to a replication job, see Specify Data Transfer Settings.

Backup Repository

The backup repository stores replica metadata that contains information on the read data blocks (such as checksums and digests). Metadata is required when Veeam Backup & Replication performs incremental replication or if you fail back from a VM replica to the original VM in the original location using quick rollback. Veeam Backup & Replication uses metadata to quickly detect changed data blocks between two replica states.

The backup repository must be deployed as close to the source backup proxy as possible and must have access to it.

For more information, see Backup Repository.

WAN Accelerators

WAN accelerators are optional components in the backup infrastructure. You can use WAN accelerators if you replicate VMs over a slow connection or over WAN.

In the replication process, WAN accelerators are responsible for global data caching. To use WAN acceleration, you must deploy two WAN accelerators in the following way:

- The source WAN accelerator must be deployed in the source side, close to the backup proxy.
- The target WAN accelerator must be deployed in the target side, close to the backup proxy.

For more information, see WAN Acceleration.

Replication Scenarios

Veeam Backup & Replication supports a number of replication scenarios that depend on the location of the target host:

- **On-site replication**. If the target host is located in the same site as the source host, use the on-site replication scenario.
- **Off-site replication**. If the target and source hosts are located in different sites, use the off-site replication scenario.

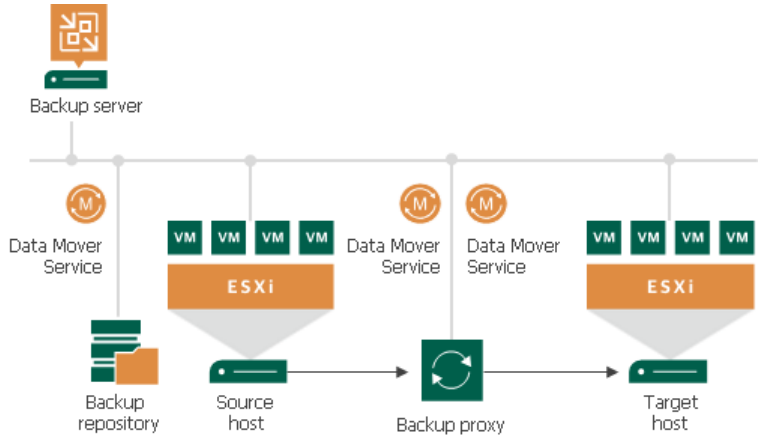
Depending on the scenario you are planning to use, different backup infrastructure components will be involved in the replication process. For more information on the components, see Backup Infrastructure for Replication.

On-Site Replication

On-site replication requires the following backup infrastructure components:

- Source and target hosts.
- Backup proxy. In the on-site replication scenario, the source Veeam Data Mover and target Veeam Data Mover are started on the same backup proxy. The backup proxy must have access to the backup server, source host, target host and backup repository that stores replica metadata.
- Backup repository for storing replica metadata.

In the on-site replication scenario, Veeam Backup & Replication does not perform data compression. Replication traffic is transferred decompressed between the two Veeam Data Mover started on the same backup proxy.



Off-Site Replication

For the off-site replication scenario, you must use at least two backup proxies:

- One backup proxy in the production site, closer to the source host.
- One backup proxy in the remote DR site, closer to the target host.

Depending on the network speed between the production and DR sites, off-site replication can run in two ways:

- **Directly over backup proxies**. Transfer data directly over backup proxies if the connection between two sites is fast and stable.
- **Over a pair of WAN accelerators**. Use WAN accelerators if the connection between the sites is slow.

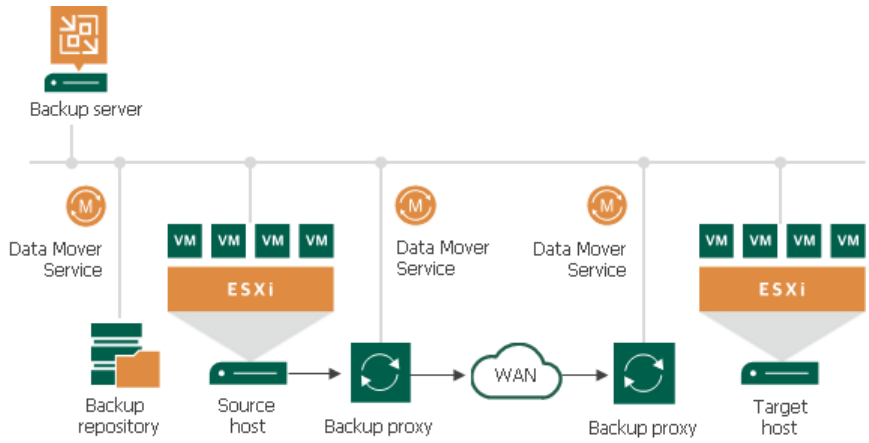
> **NOTE**
>
> With off-site replication, you can also use technologies that help reduce the amount of replication traffic and streamline replica configuration: replica seeding and mapping, network mapping and re-IP rules.

Replication Directly over Backup Proxies

Off-site replication directly over backup proxies requires the following backup infrastructure components:

- Source and target hosts.
- At least one backup proxy in the production site. The backup proxy must have access to the backup server, source host, backup proxy in the target site and backup repository that stores replica metadata.
- At least one backup proxy in the DR site. The backup proxy must have access to the backup server, target host and backup proxy in the production site.
- Backup repository for storing replica metadata.

In the off-site replication scenario, Veeam Backup & Replication uses data compression. Veeam Data Mover on the source backup proxy compresses VM data blocks and sends them to the target backup proxy in the compressed format. Veeam Data Mover on the target backup proxy decompresses VM data and stores it to a datastore in a native VMware vSphere format.



Replication over WAN Accelerators

Off-site replication over WAN accelerators requires the following backup infrastructure components:

- Source and target hosts.
- A pair of WAN accelerators at each end of the WAN link:
  - Source WAN accelerator in the production site. The source WAN accelerator must have access to the backup server, source backup proxy and target WAN accelerator.
  - Target WAN accelerator in the DR site. The target WAN accelerator must have access to the backup server, source WAN accelerator and target backup proxy.
- At least one backup proxy in the production site. The backup proxy must have access to the backup server, source host, source WAN accelerator and backup repository that stores replica metadata.
- At least one backup proxy in the DR site. The backup proxy must have access to the backup server, target host and target WAN accelerator.

- Backup repository for storing replica metadata. The backup repository must be located in the production site, closer to the source backup proxy, and must have access to it.

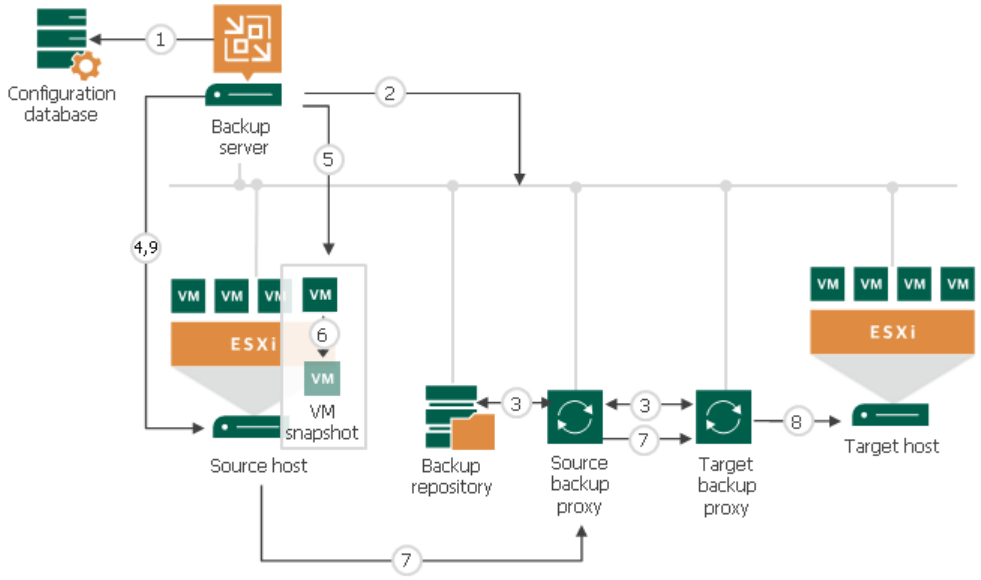
In the off-site replication scenario using WAN accelerators, Veeam Backup & Replication compresses VM data. VM data blocks are compressed on the source WAN accelerator, transported to the DR site in the compressed format and decompressed on the target WAN accelerator.

How Replication Works

Veeam Backup & Replication performs VM replication in the following way:

1. When a new replication job session starts, Veeam Backup & Replication reads job settings from the configuration database and creates a list of VMs to process. For every disk of a VM added to the job, Veeam Backup & Replication creates a new task.
2. Veeam Backup & Replication checks what backup infrastructure resources are available, and assigns backup proxies and backup repositories to process the tasks. Then Veeam Backup & Replication establishes a connection with source and target backup proxies and the backup repository, and sets a number of rules for data transfer, such as network traffic throttling rules and so on.
3. The source proxy establishes a connection with the target proxy and backup repository.
4. Veeam Backup & Replication queries information about VMs and virtualization hosts from the vCenter Server.
5. If application-aware image processing is enabled for the job, Veeam Backup & Replication connects to VM guest OSes, deploys non-persistent runtime components or uses (if necessary, deploys) persistent agent components on VM guest OSes and performs in-guest processing tasks.
6. Veeam Backup & Replication requests vCenter Server or ESXi host to create a VM snapshot. VM disks are put to the read-only state, and every virtual disk receives a delta file. All changes that the user makes to the VM during replication are written to delta files.
7. The source backup proxy reads the VM data from the read-only VM disk and copies it. During incremental job sessions, the source proxy uses changed block tracking (CBT) to retrieve only those data blocks that have changed since the previous job session. If CBT is not available, the source proxy interacts with the backup repository to obtain replica metadata, and uses this metadata to detect blocks that have changed since the previous job session.

   While copying VM data, the source proxy performs additional processing. It filters out zero data blocks, blocks of swap files and blocks of excluded VM guest OS files. The source proxy compresses VM data and transports it to the target proxy.
8. The target proxy decompresses VM data and writes the result to the destination datastore.
9. After the backup proxy finishes reading VM data, Veeam Backup & Replication requests the vCenter Server or ESXi host to commit the VM snapshot.



Replication Chain

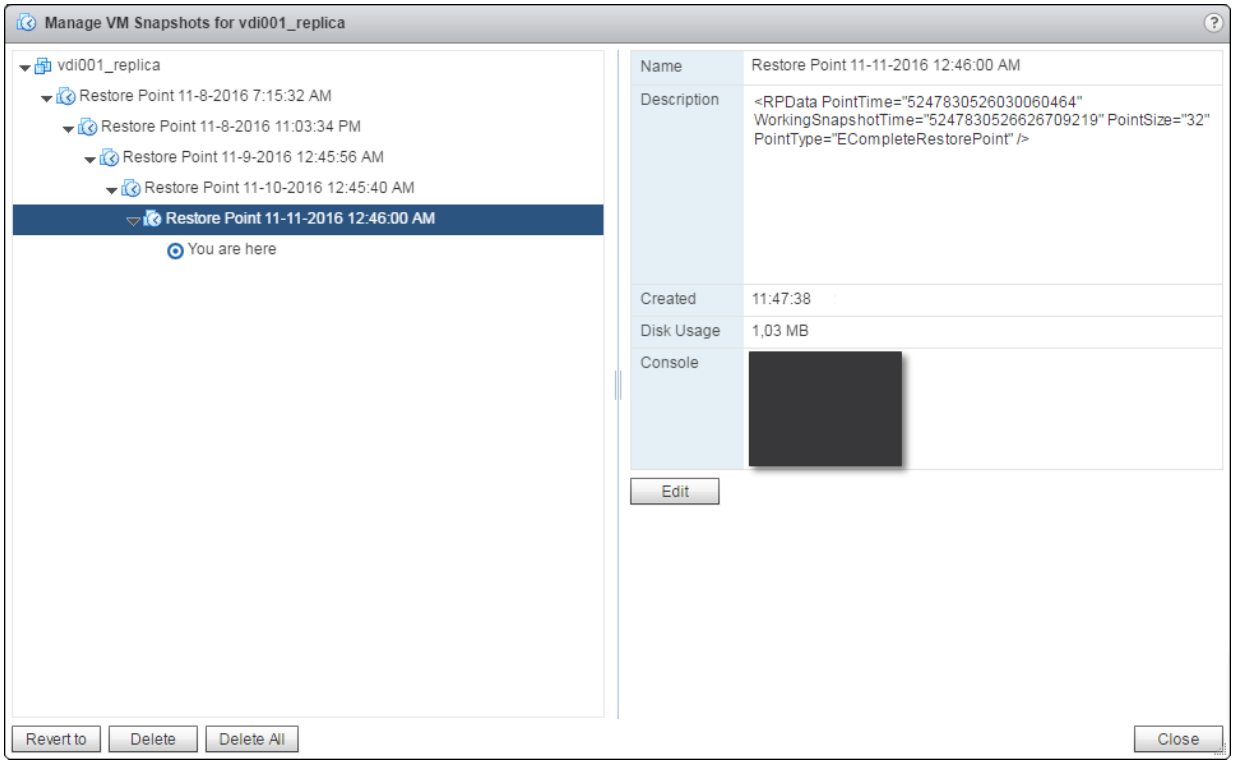
For every VM replica, Veeam Backup & Replication creates a replication chain that consists of restore points. Veeam Backup & Replication utilizes VMware ESXi snapshot capabilities to create and manage replica restore points.

Veeam Backup & Replication creates a restore point during every replication job session. During the first replication job session, Veeam Backup & Replication creates a copy of the source VM on the target host. During every subsequent replication job session, it adds a new snapshot to the replication chain for the VM replica. Blocks of data that have changed since the last job run are written to the snapshot delta file, and the snapshot delta file acts as a restore point. You can view a replication chain created for a VM using VMware vSphere client.

You can specify how many restore points you want to store in the replication chain. For this, configure retention policy settings for the replication job. For more information, see Specify Replication Job Settings.

VM replica restore points are stored in a native VMware vSphere format next to replica virtual disk files, which allows Veeam Backup & Replication to accelerate failover operations. To fail over to the necessary point of the VM replica, Veeam Backup & Replication does not need to apply rollback files. Instead, it uses a native VMware vSphere mechanism of reverting to a snapshot.

> **IMPORTANT**
>
> We recommend you against switching restore points for replicas and powering on replicas using VMware vSphere client. This may disrupt further replication operations in Veeam Backup & Replication or cause loss of important data. Instead, use Veeam Backup & Replication to perform failover operations. For more information on how to fail over to a VM replica, see Failover.



Replica Seeding and Mapping

Replica seeding and mapping are technologies that help reduce the amount of traffic sent over a network. With these technologies, Veeam Backup & Replication do not have to transfer all of VM data from the source host to the target host across the sites during the first session of a replication job (during the initial replication).

You can use seeding and mapping in the following scenarios:

- **Seeding**

  Configure replica seeding if, in a backup repository located in the disaster recovery (DR) site, you have backups of VMs that you plan to replicate. During replication, Veeam Backup & Replication will restore VMs from these backups and will synchronize the state of the restored VMs with the latest state of the original VMs. Then Veeam Backup & Replication will use these restored VMs as replicas.

  For more information on how to create backups that can be used as "seeds" for replica, see Creating Replica Seeds.

- **Mapping**

  Configure replica mapping if, on the host in the DR site, you have ready-to-use copies of the original VMs. These can be restored VMs or replicas created by other replication jobs. Veeam Backup & Replication will synchronize the state of these ready-to-use VMs with the latest state of the original VMs and will use these VMs as replicas. You can also use replica mapping if you need to reconfigure or recreate replication jobs, for example, split one replication job into several jobs.

You can also configure both replica seeding and replica mapping in the same replication job. For example, if a replication job includes 2 VMs, you can use seeding for one VM and map the other VM to an existing VM.

> **IMPORTANT**
>
> If seeding or mapping is enabled in a replication job, all VMs in the job must be covered with seeding or mapping. If a VM neither has a seed, nor is mapped to an existing VM, it will be skipped from processing.

Algorithm for Seeding

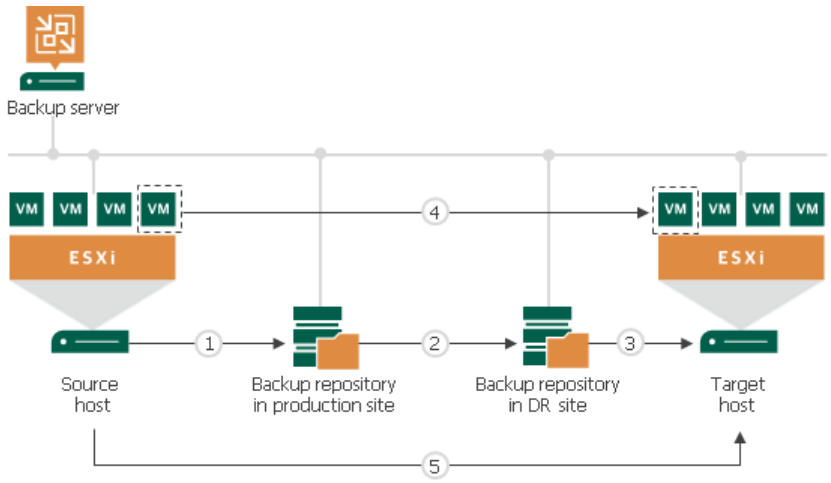
Replica seeding includes the following steps:

1. As a preparatory step for replica seeding, you need to create a backup of a VM that you plan to replicate. For more information on how to create a backup that will be used as a "seed" for replica, see Creating Replica Seeds.
2. When you create a replication job, you should point it to a backup repository in the DR site. During the initial synchronization, Veeam Backup & Replication accesses the backup repository where the replica seed is located, and restores the VM from the backup. The restored VM is registered on the target host in the DR site. Files of the restored VM are placed to the location you specify as the replica destination datastore.

   Virtual disks of a replica restored from the backup preserve their format (that is, if the original VM used thin provisioned disks, virtual disks of the VM replica are restored as thin provisioned).
3. Veeam Backup & Replication synchronizes the restored VM with the latest state of the original VM.

   After successful synchronization, in the **Home** view in the Veeam Backup & Replication console, under **Replicas** node you will see a VM replica with two restore points. One point will contain the state of the VM from the backup file; the other point will contain the latest state of the original VM you want to replicate.

4. During incremental synchronization, Veeam Backup & Replication transfers only incremental changes in a regular manner.



Replica seeding dramatically reduces traffic sent over WAN or slow connections because Veeam Backup & Replication does not send the full contents of the VM image. Instead, it transmits only differential data blocks.

> **TIP**
>
> If you add new VMs to an already existing replication job, you can enable replica seeding settings for these VMs. In this case, the newly added VMs will be seeded from the selected backups at the next pass of the replication job. VMs that have already been processed by the job by the time you add new VMs will be processed in a regular manner.

Algorithm for Mapping

Replication to a mapped VM is performed in the following way:

1. During the first run, the replication job calculates the differences between the original and mapped VM. Instead of copying and transferring all data of the original VM, the replication job transfers only incremental changes to synchronize the state of the mapped VM with the state of the original VM.

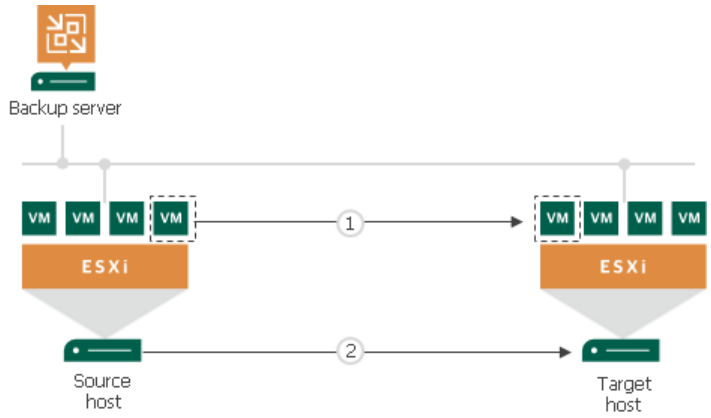
   After successful synchronization, in the **Home** view of Veeam Backup & Replication, under **Replicas** node you will see a VM replica with 2 restore points:

   - One restore point will contain the latest state of the mapped VM.
   - The other restore point will contain the latest state of the original VM on the source host.

2. All subsequent runs of the replication job will be performed in a regular manner: Veeam Backup & Replication will transfer only incremental changes to the target host.

NOTE

If a VM replica to which the original VM is mapped has any snapshots, these snapshots will be removed during the run of the replication job.



Replica from Backup

Disaster recovery plans often require that you back up and replicate the same VM for disaster recovery (DR) and high availability (HA) purposes. As a rule, this doubles the workload on the virtual infrastructure: two VM snapshots need to be created independently from one another, and VM data need to be transferred from the production site twice.

To minimize the use of compute, storage and network resources, you can use the replica from backup option. You can use this option in both on-site and off-site replication scenarios.

When you perform replication from backup, Veeam Backup & Replication does not address hosts and storage in the production environment to read VM data. As a source of data, Veeam Backup & Replication uses a backup chain that already exists in a backup repository. As a result, Veeam Backup & Replication creates only one snapshot and transfers VM data only once. Veeam Backup & Replication retrieves VM data only while a backup or backup copy job is running. The replication job re-uses retrieved data to build VM replica restore points.

Differences between Seeding and Replica from Backup

Although replica from backup might resemble replica seeding, there is difference between these options:

- Replica seeding uses a backup file only during the first run of a replication job. To further build VM replica restore points, the replication job addresses the production environment and reads VM data from the source storage.
- Replica from backup uses a backup chain in a backup repository as the only source of data. When building a new VM replica restore point, Veeam Backup & Replication always reads data from the latest restore point in the backup chain, either full or incremental. The backup chain in the backup repository may be created with a backup job or a backup copy job.

Requirements and Limitations for Replica from Backup

Consider the following requirements and limitations:

- You can perform replication only from backups of VMware vSphere virtual machines created with Veeam Backup & Replication. Replication from backups of vCloud Director virtual machines is not supported.
- Replica from backup can use the following backups only:
  - Backups that backup jobs or backup copy jobs create. These jobs must be configured on the same backup server where you configure the replication job.
  - Backups to which backup jobs or backup copy jobs are mapped. These jobs must be configured on the same backup server where you configure the replication job.
- The jobs mentioned in the previous list items must run periodically and produce new restore points. Otherwise, the replication job will have no data to retrieve and replicas will be in an outdated state.
- Replica from backup cannot use imported backups. However, there might be a situation when you need to use backups created on another backup server. In this case, use the instructions provided in Using Backups Created on Crashed Backup Server. Note that the backup server from which you import backups must not operate anymore, otherwise replica from backup may behave in an unexpected way.
- [For backups stored on scale-out backup repositories] Replica from backup ignores backups stored in a tier other than the performance tier.

How Replica from Backup Works

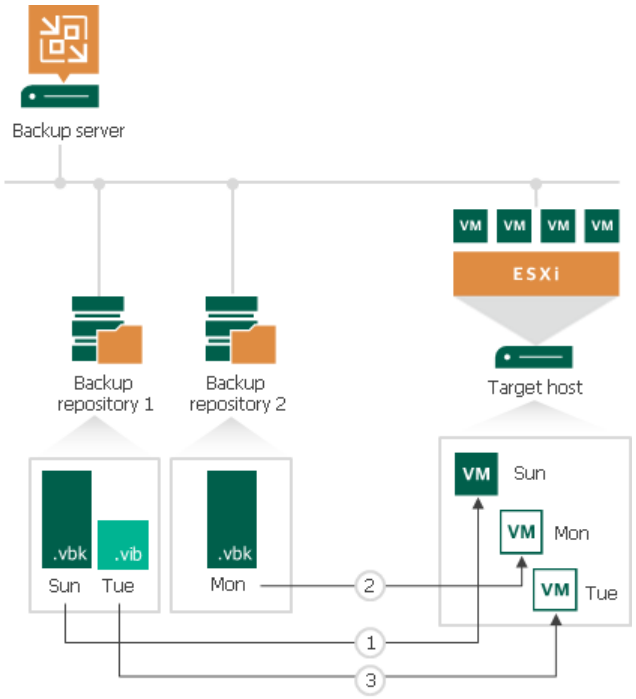
Replica from backup is performed along with a regular replication job. When you set up a replication job, you define a backup repository with VM backups as a source of data. If the backups for this VM are available in different backup repositories, you can select several backup repositories as a source. In this case, Veeam Backup & Replication will look for the latest VM restore point across these backup repositories.

For example, you configure two backup jobs that process the same VM, and target these jobs at two different backup repositories. The backup jobs create the following backup files:

- *Backup job 1* creates 2 restore points in *Backup repository 1*: full backup file on Sunday and incremental backup file on Tuesday.
- *Backup Job 2* creates 1 restore point in *Backup repository 2*: full backup file on Monday.

The replication job is configured to retrieve VM data from backups and is scheduled to run daily. In this case, the replication job retrieves VM data from backups in the following way:

1. On Sunday, the replication job retrieves VM data from the full backup file in *Backup repository 1*.
2. On Monday, the replication job retrieves VM data from the full backup file in *Backup repository 2*.
3. On Tuesday, the replication job retrieves VM data from the incremental backup file in *Backup repository 1*.
4. Till next Sunday, the replication job does not retrieve any VM data because backup files are not created.



In some situations, a new restore point in the backup repository may not be created by the time a replication job starts. In this case, Veeam Backup & Replication displays a warning notifying that the latest restore point has already been replicated. The replication job session finishes with the *Warning* status.

NOTE

When you replicate a VM over a production network, Veeam Backup & Replication retrieves VM data as of the latest VM state. When you replicate a VM from backup, Veeam Backup & Replication retrieves VM data as of the point in time when the backup was created. The VM replica restore point has the same timestamp as a corresponding VM backup restore point, not the time when the replica job session is run.

Using Backups Created on Crashed Backup Server

There might be a situation when you created backups on one backup server, the server crashed and you want to use these backups for replica from backup on another backup server.

Veeam Backup & Replication considers backups created on other backup servers as imported backups. Replica from backups cannot use imported backups, that is why you need to perform the following steps to use backups created on the crashed server:

1. Import the backups to the backup server where you create the replication job. You have several options:
   - You can connect the repository where the backups are stored to the backup server and then rescan the repository.
   - You can copy backups to a backup repository already added to the backup server and then rescan the repository.
   - You can copy backups to a backup repository already added to the backup server. Then edit repository settings and select the **Search the repository for existing backups and import them automatically check box** at the Review step of the wizard.
2. Create a new backup job or a backup copy job and map the imported backups to it. For more information on how to map backups, see Specify Backup Storage Settings.

   After you map a backup to a job, Veeam Backup & Replication stops considering the backup as imported.
3. Create a replication job. In the job settings, specify that you want to use backups as a data source and select the backup repository where the imported backups reside. For more information, see Specify Data Source.

> **IMPORTANT**
>
> Mind the following:
>
> - The backup job or backup copy job to which you map the imported backup file must run periodically and produce new restore points. Otherwise, the replication job will have no data to retrieve and replicas will be in an outdated state.
> - No other running backup servers must use the imported backups.

Creating Replication Jobs

To create VM replicas, you must configure a replication job. The replication job defines how, where and when to replicate VM data. One job can be used to process one VM or multiple VMs.

To create a replication job, use the **New Replication Job** wizard.

Before You Begin

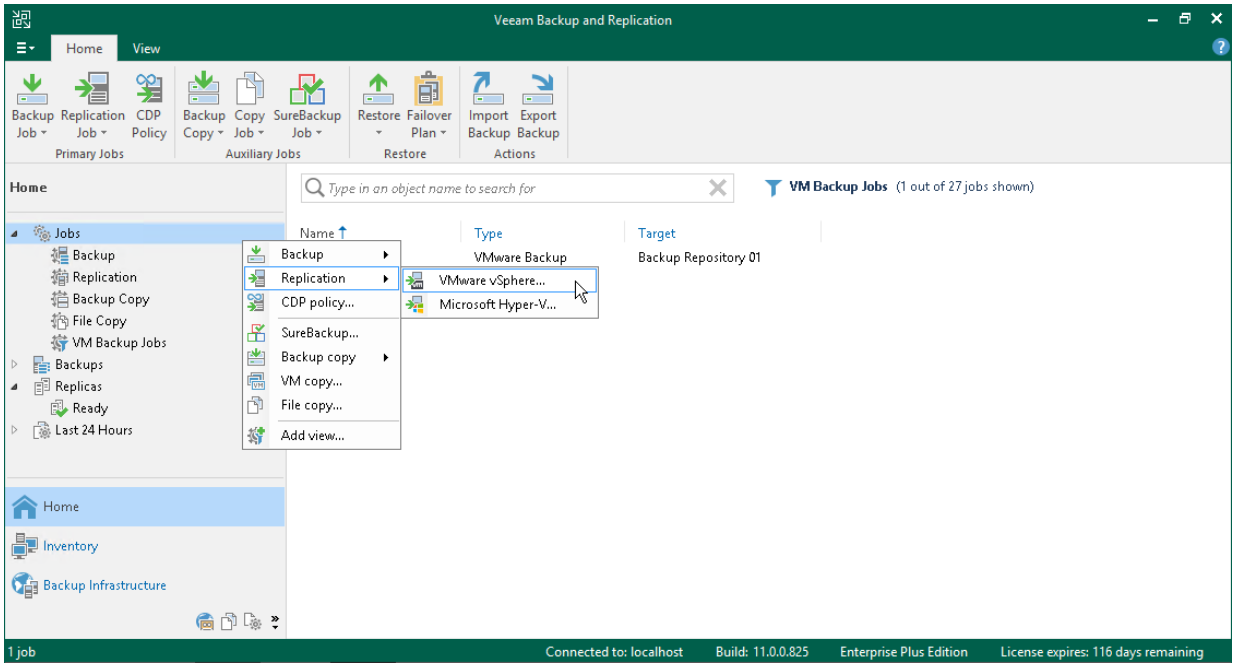
Before you create a replication job, check the following prerequisites and limitations:

- Check limitations for replication described in the Requirements and Limitations section.
- Backup infrastructure components that will take part in the replication process must be added to the backup infrastructure and properly configured. These include source and target ESXi hosts, one backup proxy for on-site replication scenario or two backup proxies for off-site replication scenario and backup repository for storing replica metadata.
- The backup server must be able to resolve short names and connect to source and target virtualization hosts.
- The target datastore must have enough free space to store disks of replicated VMs. To receive alerts about low space on the target datastore, configure global notification settings. For more information, see Specifying Other Notification Settings.
- If you plan to replicate VMs using WAN accelerators, source and target WAN accelerators must be added to the backup infrastructure and properly configured. For more information, see Adding WAN Accelerators.
- If you plan to replicate VMs using WAN accelerators, it is recommended that you pre-populate global cache on the target WAN accelerator before you start the replication job. Global cache population helps reduce the amount of traffic transferred over WAN. For more information, see Populating Global Cache.
- If you plan to replicate VMs from a backup, the backup job that you plan to use as the source must be configured beforehand. For more information, see Replica from Backup.
- If you plan to use pre-job and post-job scripts and/or pre-freeze and post-thaw scripts, you must create scripts before you configure the replication job. For the supported script format, see Pre-Freeze and Post-Thaw Scripts.

Step 1. Launch New Replication Job Wizard

To launch the **New Replication Job** wizard, do one of the following:

- On the **Home** view. On the ribbon, click **Replication Job** > **Virtual machine** > **VMware vSphere**.
- Open the **Home** view. In the inventory pane right-click the **Jobs** node and select **Replication** > **Virtual machine** > **VMware vSphere**.
- Open the **Inventory** view. In the working area, select VMs that you want to replicate and right-click one of them. Select **Add to replication job** > **New job** if you want to create a new replication job, or **Add to replication job** > *<Job Name>* if you want to add VMs to an existing replication job.



Step 2. Specify Job Name and Description

At the **Name** step of the wizard, specify a job name and description, and configure advanced settings for the replication job:

1. In the **Name** field, enter a name for the replication job.
2. In the **Description** field, provide a description for future reference.
3. If a network between your production and disaster recovery (DR) sites has low bandwidth, and you want to reduce the amount of traffic sent during the first run of the replication job, select the **Replica seeding (for low bandwidth DR sites)** check box.

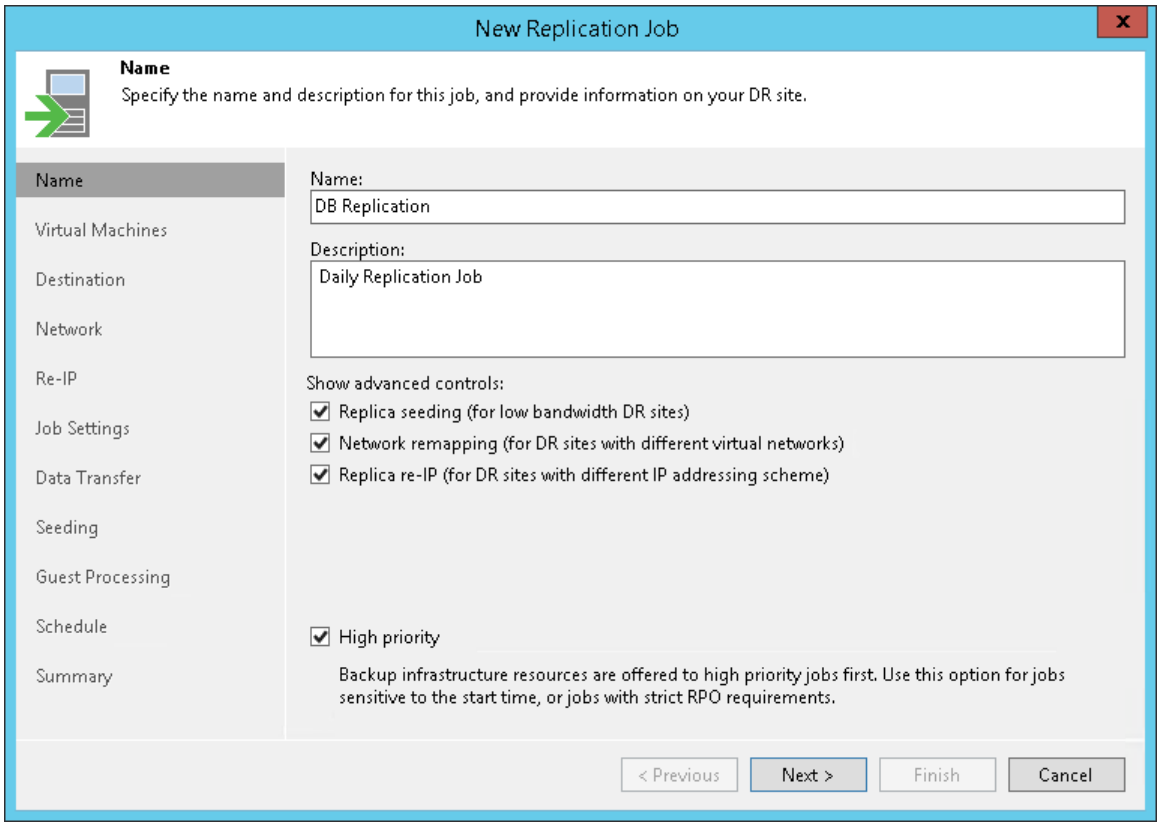
   When selected, this check box enables the **Seeding** step where you will have to configure replica seeding and mapping. For more information on seeding and mapping, see Replica Seeding and Mapping.
4. If your DR site networks do not match your production site networks, select the **Network remapping (for DR sites with different virtual networks)** check box.

   When selected, this check box enables the **Network** step where you will have to configure a network mapping table.
5. If the IP addressing scheme in your production site differs from the scheme in the DR site, select the **Replica re-IP (for DR sites with different IP addressing scheme)** check box.

   When selected, this check box enables the **Re-IP** step where you will have to configure replica re-IP rules.
6. If you want the resource scheduler of Veeam Backup & Replication to prioritize this job higher than other similar jobs and to allocate resources to it in the first place, select the **High priority** check box. For more information on job priorities, see Job Priorities.

TIP

In the UI, jobs with the **High priority** option enabled will be marked with a special flag ( ).



Step 3. Select VMs to Replicate

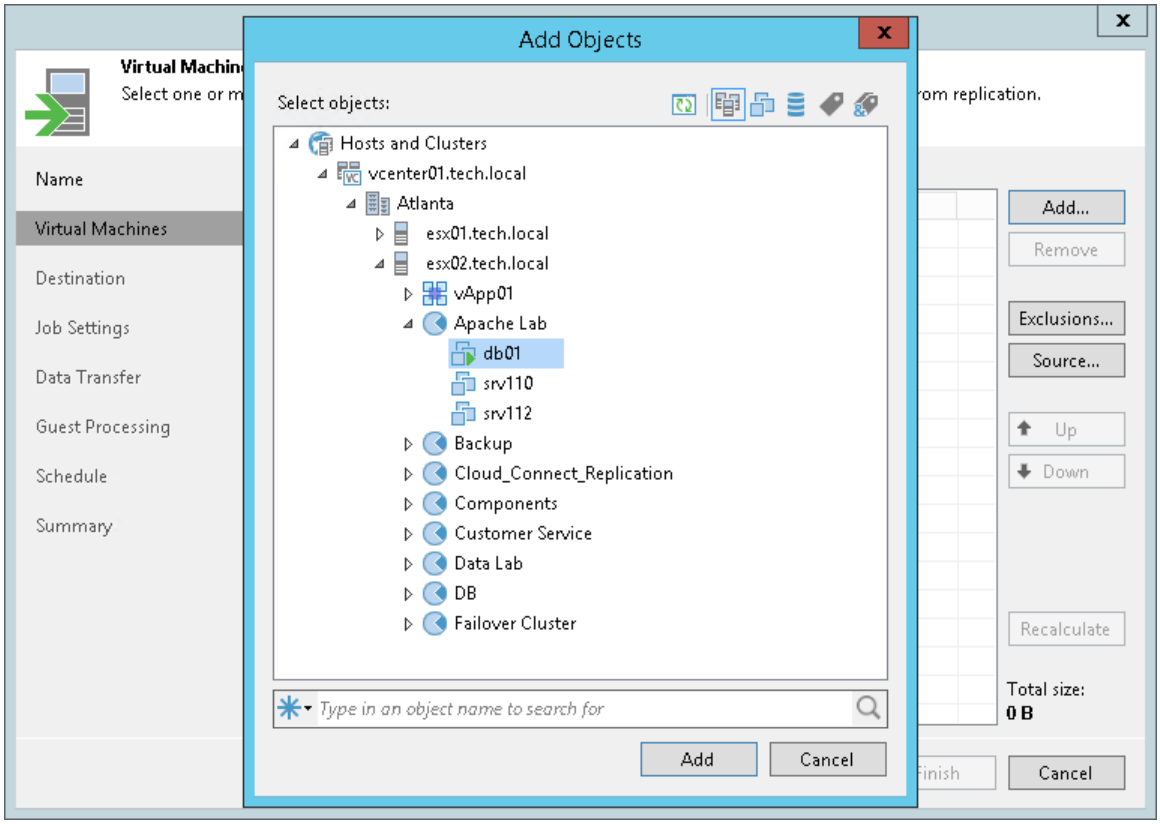
At the **Virtual Machines** step of the wizard, select VMs and VM containers (hosts, clusters, folders, resource pools, VirtualApps, datastores or tags) that you want to replicate:

1. Click **Add**.
2. In the **Add Object** window, select the necessary VMs or VM containers and click **Add**. If you select VM containers and add new VMs to this container in future, Veeam Backup & Replication will update replication job settings automatically to include these VMs.

   You can use the toolbar at the top right corner of the window to switch between views. Depending on the view you select, some objects may not be available. For example, if you select **Tags combination** view, no resource pools, hosts or clusters will be displayed in the tree. In the **Tags combination** view, you can select multiple tags and only VMs that have all the selected tags will be processed by the job.

   To quickly find the necessary VMs, you can use the search field at the bottom of the **Add Object** window. If you want to switch between types of VMs you want to search through, use the button to the left of the search field.

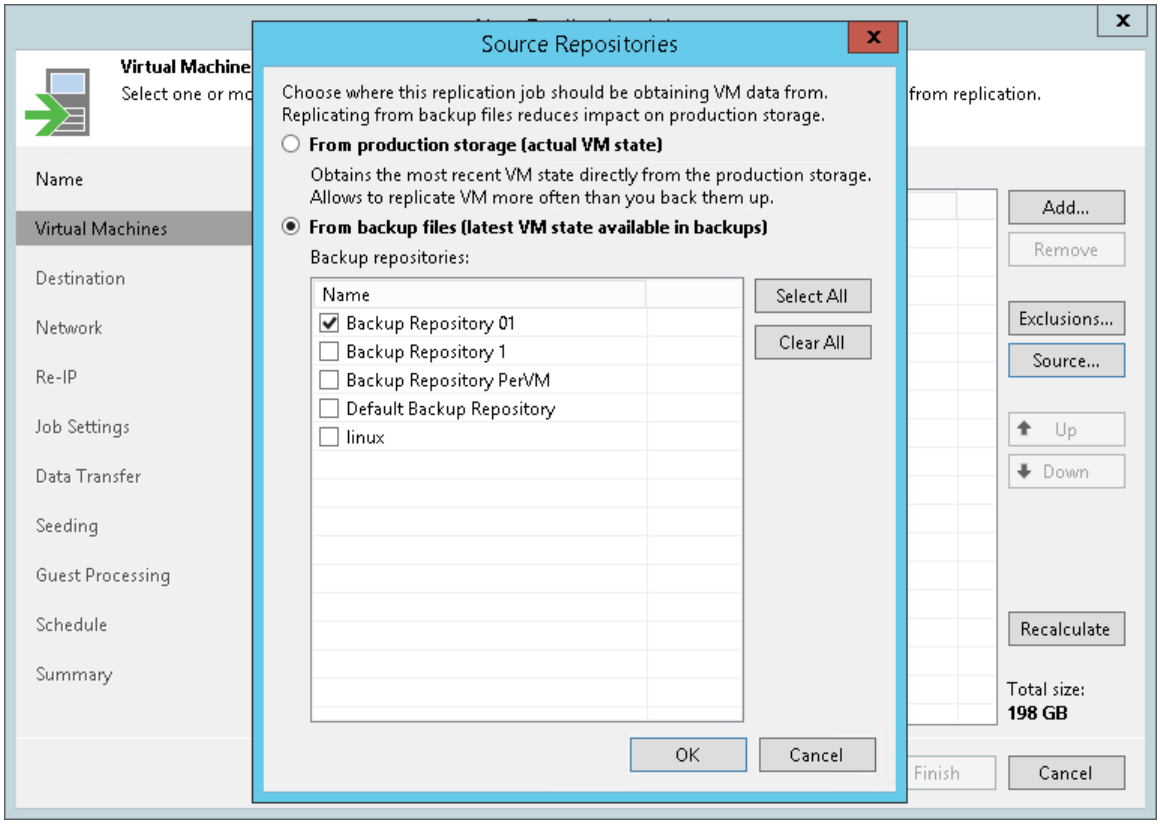
The total size of objects added to the job is displayed in the **Total size** field. Use the **Recalculate** button to refresh the total size value after you add a new object to the job.



Step 4. Specify Data Source

You can select a data source from which Veeam Backup & Replication will read VM data:

1. At the **Virtual Machines** step of the wizard, click **Source**.
2. In the **Source Repositories** window, select one of the following options:
   - **From production storage**. In this case, Veeam Backup & Replication will retrieve VM data from datastores connected to the source ESXi host.
   - **From backup files**. In this case, Veeam Backup & Replication will read VM data from the backup chain already existing in the selected backup repository. This option can be used in the replica from backup scenario. For more information, see Replica from Backup.



Step 5. Exclude Objects from Replication Job

After you have added VMs and VM containers to the replication job, you can specify which objects you want to exclude from replicas. You can exclude the following types of objects:

- VMs or VM containers
- VM disks

> **NOTE**
>
> To make the replication process faster and reduce the size of created replicas, Veeam Backup & Replication automatically excludes the following objects from replication:
>
> - VM log files
> - VM templates from VM containers

Excluding VMs or VM Containers

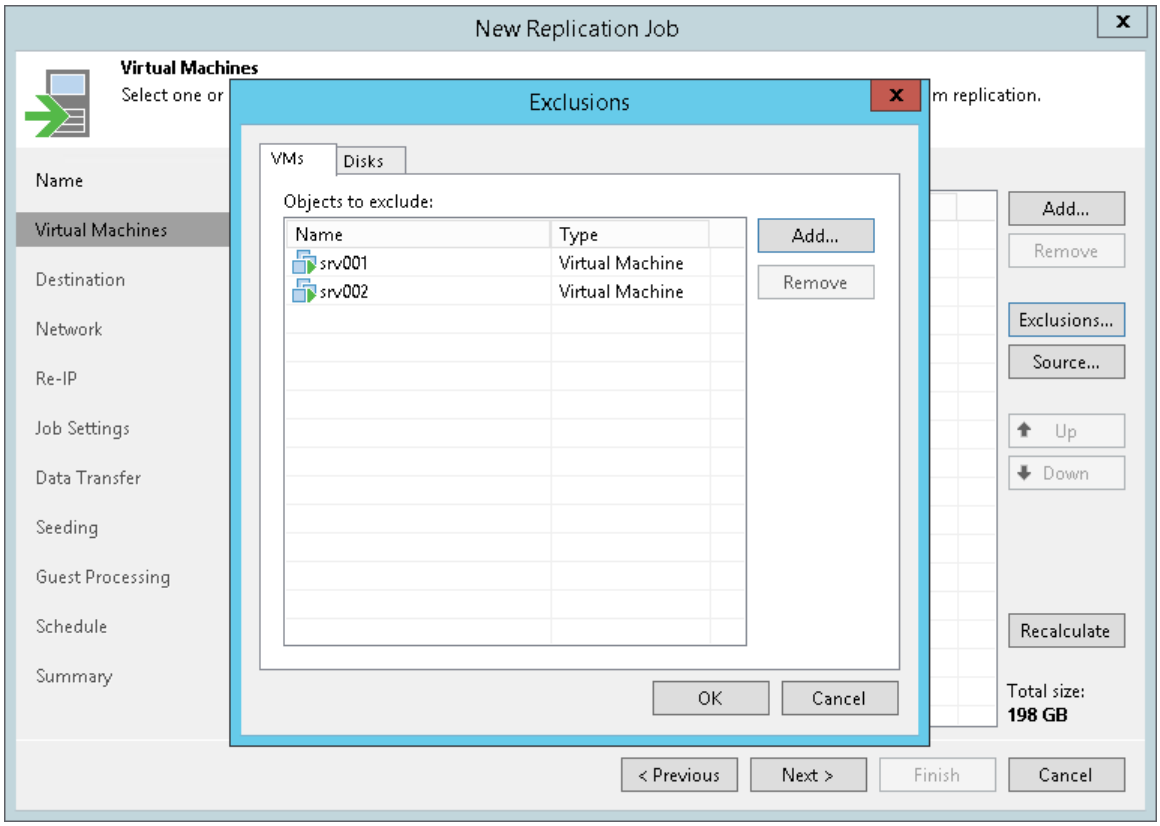
To exclude VMs from a VM container:

1. At the **Virtual Machines** step of the wizard, click **Exclusions**.
2. In the **Exclusions** window, check that the **VMs** tab is selected and click **Add**.
3. In the **Add Objects** window, select VMs or VM containers that you want to exclude from being replicated and click **Add**.

   You can use the toolbar at the top right corner of the window to switch between views. Depending on the view you select, some objects may not be available. For example, if you select **Tags combination** view, no resource pools, hosts or clusters will be displayed in the tree. In the **Tags combination** view, you can select multiple tags and only VMs that have all the selected tags will be excluded from the job.

   You can also use the **Show full hierarchy** check box to display the hierarchy of all VMware Servers added to the Veeam Backup & Replication infrastructure.

4. Click **OK**.



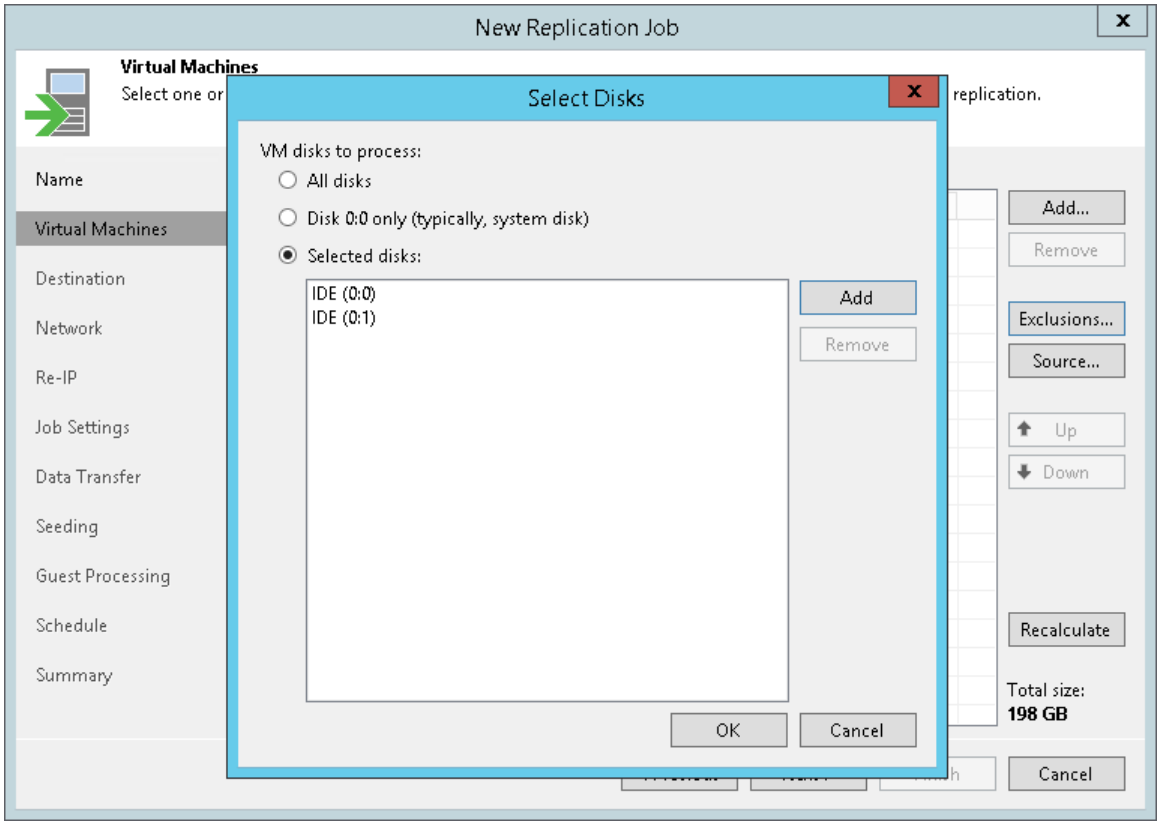
Excluding Disks

To exclude VM disks:

1. At the **Virtual Machines** step of the wizard, click **Exclusions**.
2. In the **Exclusions** window, do the following:
   a. Switch to the **Disks** tab.
   b. If you want to exclude disks of VMs that are added as a part of containers, click **Add**. In the **Add Objects** window, select the necessary VMs and click **Add**. Veeam Backup & Replication will include these VMs in the list as standalone objects.
   c. In the **Disks to process** list, select VMs or VM containers whose disks you want to exclude.
   d. Click **Edit**.
3. In the **Select Disks** window, select disks that you want to replicate: all disks, 0:0 disks (as a rule, system disks) or specific IDE, SCSI, SATA or NVMe disks. Disks that you do not select will be excluded from processing. Click **OK**.
4. In the **Exclusions** window, click **OK**.

NOTE

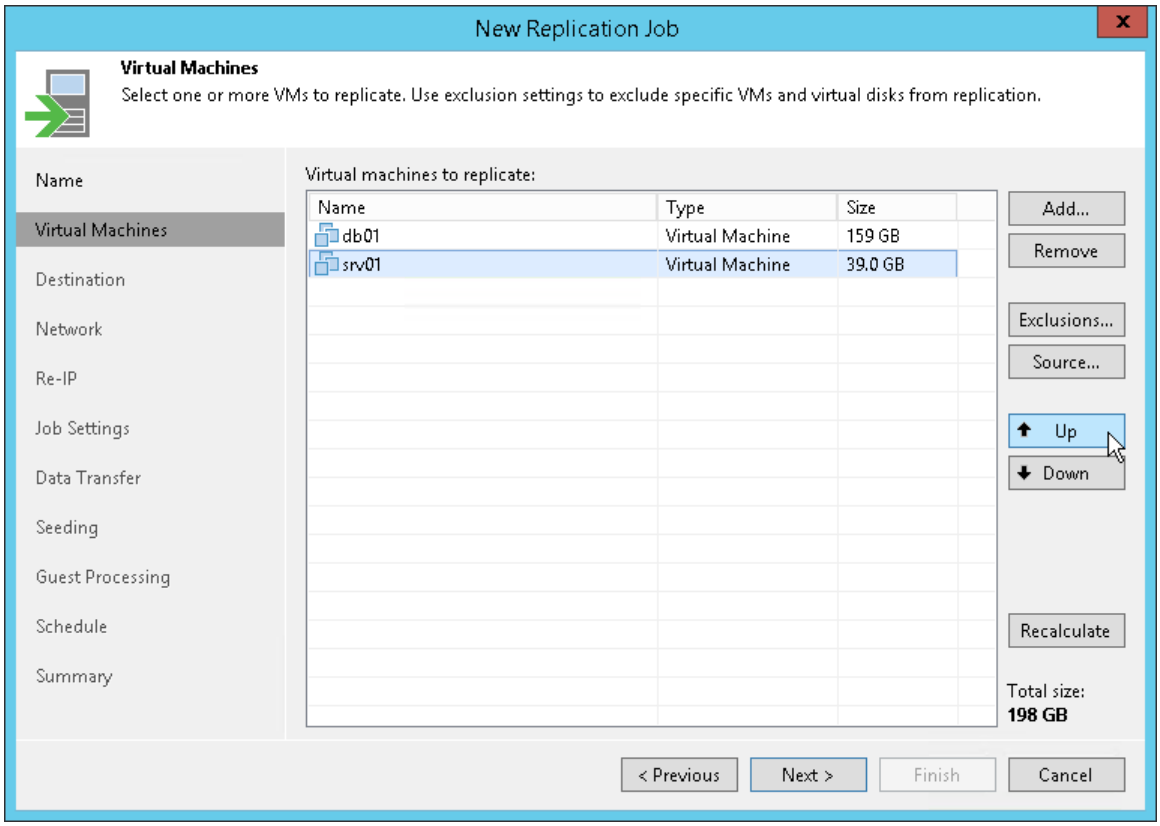
If you exclude disks from being replicated and enable application-aware processing, Veeam Backup & Replication will still perform application-aware processing for the excluded disks. This means that VSS will process disk data.



Step 6. Define VM Processing Order

At the **Virtual Machines** step of the wizard, click **Up** and **Down** to change the processing order. VMs at the top of the list have a higher priority and will be processed first.

> **NOTE**
>
> Mind the following:
>
> - VMs inside a VM container are processed at random. To ensure that VMs are processed in the defined order, you must add them as standalone VMs, not as a part of containers.
> - The processing order may differ from the order that you have defined. For example, if resources of a VM that is higher in the priority are not available, and resources of a VM that is lower in the priority are available, Veeam Backup & Replication will process the VM with the lower priority first.



Step 7. Specify Replica Destination

At the **Destination** step of the wizard, select a target host or cluster, resource pool, folder and datastore for replicas, and types of replica disks:

1. Next to the **Host or cluster** field, click **Choose** and select a host or cluster where replicas must be registered. If you select a cluster or vCenter Server, the replication process will become more sustainable — the replication process will not fail if there is at least one available host in the cluster.

   If you select a cluster as a destination, Veeam Backup & Replication will request VMware to send the list of available hosts, and will select the first host in this list as the destination for the replicas. These replicas will be stored on the datastore with the most free disk space.

2. Next to the **Resource pool** field, click **Choose** and select a resource pool to which replicas will be added.

   If you have selected to replicate multiple VMs and want to add individual replicas to other resource pools:

   a. Click the **Pick resource pool for selected replicas** link.

   b. In the **Choose Resource Pool** window, click **Add VM**.

   c. In the **Add Objects** window, select the necessary VMs and click **Add**.

   d. In the **Choose Resource Pool** window, select the necessary VMs in the **Replica VM resource pool** list. At the bottom of the window, click **Resource Pool**.

   e. In the **Select Resource Pool** window, select the necessary resource pool and click **OK**.

3. Next to the **VM folder** field, click **Choose** and select a folder where all VM files will be stored. Note that the **VM folder** section is disabled if you have selected a standalone ESXi host as the target for replicas.

   If you have selected to replicate multiple VMs and want to place individual replicas to other folders:

   a. Click the **Pick VM folder for selected replicas** link.

   b. In the **Choose Folder** window, click **Add VM**.

   c. In the **Add Objects** window, select the necessary VMs and click **Add**.

   d. In the **Choose Folder** window, select the necessary VMs in the **Replica VM folder** list. At the bottom of the window, click **VM Folder**.

   e. In the **Select Folder** window, select the necessary folder.

4. Next to the **Datastore** field, click **Choose** and select a datastore where replica files will be stored. Note that if you have chosen to replicate VMs to a cluster, Veeam Backup & Replication displays only shared datastores.

   If you have selected to replicate multiple VMs and want to place individual replicas to other datastores:

   a. Click the **Pick datastore for selected virtual disks** link.

   b. In the **Choose VM Files Location** window, click **Add VM**.

   c. In the **Add Objects** window, select the necessary VMs and click **Add**.

   d. In the **Choose VM Files Location** window, select the necessary VMs in the **Files location** list. At the bottom of the window, click **Datastore**.

   e. In the **Select Datastore** window, select the necessary datastore.

5. If you want to store replica configuration files and disk files in different datastores:
   a. Click the **Pick datastore for selected virtual disks** link.
   b. In the **Choose VM Files Location** window, click **Add VM**.
   c. In the **Add Objects** window, select the necessary VMs and click **Add**.
   d. In the **Choose VM Files Location** window, expand the necessary VMs in the **Files location** list, and select the necessary files. At the bottom of the window, click **Datastore**.
   e. In the **Select Datastore** window, select the destination for the selected type of files.

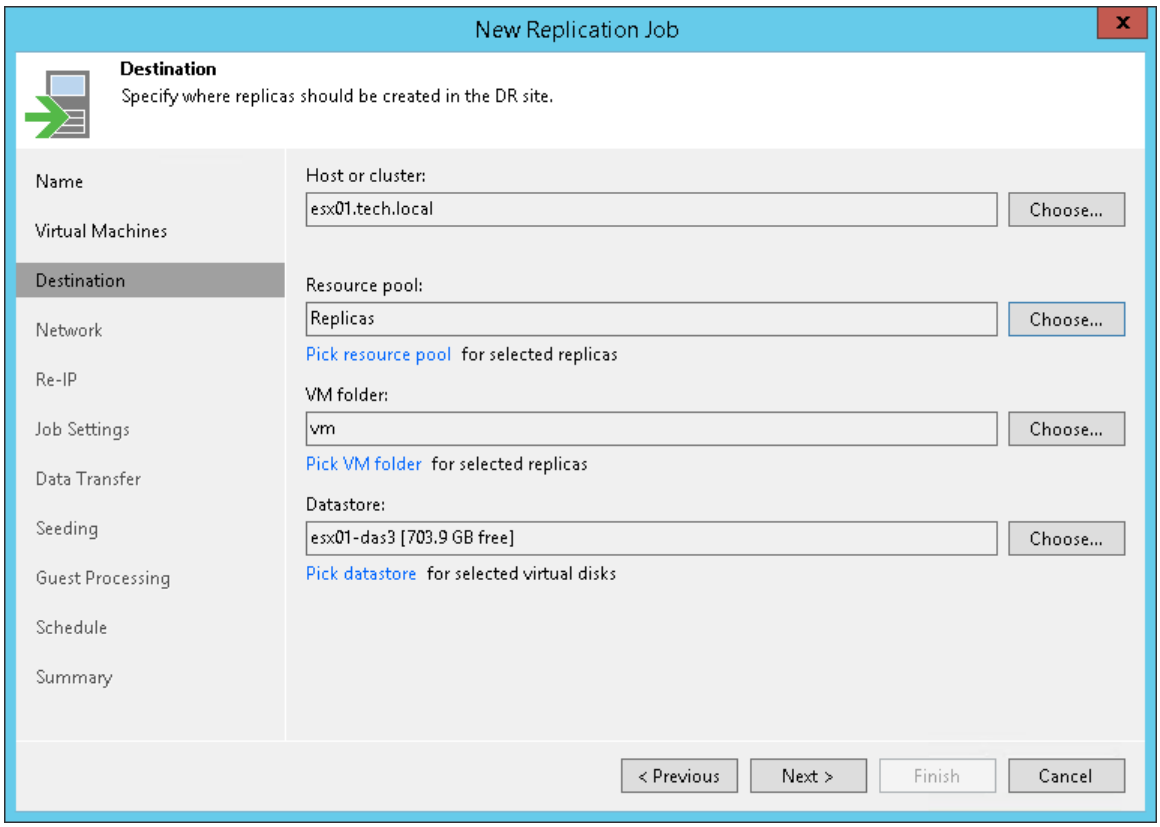
> **NOTE**
>
> After the first run of the replication job, you can change the target location for replicated files. However, the target will be changed only for new files that were created on the source VM after the first run. The target for old files will not be changed.

6. You can change types of replica disks. By default, Veeam Backup & Replication saves disks in the thin type.

   To change replica disk types:
   a. Click the **Pick datastore for selected virtual disks** link.
   b. In the **Choose VM Files Location** window, click **Add VM**.
   c. In the **Add Objects** window, select VMs whose disk type you want to change and click **Add**.
   d. In the **Choose VM Files Location** window, select the necessary VMs in the **Files location** list. At the bottom of the window, click **Disk type**.
   e. In the **Disk Type Settings** window, select a type that will be used to restore replica disk files: same as source, thin, thick lazy zeroed or thick eager zeroed.

      For more information about disk types, see VMware Docs.

> **NOTE**
>
> Disk type change is available only for VMs that use virtual hardware version 7 or later.



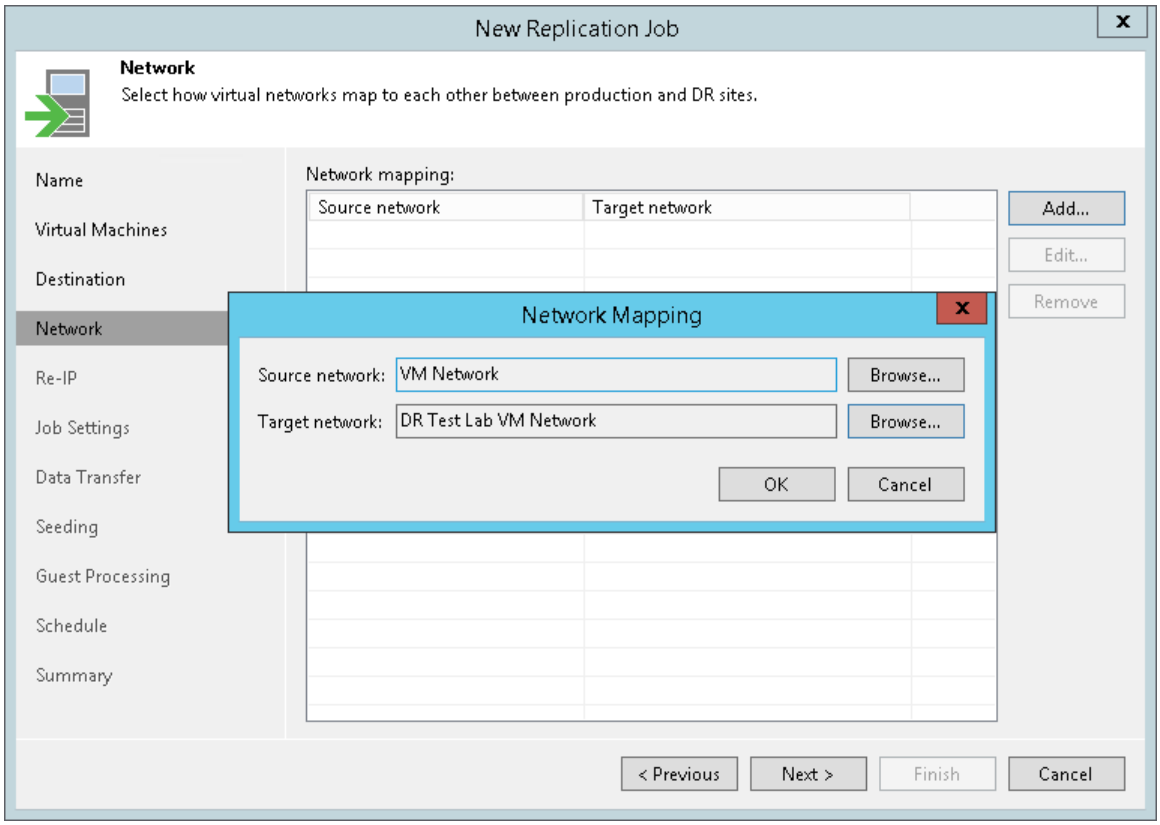
Step 8. Create Network Map Table

The **Network** step of the wizard is available if you have selected the **Network remapping** option at the Name step of the wizard.

At the **Network** step of the wizard, configure a network mapping table. This table maps networks in the production site to networks in the disaster recovery (DR) site. When the job starts, Veeam Backup & Replication will check the network mapping table. Then Veeam Backup & Replication will update replica configuration files to replace the production networks with the specified networks in the DR site. As a result, you will not have to re-configure network settings manually.

To add a row to a network mapping table:

1. Click **Add**.
2. In the **Network Mapping** window, click **Browse** next to the **Source network** field.
3. In the **Select Network** window, select the production network to which the original VMs are connected and click **OK**.
4. In the **Network Mapping** window, click **Browse** next to the **Target network** field.
5. In the **Select Network** window, select a network in the DR site to which replicas will be connected and click **OK**.
6. In the **Network Mapping** window, click **OK**.



Step 9. Configure Re-IP Rules

The **Re-IP** step of the wizard is available if you have selected the **Replica re-IP** option at the Name step of the wizard. This step applies only to VMs with Microsoft Windows OSes.

At the **Re-IP** step of the wizard, configure re-IP rules. These rules map IPs in the production site to IPs in the disaster recovery (DR) site. When you perform failover, Veeam Backup & Replication will check the configured re-IP rules and will change replica IPs if the rules apply. VM replicas will get new IP addresses according to the network masks specified in the rules, so that you will be able to reach replicas in the DR site.

To configure a re-IP rule:

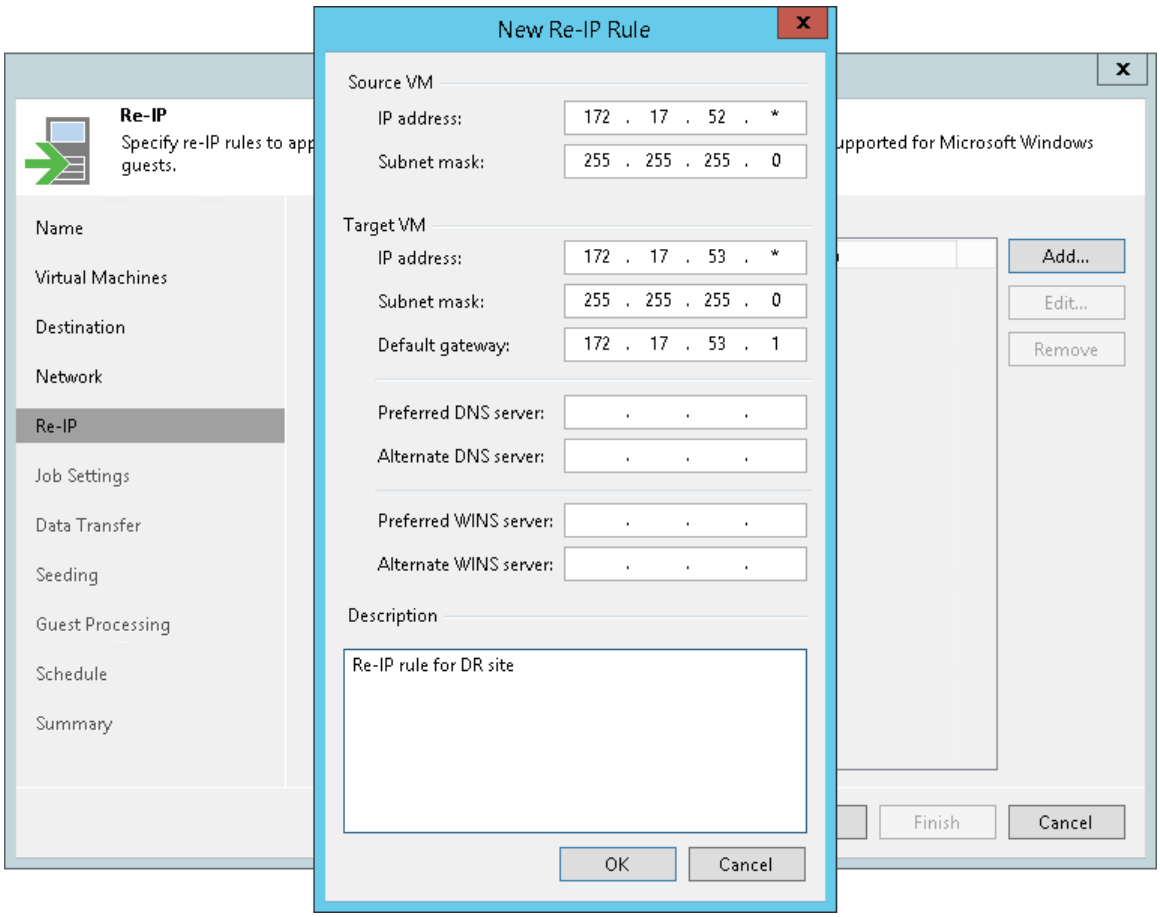
1. Click **Add**.
2. In the **Source VM** section, specify an IP numbering scheme used in the production site.

   To facilitate the configuration, Veeam Backup & Replication detects an IP address and subnet mask for the backup server and pre-populates values in the **Source VM** section.
3. In the **Target VM** section, specify an IP address, subnet mask and default gateway that will be used for replicas in the DR site. If required, specify the DNS and WINS server addresses.
4. In the **Description** field, provide a description the rule.
5. Click **OK**.

NOTE

Mind the following:

- Do not use 0 to specify IP address ranges. In Veeam Backup & Replication, value 172.16.17.0 means a regular IP address 172.16.17.0, not an IP address range. To specify a range, use the asterisk character (*).
- Replica re-IP works only if you perform replica failover using Veeam Backup & Replication. If you power on a VM replica in some other way, for example, manually using vSphere Client, re-IP rules will not be applied to it.
- The backup server OS must support mounting of the system disks of VMs that will be replicated.



Step 10. Specify Replication Job Settings

At the **Job Settings** step of the wizard, specify a backup repository for storing replica metadata, replica name and number of restore points to keep:

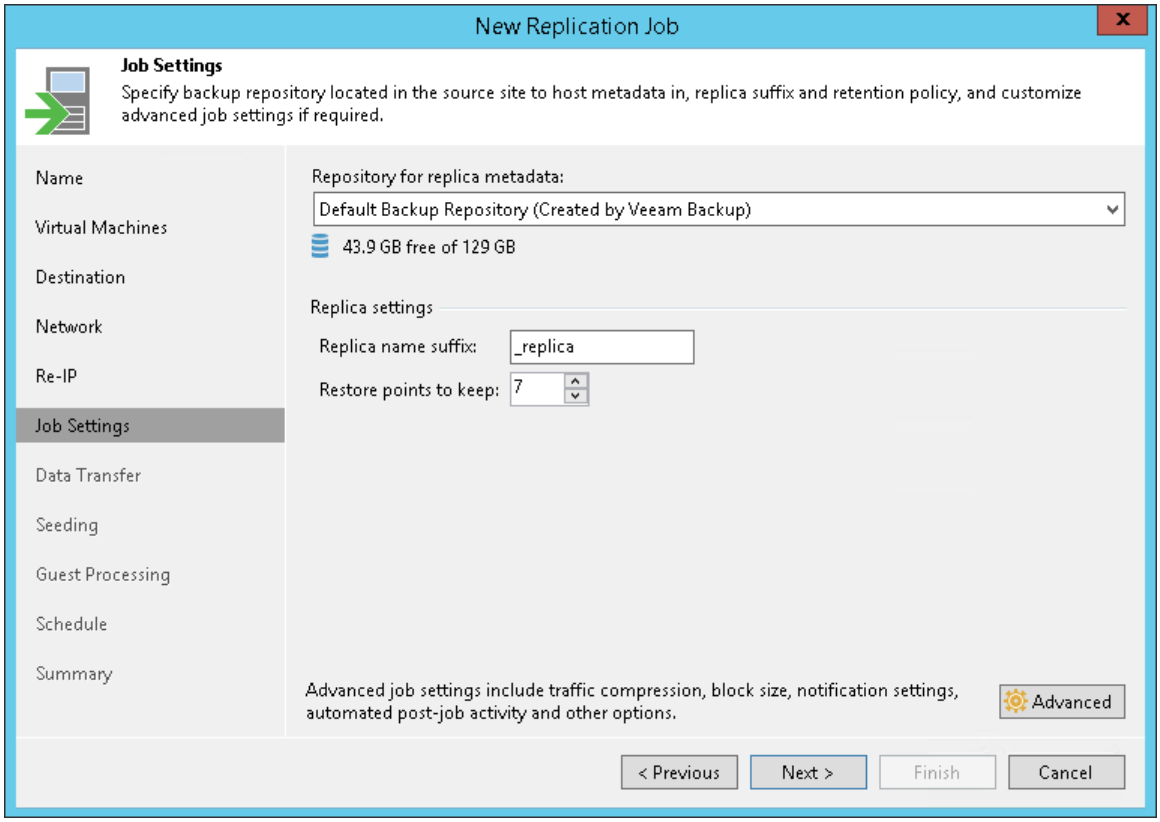
1. From the **Repository for replica metadata** list, select a backup repository that will store metadata for VM replicas.

   > **IMPORTANT**
   >
   > Mind the following:
   >
   > - You cannot store VM replica metadata on deduplicating storage appliances. During replication jobs, Veeam Backup & Replication frequently reads and writes small portions of metadata from/to the backup repository. Frequent access to metadata causes low performance of deduplicating storage appliances, which may result in low performance of replication jobs.
   > - You cannot store replica metadata in a scale-out backup repository.

2. In the **Replica name suffix** field, enter a suffix that will be added to the original VM names.

   To register a VM replica on the target host, Veeam Backup & Replication appends the specified suffix to the name of the source VMs. Files of the VM replica are placed to the *VMname_suffix* folder on the selected datastore.

3. In the **Restore points to keep** field, specify the number of restore points that the replication job must maintain. Due to VMware restrictions on the number of VM snapshots, the maximum number of restore points for VM replicas is limited to 28. When the specified number is exceeded, Veeam Backup & Replication removes the earliest restore points.



Step 11. Specify Advanced Replica Settings

At the **Job settings** step of the wizard, you can specify the following settings for the replication job:

- Traffic settings
- Notification settings
- vSphere settings
- Integration settings
- Script settings

> **TIP**
>
> After you specify necessary settings for the replication job, you can save them as default settings. To do this, click **Save as Default** at the bottom left corner of the **Advanced Settings** window. When you create a new replication job, Veeam Backup & Replication will automatically apply the default settings to the new job.

Traffic Settings

You can optimize data traffic sent over network by specifying which data you want to replicate, data compression level and optimize the job performance and storage usage:

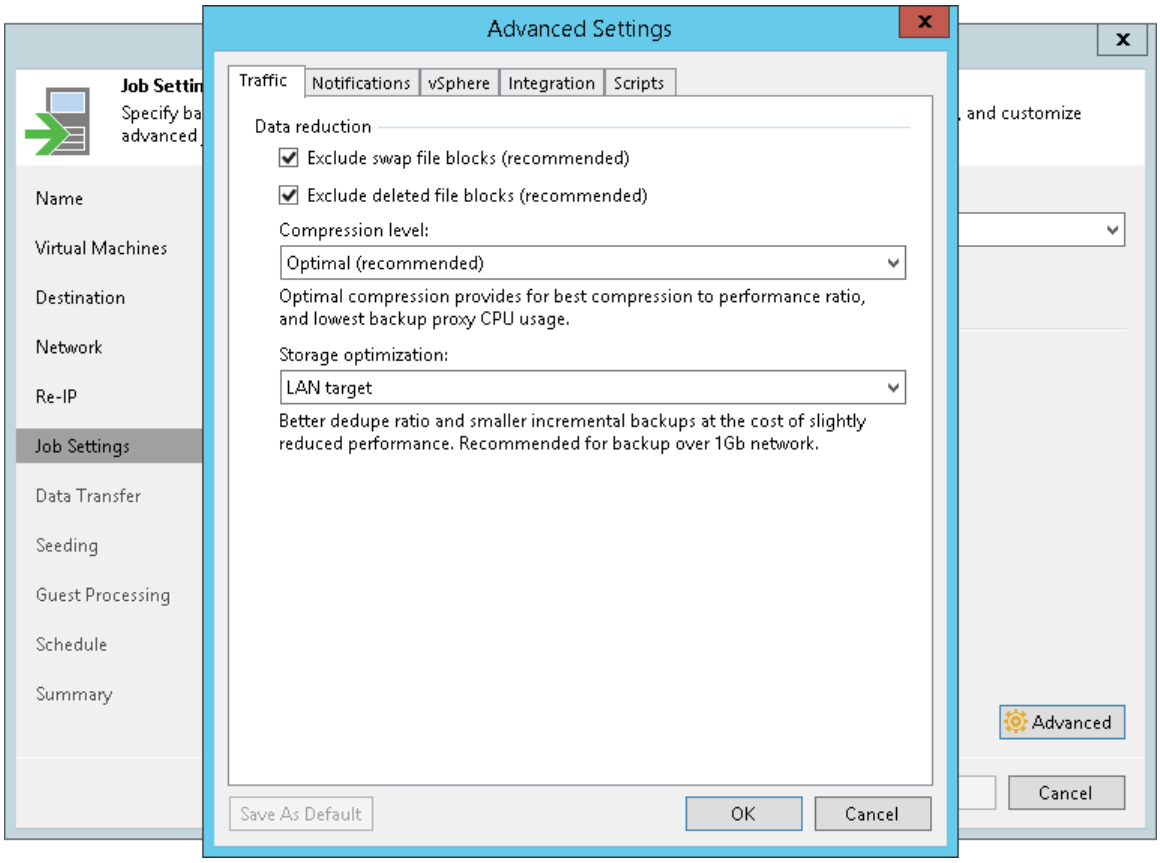
1. At the **Job Settings** step of the wizard, click **Advanced**.
2. In the **Advanced Settings** window, check that the **Traffic** tab is selected.
3. [For Microsoft Windows NTFS] By default, Veeam Backup & Replication excludes data blocks of the `hiberfil.sys` and `pagefile.sys` system files from replicas. For more information on how Veeam Backup & Replication excludes data blocks of these system files, see Swap Files.

   If you want to include data blocks of the `hiberfil.sys` and `pagefile.sys` system files into replicas, clear the **Exclude swap file blocks** check box. Note that including these files into replicas will increase their size.
4. By default, Veeam Backup & Replication does not copy deleted file blocks ("dirty" blocks on the VM guest OS) to the target location. For more information, see Deleted File Blocks.

   If you want to include dirty data blocks into VM replicas, clear the **Exclude deleted file blocks** check box. Note that including these files into replicas will increase their size.
5. From the **Compression level** list, select a compression level for VM replicas. For more information on data compression and compressions levels, see Data Compression.

   > **NOTE**
   >
   > Compression is applicable only if VM data is transferred between two backup proxies. If one backup proxy acts as the source and target backup proxy, VM data is not compressed at all.

6. In the **Storage optimization** section, select what type of target you plan to use for replication. Depending on the selected type, Veeam Backup & Replication will use different data block sizes to process VMs. For more information on the available storage optimization options, the used data blocks and how they affect performance, see Storage Optimization.



Notification Settings

To specify notification settings for the replication job:

1. At the **Job Settings** step of the wizard, click **Advanced**.
2. In the **Advanced Settings** window, click the **Notifications** tab.
3. To receive SNMP traps when the job completes successfully, select the **Send SNMP notifications for this job** check box.

   SNMP traps will be sent if you configure global SNMP settings in Veeam Backup & Replication and configure software on recipient machine to receive SNMP traps. For more information, see Specifying SNMP Settings.

4. To receive notifications by email in case of job failure, success or warning, select the **Send email notifications to the following recipients** check box. Then configure notification settings:

   a. Check that you have configured global email notification settings as described in the Configuring Global Email Notification Settings section.

   b. In the text field, specify a recipient email address. If you want to specify multiple addresses, separate them by a semicolon.

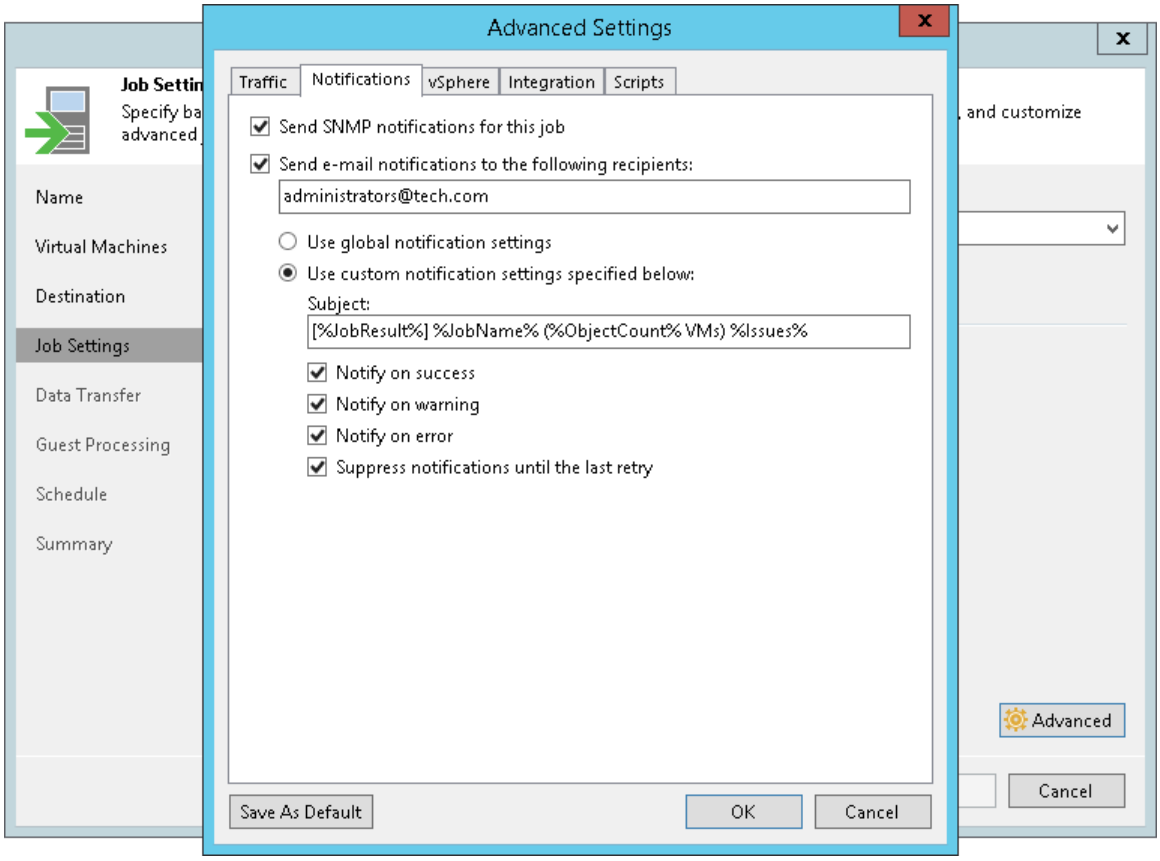
   c. To use global notification settings, select **Use global notification settings**.

d. To specify a custom notification subject and redefine at which time notifications must be sent, select **Use custom notification settings specified below.** Then specify the following settings:

   i. In the **Subject** field, specify a notification subject. You can use the following variables in the subject: *%JobResult%*, *%JobName%*, *%ObjectCount%* (number of VMs in the job) and *%Issues%* (number of VMs in the job that have been processed with the *Warning* or *Failed* status).

   ii. Select the **Notify on success**, **Notify on error** or **Notify on warning** check boxes to receive an email notification if the job gets the *Warning*, *Success* or *Error* status.

   iii. Select the **Suppress notifications until the last retry** check box to receive the notification about the final job status. If you do not enable this option, Veeam Backup & Replication will send one notification per every job retry.



vSphere Settings

To specify vSphere settings for the replication job:

1. At the **Job Settings** step of the wizard, click **Advanced**.
2. Click the **vSphere** tab.
3. Select the **Enable VMware tools quiescence** check box to freeze the file system of processed VMs during replication.

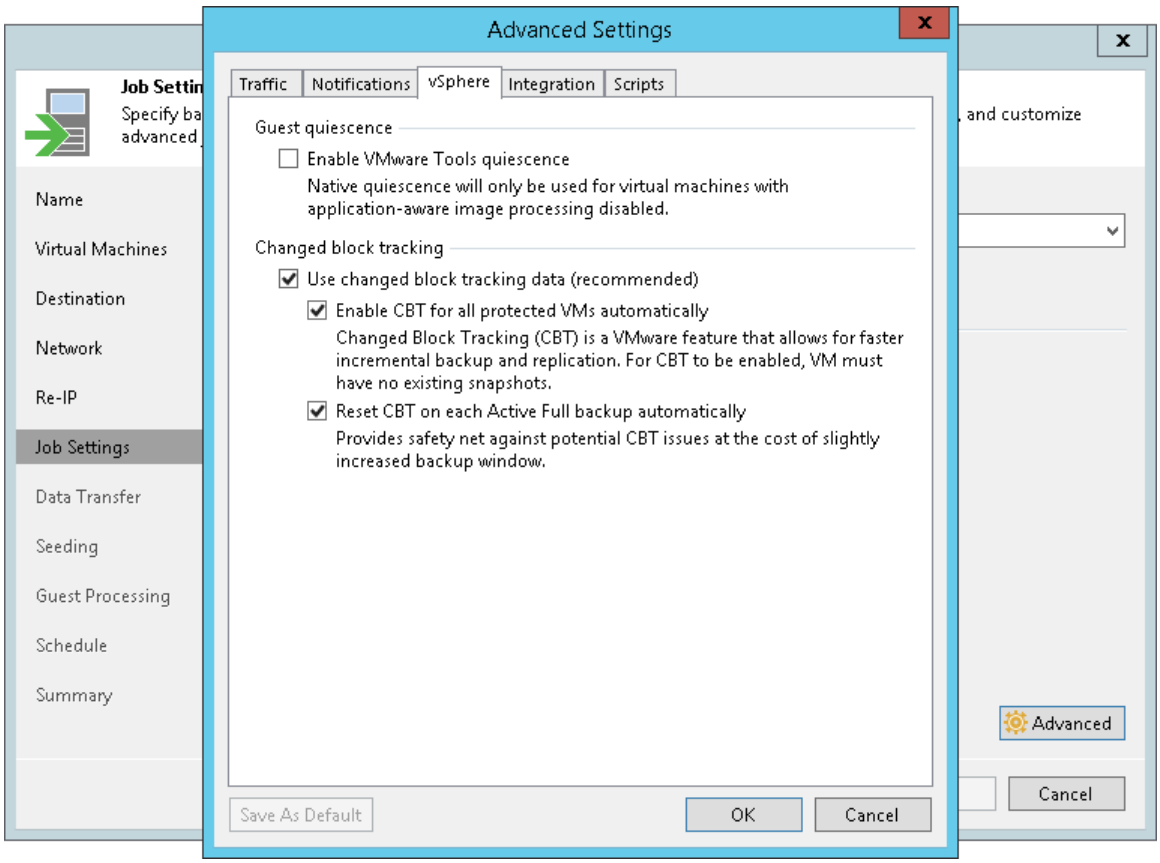
   Depending on the VM version, Veeam Backup & Replication will use the VMware FileSystem Sync Driver (vmsync) driver or VMware VSS component in VMware Tools for VM snapshot creation. These tools are responsible for quiescing the VM file system and bringing the VM to a consistent state suitable for backup. For more information, see VMware Tools Quiescence.

4. In the **Changed block tracking** section, configure VMware vSphere changed block tracking (CBT):
   a. To enable CBT, make sure that the **Use changed block tracking data** check box is selected.
   b. To force using CBT even if CBT is disabled in VM configuration, make sure that the **Enable CBT for all processed VMs automatically** check box is selected.
   c. To reset CBT after the replication job starts for the first time, make sure that the **Reset CBT on each Active Full backup automatically** check box is selected.

      CBT reset helps avoid issues, for example, when CBT returns incorrect changed data.

   For more information on CBT, see Changed Block Tracking.

> **IMPORTANT**
>
> You can use CBT for VMs with virtual hardware version 7 or later. These VMs must not have existing snapshots.



Integration Settings

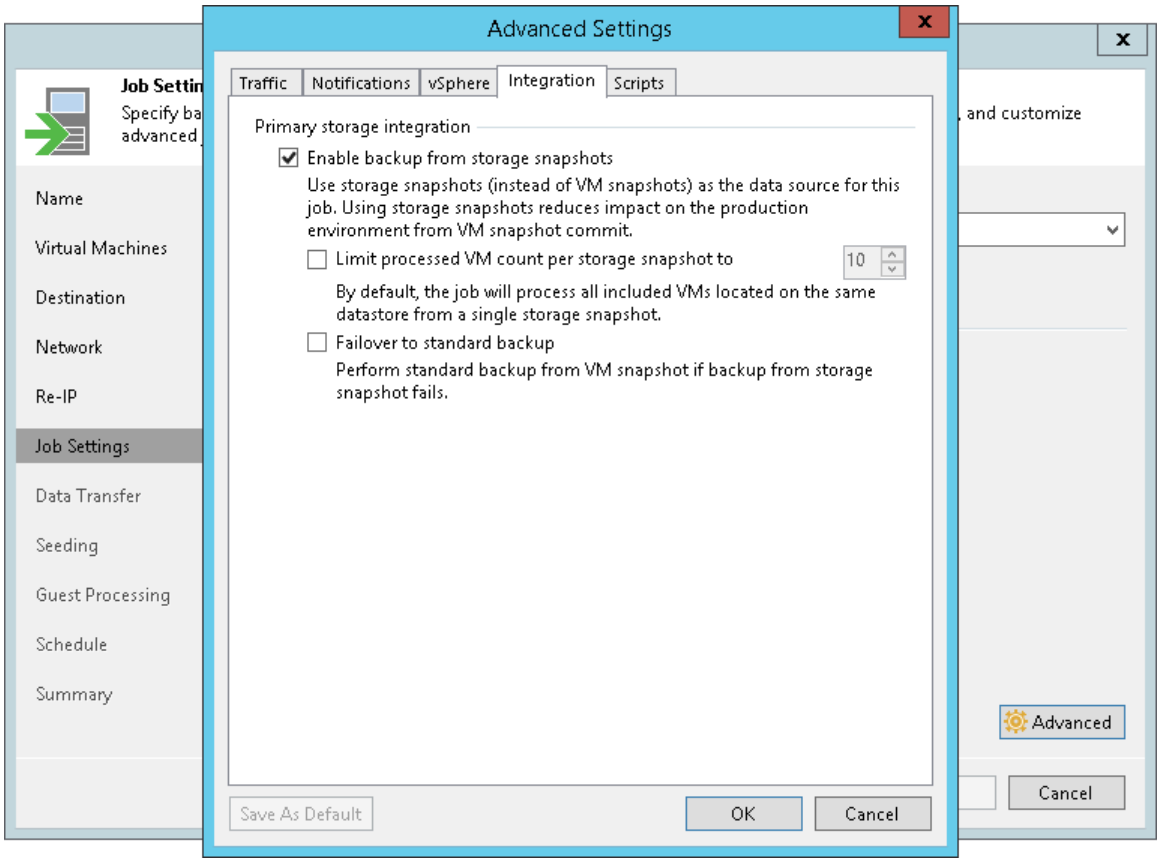
On the **Integration** tab, define whether you want to use the Backup from Storage Snapshots technology to create a VM replica. Backup from Storage Snapshots lets you leverage storage snapshots for VM data processing. The technology improves RPOs and reduces the impact of replication activities on the production environment. For more information, see the Backup from Storage Snapshot section in the Veeam Backup & Replication User Guide.

To specify storage integration settings for the replication job:

1. Check prerequisites. For more information, see the Configuring Backup from Storage Snapshots section in the Veeam Backup & Replication User Guide.

2. At the **Job Settings** step of the wizard, click **Advanced**.
3. Click the **Integration** tab.
4. By default, Backup from Storage Snapshots functionality is enabled. If you do not want to use it, clear the **Enable backup from storage snapshots** check box.
5. If you add to the job many VMs whose disks are located on the same volume or LUN, select the **Limit processed VM count per storage snapshot to** check box and specify the number of VMs for which one storage snapshot must be created.

   In a regular job processing course, Veeam Backup & Replication creates a VMware snapshot for every VM added to the job and then triggers one storage snapshot for all VMs. In some situations, creating VMware snapshots for all VMs may require a lot of time. If you limit the number of VMs per storage snapshot, Veeam Backup & Replication will divide VMs into several groups, trigger a separate storage snapshot for every VM group and read VM data from these snapshots. As a result, the job performance will increase. For more information, see the Limitation on Number of VMs per Snapshot section in the Veeam Backup & Replication User Guide.
6. If Veeam Backup & Replication fails to create a storage snapshot, VMs whose disks are hosted on the storage system will not be processed by the job. To fail over to the regular data processing mode and replicate such VMs, select the **Failover to standard backup** check box.



Script Settings

To specify script settings for the replication job:

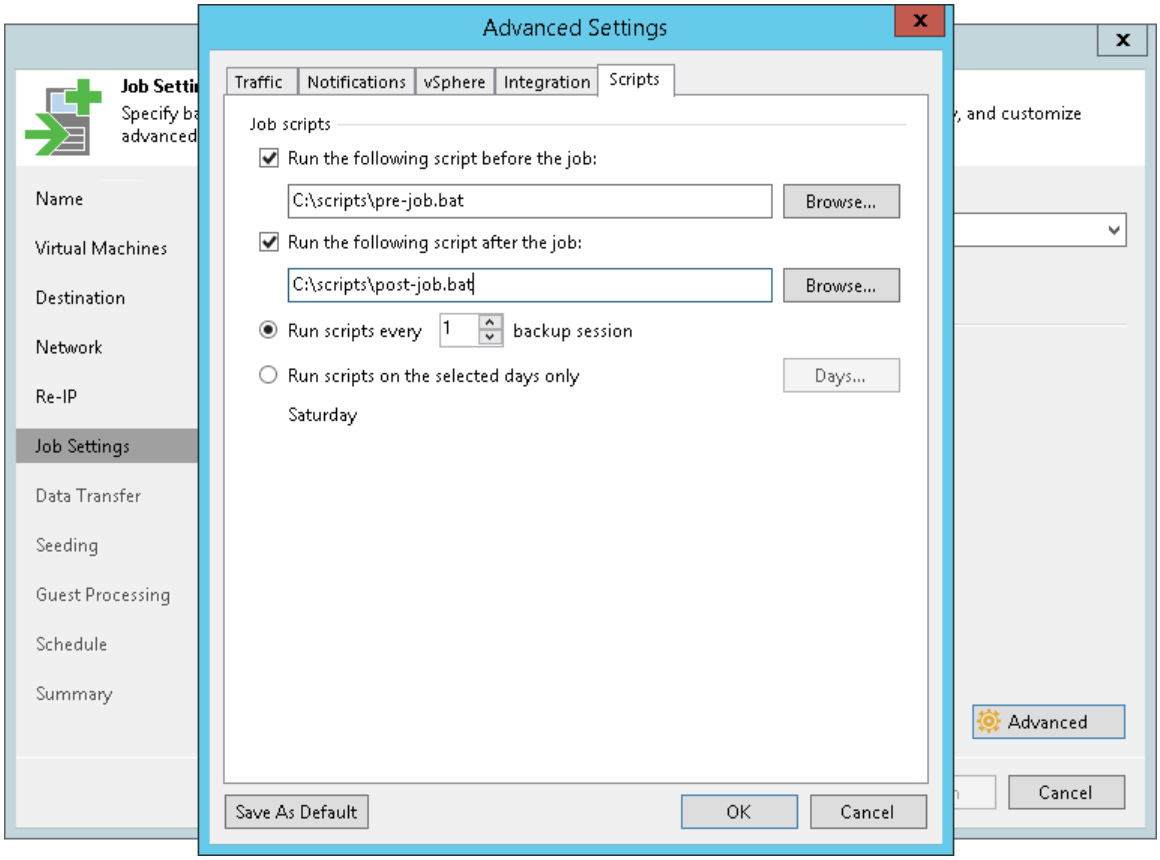
1. At the **Job Settings** step of the wizard, click **Advanced**.
2. In the **Advanced Settings** window, click the **Scripts** tab.

3. If you want to execute custom scripts before or after the replication job, select the **Run the following script before the job** or **Run the following script after the job** check boxes. Click **Browse** to choose executable files from a local folder on the backup server. The scripts will be executed on the backup server.
4. Configure when to execute pre- and post-replication scripts:
   - To execute scripts after a number of job sessions, select the **Run scripts every... backup session** option and then specify the number of the replication job sessions.
   - To execute scripts on specific week days, select the **Run scripts on selected days only** option. Click **Days** and specify week days on which scripts must be executed.

NOTE

Mind the following:

- Custom scripts that you define in the advanced job settings relate to the job itself, not the VM quiescence process. To add pre-freeze and post-thaw scripts for VM image quiescence, use the **Guest Processing** step of the wizard.
- If you select the **Run scripts on the selected days only** option, Veeam Backup & Replication executes scripts only once on each selected day — when the job runs for the first time. During subsequent job runs, scripts are not executed.
- To run the script, Veeam Backup & Replication uses the Service Account under which the Veeam Backup Service is running.



Step 12. Specify Data Transfer Settings

At the **Data Transfer** step of the wizard, select backup infrastructure components that must be used for the replication process and choose a path for VM data transfer:

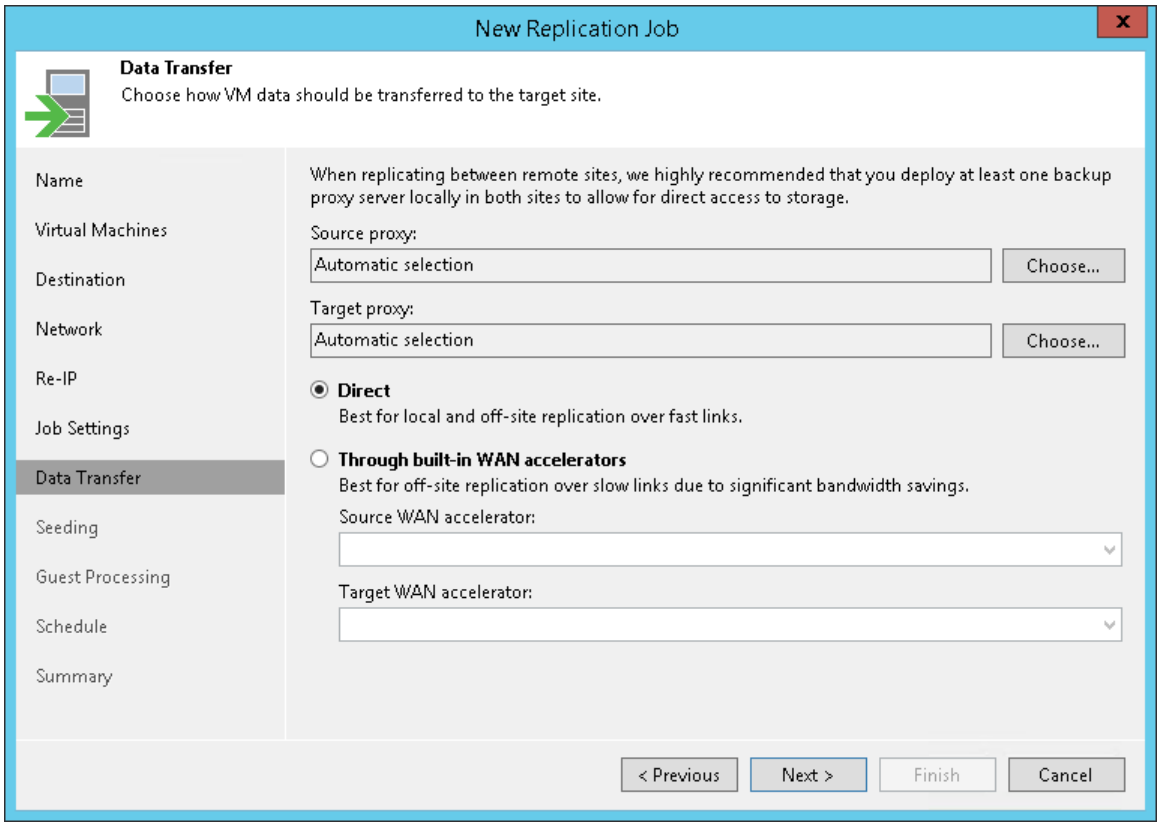
1. Specify which backup proxies you want to use:
   - If you want Veeam Backup & Replication to select proxies automatically, leave **Automatic selection** in the **Source proxy** and **Target proxy** fields.

     Veeam Backup & Replication will assign backup proxies for VM processing one by one. Before processing a new VM from the list, Veeam Backup & Replication will check available backup proxies. If more than one backup proxy is available, Veeam Backup & Replication analyzes transport modes that the backup proxies can use and the current workload on the backup proxies to select the most appropriate backup proxy for VM processing.
   - If you want to select backup proxies manually, do the following:
     i. Click **Choose** next to the **Source proxy** field if you want to select VMware CDP proxies in the production site, or next to the **Target proxy** field if you want to select VMware CDP proxies in the disaster recovery site.
     ii. In the **Backup Proxy** window, click **Use the selected backup proxy servers only**. Select proxies that you want to use and click **OK**.

> **NOTE**
>
> Mind the following:
>
> - If you plan to replicate VM data within one site, the same backup proxy can act as the source and target backup proxy. For off-site replication, you must deploy at least one backup proxy in each site to establish a stable connection for VM data transfer across sites.
> - We recommend you to select at least two backup proxies to ensure that the job will be performed if one of backup proxies fails or loses its connectivity to the source datastore.

2. Select a path for VM data transfer:
   - To transport VM data directly using backup proxies to the target datastore, select **Direct**.
   - To transport VM data using WAN accelerators, select **Through built-in WAN accelerators**. From the **Source WAN accelerator** list, select the WAN accelerator configured in the source site. From the **Target WAN accelerator** list, select the WAN accelerator configured in the target site.

> **NOTE**
>
> You should not assign one source WAN accelerator to several replication jobs that you plan to run simultaneously. The source WAN accelerator requires a lot of CPU and RAM resources, and does not process multiple replication tasks in parallel. As an alternative, you can create one replication job for all VMs you plan to process over one source WAN accelerator.
>
> The target WAN accelerator, however, can be assigned to several replication jobs. For more information, see Adding WAN Accelerators.



Step 13. Define Seeding and Mapping Settings

The **Seeding** step is available if you have selected the **Replica seeding** check box at the Name step of the wizard.

At the **Seeding** step of the wizard, configure replica seeding and mapping. Seeding and mapping help reduce the amount of traffic sent during the initial replica synchronization. For more information on when to use seeding and mapping, see Replica Seeding and Mapping.

If you use replica seeding or mapping, make sure that you select correct backup infrastructure components for the job: source-side backup repository for metadata and backup proxies. It is recommended that you explicitly assign backup proxies in the production site and disaster recovery (DR) site. For more information, see Specify Data Transfer Settings.

> **IMPORTANT**
>
> If the **Replica seeding** check box is enabled in a replication job, all VMs in the job must be covered with seeding or mapping. If a VM is neither has a seed, nor is mapped to an existing VM, it will be skipped from processing.

Configuring Replica Seeding

To configure replica seeding:

1. Make sure that you have backups of replicated VMs in a backup repository in the DR site. If you do not have the backups, create them as described in the Creating Replica Seeds section.

   > **IMPORTANT**
   >
   > Mind the following:
   >
   > - Backups must be created by Veeam Backup & Replication.
   > - Backups must not reside in a scale-out backup repository.

2. In the **Initial seeding** section, select the **Get seed from the following backup repository** check box.
3. From the list of available backup repositories, select the repository where your replica seeds are stored.

> **NOTE**
>
> If a VM has a seed and is mapped to an existing replica, replication will be performed using replica mapping because mapping has a higher priority.

Configuring Replica Mapping

To configure replica mapping:

1. Select the **Map replicas to existing VMs** check box.
2. If you want Veeam Backup & Replication to scan the DR site to detect existing copies of VMs that you plan to replicate, click **Detect**.

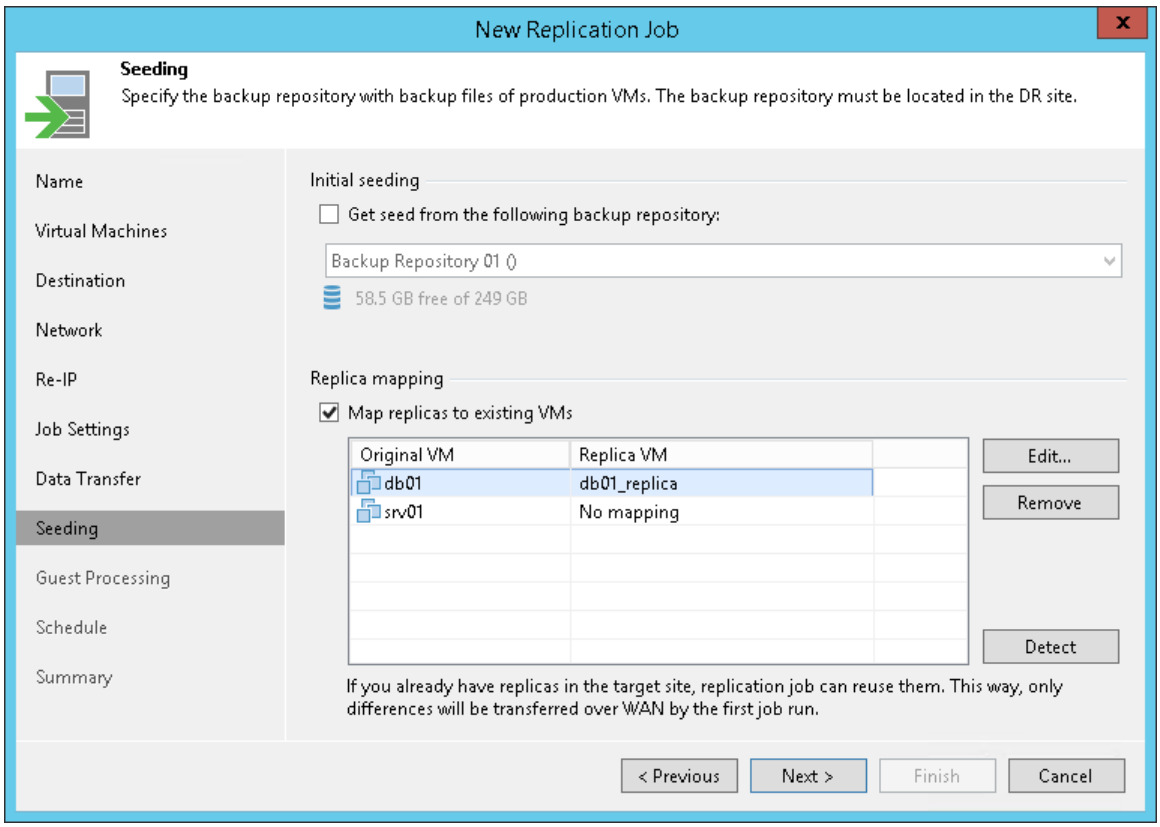
   If any matches are found, Veeam Backup & Replication will populate the mapping table. If Veeam Backup & Replication does not find a match, you can map a VM to its copy manually.

3. If you want to map a VM manually, select an original VM from the list, click **Edit** and select the copy of this VM on the target host in the DR site.

If there is no existing VM replica in the DR site, you can restore a VM from the backup and map it to the original VM.

To remove a mapping association, select a VM in the list and click **Remove**.

> **NOTE**
>
> The mapping list does not display VMs added to the list of exclusions. For more information, see Exclude Objects from Replication Job.



Step 14. Specify Guest Processing Settings

At the **Guest Processing** step of the wizard, enable guest OS processing that allows you to create transactionally consistent replicas, exclude files from processing, and use pre-freeze and post-thaw scripts. For more information on guest OS processing, see the Guest Processing section.

To enable guest OS processing:

1. Select **Enable application-aware processing**.
2. If you have added Microsoft Windows VMs to the job, specify which guest interaction proxy Veeam Backup & Replication can use to perform different guest processing tasks:
   - If you want Veeam Backup & Replication to select the guest interaction proxy automatically, leave **Automatic selection** on the **Guest interaction proxy** field.
   - If you want to explicitly specify which servers will perform the guest interaction proxy role, click **Choose**. In the **Guest Interaction Proxy** window, click **Prefer the following guest interaction proxy server**, and select the necessary proxies.

   For more information on the guest interaction proxy,requirements and limitations for it, see Guest Interaction Proxy.
3. From the **Guest OS credentials** list, select a user account that has enough permissions. For more information on the permissions and requirements for the user account, see Required Permissions for Guest Processing.

   Veeam Backup & Replication will use the account to connect to guest OSes and deploy the non-persistent runtime components or use (if necessary, deploy) persistent agent. For more information on guest agent and non-persistent components, see Non-Persistent Runtime Components and Persistent Agent Components.

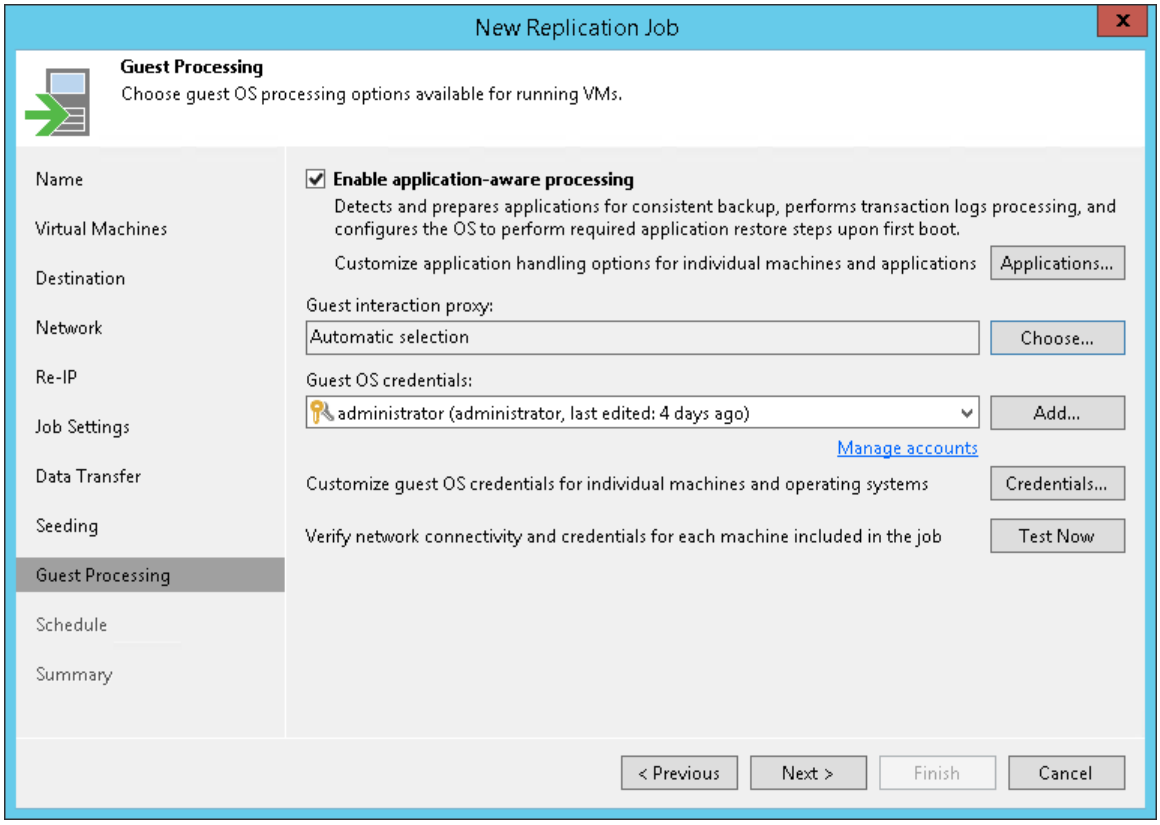
   If you have not set up credentials beforehand, click the **Manage accounts** link or click the **Add** button to add credentials. For more information on adding credentials, see the Credentials Manager section.

   > **NOTE**
   >
   > If you plan to use Kerberos authentication, check limitations and requirements listed in the Guest Processing section.
4. To specify other credentials for individual VMs, click **Credentials**. Then select the necessary VM and set user credentials for it.
5. To check whether Veeam Backup & Replication can connect to VMs using the specified guest OS credentials and can deploy the non-persistent runtime components or connect to persistent agent components on the guest OSes, click **Test Now**.

After you have enabled guest OS processing, configure the following guest processing options:

- Application-aware processing
- Microsoft SQL Server transaction log settings
- Oracle archived log settings
- VM guest OS file exclusion

- Pre-freeze and post-thaw scripts



Application-Aware Processing

If VMs run Microsoft Active Directory, Microsoft SQL Server, Microsoft SharePoint, Microsoft Exchange or Oracle, you can enable application-aware processing to create transactionally consistent replicas. The transactionally consistent replicas guarantee proper recovery of applications without data loss.

To configure application-aware processing:

1. Check that you have configured guest OS processing.
2. At the **Guest Processing** step of the wizard, click **Applications**.
3. In the **Application-Aware Processing Options** list, select a VM for which you want to configure application-aware processing, and click **Edit**.

   To define custom settings for a VM added as a part of a VM container, you must include the VM in the list as a standalone object. To do this, click **Add** and choose the necessary VM. Then select the VM in the list and define the necessary settings.
4. On the **General** tab, in the **Applications** section, specify the behavior scenario for application-aware processing:
   - Select **Require successful processing** if you want Veeam Backup & Replication to stop the replication process if any error occurs during application-aware processing.
   - Select **Try application processing, but ignore failures** if you want to continue the replication process even if an error occurs during application-aware processing. This option guarantees that the replication job will continue working. However, the resulting replica will be crash consistent, not transactionally consistent.
   - Select **Disable application processing** if you want to disable application-aware processing for the VM.

5. [For Microsoft Exchange, Microsoft SQL and Oracle] In the **Transaction logs** section, specify if Veeam Backup & Replication must process transaction logs or create copy-only replicas:

   a. Select **Process transaction logs with this job** if you want Veeam Backup & Replication to process transaction logs.

      [For Microsoft Exchange VMs] With this option selected, the non-persistent runtime components or persistent components running on the VM guest OS will wait for replication to complete successfully and then trigger truncation of transaction logs. If the replication job fails, the logs will remain untouched on the VM guest OS until the next start of the non-persistent runtime components or persistent components.

      [For Microsoft SQL Server VMs and Oracle VMs] You will need to configure how to process transaction logs. For more information on how to configure log handling, see Microsoft SQL Server Transaction Log Settings and Oracle Archived Log Settings.

   b. Select **Perform copy only** if you use another replication tool to perform guest level replication, and this tool maintains consistency of the database state. Veeam Backup & Replication will create a copy-only replica for the selected VM. The copy only replica preserves the chain of full and differential files and transaction logs on the VM. For more information, see Microsoft Docs.

6. In the **Persistent guest agent** section, select the **Use persistent guest agent** check box to use for application-aware processing persistent guest agents on each protected VM.

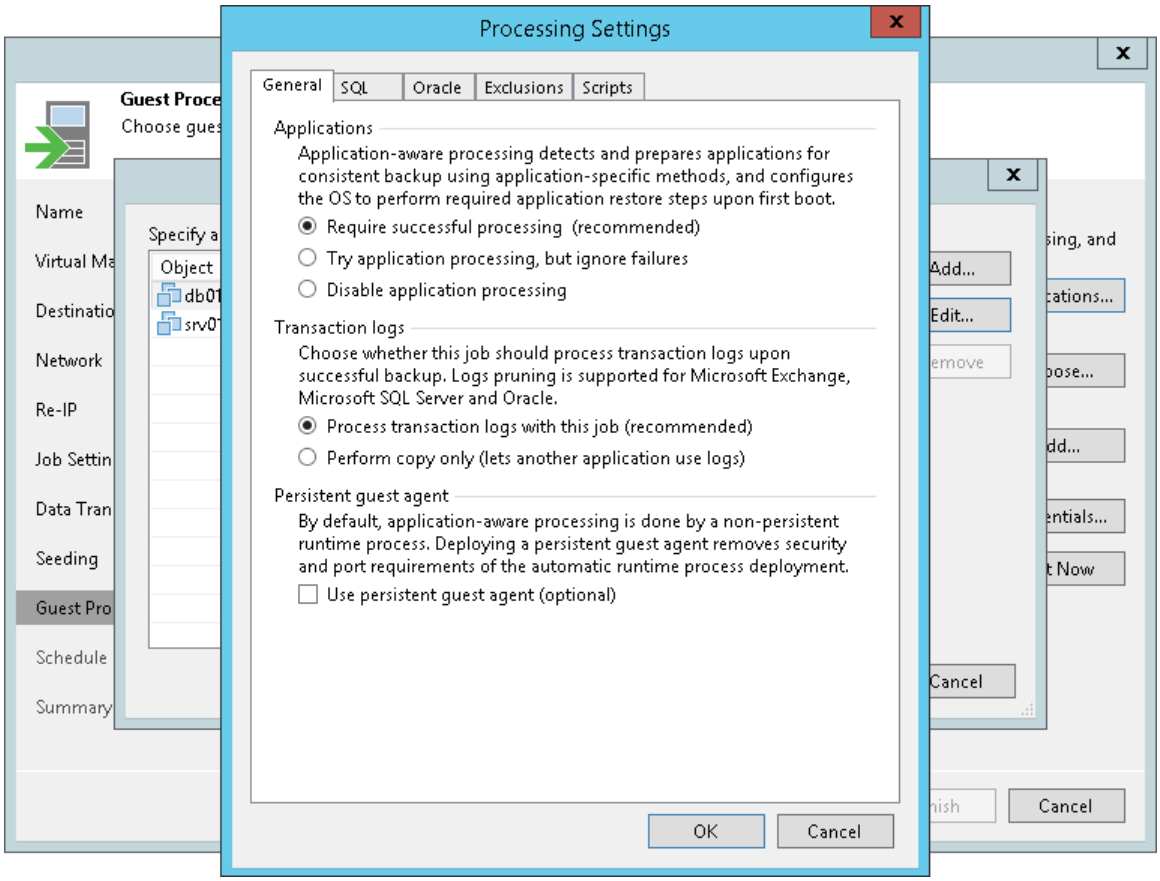
   By default, Veeam Backup & Replication uses non-persistent runtime components.
   Veeam Backup & Replication deploys runtime components on each protected VM when the replication job starts, and removes the runtime components as soon as the replication job finishes.

   For more information on guest agent and non-persistent components, see Non-Persistent Runtime Components and Persistent Agent Components.

To configure application-aware processing for multiple VMs at once, do the following:

1. At the **Guest Processing** step of the wizard, click **Applications.**
2. In the **Application-Aware Processing Options** list, select multiple VMs.
3. Click **Edit** > **General**.

4. Configure options described in the previous list.



Microsoft SQL Server Transaction Log Settings

The **SQL** tab applies to VMs that run Microsoft SQL Server and if you have selected **Process transaction logs with this job** when configuring application-aware processing.

To configure how Veeam Backup & Replication must process transaction logs of an Microsoft SQL Server:

1. At the **Guest Processing** step of the wizard, click **Applications**.
2. In the **Application-Aware Processing Options** list, select a VM that runs the Microsoft SQL Server.
3. Click **Edit**.
4. In the **Processing Settings** window, switch to the **SQL** tab and specify how transaction logs must be processed:
   - If you want Veeam Backup & Replication to trigger truncation of transaction logs only after the job completes successfully, select **Truncate logs**.

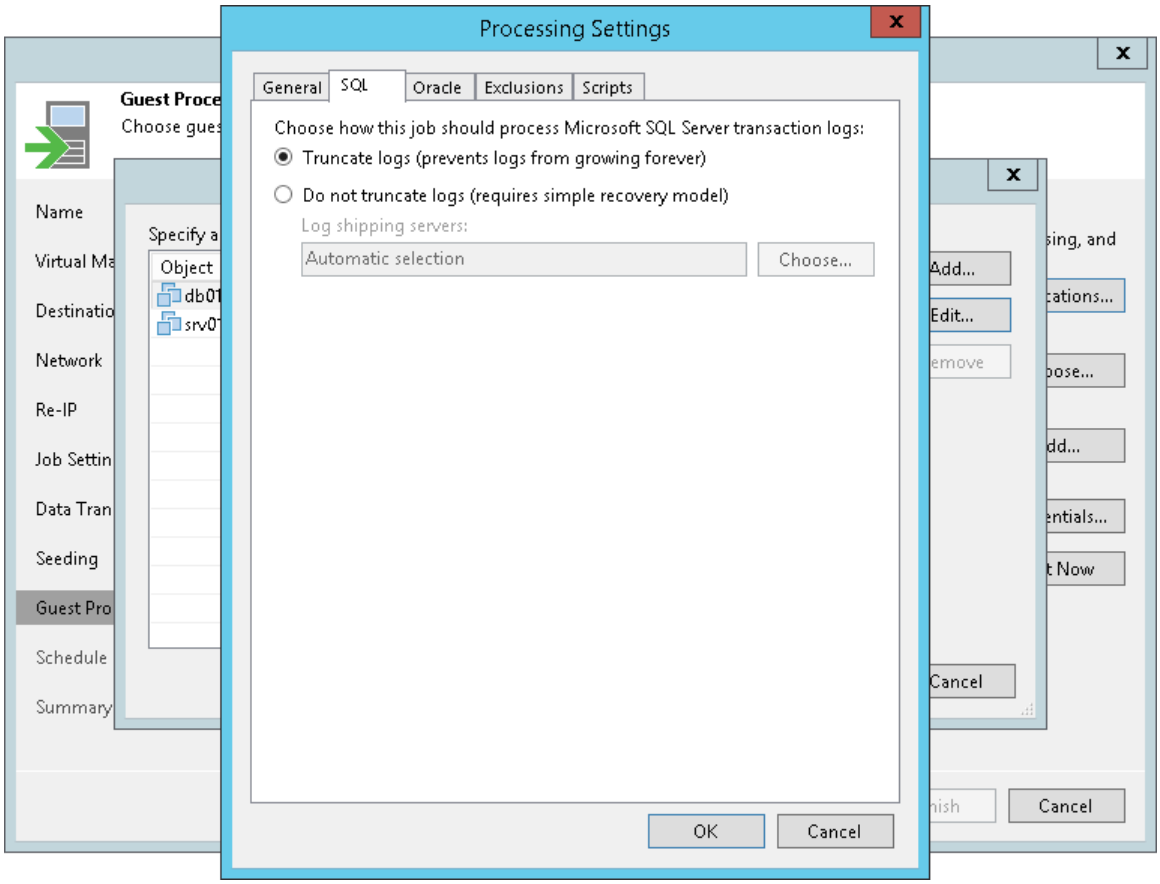
     In this case, the non-persistent runtime components or persistent components will wait for VM replication to complete and then trigger truncation of transaction logs. If the replication job fails, the logs will remain untouched on the VM guest OS until the next start of the non-persistent runtime components or persistent components.

- If you do not want Veeam Backup & Replication to truncate logs at all, select **Do not truncate logs**.

  This option is recommended if you are using another backup tool to perform VM guest-level backup or replication, and this tool maintains consistency of the database state. In such scenario, Veeam Backup & Replication will not trigger transaction log truncation. After you fail over to the necessary restore point of the VM replica, you will be able to apply transaction logs to get the database system to the necessary point in time between replication job sessions.

To configure transaction log handling for multiple VMs at once, do the following:

1. At the **Guest Processing** step of the wizard, click **Applications.**
2. In the **Application-Aware Processing Options** list, select multiple VMs.
3. Click **Edit** > **Microsoft SQL Server**.
4. Configure options described in the previous list.



Oracle Archived Log Settings

The **Oracle** tab applies to VMs that run Oracle and if you have selected **Process transaction logs with this job** when configuring application-aware processing.

To configure how Veeam Backup & Replication must process archive logs of an Oracle server:

1. At the **Guest Processing** step of the wizard, click **Applications**.
2. In the **Application-Aware Processing Options** list, select a VM that runs Oracle.
3. Click **Edit**.
4. In the **Processing Settings** window, switch to the **Oracle** tab.

5. From the **Specify Oracle account with SYSDBA privileges** drop-down list, select a user account that Veeam Backup & Replication will use to connect to the Oracle databases. The account must have SYSDBA rights on the Oracle database.

   You can select **Use guest credentials** from the list of user accounts. In this case, Veeam Backup & Replication will use the account specified at the **Guest Processing** step of the wizard to access the VM guest OS and connect to the Oracle databases. Note that this account must be a member of specific groups or granted specific privileges. For more information, see Required Permissions.

6. In the **Archived logs** section, specify how to process archived logs:

   - If you want to preserve archived logs on the VM guest OS, select **Do not delete archived logs**. When the replication job completes, the non-persistent runtime components or persistent components will not truncate transaction logs.

     It is recommended that you select this option for databases where the ARCHIVELOG mode is turned off. If the ARCHIVELOG mode is turned on, archived logs may grow large and consume all disk space.

   - If you want to delete archived logs older than <N> hours, select **Delete logs older than <N> hours** and specify the number of hours.

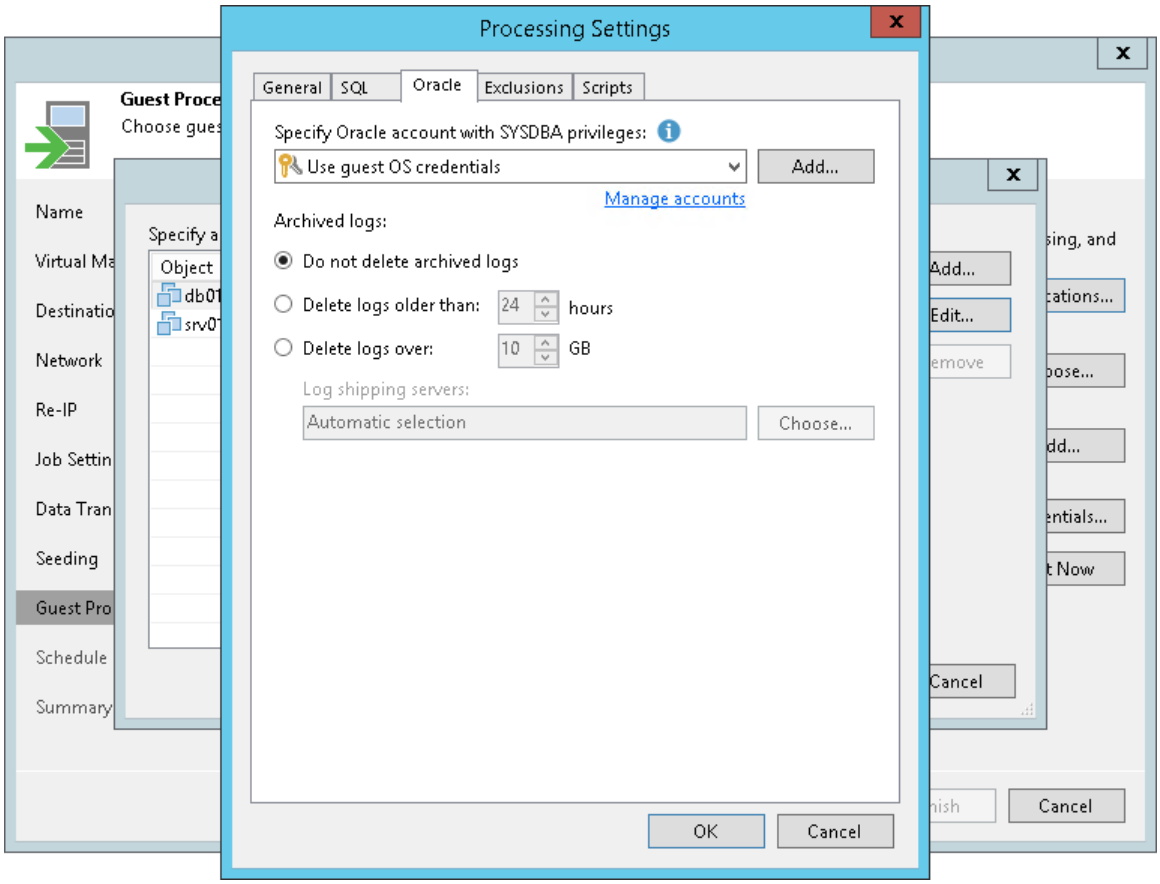
   - If you want to delete archived logs larger than <N> GB, select **Delete logs over <N> GB** and specify the size. The specified size refers to the log size of each database, not all databases on the selected Oracle server.

   The non-persistent runtime components or persistent components running on the VM guest OS will wait for the replication job to complete successfully and then trigger transaction logs truncation using Oracle Call Interface (OCI). If the job does not manage to replicate the Oracle VM, the logs will remain untouched on the VM guest OS until the next start of the non-persistent runtime components or persistent components.

To configure archive log handling for multiple VMs at once, do the following:

1. At the **Guest Processing** step of the wizard, click **Applications.**
2. In the **Application-Aware Processing Options** list, select multiple VMs.
3. Click **Edit** > **Oracle**.

4. Configure options described in the previous list.



VM Guest OS File Exclusion Settings

These settings apply only to Microsoft Windows VMs.

To exclude guest OS files and folders from being replicated:

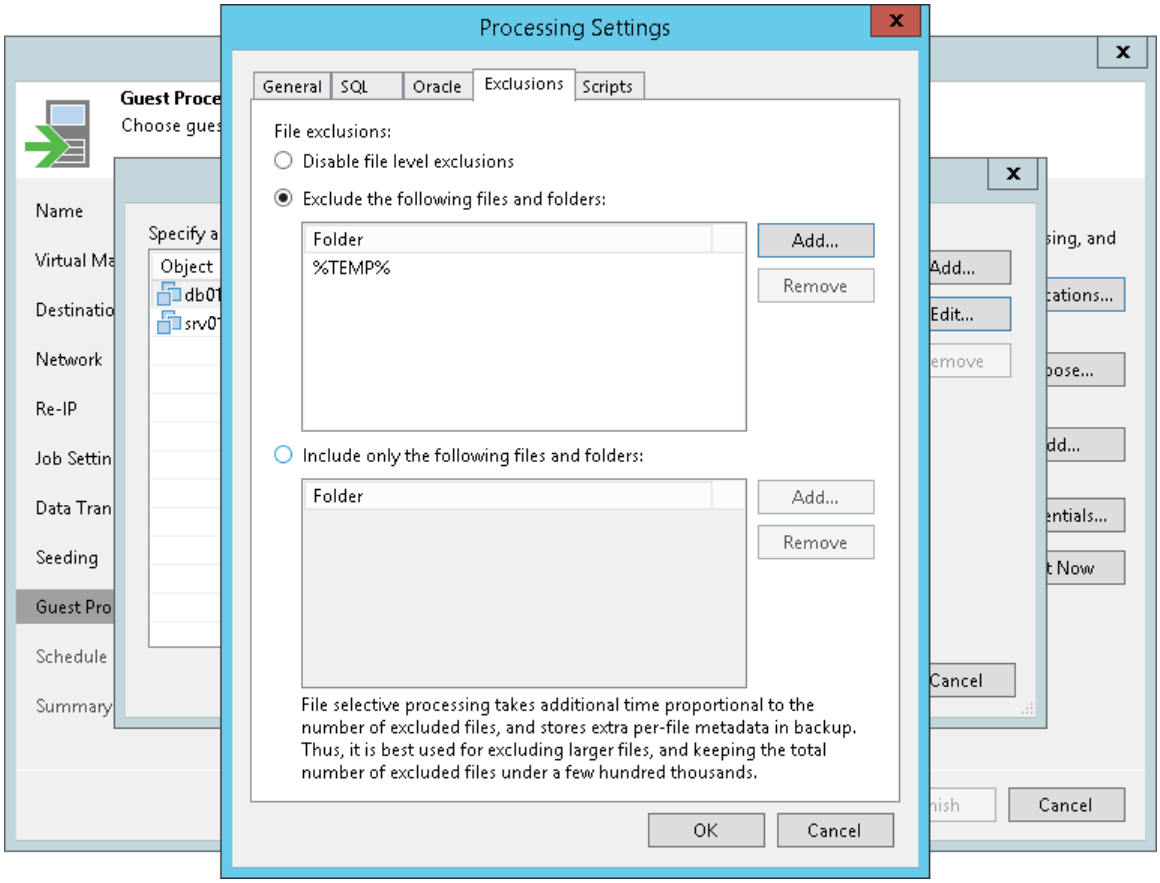
1. Check that you have enabled guest OS processing.
2. At the **Guest Processing** step of the wizard, click **Applications**.
3. In the **Application-Aware Processing Options** list, select a VM for which you want to exclude files, and click **Edit**.

   To define custom settings for a VM added as a part of a VM container, you must include the VM in the list as a standalone object. To do this, click **Add** and choose the necessary VM. Then select the VM in the list and define the necessary settings.
4. In the **Processing Settings** window, switch to the **Exclusions** tab and specify whether you want to exclude or include files and folders:
   - To remove individual files and folders from replicas, select **Exclude the following files and folders** and click **Add**.
   - To include only the specified files and folders in replicas, select **Include only the following files and folders** and click **Add**.
5. In the **Specify Folder** window, specify which files and folders you want to include or exclude. For the methods that you can use to specify the list of exclusions or inclusions, see VM Guest OS Files.

> **NOTE**
>
> When you select files to be included or excluded, mind requirements and limitations that are listed in the Requirements and Limitations for VM Guest OS File Exclusion section.

To configure file exclusions for multiple VMs at once, do the following:

1. At the **Guest Processing** step of the wizard, click **Applications.**
2. In the **Application-Aware Processing Options** list, select multiple VMs.
3. Click **Edit** > **File exclusions**.
4. Configure options described in the previous list.



Pre-Freeze and Post-Thaw Script Settings

If you plan to replicate VMs running applications that do not support VSS, you can instruct Veeam Backup & Replication to run custom pre-freeze and post-thaw scripts for these VMs. The pre-freeze script quiesces the VM file system and application data to bring the VM to a consistent state before Veeam Backup & Replication triggers a VM snapshot. After the VM snapshot is created, the post-thaw script brings the VM and applications to their initial state.

To specify pre-freeze and post-thaw scripts for the job:

1. Check that you have enabled guest OS processing.
2. At the **Guest Processing** step of the wizard, click **Applications**.

3. In the **Application-Aware Processing Options** list, select a VM for which you want to configure scripts, and click **Edit**.

   To define custom settings for a VM added as a part of a VM container, you must include the VM in the list as a standalone object. To do this, click **Add** and choose the necessary VM. Then select the VM in the list and define the necessary settings.

4. Click the **Scripts** tab.
5. In the **Script processing mode** section, select a scenario for script execution:
   - Select **Require successful script execution** if you want Veeam Backup & Replication to stop the replication process if scripts fail.
   - Select **Ignore script execution failures** if you want to continue the replication process even if script errors occur.
   - Select **Disable script execution** if you do not want to run scripts for the VM.
5. In the **Windows scripts** section, specify paths to scripts for Microsoft Windows VMs. For the list of supported script formats, see Pre-Freeze and Post-Thaw Scripts.
6. In the **Linux scripts** section, specify paths to scripts for Linux VMs. For the list of supported script formats, see Pre-Freeze and Post-Thaw Scripts.

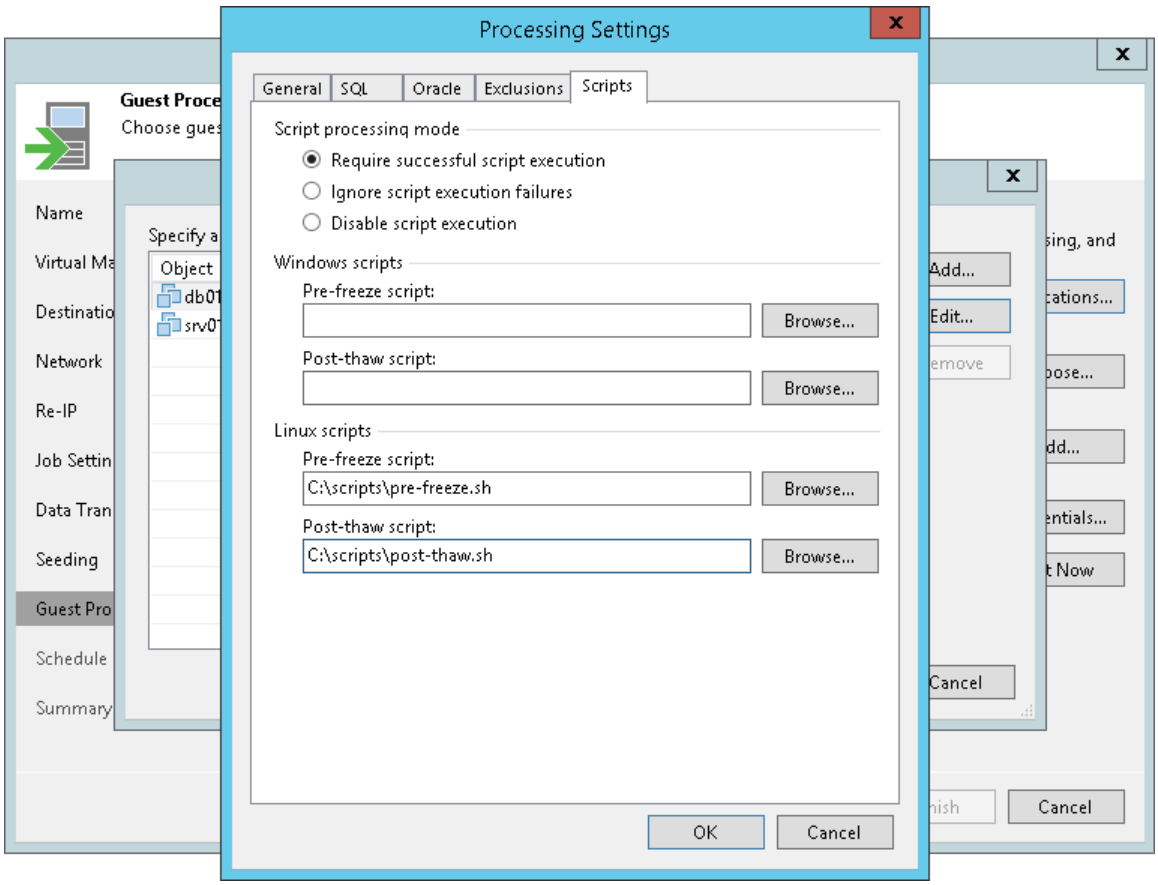
   If you have added to the replication job a VM container with Microsoft Windows and Linux VMs, you can select to execute both Microsoft Windows and Linux scripts. When the job starts, Veeam Backup & Replication will automatically determine which OS type is installed on the VM and use corresponding scripts for this VM.

To configure scripts for multiple VMs at once, do the following:

1. At the **Guest Processing** step of the wizard, click **Applications.**
2. In the **Application-Aware Processing Options** list, select multiple VMs.
3. Click **Edit** > **Scripts**.
4. Configure options described in the previous list.

TIP

Beside pre-freeze and post-thaw scripts for VM quiescence, you can instruct Veeam Backup & Replication to run custom scripts before the job starts and after the job completes. For more information, see Advanced Settings.



Step 15. Define Job Schedule

At the **Schedule** step of the wizard, select to run the replication job manually or schedule the job to run on a regular basis:

1. To run the replication job automatically, select the **Run the job automatically** check box. If you do not select this check box, you will have to start the job manually.
2. Define scheduling settings for the job:
   - To run the job at specific time daily, on weekdays or with specific periodicity, select **Daily at this time**. In the fields on the right of the check box, specify the time and required days.
   - To run the job once a month on specific days, select **Monthly at this time**. In the fields on the right of the check box, specify the necessary days.

     > **NOTE**
     >
     > When you configure the job schedule, keep in mind possible date and time changes (for example, related to daylight saving time transition).

   - To run the job repeatedly throughout a day with a set time interval, do the following:
     i. Select **Periodically every**. In the field on the right, select the necessary time unit: *Hours* or *Minutes.*
     ii. If you want to specify the permitted time window for the job, click **Schedule**. In the **Time Periods** window, specify the schedule.

     If you want to shift the schedule, specify the offset in the **Start time within an hour** field. For example, you schedule the prohibited hours from 08:00 AM to 10:00 AM, and set the offset value to 25. The schedule will be shifted forward, and the prohibited hours will be from 8:25 AM and to 10:25 AM.

     A repeatedly run job is started by the following rules:

     - Veeam Backup & Replication always starts counting defined intervals from 12:00 AM. For example, if you configure to run a job with a 4-hour interval, the job will start at 12:00 AM, 4:00 AM, 8:00 AM, 12:00 PM, 4:00 PM and so on.
     - If you define permitted hours for the job, after the denied interval is over, Veeam Backup & Replication will immediately start the job and then run the job by the defined schedule.

     For example, you have configured a job to run with a 2-hour interval and defined permitted hours from 9:00 AM to 5:00 PM. According to the rules above, the job will first run at 9:00 AM, when the denied period is over. After that, the job will run at 10:00 AM, 12:00 PM, 2:00 PM and 4:00 PM.

   - To run the job continuously, select the **Periodically every** option and choose **Continuously** from the list on the right.
   - To chain jobs, use the **After this job** field. For the first job in the chain, you must specify an automatic time schedule, otherwise the chained jobs will not be started. For the chained jobs, select the **After this job option** and choose the preceding job from the list. For more information on job chaining and recommendations for it, see Chained Jobs.

3. In the **Automatic retry** section, select the **Retry failed VMs processing** if Veeam Backup & Replication must attempt to run the job again for VMs whose processing failed for some reason. Enter the number of attempts to run the job and define time spans between them.

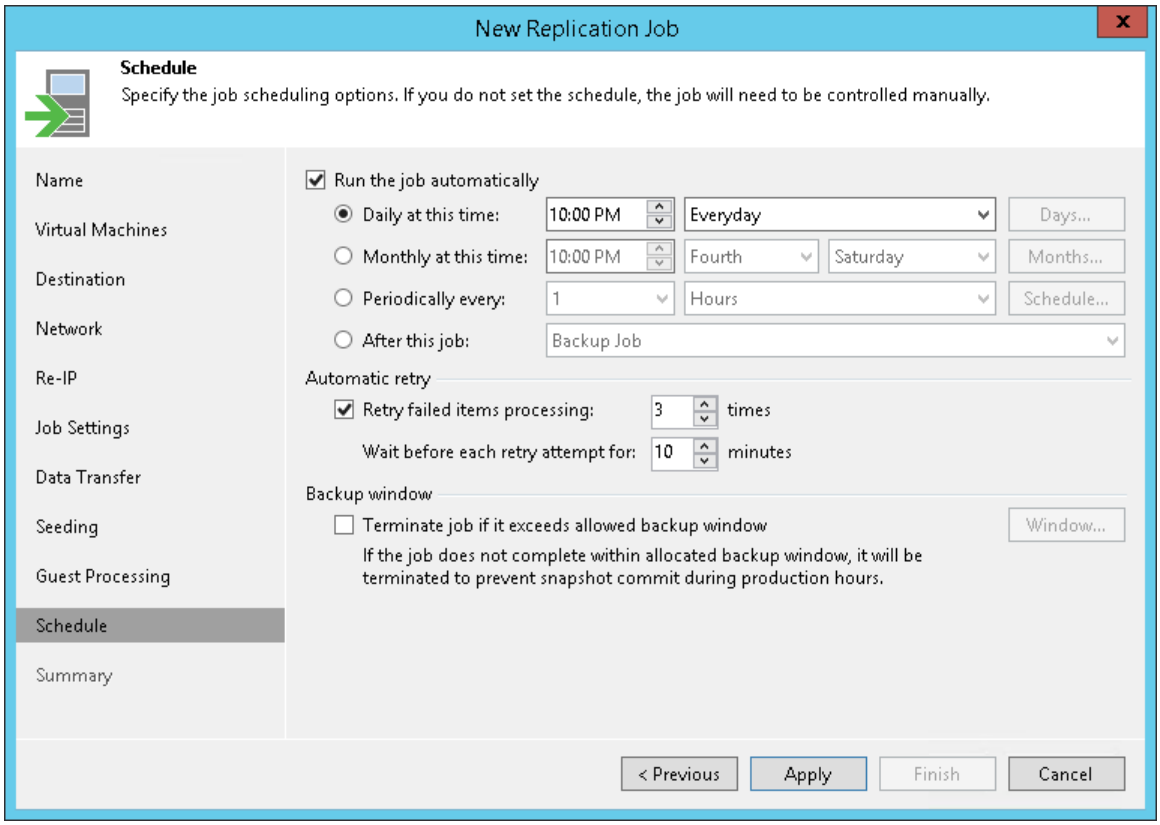
   If you select continuous schedule for the job, Veeam Backup & Replication will retry the job for the defined number of times without any time intervals between the job sessions.

4. In the **Backup window** section, specify a time interval within which the job must be completed. The backup window prevents the job from overlapping with production hours and ensures the job does not provide unwanted overhead on your production environment.

   To set up a backup window for the job:

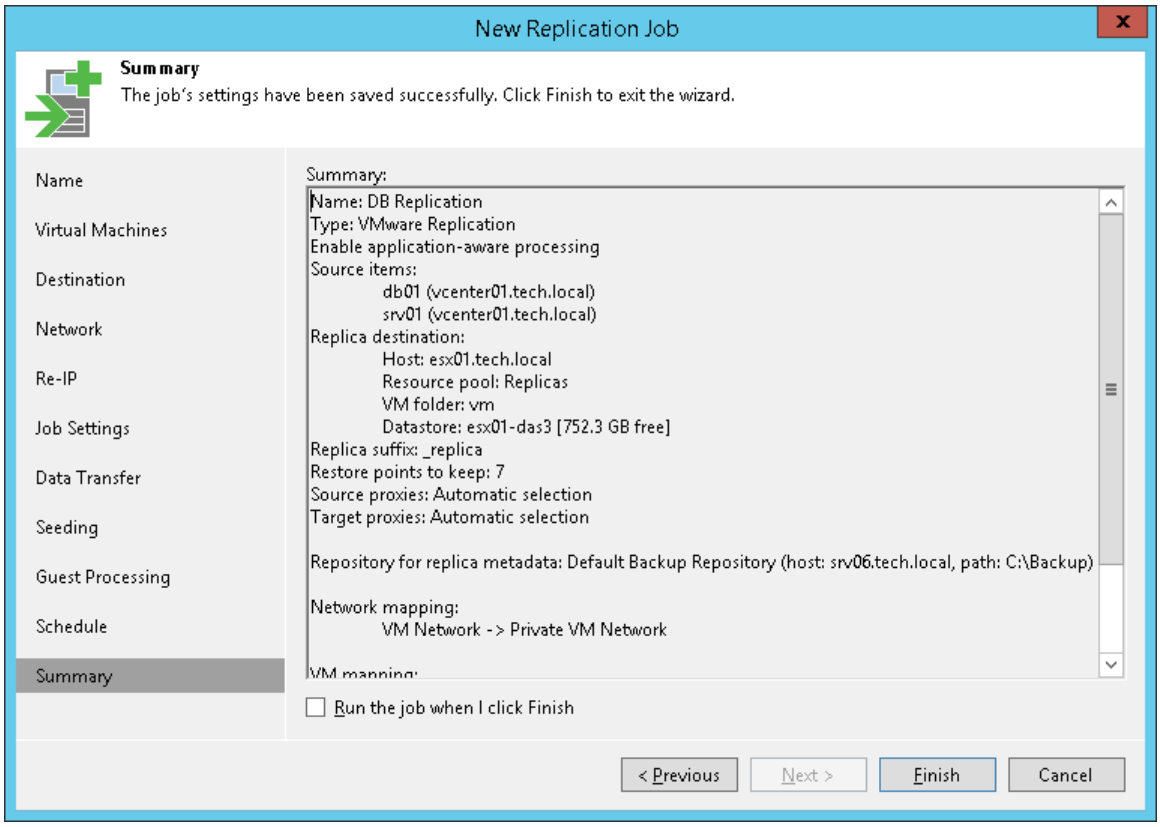
   a. Select the **Terminate job if it exceeds allowed backup window** check box and click **Window**.

   b. In the **Time Periods** window, define the allowed hours and prohibited hours for VM replication. If the job exceeds the allowed window, the job will be automatically terminated.



Step 16. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the replication job. If you want to start the job right after you close the wizard, select the **Run the job when I click Finish** check box, otherwise leave the check box unselected. Then click **Finish** to close the wizard.



Creating Replica Seeds

To use replica seeding in a replication job, you must have backups of replicated VMs in a backup repository in the disaster recovery (DR) site. These backups are known as replica seeds. For more information on seeding and when to use it, see Replica Seeding and Mapping.

If you do not have replica seeds in the DR site, do the following:

1. Create a backup of VMs that you plan to replicate as described in the Creating Backup Jobs section. As the target repository for this job, select a backup repository in the production site. Then run the job.

   If you already have backups containing the necessary VMs, there is no need to configure and run a new backup job. For seeding, you can use any existing backups created with Veeam Backup & Replication. The backup must include VBK and VBM files. If you have a full backup and a chain of forward increments, you can use VIB files together with the VBK and VBM files. In this case, Veeam Backup & Replication will restore VMs from the seed to the latest available restore point.

2. Copy the backup from the backup repository in the production site to a backup repository in the DR site.

   You can move the backup using a file copy job or any other appropriate method, for example, copy the backup to a removable storage device, ship the device to the DR site and copy backups to the backup repository in the DR site.

   If you do not have a backup repository in the DR site, you need to create the repository as described in Backup Repository.

   > **IMPORTANT**
   >
   > You cannot copy backups to a scale-out backup repository in the DR site.

3. After the backup is copied to the backup repository in the DR site, perform rescan of this backup repository as described in the Rescanning Backup Repositories section. Otherwise, Veeam Backup & Replication will not be able to detect the copied backup.

Managing Replicas

You can perform the following operations with replicas:

- View replica properties
- Rescanning Replicas
- Remove a replica from configuration
- Delete a replica from disks

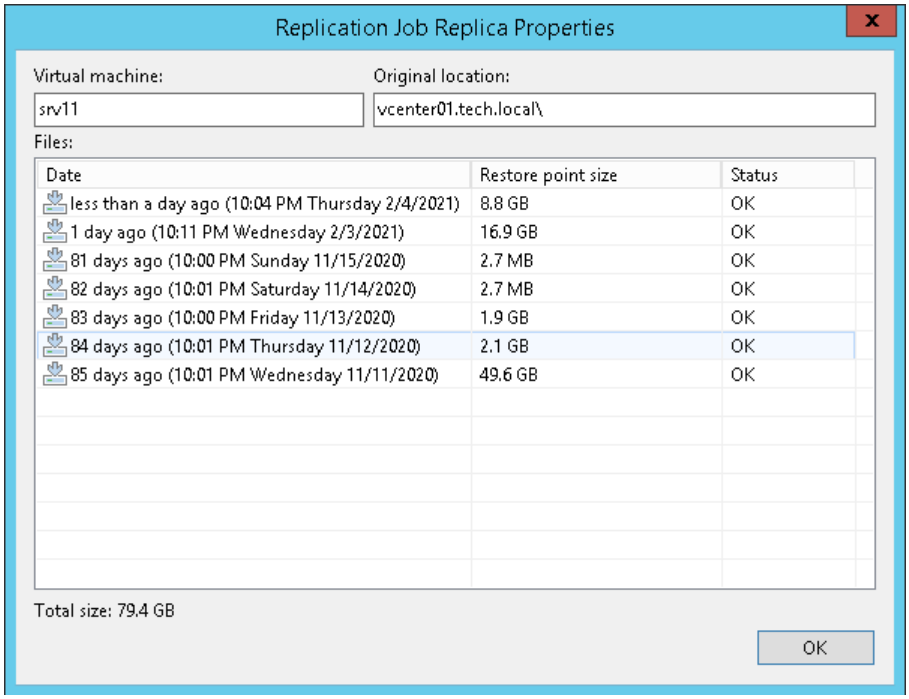
Viewing Replica Properties

You can view summary information about created replicas. The summary information provides the following data:

- Available restore points
- Date of restore points creation
- Data size and replica status

To view summary information for replicas:

1. Open the **Home** view.
2. In the inventory pane, select **Replicas**.
3. In the working area, right-click the replica and select **Properties**.



Rescanning Replicas

You may need to perform replica rescan in the following cases:

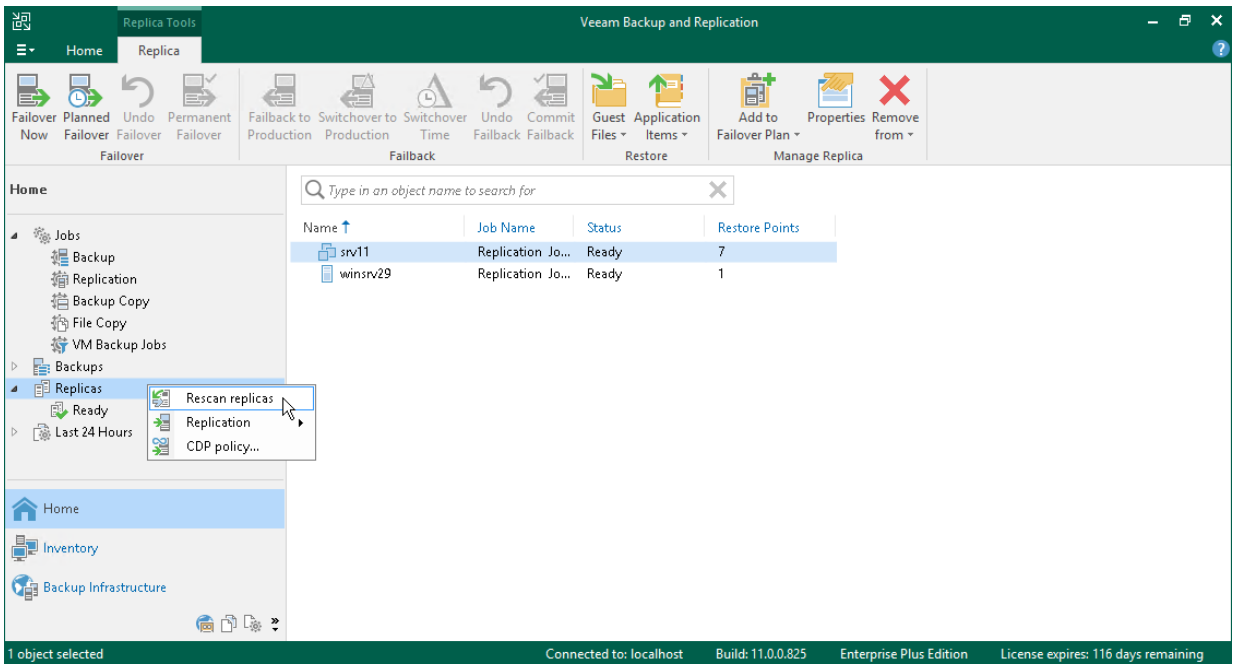
- After you delete restore points for one or more replicas. For details, see this Veeam KB article.
- After you restore the configuration database, and the session results show that some hosts used to register replicas were unavailable during the session.

  To check whether any errors occurred during the database restore session, open the **Home** view and select **System** in the inventory pane. In the working area, right-click the **Configuration Database Resynchronize** job and select **Statistics**.

During the rescan process, Veeam Backup & Replication gathers information on replicas that are currently available in backup repositories and updates the list of replicas in the configuration database.

To rescan replicas, do the following:

1. Open the **Home** view.
2. In the inventory pane, right-click the **Replicas** node and select **Rescan Replicas**.



Removing Replicas from Configuration

If you want to remove records about replicas from the Veeam Backup & Replication console and configuration database, you can use the **Remove from configuration** operation.

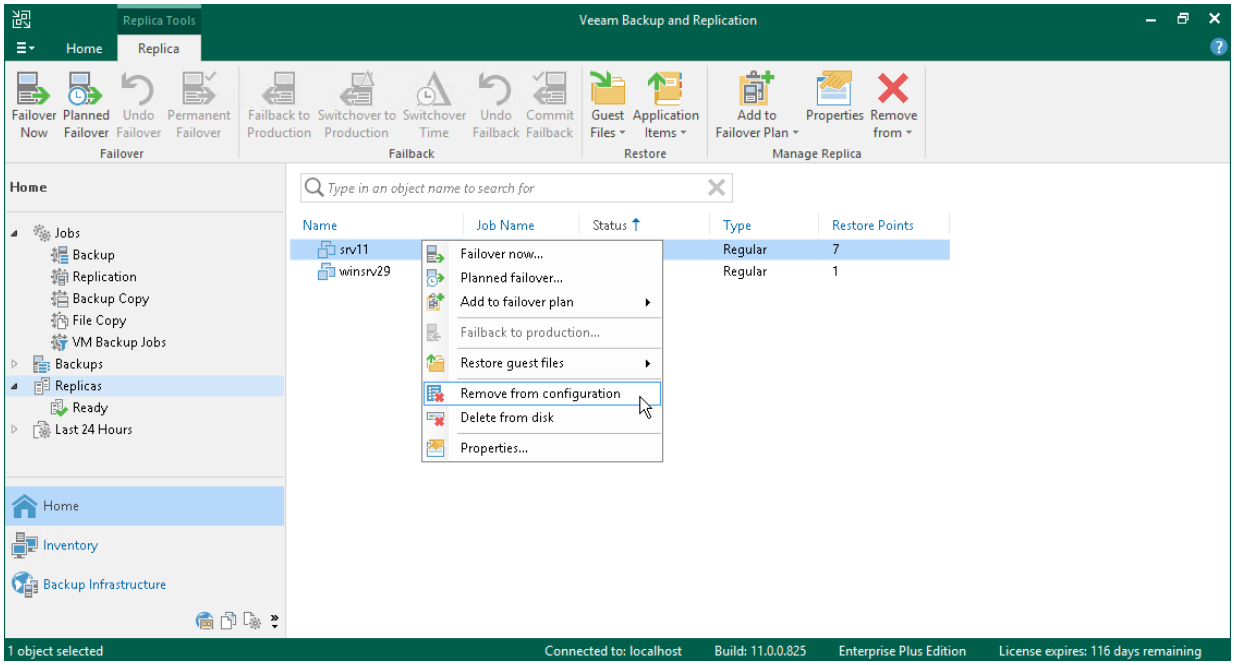
Replicated VMs remain on target hosts. If necessary, you can start them manually after the **Remove from configuration** operation is performed.

Mind the following:

- The **Remove from configuration** operation can be performed only for VM replicas in the *Ready* state. If the VM replica is in the *Failover* or *Failback* state, this option is disabled.
- When you perform the **Remove from configuration** operation for a VM that is replicated as a standalone object, Veeam Backup & Replication removes this VM from the initial replication job. When you perform the **Remove from configuration** operation for a VM that is replicated as part of a VM container (host, cluster, folder, resource pool, VirtualApp, datastore or tag), Veeam Backup & Replication adds this VM to the list of exclusions in the initial replication job. For more information, see Exclude Objects from Replication Job.

To remove records about VM replicas from the Veeam Backup & Replication console and configuration database:

1. Open the **Home** view.
2. In the inventory pane, select **Replicas**.
3. In the working area, select the replica and click **Remove from** > **Configuration** on the ribbon. You can also right-click the replica and select **Remove from configuration**.



Deleting Replicas from Disk

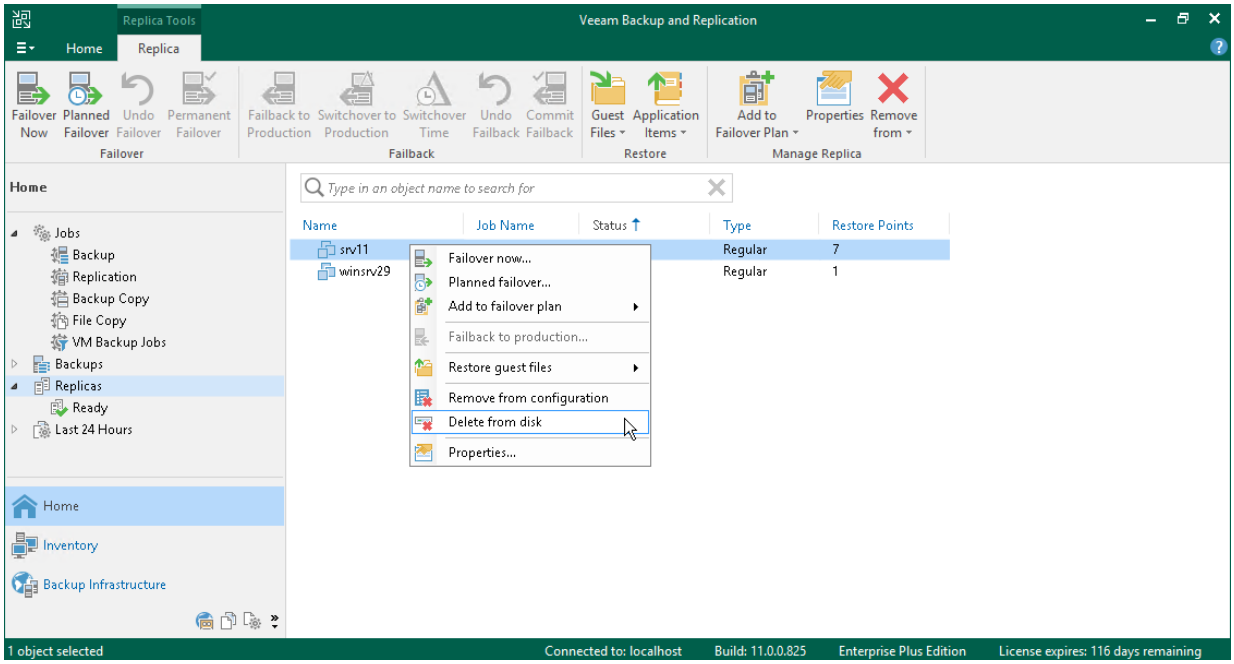
If you want to delete records about replicas from the Veeam Backup & Replication console and configuration database and, additionally, delete replica files from the destination storage, you can use the **Delete from disk** operation.

Mind the following:

- Do not delete replica files from the destination storage manually. Use the **Delete from disk** option instead. If you delete replica files manually, subsequent replication job sessions will fail.
- The **Delete from disk** operation can be performed only for VM replicas in the *Ready* state. If the VM replica is in the *Failover* or *Failback* state, this option is disabled.

To delete replica files from disk:

1. Open the **Home** view.
2. In the inventory pane, select **Replicas**.
3. In the working area, select the VM replica and click **Remove from** > **Disk** on the ribbon. You can also right-click the VM replica and select **Delete from disk**.



Failover and Failback for Replication

Failover and failback operations help you ensure that your business will function even if a disaster strikes your production site. Failover is a process of switching from the original VM on the source host to its VM replica on a host in the disaster recovery site. Failback is a process of returning from the VM replica to the original VM.

Veeam Backup & Replication provides the following failover and failback operations:

- **Perform failover**

  When you perform failover, you shift all processes from the original VM in the production site to the VM replica in the disaster recovery site. Failover is an intermediate step that needs to be finalized: you can undo failover, perform permanent failover or perform failback.

  For more information on how failover is performed, see Failover.

- **Perform planned failover**

  When you perform planned failover, you shift all processes from the original VM to its replica. Planned failover is helpful when you know that the original VM is about to go offline, for example, you plan to perform datacenter maintenance, and you want to proactively switch the workload to the replica. The procedure is designed to transfer the current workload, that is why it does not suggest to select a restore point.

  For more information on how planned failover is performed, see Planned Failover.

- **Create failover plan**

  When you create a failover plan, you define the order in which Veeam Backup & Replication must perform failover for VMs, and an interval of time for which Veeam Backup & Replication must wait before starting the failover operation for the next VM in the list.

  For more information on failover plans, see Failover Plan.

- **Perform permanent failover**

  When you perform permanent failover, you permanently switch from the original VM to a VM replica and use this replica as the original VM. You can use this scenario if the original VM and VM replica are located in the same site and are nearly equal in terms of resources. Otherwise, perform failback.

  For more information on how permanent failover is performed, see Permanent Failover.

- **Undo failover**

  When you undo failover, you switch back to the original VM and discard all changes made to the VM replica while it was running. You can use the undo failover scenario if you have failed over to the VM replica for testing and troubleshooting purposes, and you do not need to synchronize the original VM state with the current state of the replica.

  For more information on how failover undo is performed, see Failover Undo.

- **Perform failback**

  When you perform failback, you switch back to the original VM and send to the original VM all changes that took place while the VM replica was running. If the source host is not available, you can recover a VM with the same configuration as the original VM and switch to it. For more information on how failback is performed, see Failback.

When you perform failback, changes are only sent to the original/recovered VM but not published. You must test whether the original/recovered VM works with these changes. Depending on the test results, you can do the following:

- **Commit failback**. When you commit failback, you confirm that the original/recovered VM works as expected and you want to get back to it.

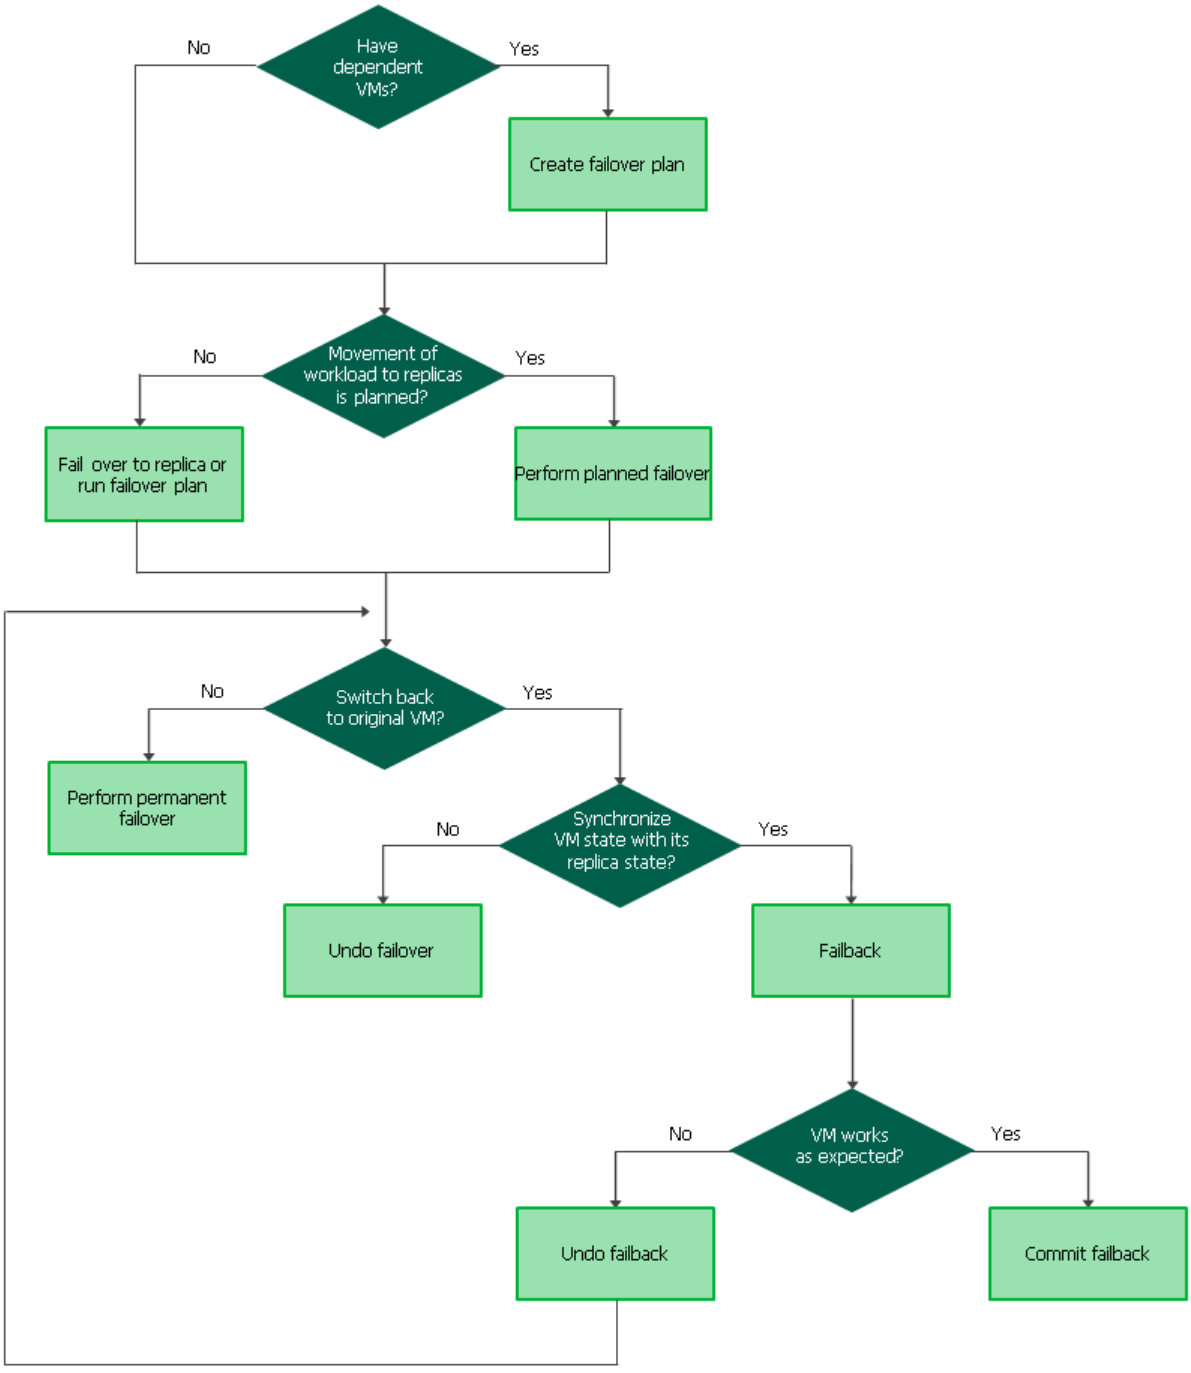
  For more information on how failback commit is performed, see Failback Commit.

- **Undo failback**. When you undo failback, you confirm that the original/recovered VM is not working as expected and you want to get back to the VM replica.

  For more information on how failback undo is performed, see Failback Undo.

Veeam Backup & Replication supports failover and failback operations for one VM and for several VMs. In case one or several hosts fail, you can use batch processing to restore operations with minimum downtime.

The following scheme can help you decide at which moment which operations are preferable.



Failover Plans

A failover plan helps you perform failover for dependent VMs one by one, as a group. To do this automatically, you can prepare a failover plan.

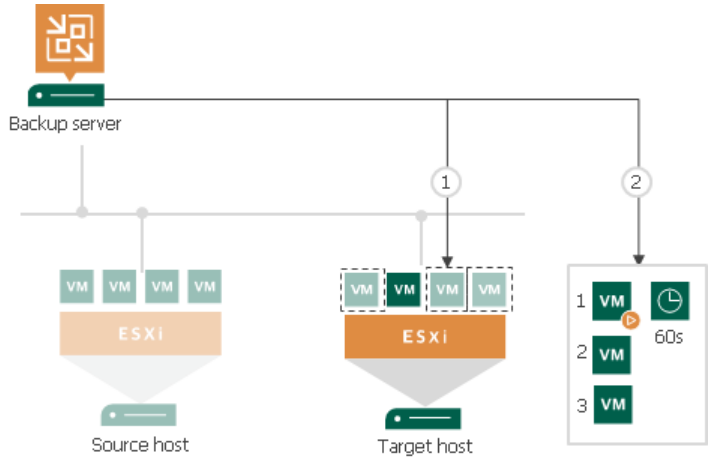
In the failover plan, you define the order in which VMs must be processed and an interval of time for which Veeam Backup & Replication must wait before starting the failover operation for the next VM in the list. The failover plan helps ensure that some VMs, such as a DNS server, are already running at the time the dependent VMs start.

> **IMPORTANT**
>
> The failover plan must be created in advance.

In case the primary VM group goes offline, you can start the failover plan manually. When you start the plan, you can choose to fail over to the latest state or select the point in time to which VM replicas must be started. Veeam Backup & Replication will look for the closest restore points to this point in time and use them to start VM replicas. The source VMs will not be powered off.

The failover process is performed in the following way:

1. For each VM, Veeam Backup & Replication detects its replica. The VMs whose replicas are already in the *Failover* or *Failback* state are skipped from processing.
2. The replica VMs are started in the order they appear in the failover plan within the set time intervals.



Limitations for Failover Plans

The maximum number of VMs that can be started simultaneously when you run a failover plan is 10. If you have added more VMs to the failover plan and scheduled them to start simultaneously, Veeam Backup & Replication will wait for the first VMs in the list to fail over and then start the failover operation for subsequent VMs. This limitation helps reduce the workload on the production infrastructure and backup server.

For example, if you have added 14 VMs to the failover plan and scheduled them to start at the same time, Veeam Backup & Replication will start the failover operation for the first 10 VMs in the list. After the 1st VM is processed, Veeam Backup & Replication will start the failover operation for the 11th VM in the list, then for the 12th VM and so on.

Finalizing Failover Plans

Failover is a temporary intermediate step that needs to be finalized. You can finalize group failover in the same ways as regular failover: undo failover, perform permanent failover or failback.

When you perform failback or permanent failover, you need to process each VM individually. For more information, see Performing Failback and Performing Permanent Failover. When you undo failover, you can process the whole group. For more information, see Undoing Failover by Failover Plans.

Creating Failover Plans

To create a failover plan, use the **New Failover Plan** wizard.

Before You Begin

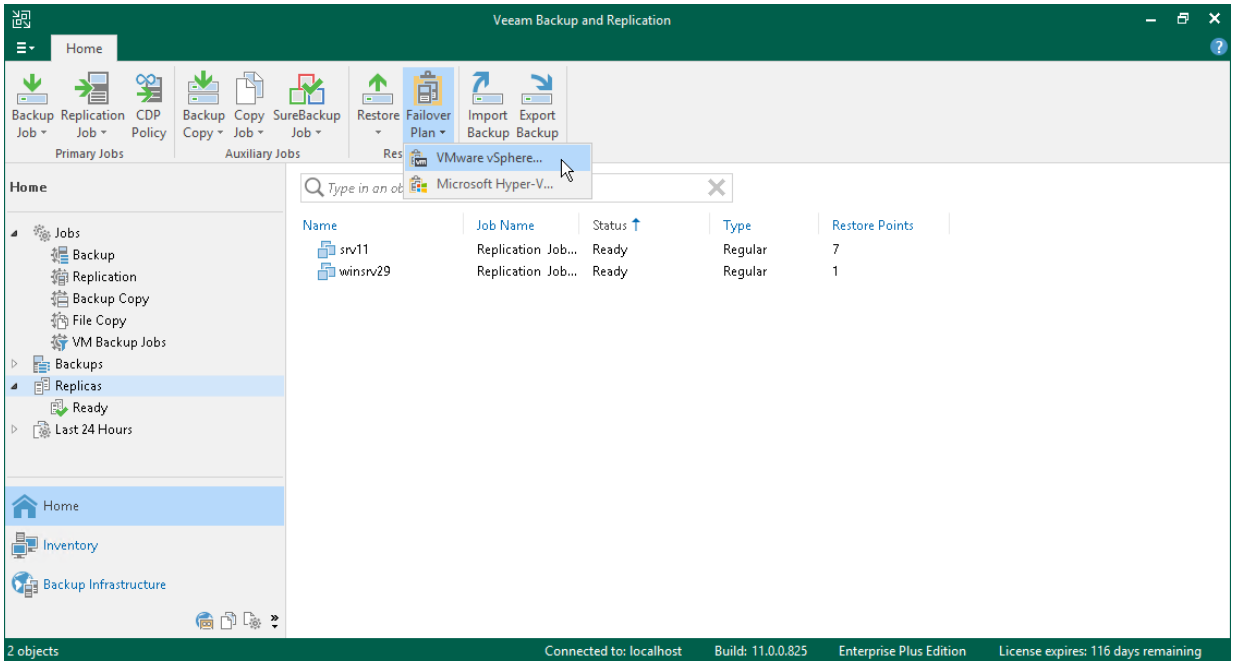
Before you create a failover plan, check the following prerequisites:

- [If you plan to select VMs from replication jobs] VMs that you plan to include in the failover plan must be successfully replicated at least once.
- [If you plan to select VMs from replication jobs] VM replicas must be in the *Ready* state.
- If you plan to use pre-failover and/or post-failover scripts for the failover plan, you must create scripts before you configure the failover plan.

Step 1. Launch New Failover Plan Wizard

To launch the **New Failover Plan** wizard, do one of the following:

- On the **Home** tab, click **Failover Plan** and select **VMware vSphere**.
- Open the **Home** view. In the working area, select VMs that you want to add to a failover plan. On the ribbon, click **Add to Failover Plan** > **New failover plan** if you want to create a new failover plan, or **Add to Failover Plan** > *<Plan Name>* if you want to add VMs to an existing failover plan.
- Open the **Home** view. In the working area, select VMs that you want to add to a failover plan and right-click one of them. Select **Add to failover plan** > **New failover plan** if you want to create a new failover plan, or **Add to failover plan** > *<Plan Name>* if you want to add VMs to an existing failover plan.

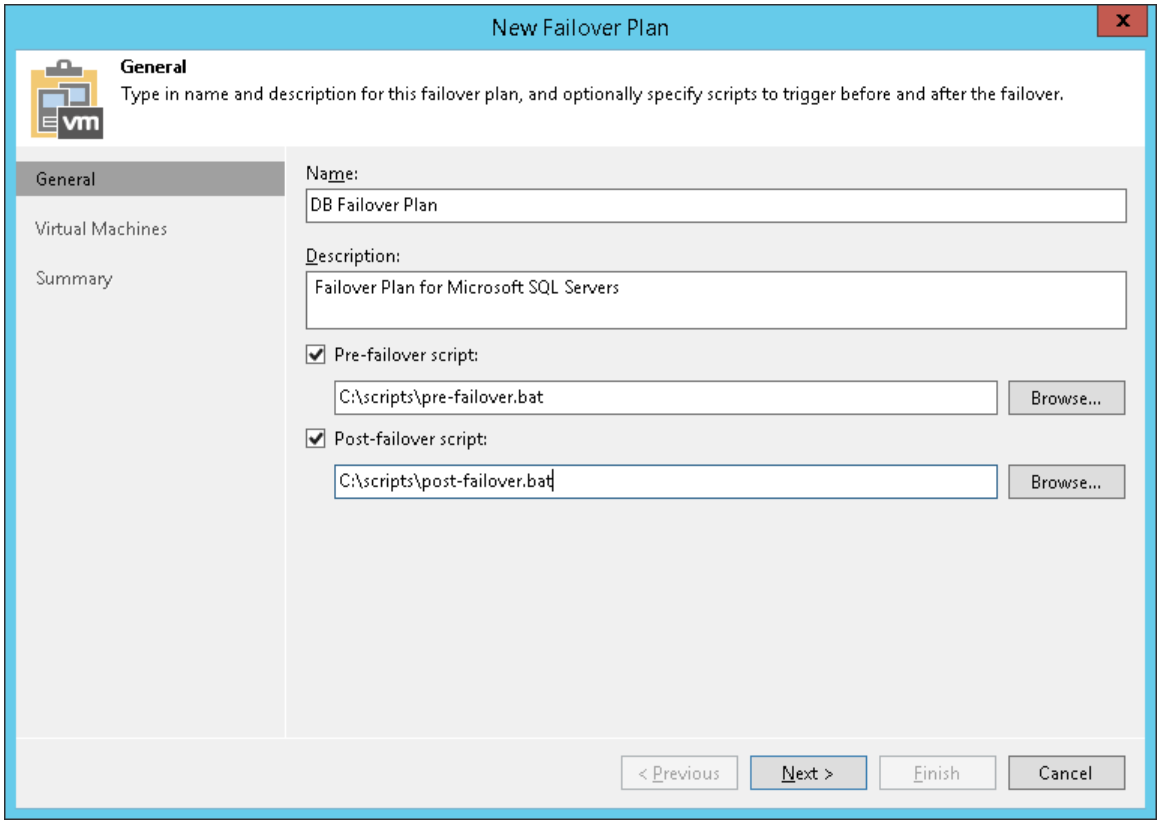


Step 2. Specify Failover Plan Name and Description

At the **General** step of the wizard, specify a name and description for the failover plan.

If you want to execute custom scripts before or after the failover plan, select the **Pre-failover script** and **Post-failover** script check boxes and click **Browse** to choose executable files. For example, you may want stop some applications on production VMs before the failover plan starts or send an email to backup administrators after the failover plan finishes.

The scripts will be executed on the backup server. Veeam Backup & Replication supports the script files in the following formats: BAT, CMD, EXE and PS1.



Step 3. Select VMs

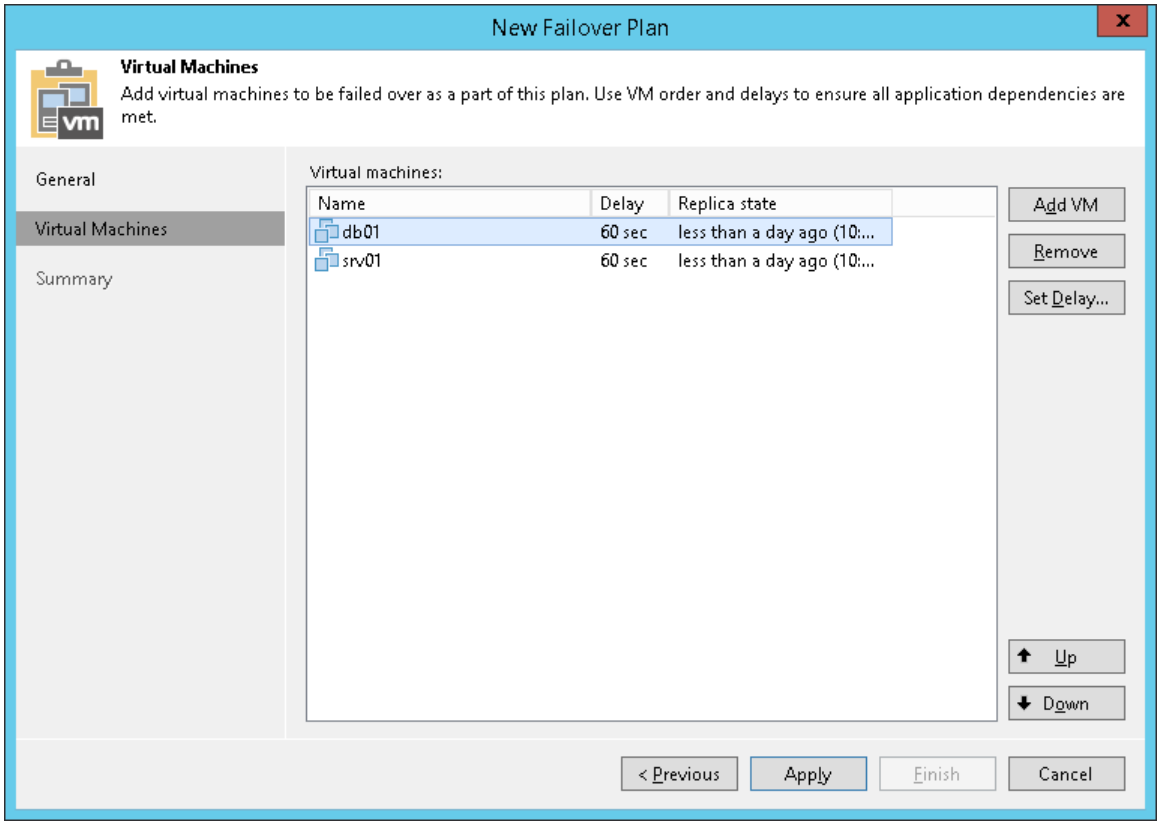
At the **Virtual Machines** step of the wizard, select VMs that you want to add to the failover plan. You can add separate VMs and whole VM containers (hosts, clusters, folders, resource pools, VirtualApps, datastores or tags).

To add VMs and VM containers:

1. Click **Add VM**.
2. Select where to browse for VMs and VM containers:
   - **From infrastructure** – browse the virtual environment and select VMs or VM containers. If you choose a VM container, Veeam Backup & Replication will expand it to a plain VM list.
   - **From replicas** – browse existing replication jobs and select all VMs or specific VMs from replication jobs.

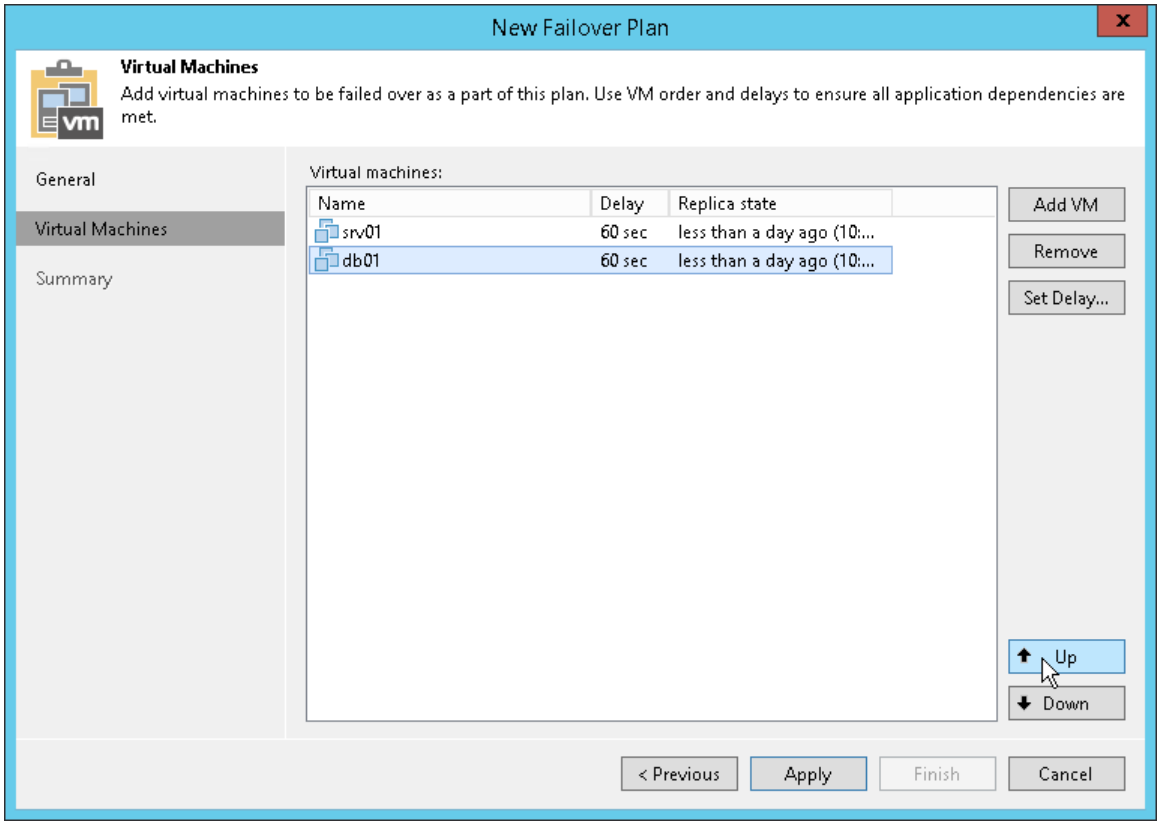
To quickly find the necessary VMs or VM containers, you can use the search fields.

> **NOTE**
>
> A source from which you add a VM to a failover plan does not affect whether you fail over to the latest or specific restore point. It is the command that you select when starting a failover plan that defines the restore point. For more information, see Running Failover Plans.



Step 4. Define VM Failover Order

At the **Virtual Machines** step of the wizard, click **Up** and **Down** to change the processing order. VMs at the top of the list have a higher priority and will be started first. If some VMs provide environment for other dependent VMs, make sure that they are started first.



Step 5. Set Time Delay

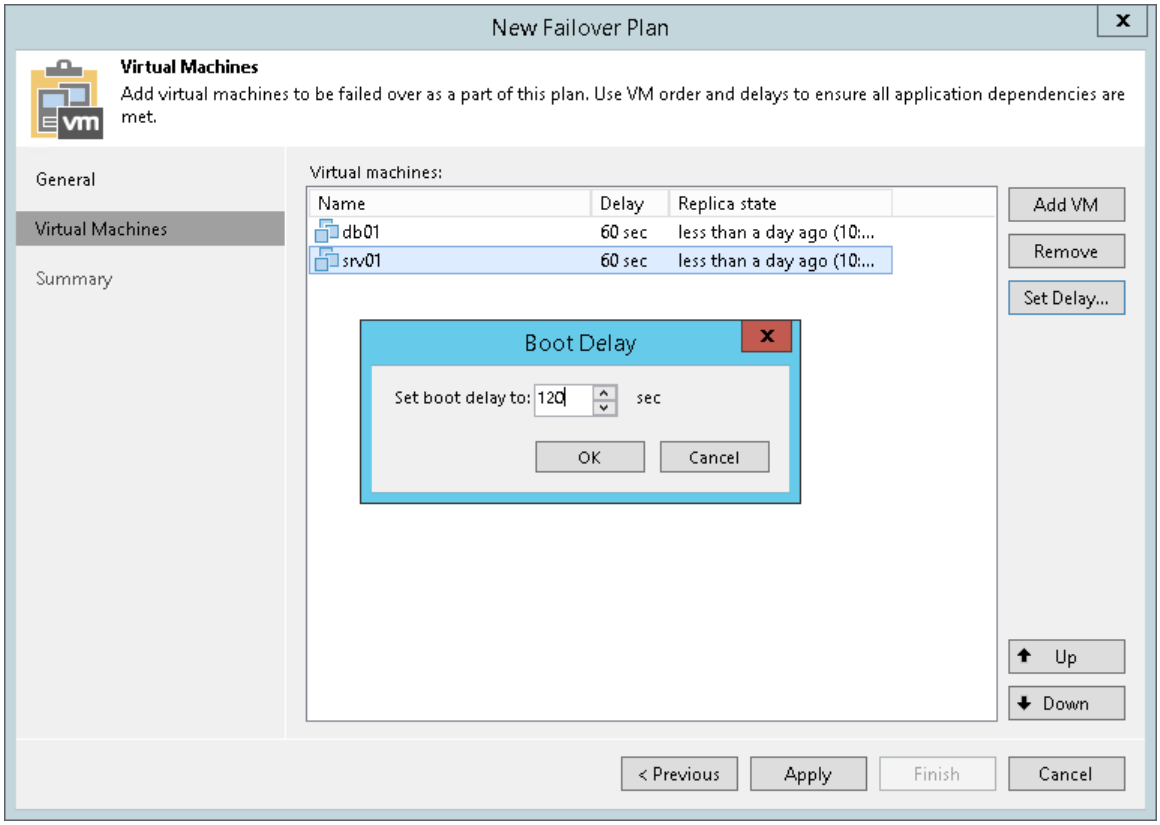
After you have set the order for VMs in the failover plan, you need to set a time delay for VMs. The delay time defines for how long Veeam Backup & Replication must wait before starting the failover operation for the next VM in the list. You can use time delays to make sure that some VMs are already running at the moment dependent VMs start.

For example, you have added 2 VMs to the failover plan and set a time delay to 60 seconds for the first VM in the list. Veeam Backup & Replication will perform failover in the following manner: Veeam Backup & Replication will start the failover operation for the first VM in the list, then wait for 60 seconds and start the failover operation for the second VM in the list.

> **NOTE**
>
> Time delays can be specified for all VMs in the list except the last one. If you do not specify time delays, VMs will be started simultaneously.

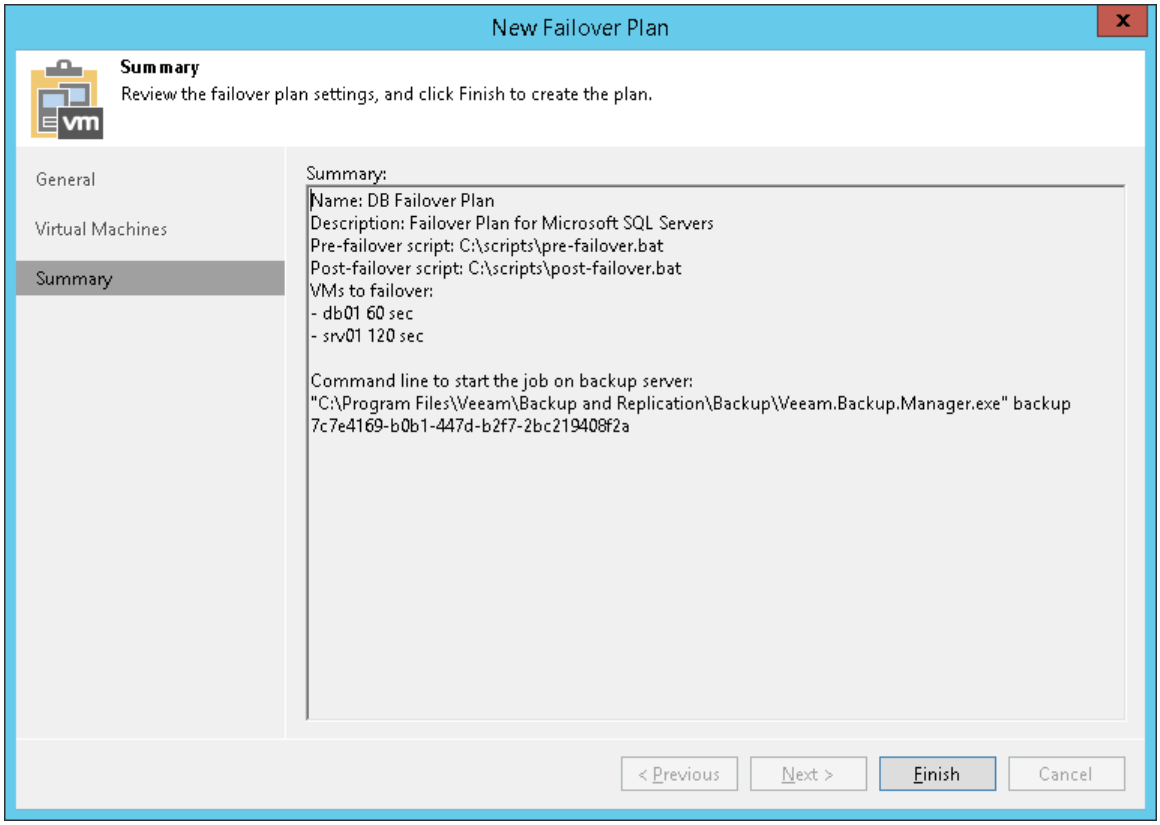
To set the time delay for a VM:

1. Select it and click **Set Delay** on the right or double-click the VM in the list.
2. Enter the time interval that you consider sufficient for this VM to boot.



Step 6. Finish Working with Wizard

At the **Summary** step of the wizard, review details for the configured failover plan and click **Finish** to create the plan.



Running Failover Plans

You have the following options to run a failover plan:

- You can fail over to latest restore points of VM replicas.

  In this case, Veeam Backup & Replication searches for the latest restore point of VM replicas across all replication jobs configured on the backup server. For example, you have 2 jobs that replicate the same VM: *Job 1* has created the most recent point at 2:00 AM and *Job 2* has created the most recent restore point at 3:00 AM. When you run the failover plan using the **Start** command, Veeam Backup & Replication will pick the restore point created at 3:00 AM with *Job 2*.

- You can fail over to specific restore points of VM replicas.

Failing Over to Latest Restore Points

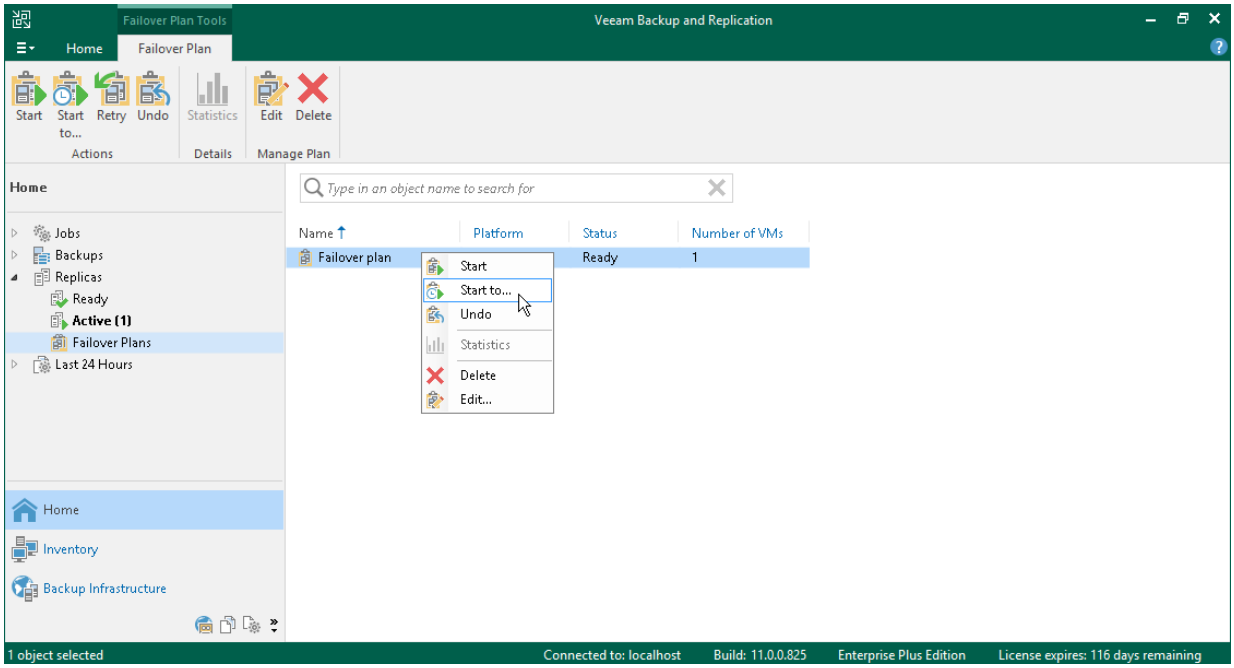
To fail over to the latest restore points of VM replicas:

1. Open the **Home** view.
2. Expand the **Replicas** node.
3. Select **Failover Plans**.
4. In the working area, right-click the failover plan and select **Start**.

Failing Over to Specific Restore Points

To fail over to specific restore points of VM replicas:

1. Open the **Home** view.
2. Expand the **Replicas** node.
3. Select **Failover Plans**.
4. In the working area, right-click the failover plan and select **Start to**.
5. In the displayed window, select the backup date and time. Veeam Backup & Replication will find the closest restore point prior to the entered value for each VM and will fail over to it.



Undoing Failover by Failover Plans

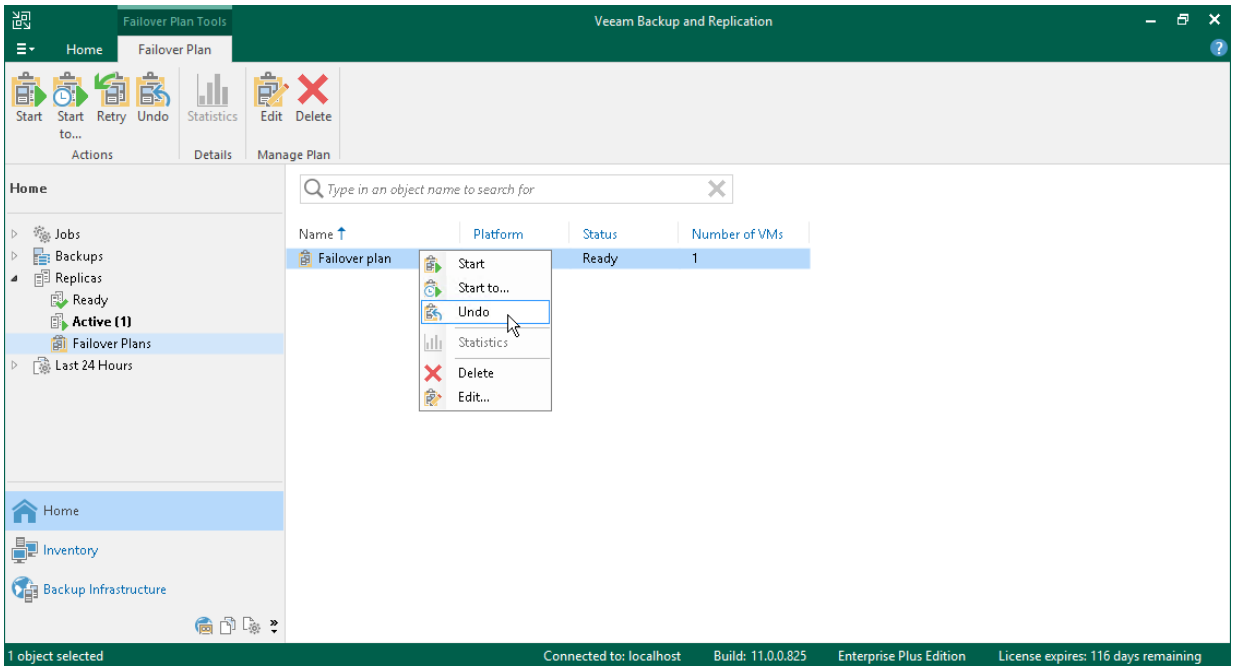
You can undo failover for all VMs added to the failover plan at once. When you undo failover, you switch the workload back to original VMs and discard all changes that were made to VM replicas during failover. If some of the VMs were already failed back, for example manually by the user, they are skipped from processing.

Veeam Backup & Replication starts the undo failover operation for a group of 5 VMs at the same time. The time interval between the operation starts is 10 seconds. For example, if you have added 10 VMs to the failover plan, Veeam Backup & Replication will undo failover for the first 5 VMs in the list, then will wait for 10 seconds and undo failover for the remaining 5 VMs in the list. Time intervals between the operation starts help Veeam Backup & Replication reduce the workload on the production environment and the backup server.

To undo failover by a failover plan:

1. Open the **Home** view.
2. Expand the **Replicas** node.
3. Select **Failover Plans**.
4. In the working area, right-click the failover plan and select **Undo**.

5. In the displayed dialog box, click **Yes** to confirm the operation.



Failover

Failover is a process of switching from the original VM in the production site to its VM replica in the disaster recovery site. During failover, Veeam Backup & Replication recovers the VM replica to the required restore point and shifts all I/O processes from the original VM to its replica. As a result, you have a fully functional VM within a couple of seconds, and your users can access services and applications with minimum disruption.

You can fail over to replicas not only when a disaster strikes the production site, but also to test replicas for recoverability. You can perform failover while the original VM is running. After all the necessary tests, you can undo failover and get back to the normal mode of operation. If the original VMs and VM replicas are located in the same network, consider temporary disconnecting the original VMs from the network to avoid IP address or machine name conflicts. As an alternative way of testing, Veeam Backup & Replication also provides the SureReplica technology. For more details, see SureReplica.

> **IMPORTANT**
>
> Use Veeam Backup & Replication to perform failover operations. Avoid powering on a replica manually — this may disrupt further replication operations or cause loss of important data.

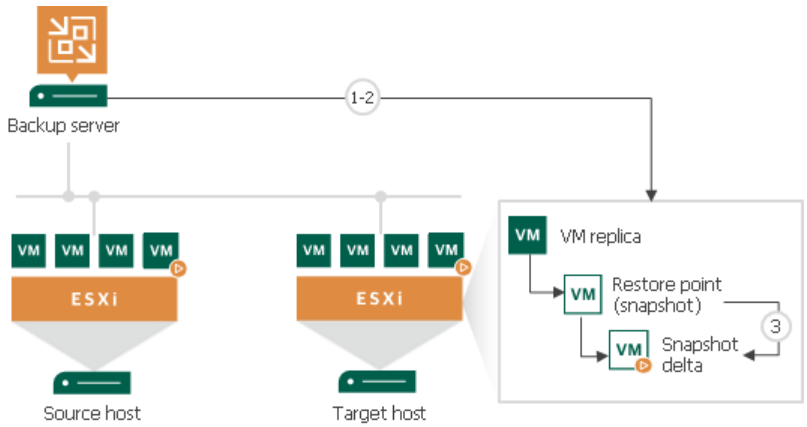
The failover operation is performed in the following way:

1. Veeam Backup & Replication rolls back the VM replica to the required restore point. To do this, it reverts the VM replica to the necessary snapshot in the replica chain.
2. Veeam Backup & Replication powers on the VM replica. The state of the VM replica is changed from *Ready* to *Failover*.

   If you perform failover for testing or disaster recovery (DR) simulation purposes, and the original VM still exists and is running, the original VM remains powered on.

   > **NOTE**
   >
   > Veeam Backup & Replication stops all replication activities for the original VM until its replica is returned to the *Ready* state.

3. All changes made to the VM replica while it is running in the *Failover* state are written to the delta file of the snapshot, or restore point, to which you have selected to roll back.



Finalizing Failover

Failover is an intermediate step that needs to be finalized. You can use one of the following operations:

- Undo failover
- Perform permanent failover
- Perform failback

Performing Failover

To perform failover, use the **VMware Failover** wizard.

Before You Begin

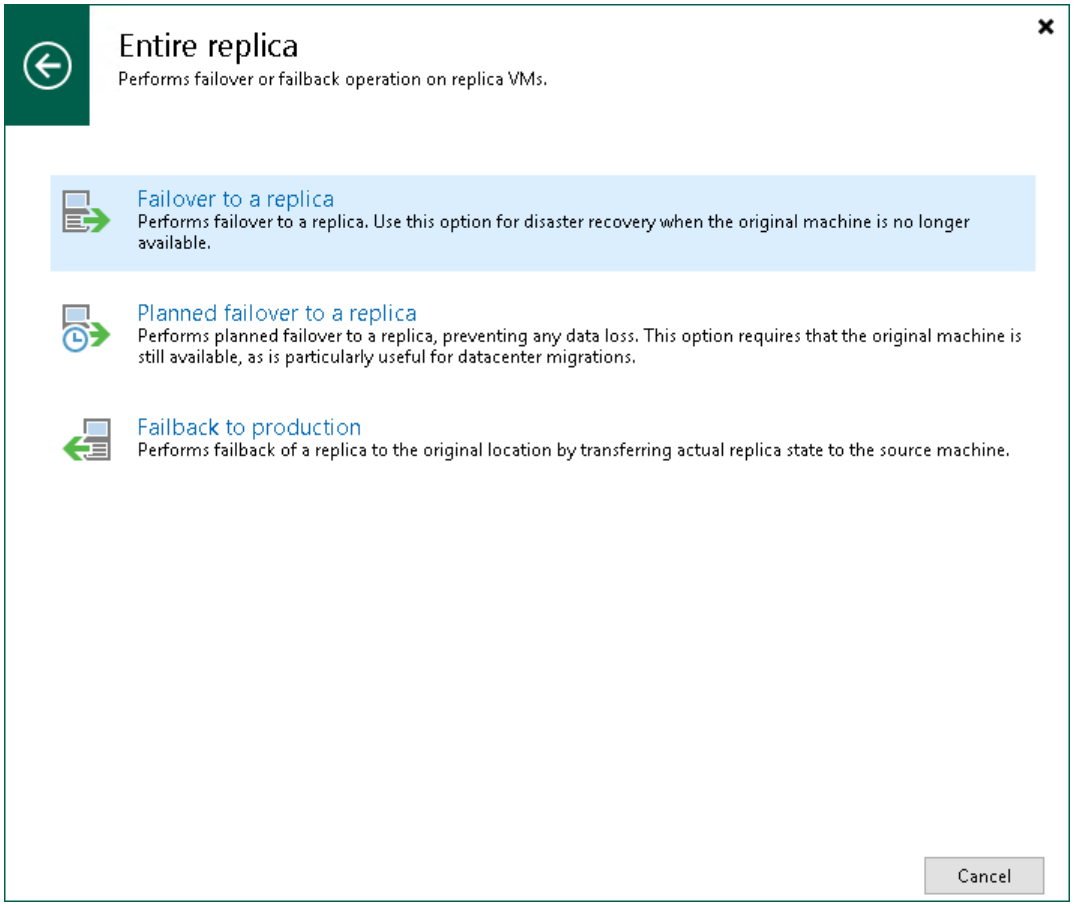
Before you fail over to a VM replica, check the following prerequisites:

- The failover operation can be performed for VMs that have been successfully replicated at least once.
- VM replicas must be in the *Ready* state.

Step 1. Launch Failover Wizard

To launch the **Failover** wizard, do one of the following:

- On the **Home** tab, click **Restore** > **VMware vSphere** > **Restore from replica** > **Entire replica** > **Failover to a replica**.
- Open the **Home** view. In the inventory pane select **Replicas** > **Ready**. In the working area, select the necessary replica and click **Failover Now** on the ribbon.
- Open the **Home** view. In the inventory pane select **Replicas** > **Ready**. In the working area, right-click the necessary replica and select **Failover Now**.



Step 2. Select VMs

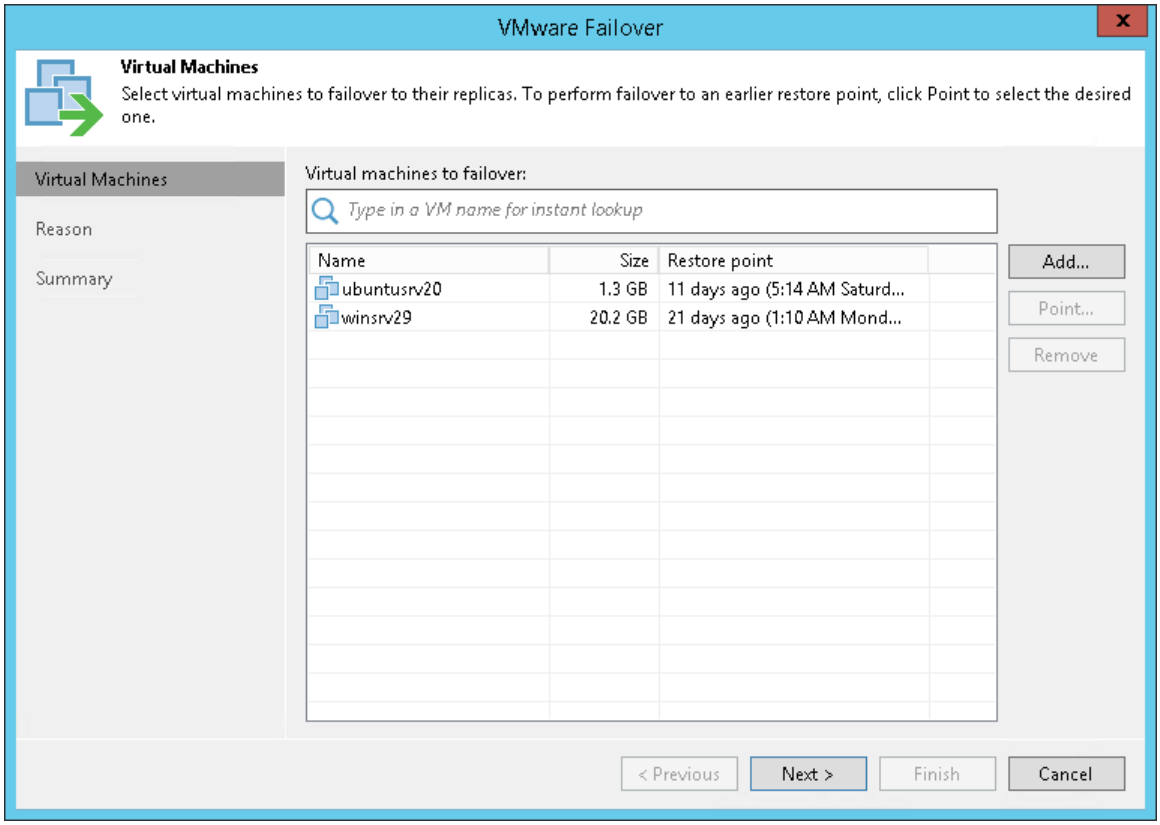
At the **Virtual Machines** step of the wizard, select one or more VMs for which you want to perform failover. You can perform failover for separate VMs and whole VM containers (hosts, clusters, folders, resource pools, VirtualApps, datastores or tags).

To select VMs and VM containers:

1. Click **Add VM**.
2. Select where to browse for VMs and VM containers:
   - **From infrastructure** – browse the virtual environment and select VMs or VM containers. If you choose a VM container, Veeam Backup & Replication will expand it to a plain VM list.
   - **From replicas** – browse existing replicas and select VMs or VM containers under replication jobs.

To quickly find the necessary VMs within existing replicas, you can use the search field at the top of the window.

> **NOTE**
>
> Make sure that VMs you select from the virtual environment have been successfully replicated at least once.

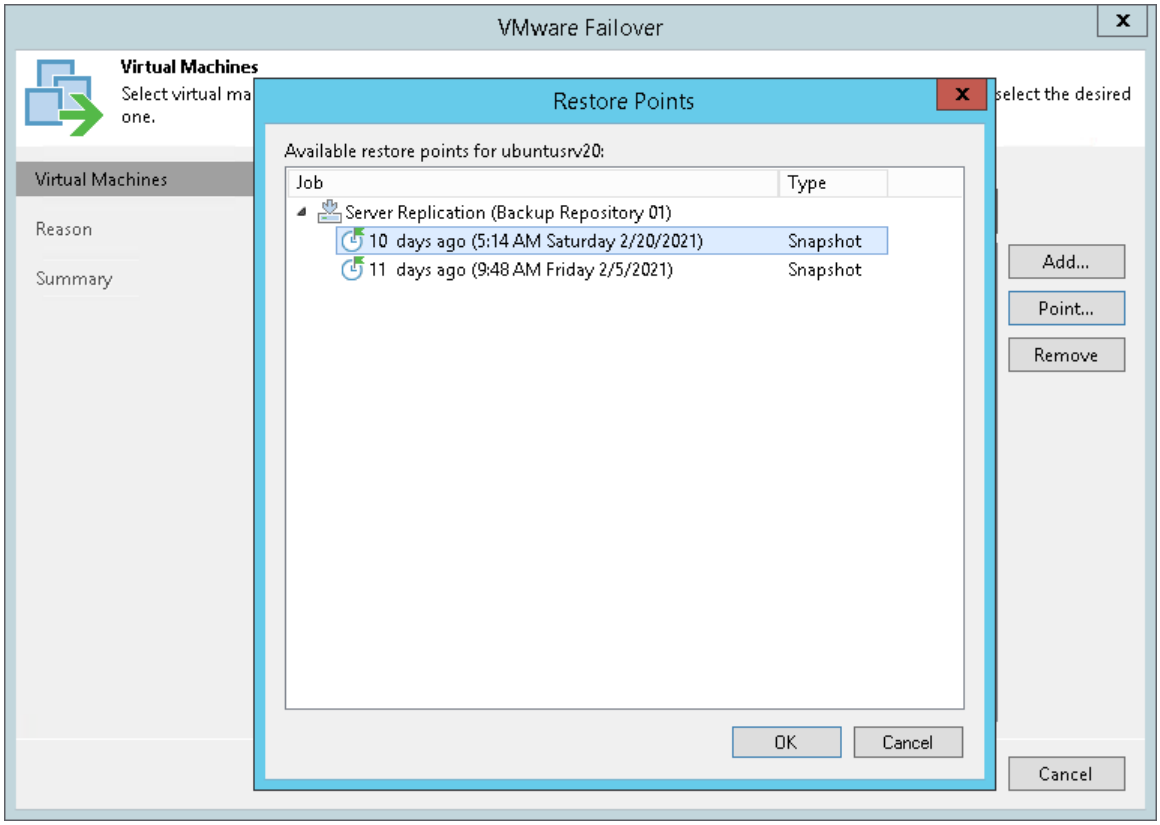


Step 3. Select Restore Points

By default, Veeam Backup & Replication uses the latest valid restore point of the VM replica. However, you can fail over to an earlier state of the VM. If you have chosen to perform failover for several VMs, you can select the necessary restore point for every VM in the list.

To select a restore point for a VM:

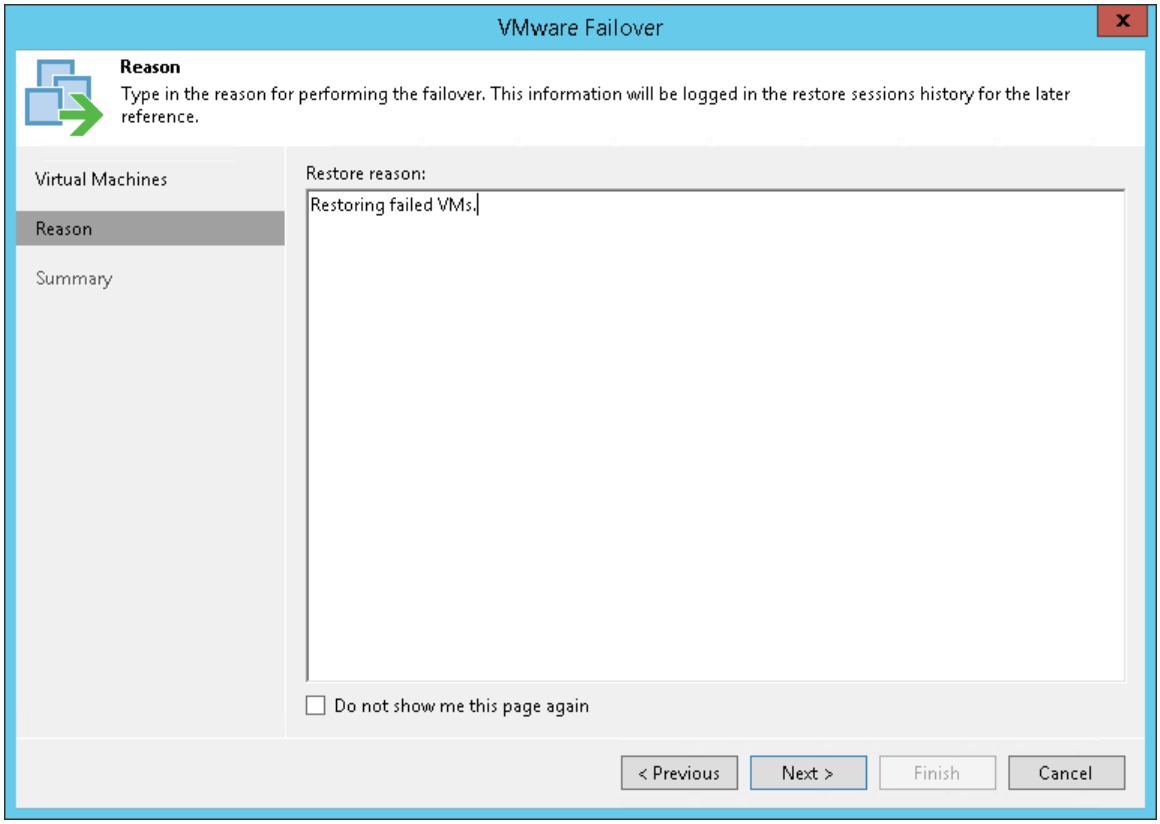
1. In the **Virtual machines to failover** list, select the necessary VM and click **Point**.
2. In the **Restore Points** window, select the necessary restore point.



Step 4. Specify Failover Reason

At the **Reason** step of the wizard, enter a reason for failing over to the VM replicas. The information you provide will be saved in the session history and you can reference it later.

> **TIP**
>
> If you do not want to display the **Reason** step of the wizard in future, select the **Do not show me this page again** check box.



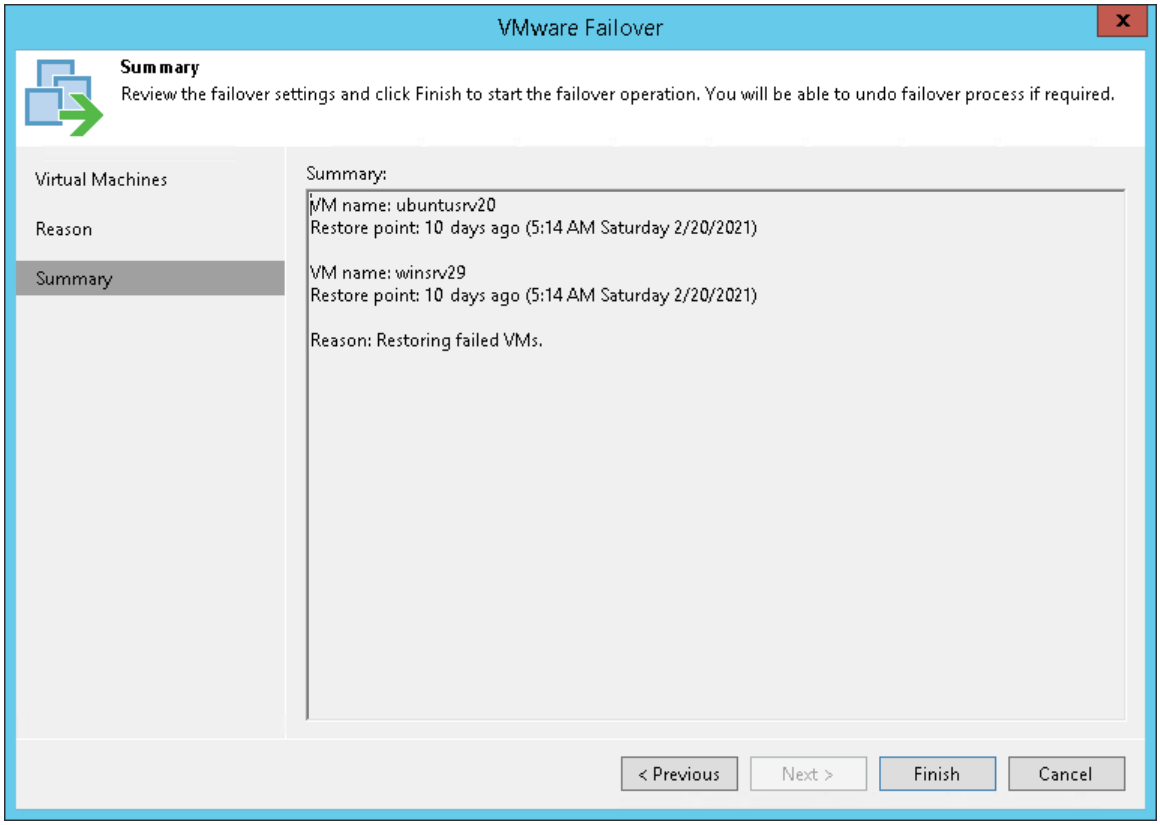
Step 5. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the failover task and click **Finish** to exit the wizard. When the failover process is complete, the VM replicas will be started on the target host.

What You Do Next

Failover is an intermediate step that needs to be finalized. You can finalize failover in the following ways:

- Perform permanent failover
- Undo failover
- Perform failback



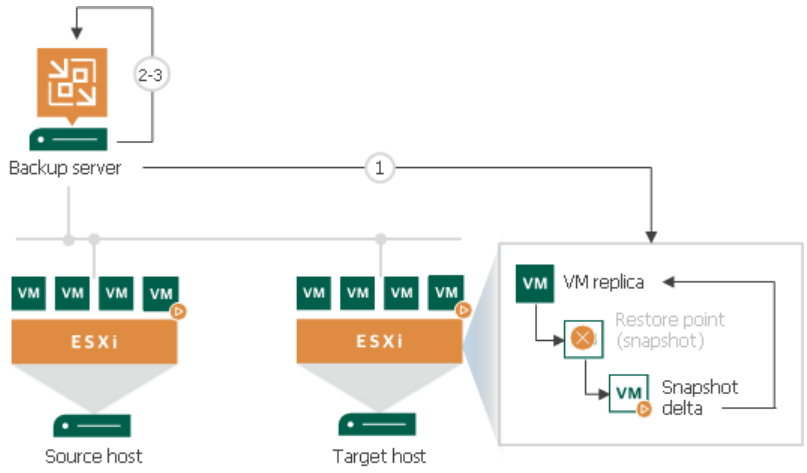
Permanent Failover

Permanent failover is one of the ways to finalize failover. When you perform permanent failover, you permanently switch from the original VM to it replica. As a result of permanent failover, the VM replica stops acting as a replica and starts acting as the production VM.

> **NOTE**
>
> We recommend you to perform permanent failover only if the original VM and its replica are located in the same site and are nearly equal in terms of resources. In this case, users will not experience any latency in ongoing operations. Otherwise, perform failback.

The permanent failover operation is performed in the following way:

1. Veeam Backup & Replication removes snapshots (restore points) of the VM replica from the snapshot chain and deletes associated files from the datastore. Changes that were written to the snapshot delta file are committed to the VM replica disk files to bring the VM replica to the most recent state.
2. Veeam Backup & Replication removes the VM replica from the list of replicas in the Veeam Backup & Replication console.
3. To protect the VM replica from corruption after permanent failover is complete, Veeam Backup & Replication reconfigures the replication job: adds the VM replica in the list of processed VMs and adds the original VM to the list of exclusions. When the replication job starts, the original VM is skipped from processing. As a result, no data is written to the working VM replica.

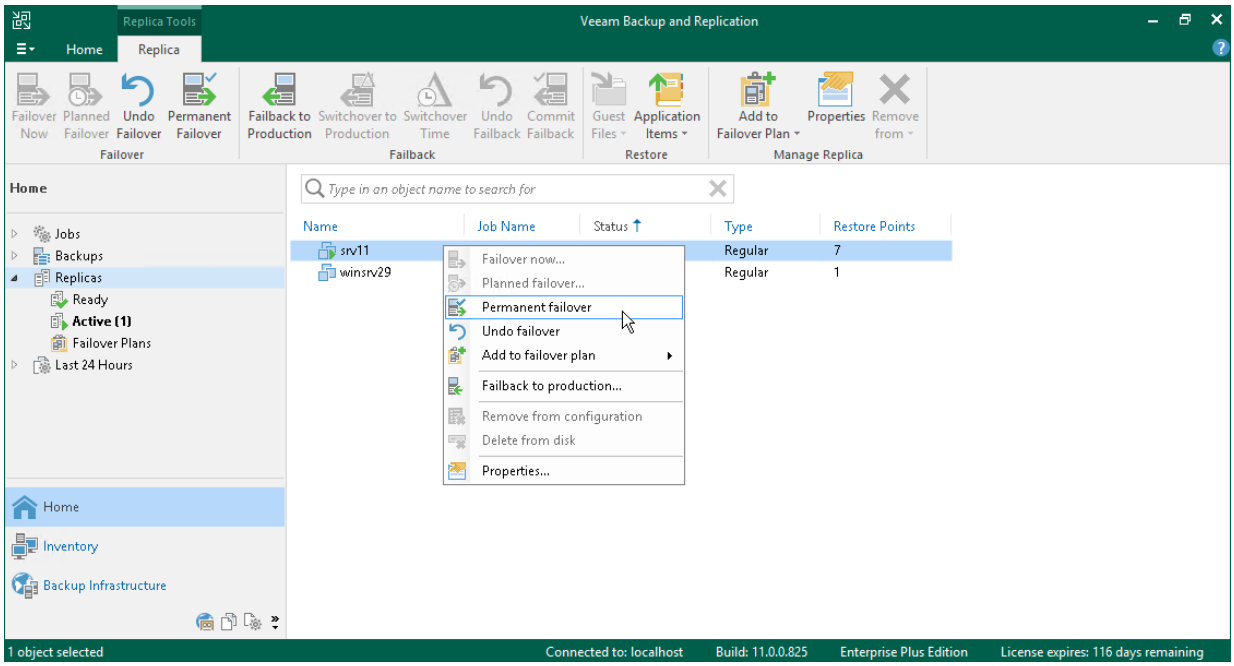


Performing Permanent Failover

To perform permanent failover, do one of the following:

- Open the **Home** view. In the inventory pane navigate to the **Replicas** > **Active** node. In the working area, select the necessary replica and click **Permanent Failover** on the ribbon.

- Open the **Home** view. In the inventory pane navigate to the **Replicas** > **Active** node. In the working area, right-click the necessary replica and select **Permanent failover**.



Planned Failover

Planned failover is smooth manual switching from a primary VM to its replica with minimum interrupting in operation. Planned failover is helpful when you know that your primary VMs are about to go offline and you need to proactively switch the workload from original VMs to their replicas. You can use the planned failover, for example, if you plan to perform datacenter migration, maintenance or software upgrade of the primary VMs. You can also perform planned failover if you have noticed some signs of the approaching disaster.

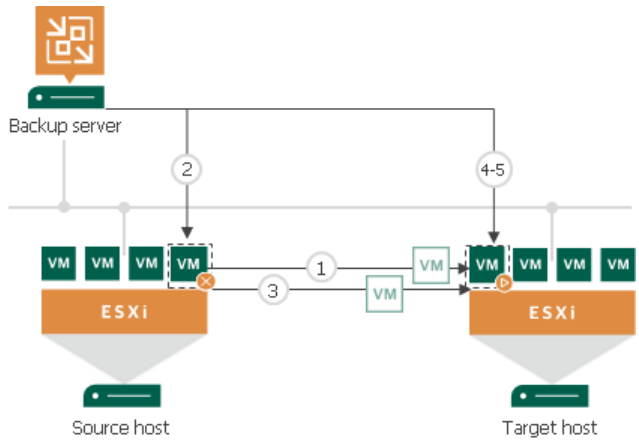
As the procedure is designed to transfer the current workload to the replica, it does not suggest selecting a restore point to switch.

When you start the planned failover, Veeam Backup & Replication performs the following operations:

1. The failover process triggers the replication job to perform an incremental replication run and copy the un-replicated changes to the replica.
2. The guest OS of the VM is shut down or the VM is powered off.

   If VMware Tools are installed on the VM, Veeam Backup & Replication tries to shut down the VM guest OS. If nothing happens after 15 minutes, Veeam Backup & Replication powers off the VM. If VMware Tools are not installed on the VM or the VM is suspended, Veeam Backup & Replication powers off the VM.
3. The failover process triggers the replication job to perform another incremental replication run and copy the portion of last-minute changes to the replica. The replica becomes fully synchronized with the source VM.
4. The VM is failed over to its replica.
5. The VM replica is powered on.

During the planned failover, Veeam Backup & Replication creates two helper restore points (steps 1 and 3) that are not deleted afterwards. You can see these restore points in the list of restore points for the VM. You can use the restore points later to roll back to the necessary VM replica state.



NOTE:

During planned failover, Veeam Backup & Replication always retrieves VM data from the production infrastructure, even if the replication job uses the backup as a data source. This approach helps Veeam Backup & Replication synchronize the VM replica to the latest state of the production VM.

Finalizing Planned Failover

When your primary host is online again, you can switch back to it. You can finalize planned failover in the same ways as regular failover: undo failover, perform permanent failover or failback.

Limitations for Planned Failover

Planned failover has the following limitations:

- If you start planned failover for several VMs that are replicated with one replication job, these VMs will be processed one by one, not in parallel.
- Each planned failover task for each VM is processed as a separate replica job session. If a backup proxy is not available and the session has to wait for resources, job sessions for other VMs in the same task cannot be started before the current session is finished.
- The user account under which you launch the planned failover operation must have the *Veeam Backup Administrator* role or *Veeam Backup Operator* and *Veeam Restore Operator* roles in Veeam Backup & Replication. For more information, see Roles and Users.

Performing Planned Failover

To perform planned failover, use the **Planned Failover** wizard.

Before You Begin

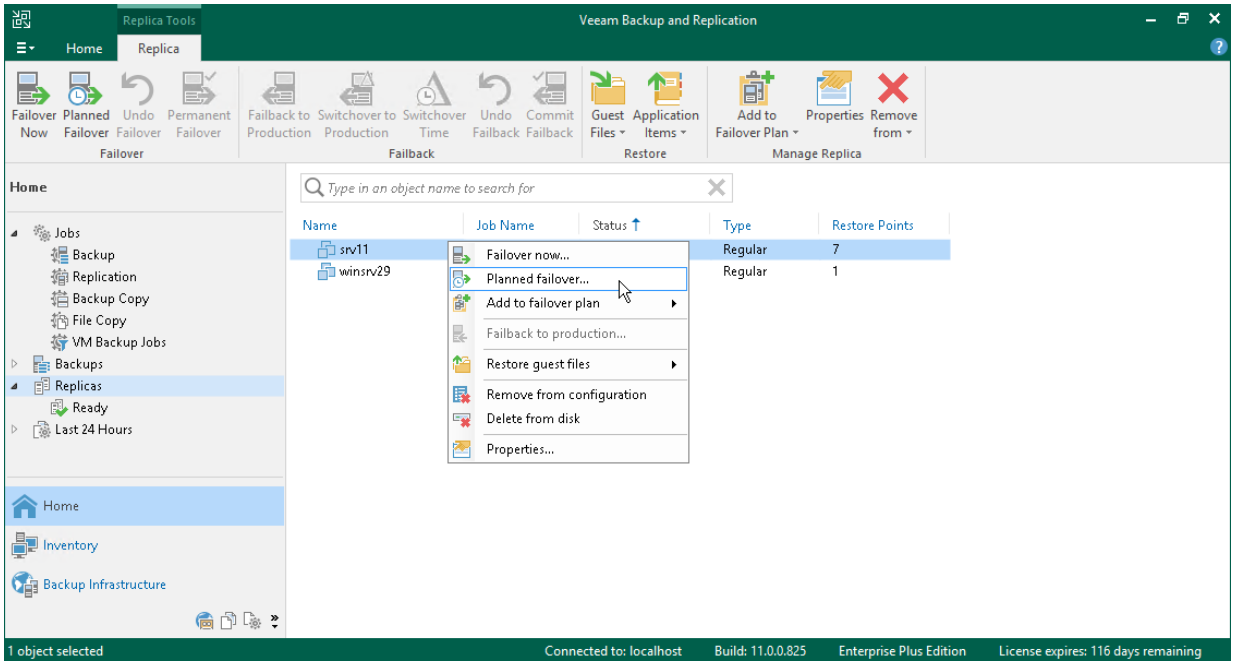
Before you perform planned failover, check the following prerequisites:

- VMs for which you plan to perform planned failover must be successfully replicated at least once.
- VM replicas must be in the *Ready* state.

Step 1. Launch Planned Failover Wizard

To launch the **Planned Failover** wizard, do one of the following:

- On the **Home** tab, click **Restore** and select **VMware vSphere** > **Restore from replica** > **Entire replica** > **Planned failover to a replica**.
- Open the **Home** view, expand the **Replicas** node. In the working area, select one or more VMs and click **Planned Failover** on the ribbon. You can also right-click one of the selected VMs and click **Planned Failover**.
- Open the **Inventory** view. In the working, select one or more VMs and right-click one of the selected VMs and click **Restore** > **Planned Failover**.



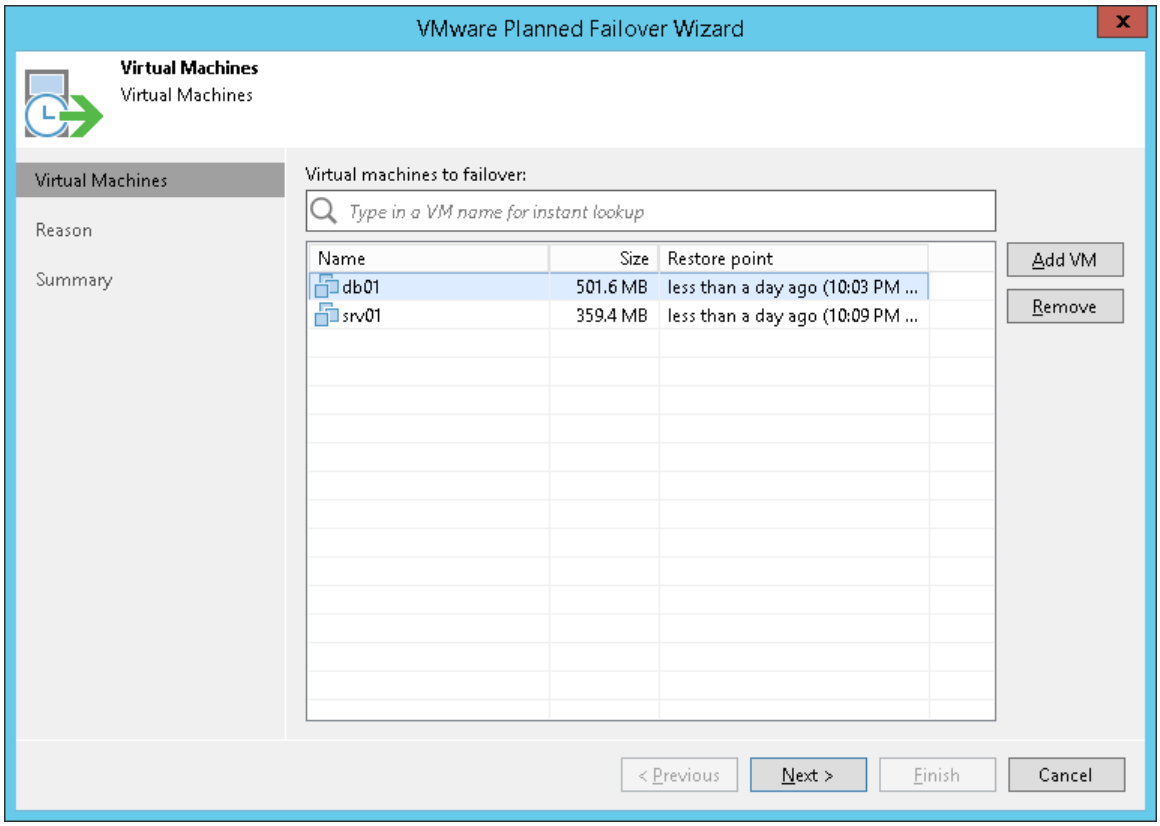
Step 2. Select VMs

At the **Virtual Machines** step of the wizard, select one or more VMs for which you want to perform failover. You can perform failover for separate VMs and whole VM containers (hosts, clusters, folders, resource pools, VirtualApps, datastores or tags).

To select VMs and VM containers:

1. Click **Add VM**.
2. Select where to browse for VMs and VM containers:
   - **From infrastructure** – browse the virtual environment and select VMs or VM containers. If you choose a VM container, Veeam Backup & Replication will expand it to a plain VM list.
   - **From replicas** – browse existing replication jobs and select all VMs or specific VMs from replication jobs.

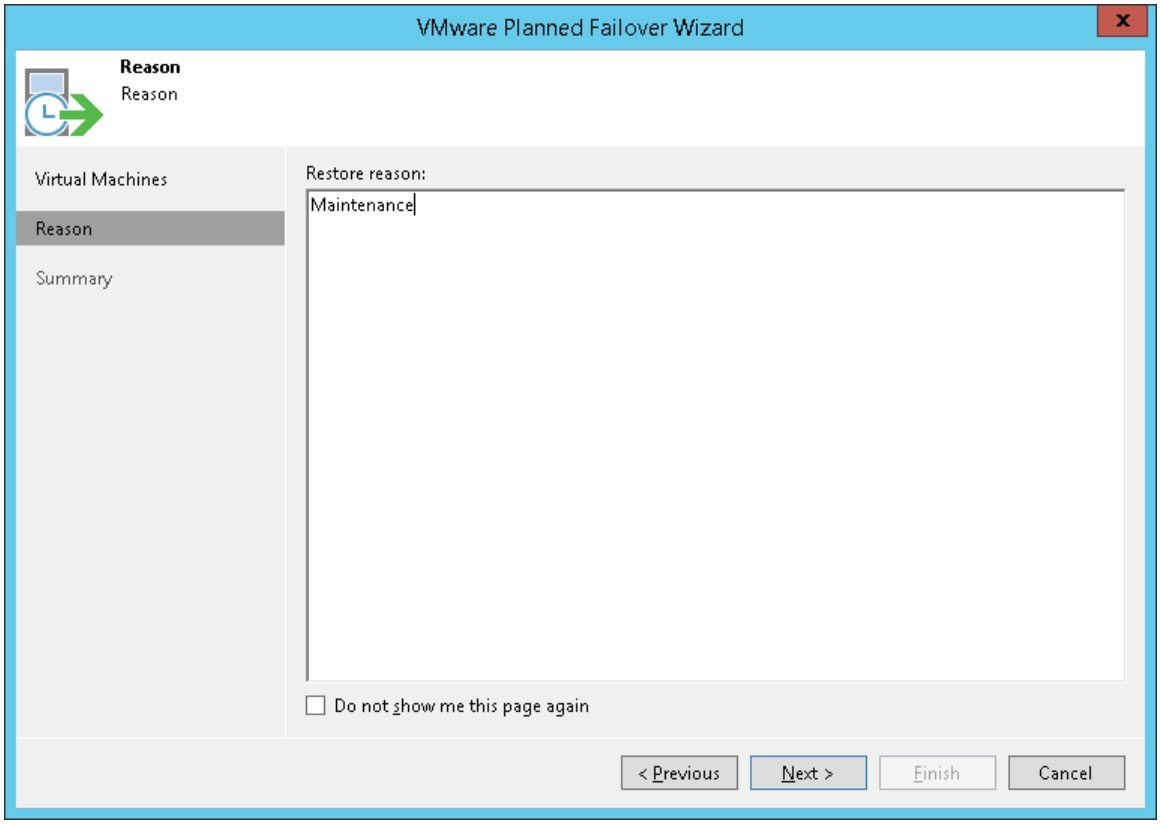
To quickly find the necessary VMs or VM containers, you can use the search fields.



Step 3. Specify Failover Reason

At the **Reason** step of the wizard, enter a reason for failing over to VM replicas. The information you provide will be saved in the session history and you can reference it later.

> **TIP**
>
> If you do not want to display the **Reason** step of the wizard in future, select the **Do not show me this page again** check box.



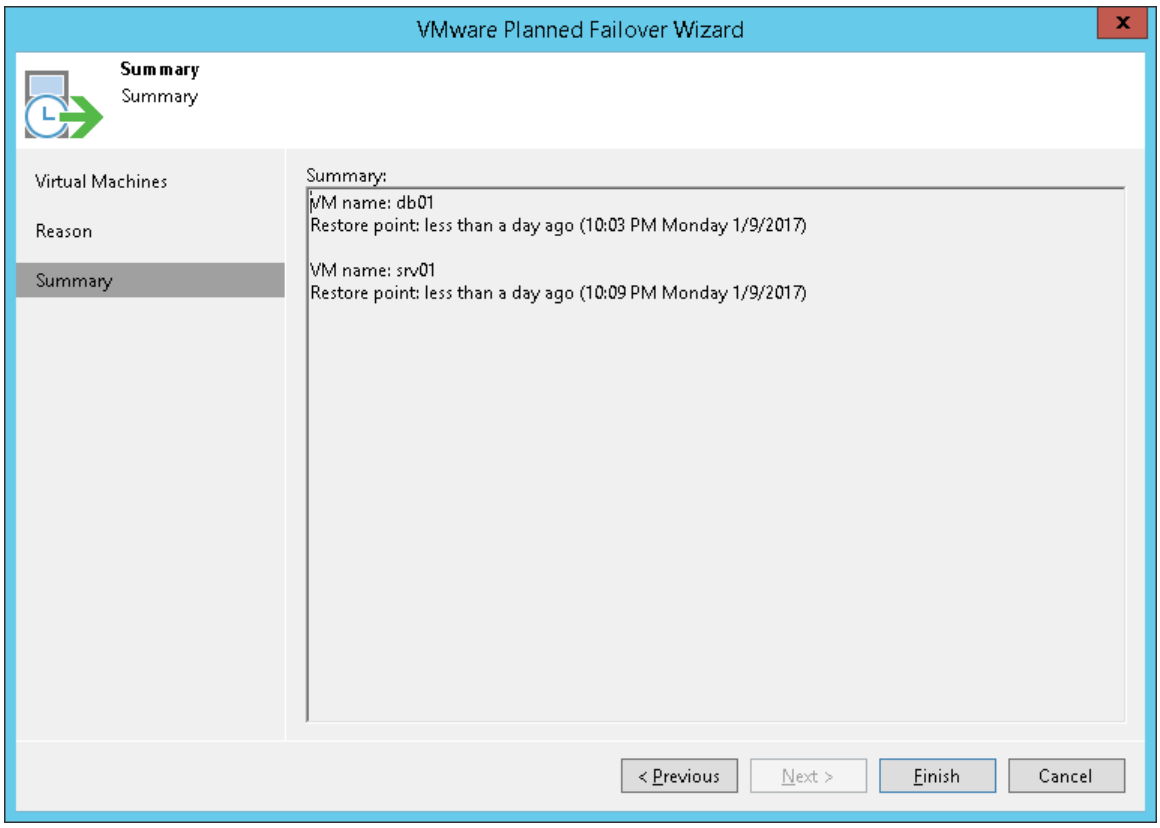
Step 4. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the failover task and click **Finish** to start the failover process. Once planned failover is complete, VM replicas will be started on the target host.

What You Do Next

Planned failover is an intermediate step that needs to be finalized. You can finalize failover in the following ways:

- Perform permanent failover
- Undo failover
- Perform failback

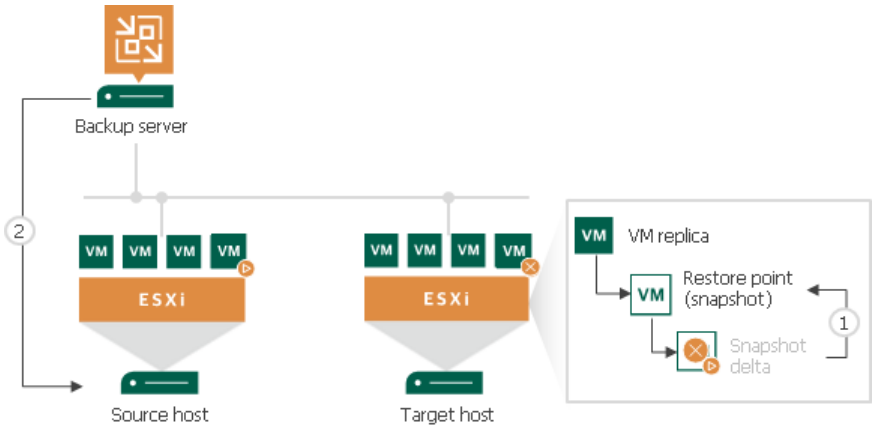


Failover Undo

Failover undo is one of the ways to finalize failover. When you undo failover, you switch back from a VM replica to the original VM. Veeam Backup & Replication discards all changes made to the VM replica while it was in the *Failover* state. You can use the undo failover scenario if you have failed over to the VM replica for testing and troubleshooting purposes and want to get back to the normal operation mode.

The undo failover operation is performed in the following way:

1. Veeam Backup & Replication reverts the VM replica to its pre-failover state. To do this, Veeam Backup & Replication powers off the VM replica and gets it back to the state of the latest snapshot in the snapshot chain. Changes that were written to the snapshot delta file while the VM replica was in the *Failover* state are discarded.
2. The state of the VM replica gets back to *Ready*, and Veeam Backup & Replication resumes replication activities for the original VM on the source host.



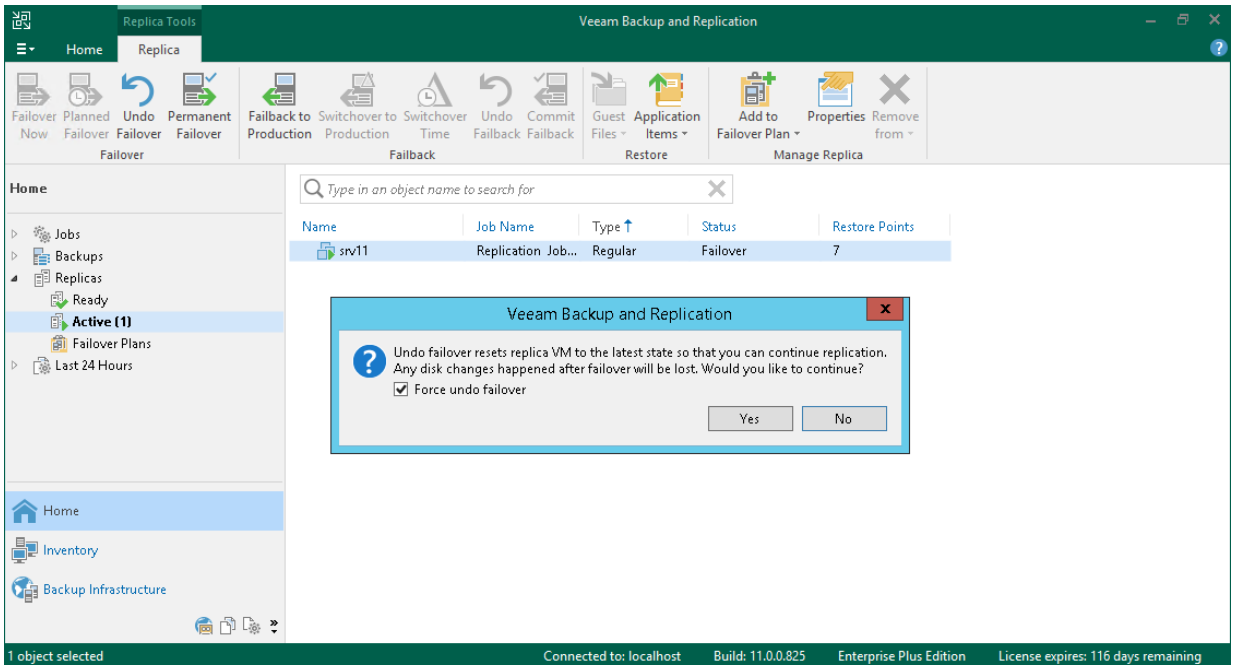
Undoing Failover

To undo failover:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Replicas > Active** node.
3. In the working area, select the necessary replica and click **Undo Failover** on the ribbon. As an alternative, right-click one of the selected replicas and click **Undo failover**.
4. In the displayed window, do the following:
   a. If you want to force failover undo, select the **Force undo failover** check box.

      When you force failover, Veeam Backup & Replication attempts to perform the undo failover operation in a regular way. If the host on which the VM replica resides is unavailable, Veeam Backup & Replication changes the VM replica state to *Ready* in the configuration database and console. This helps avoid failure of the undo failover operation.

b. Click **Yes**.


Replica Tools
Veeam Backup and Replication
Home
Replica
Failover Now
Planned Failover
Undo Failover
Permanent Failover
Failover
Failback to Production
Switchover to Production
Switchover Time
Undo Failback
Commit Failback
Failback
Guest Files
Application Items
Restore
Add to Failover Plan
Properties
Remove from
Manage Replica
Home
Type in an object name to search for
Name
Job Name
Type
Status
Restore Points
srv11
Replication Job...
Regular
Failover
7
Jobs
Backups
Replicas
Ready
Active (1)
Failover Plans
Last 24 Hours
Veeam Backup and Replication
Undo failover resets replica VM to the latest state so that you can continue replication.
Any disk changes happened after failover will be lost. Would you like to continue?
Force undo failover
Yes
No
Home
Inventory
Backup Infrastructure
1 object selected
Connected to: localhost
Build: 11.0.0.825
Enterprise Plus Edition
License expires: 116 days remaining

Failback

Failback is one of the ways to finalize failover. When you perform failback, you switch back to the production VM from a VM replica, shift I/O processes from the disaster recovery site to the production site.

Veeam Backup & Replication provides you the following options to perform failback:

- You can fail back to the original VM in the original location.
- You can fail back to a VM already recovered to a new location. This VM must be recovered before you perform failback. For example, you can recover the VM from a backup.
- You can fail back to a VM recovered from a replica to a new location, or to any location but with different settings. The VM will be recovered from the replica during the failback process.

The first two options help you decrease recovery time and the use of the network traffic because Veeam Backup & Replication needs to transfer only differences between the original/recovered VM and VM replica. For the third option, Veeam Backup & Replication needs to transfer the whole VM data, including its configuration and virtual disk content. Use the third option if there is no way to use the original VM or restore it from a backup.

Veeam Backup & Replication performs failback in two phases:

- **First phase**: Veeam Backup & Replication synchronizes the state of the production VM (the original VM, an already recovered VM or a VM that will be recovered from the replica) with the current state of its replica. This phase may take a lot of time especially if VM is large. While Veeam Backup & Replication performs the first phase of failback, VM replicas are still up and running, users can access these VMs and perform daily routine tasks as normal.
- **Second phase**: Veeam Backup & Replication switches all processes from the VM replica to the production VM, turns off the replica and also sends to the production VM changes made to the VM replica since the end of the first phase.

  The time when the second phase starts depends on how you want to switch from the replica to the production VM. You can switch to the production VM automatically, at the scheduled time or manually. If you select to switch automatically, the second phase will start right after the first phase finishes. If you select to switch at the scheduled time or manually, the second phase will start at the time you want.

The process of failing back to the original VM or an already recovered VM differs from the process of failing back to a VM recovered from a replica:

- How failback to the original VM and already recovered VM works
- How failback to a VM recovered from a replica works

How Failback to Original VM or Already Recovered VM Works

When you fail back to the original VM or an already recovered VM, Veeam Backup & Replication performs the following operations during the first phase:

1. If the production VM is running, Veeam Backup & Replication powers it off.
2. Veeam Backup & Replication creates a working failback snapshot for the production VM.
3. Veeam Backup & Replication creates a failback protective snapshot for the VM replica. You can use this snapshot to return to the pre-failback state of the VM replica afterwards.

4. Veeam Backup & Replication calculates the difference between disks of the production VM and disks of the VM replica in the *Failover* state. Difference calculation helps Veeam Backup & Replication understand what data needs to be transferred to the production VM to synchronize its state with the state of the VM replica.

   If you fail back to the original VM in the original location and you have enabled the **Quick rollback** option, difference calculation can be performed much faster than without this option enabled. For more information on quick rollback, see Quick Rollback.

5. Veeam Backup & Replication transfers the data that was detected at the previous step to the production VM. The transferred data is written to the delta file of the working failback snapshot on the production VM.

6. Veeam Backup & Replication removes the working failback snapshot from the production VM.

7. Veeam Backup & Replication changes the state of the VM replica from *Failover* to *Ready to switch*.

During the second phase, Veeam Backup & Replication performs the following operations:

1. Veeam Backup & Replication creates a working failback snapshot on the production VM.

2. The guest OS of the VM replica is shut down or the VM replica is powered off.

   If VMware Tools are installed on the VM replica, Veeam Backup & Replication tries to shut down the replica guest OS. If nothing happens after 15 minutes, Veeam Backup & Replication powers off the VM replica. If VMware Tools are not installed on the VM or the VM is suspended, Veeam Backup & Replication powers off the VM. The VM replica remains powered off until you commit failback or undo failback.

3. Veeam Backup & Replication creates a failback protective snapshot for the VM replica. The snapshot acts as a new restore point and saves the pre-failback state of the VM replica. You can use this snapshot to return to the pre-failback state of the VM replica afterwards.

4. Sends data changed on the VM replica while it was in the *Ready to switch* state to the working failback snapshot on the production VM.

5. Veeam Backup & Replication removes the protective snapshot from the VM replica.

6. Veeam Backup & Replication removes the working failback snapshot from the production VM. Changes written to the delta file of this snapshot are committed to the production VM disks.

7. The state of the VM replica is changed from *Ready to switch* to *Failback*. Veeam Backup & Replication temporarily puts replication activities for the production VM on hold.

8. [If you fail back to a VM already recovered to a new location] Veeam Backup & Replication updates the ID of the original VM in the Veeam Backup & Replication configuration database. The ID of the original VM is replaced with the ID of the recovered VM.

9. If you have selected to power on the production VM after failback, Veeam Backup & Replication powers on the production VM on the host.

How Failback to VM Recovered from Replica Works

When you fail back to a VM recovered from a replica, Veeam Backup & Replication performs the following operations during the first phase:

1. Veeam Backup & Replication requests vCenter Server to create on the target host an empty VM with the same configuration as the VM replica. vCenter Server registers the created production VM.

2. Veeam Backup & Replication creates a working failback snapshot for the production VM.

3. Veeam Backup & Replication creates a failback protective snapshot for the VM replica. You can use this snapshot to return to the pre-failback state of the VM replica afterwards.
4. Veeam Backup & Replication transfers data of the VM replica to the production VM to update the production VM state to the VM replica state.
5. Veeam Backup & Replication removes the working failback snapshot from the production VM.
6. Veeam Backup & Replication changes the state of the VM replica from *Failover* to *Ready to switch*.

During the second phase, Veeam Backup & Replication performs the same operations as described in How Failback to Original VM or Already Recovered VM Works.

Finalizing Failback

Failback is an intermediate step that needs to be finalized. If the production VM works as expected and you want to get back to it, commit failback. If the VM does not work as expected, undo failback.

Failback on VSAN

Due to specifics of VSAN data storage organization, Veeam Backup & Replication cannot get the difference between disks of a VM replica located on VSAN and disks of the original VM in a regular manner. Veeam Backup & Replication needs to read VM disks data anew in every failback process phase. As a result, failback for VMs replicas on VSAN slightly differs from the regular failback course.

Before Veeam Backup & Replication starts the failback process, it checks the location of VM replica disks. If at least one disk is located on VSAN, Veeam Backup & Replication performs failback in the following way:

1. Veeam Backup & Replication creates a working failback snapshot for the original VM.
2. For every VM disk, Veeam Backup & Replication performs the following actions:
    a. If you fail back to the original VM location, Veeam Backup & Replication calculates the difference between the VM replica disk and the original VM disk. To do this, Veeam Backup & Replication reads the whole amount of disk data from VSAN, and transfers only changed data to the original VM side.
    b. If you fail back to a new location, Veeam Backup & Replication transfers the whole disk without calculating the difference.
3. The VM replica is powered off.
4. Veeam Backup & Replication creates a protective failback snapshot for the VM replica. Using the protective failback snapshot, Veeam Backup & Replication detects what changes took place on the VM replica while VM disk data was being transported. As well as before, Veeam Backup & Replication reads the whole amount of VM disks data but transports only those data blocks that have changed since the VM disks transfer.

The rest of the failback process does not differ from the regular failback process.

Quick Rollback

If you fail back from a VM replica to the original VM in the original location, you can instruct Veeam Backup & Replication to perform quick rollback. Quick rollback significantly reduces the failback time and has little impact on the production environment.

During failback with the quick rollback option enabled, Veeam Backup & Replication does not calculate digests for entire VM replica disks to get the difference between the original VM and VM replica. Instead, it queries CBT to get information about disk sectors that have changed, and calculates digests only for these disk sectors. As a result, digest calculation is performed much faster. After that, Veeam Backup & Replication performs failback in a regular way: transport changed blocks to the original VM, powers off the VM replica and synchronizes the original VM with the VM replica once again.

It is recommended that you use quick rollback if you fail back to the original VM after a problem that has occurred at the level of the guest OS of the VM replica — for example, there has been an application error or a user has accidentally deleted a file on the VM replica guest OS. Do not use quick rollback if the problem has occurred at the VM hardware level, storage level or due to a power loss.

Requirements for Quick Rollback

To perform quick rollback, make sure that the following requirements are met:

- You must perform failback to the original VM in the original location.
- CBT must be enabled for the original VM.
- The VM replica must be created with the **Use changed block tracking data** option enabled.

Limitations for Quick Rollback

The following limitations apply to quick rollback:

- During the first replication job session after failback with quick rollback, CBT on the original VM is reset. Due to that Veeam Backup & Replication will read data of the entire VM.
- Quick rollback can be performed in the Direct NFS access, Virtual appliance, Network transport mode. The Direct SAN access transport mode cannot be used for quick rollback due to VMware limitations.

Performing Failback

To switch back to the original VM, use the **Failback** wizard.

Before You Begin

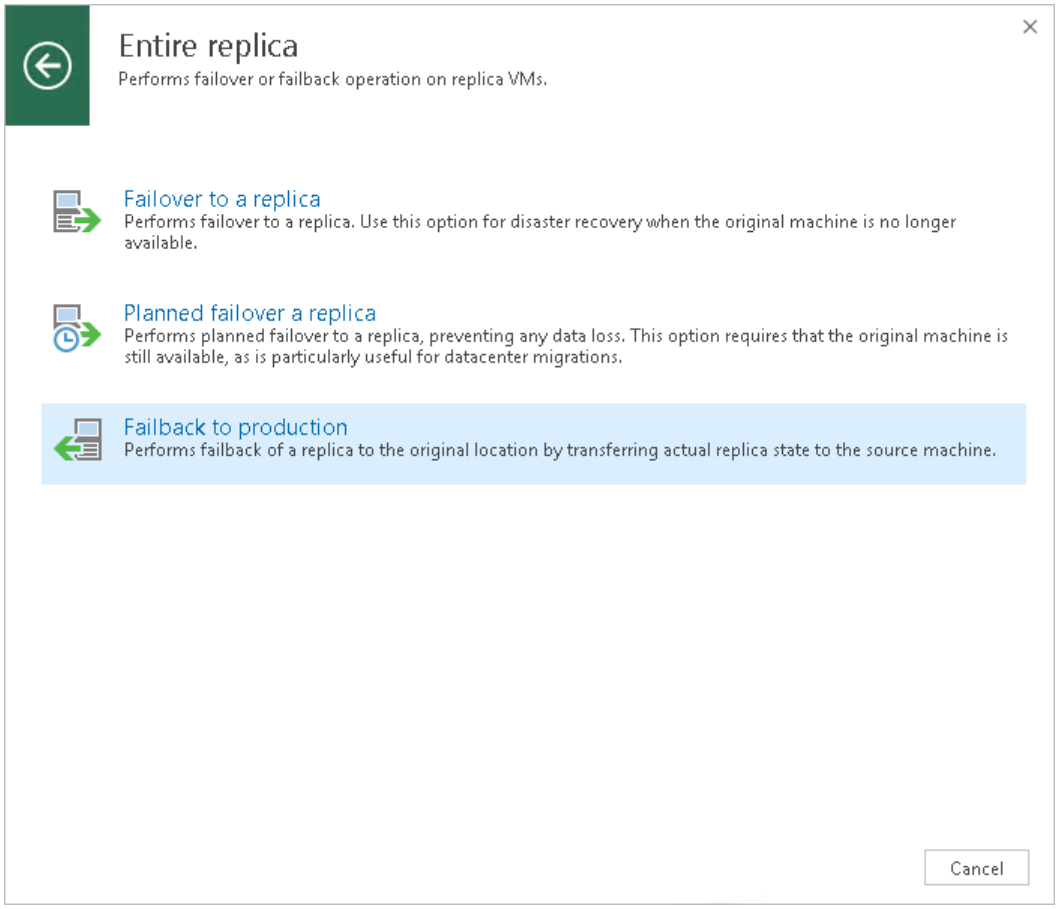
Before you perform failback, check the following prerequisites:

- VMs for which you plan to perform failback must be successfully replicated at least once.
- VM replicas must be in the *Failover* state.

Step 1. Launch Failback Wizard

To launch the **Failback** wizard, do one of the following:

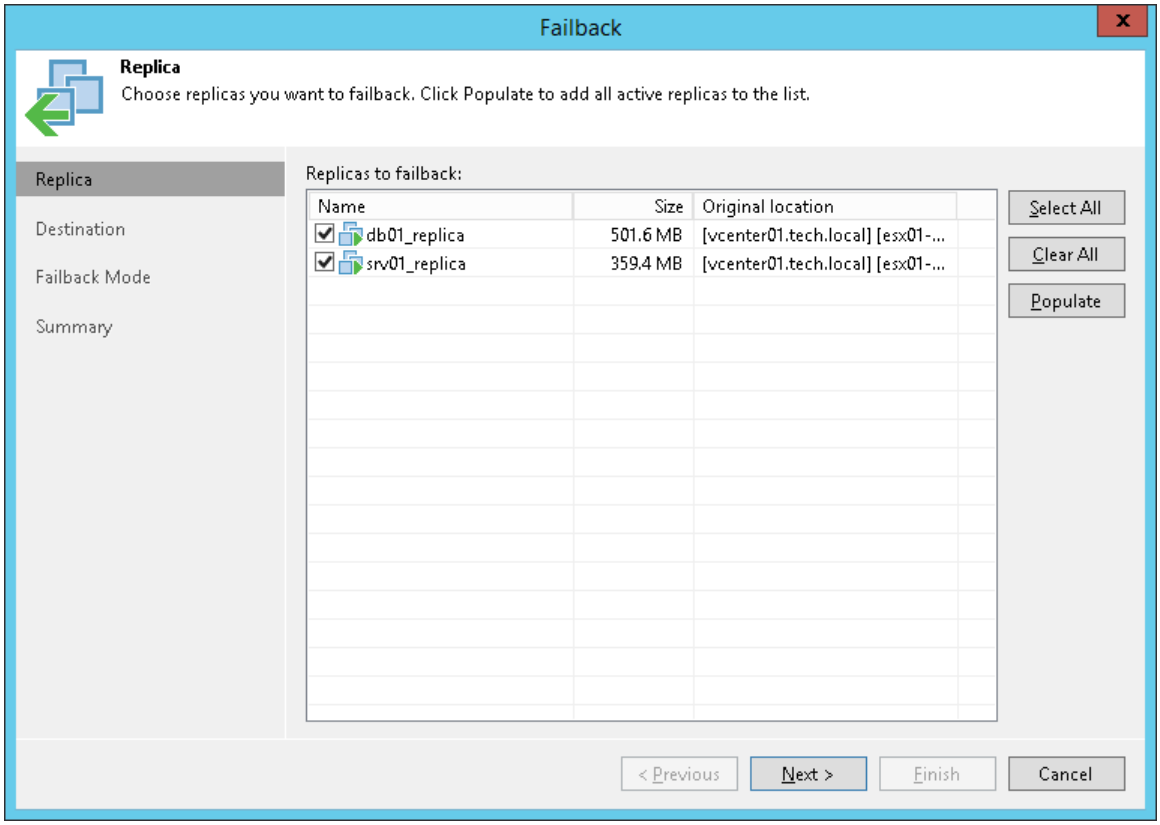
- On the **Home** tab, click **Restore** > **VMware vSphere** > **Restore from replica** > **Entire replica** > **Failback to production**.
- Open the **Home** view. In the inventory pane navigate to the **Replicas** > **Active** node. In the working area, select the necessary replica and click **Failback to Production** on the ribbon.
- Open the **Home** view. In the inventory pane navigate to the **Replicas** > **Active** node. In the working area, right-click the necessary replica and select **Failback to production**.



Step 2. Select VM Replicas to Fail Back

At the **Replica** step of the wizard, select replicas from which you want to fail back.

To update the list of replicas that are ready for failback (replicas in the *Failover* state), click **Populate**.



Step 3. Select Failback Destination

At the **Destination** step of the wizard, select a failback destination and backup proxies for data transfer during failback:

1. Select a destination for failback. Veeam Backup & Replication supports the following options:

   - **Failback to the original VM** — select this option if you want to fail back to the original VMs that reside on the source hosts. Veeam Backup & Replication will synchronize the state of the original VMs with the current state of their replicas to apply any changes that occurred to the VM replicas while running in the DR site.

     If you select this option, you will proceed to the Summary step of the wizard.

   - **Failback to the original VM restored in a different location** — select this option if the original VMs have already been recovered to a new location, and you want to switch to the recovered VMs from their replicas. Veeam Backup & Replication will synchronize the state of the recovered VMs with the current state of the VM replicas to apply any changes that occurred to the replicas while running in the DR site.

     If you select this option, you will proceed to the Target VM step of the wizard.

   - **Failback to the specified location** — select this option if you want to recover VMs from replicas. You can recover VMs to a new location, or to any location but with different settings (such as network settings, virtual disk type, configuration file path and so on). Select this option if there is no way to fail back to the original VM or an already recovered VM.

     If you select this option, the wizard will include additional steps.

   If you select one of the first two options, Veeam Backup & Replication will send to the original/recovered VMs only differences between the existing virtual disks. Veeam Backup & Replication will not send replica configuration changes such as different IP address or network settings (if replica Re-IP and network mapping were applied), new hardware or virtual disks added while the replicas were in the *Failover* state.

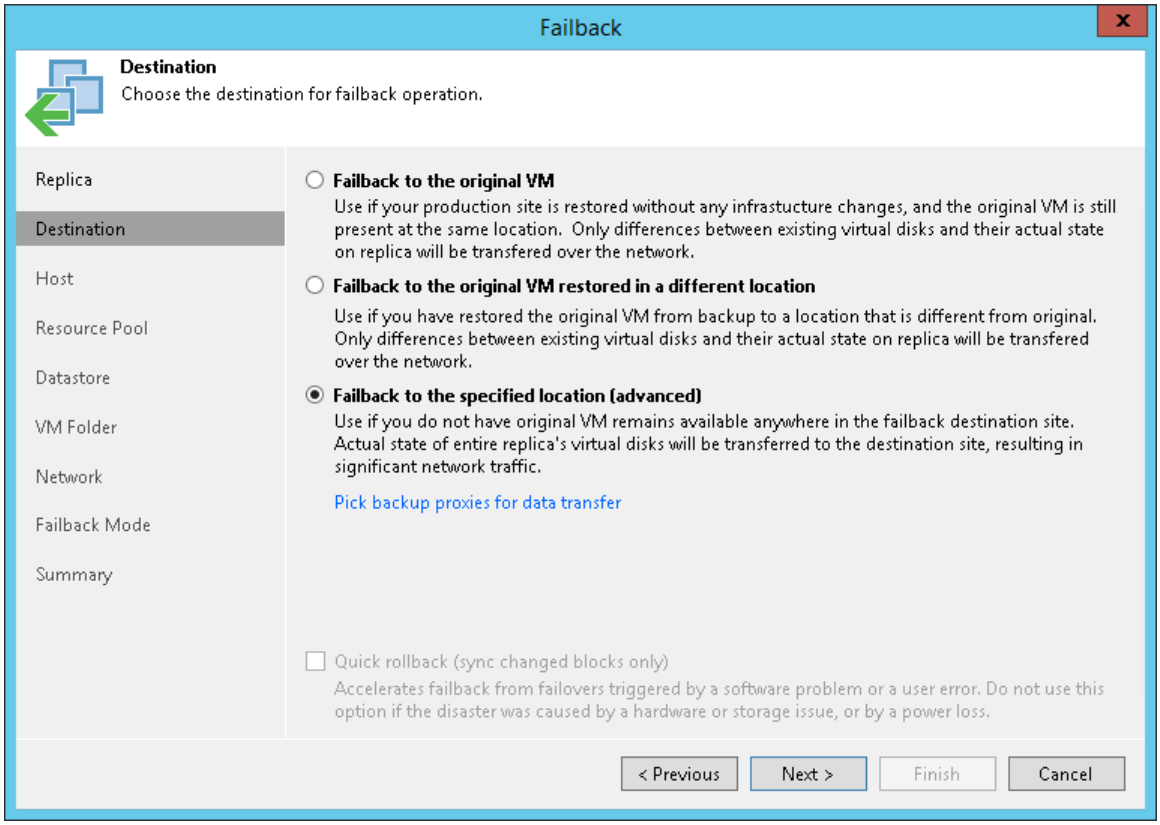
   If you select **Failback to the specified location**, Veeam Backup & Replication will send to the specified location whole VM data, including VM configurations and virtual disk content.

2. To select which backup proxies will be used for data transfer, click **Pick backup proxies for data transfer**. By default, Veeam Backup & Replication selects proxies automatically.

   If you leave automatic proxy selection, Veeam Backup & Replication will check available backup proxies before processing each VM from the VM list. If more than one backup proxy is available, Veeam Backup & Replication analyzes transport modes that the backup proxies can use, the current workload on the backup proxies to select the most appropriate resource for VM processing.

   If you want to assign proxies manually, use the following instructions. If VMs and their replicas reside in different sites, select at least one backup proxy in the production site and one proxy in the disaster recovery site. If VMs and replicas reside in the same site, you can use the same backup proxy as the source and target one. We recommend you to select at least two backup proxies in each site to ensure that failback will be performed in case one proxy fails or looses the network connection.

3. [For failback to the original VMs] If you want to fasten failback, and the original VMs had problems at the guest OS level, select the **Quick rollback** check box. For more information on quick rollback, its requirements and limitations, see Quick Rollback.



Restoring Storage Policies

If the replicated VM was associated with the storage policy, in the failback to original location scenario, Veeam Backup & Replication will associate the restored VM with this storage policy.

When you click **Next**, Veeam Backup & Replication will check storage policies in the virtual environment and compare this information with the information about the replica storage policy. If the original storage policy has been changed or deleted, Veeam Backup & Replication will display a warning. You can select one of the following options:

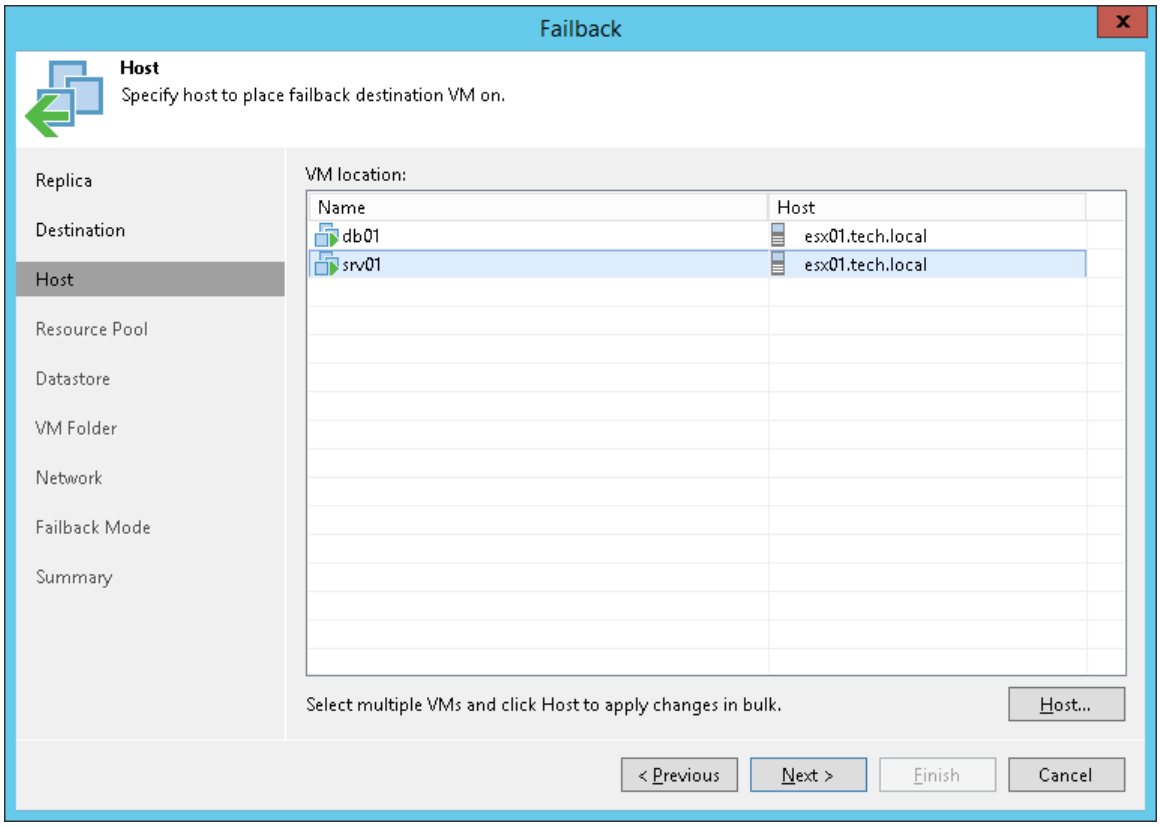
- **Current** – the restored VM will be associated with the profile with which the original VM in the production environment is currently associated.
- **Default** – the restored VM will be associated with the profile that is set as default for the target datastore.
- **Stored** – the restored VM will be associated with the profile that was assigned to the original VM at the moment of replication.

For more information, see Storage Policy Restore.

Step 4. Select Target Host

The **Host** step is available if you have selected the **Failback to the specified location** option at the Destination step.

At the **Host** step of the wizard, specify hosts or clusters where the recovered VMs will be registered. To do this, select the necessary VMs, click **Host** and select a host or cluster.

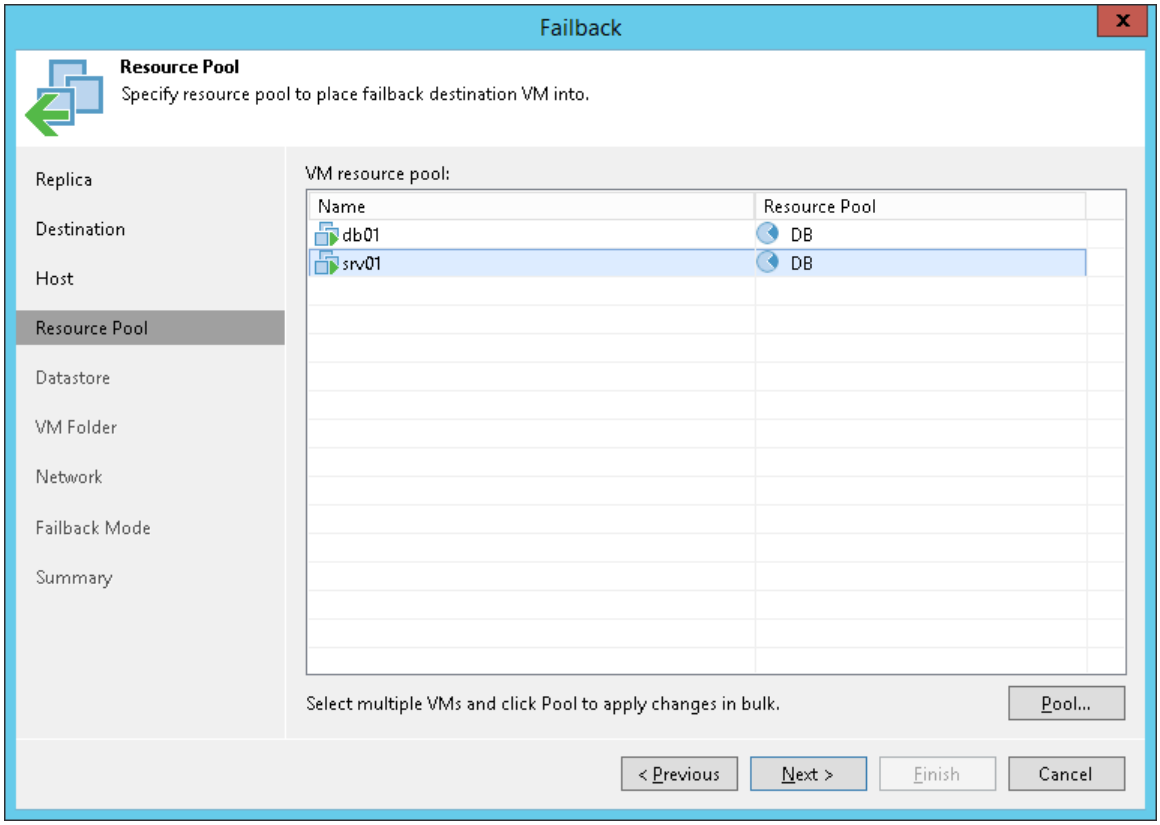


Step 5. Select Target Resource Pool

The **Resource Pool** step is available if you have selected the **Failback to the specified location** option at the Destination step.

At the **Resource Pool** step of the wizard, select resource pools to which the recovered VMs will be added. To do this, select VMs that you want to add to the same resource pool, click **Pool** and select the necessary resource pool in the **Select Resource Pool** window.

As an alternative, you can select a vAPP to which the restored VM will be included. To find the necessary vAPP, at the left bottom corner of the **Select Resource Pool** window, click the resource pool icon ( ) and select *VirtualApp*.



Step 6. Select Target Datastore

The **Datastore** step is available if you have selected the **Failback to the specified location** option at the Destination step.

At the **Datastore** step of the wizard, specify datastores where you want to store configuration files and disk files of VMs that will be recovered. Also, you can change disk types.

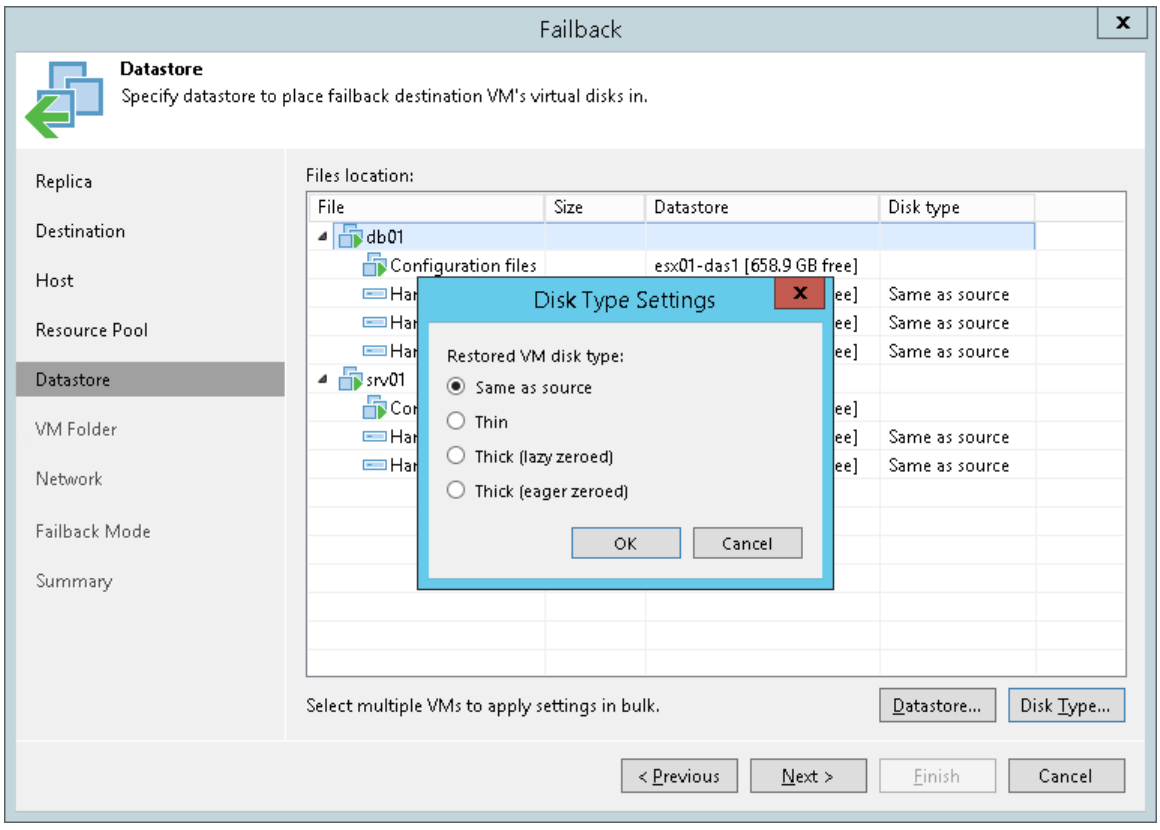
1. To change a datastore where VM files will be stored, select the necessary VMs and click **Datastore**. From the list of available datastores, select the necessary datastore.

   If configuration and disk files of VMs must be placed to different datastores, select files of the necessary type, click **Datastore** and select the necessary datastore.

2. To change disk type settings, select the necessary disk files and click **Disk Type**. In the **Disk Type Settings** window, select the necessary disk type. For more information about disk types, see VMware docs.

   By default, Veeam Backup & Replication preserves disk types of the original VMs.

   > **NOTE:**
   >
   > You can change disk types only for VMs with Virtual Hardware version 7 or later.



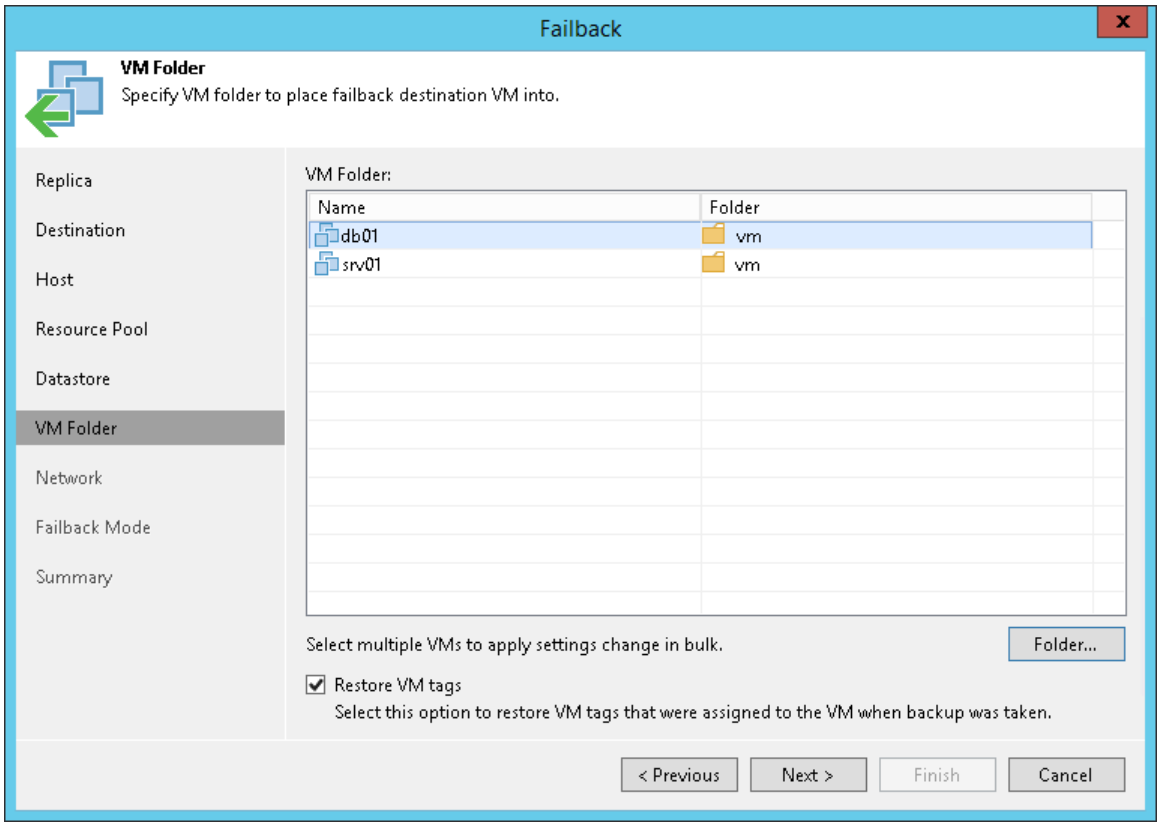
Step 7. Select Target Folder

The **VM Folder** step is available if you have selected the **Failback to the specified location** option at the Destination step.

At the **VM Folder** step of the wizard, specify folders in the target datastores where all files of the recovered VMs will be stored.

If you want the recovered VMs to have the same tags as the source VMs, select the **Restore VM Tags** check box.

> **NOTE**
>
> Mind the following:
>
> - You can select destination folders only if you recover VMs to destinations other than standalone hosts.
> - You can recover VM tags only if you recover VMs to their original locations, and the original VM tags are still available on the source vCenter Server.



Step 8. Select Target Network

The **Network** step is available if you have selected the **Failback to the specified location** option at the Destination step. This step applies if you fail back to VMs recovered to new locations, and if networks in those locations differ from networks in the disaster recovery (DR) site.

At the **Network** step of the wizard create a network mapping table. This table maps networks in the DR site to networks in the site where the recovered VMs will reside. Veeam Backup & Replication will use the network mapping table to update configuration files of VMs on the fly, during the failback process.

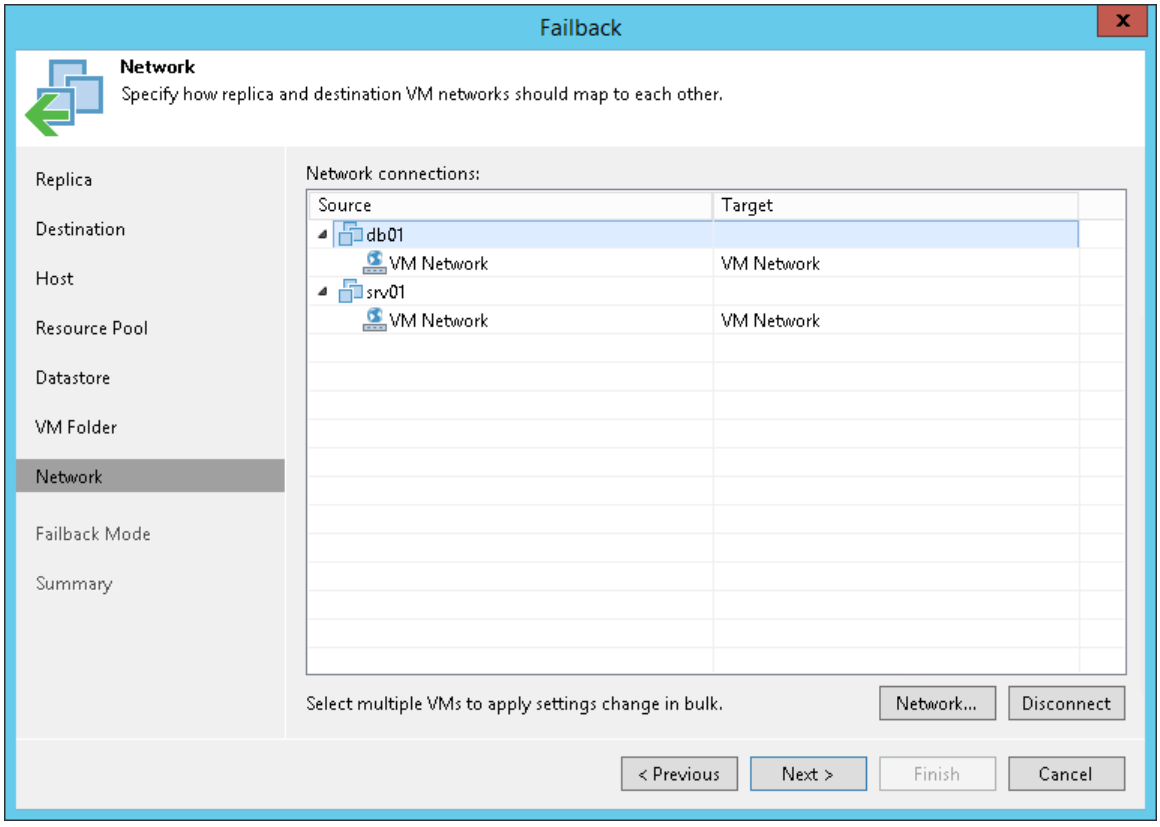
To change networks to which the restored VMs will be connected:

1. In the **Network connections** list, select the necessary VMs and click **Network**.

   If VMs are connected to multiple networks, select networks which you want to map.

2. In the list of available networks, select a network to which the recovered VMs will be connected.

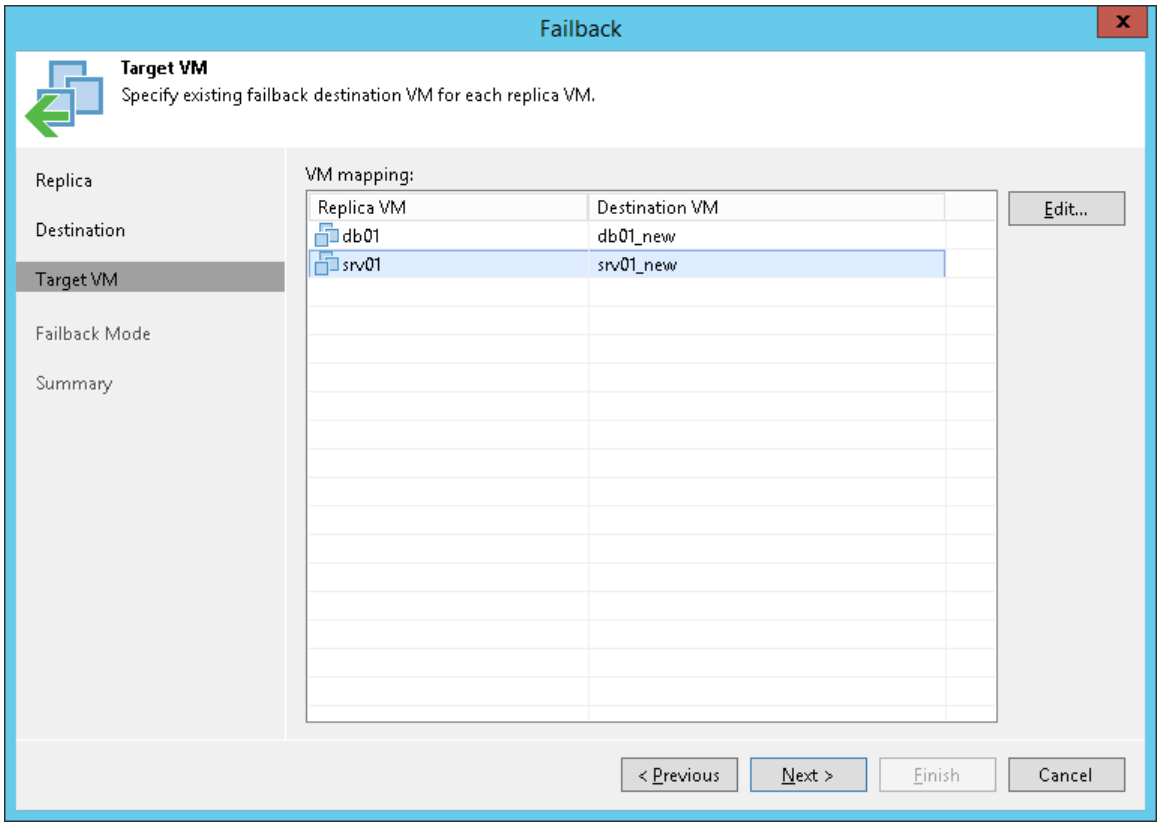
If you do not want to connect the recovered VMs to any virtual network, select the necessary VMs and click **Disconnect**.



Step 9. Map Replicas to Restored VMs

The **Target VM** step is available if you have selected the **Failback to the original VM restored in a different location** option at the Destination step.

At the **Target VM** step of the wizard, specify to which VMs you want to fail back from replicas. These VMs must be already restored from backups in the required location.

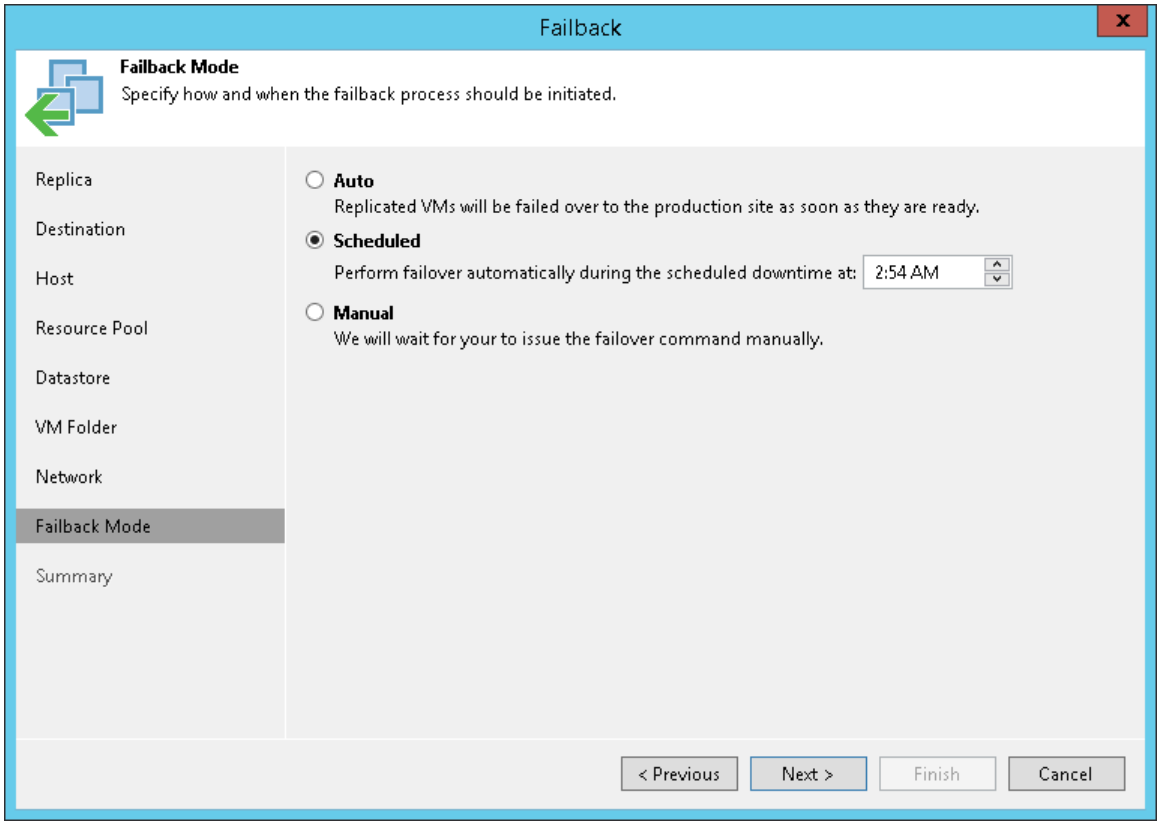


Step 10. Schedule Switch to Production VMs

At the **Failback Mode** step of the wizard, specify when switch from replicas to production VMs must be performed:

- Select **Auto** if you want Veeam Backup & Replication to perform the switch automatically right after the state of the production VMs is synchronized with the state of their replicas.
- Select **Scheduled** if you want Veeam Backup & Replication to perform the switch at a specific time.
- Select **Manual** if you want to perform the switch manually.

If you select the **Scheduled** or **Manual** option, you can further reset/set the scheduled time or switch to the production VM manually. For more information, see Changing Switching Time and Switching Replicas to Production VMs Manually.



Step 11. Review Summary and Finish Working with Wizard

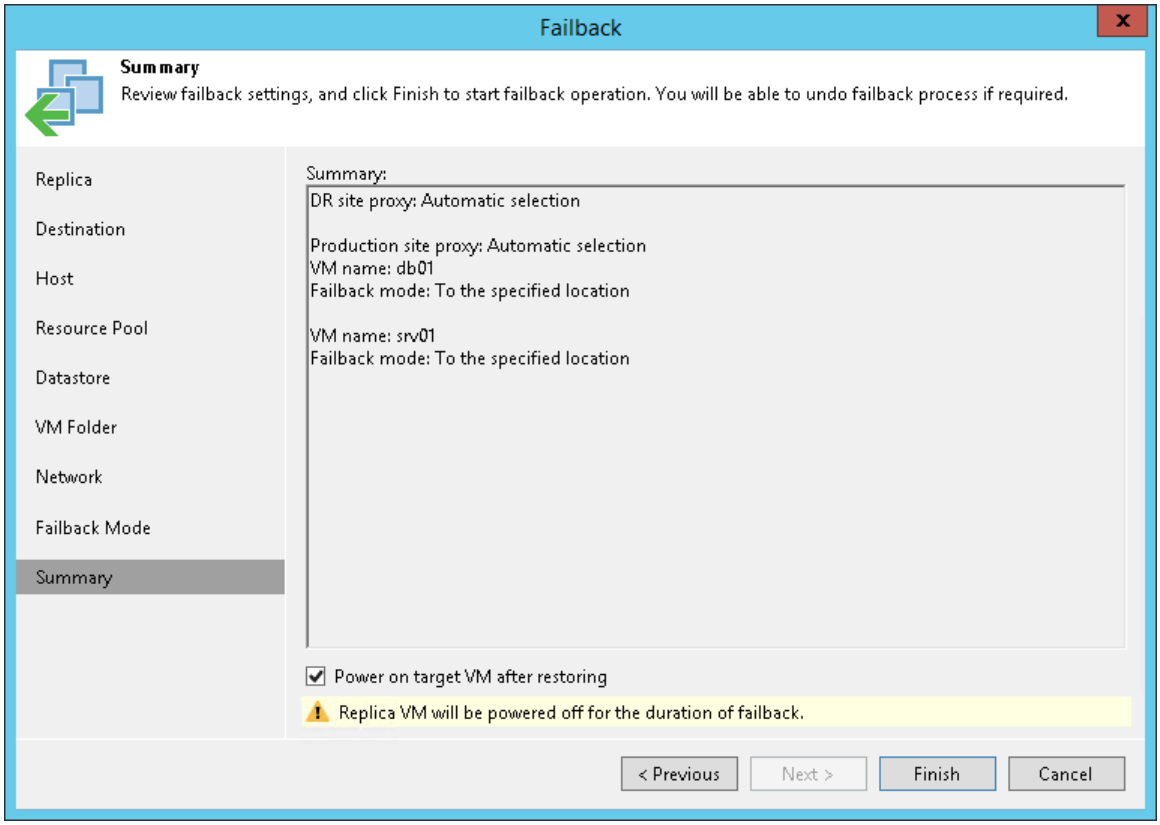
At the **Summary** step of the wizard, review the configured failback settings and click **Finish**.

If you want to power on the production VMs right after the switch to production operation is performed, select the **Power on target VM after restoring** check box.

What You Do Next

Failback is an intermediate step that needs to be finalized. You can finalize failback in the following ways:

- Commit failback
- Undo failback



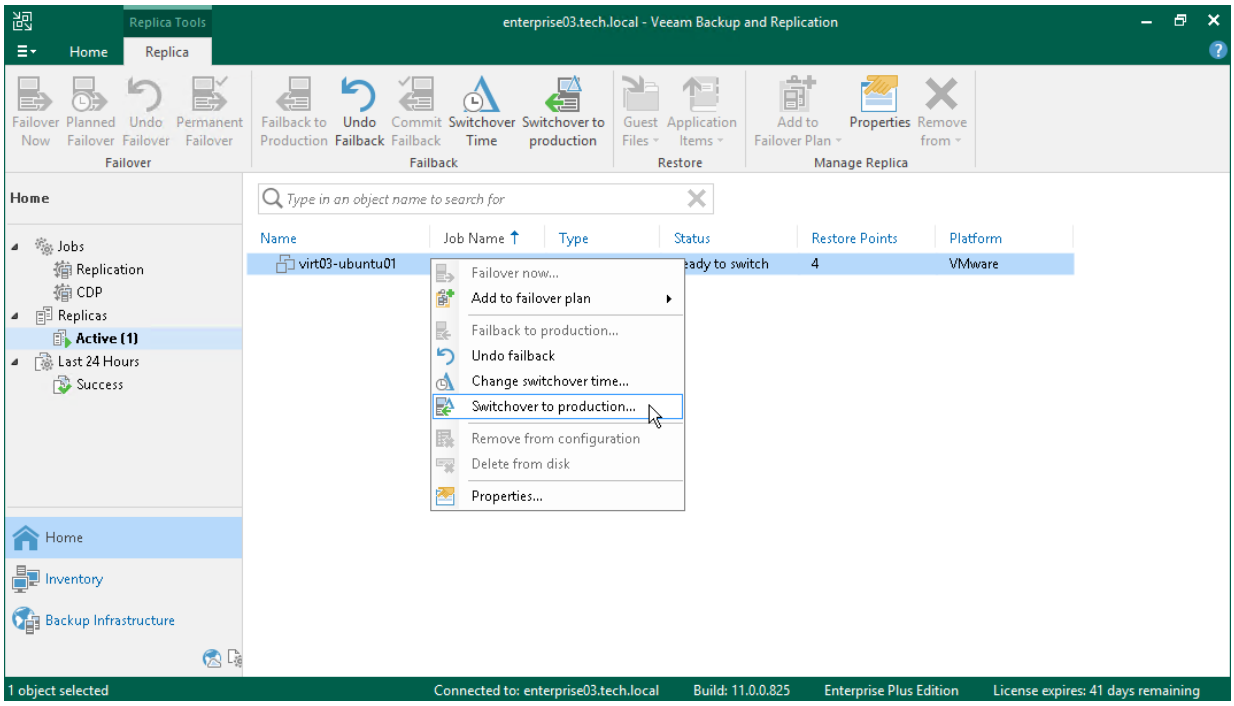
Switching to Production VMs Manually

The following instructions apply if you have selected to switch from replicas to production VMs manually or at the scheduled time at the **Failback Mode** step of the **Failback** wizard.

To switch to a production VM from its replica, do the following:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Replicas** > **Active** node.

3. Right-click a replica in the *Ready to switch* state and select **Switch to production**.



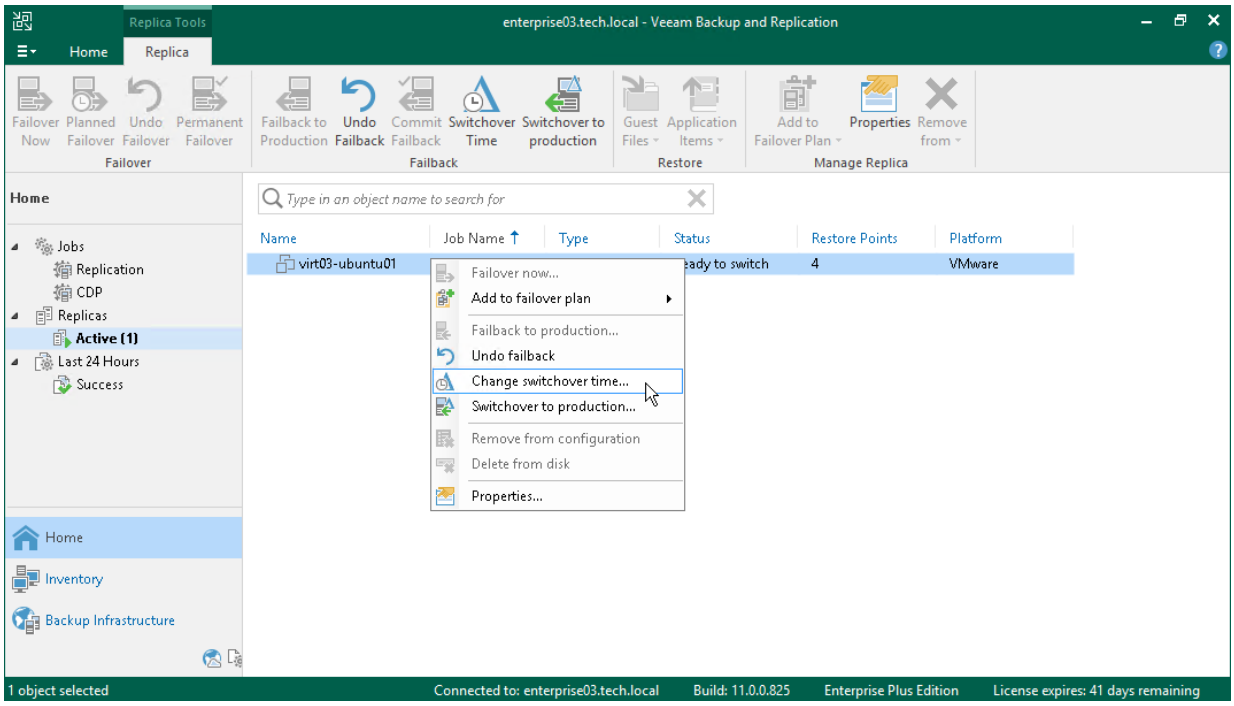
Changing Switching Time

The following instructions apply if you have selected to switch from replicas to production VMs manually or at the scheduled time at the **Failback Mode** step of the **Failback** wizard.

To change the time when Veeam Backup & Replication will switch from replicas to production VMs:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Replicas** > **Active** node.

3. Right-click a replica in the *Ready to switch* state and select **Change switching time.**



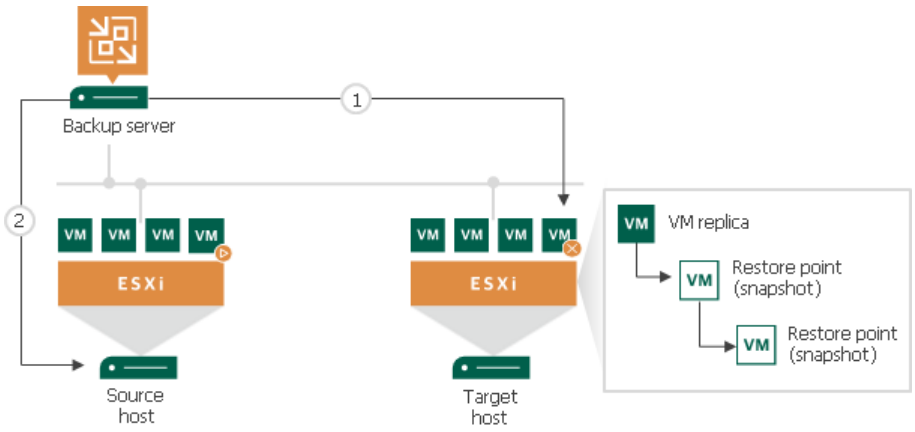
Failback Commit

Failback commit is one of the ways to finalize failback. When you commit failback, you confirm that the VM to which you failed back (the production VM) works as expected. After the commit operation, Veeam Backup & Replication resumes replication activities for the production VM.

The commit failback operation is performed in the following way:

1. Veeam Backup & Replication changes the state of the replica from *Failback* to *Ready*.
2. Further operations depend on whether you have failed back to the original VM or recovered VM:
   - If you have failed back to a VM recovered from a backup or replica, Veeam Backup & Replication reconfigures all existing jobs where the original VM is present and adds the original VM to the list of exclusions. The recovered VM takes the role of the original VM and is included into all jobs instead of the excluded VM. When the replication job starts, Veeam Backup & Replication processes the recovered VM instead of the former original VM.
   - If you have failed back to the original VM, the replication job is not reconfigured. When the replication job starts, Veeam Backup & Replication still processes the original VM.

During failback commit, the failback protective snapshot that saves the pre-failback state of a VM replica is not deleted. Veeam Backup & Replication uses this snapshot as an additional restore point for VM replica. With the pre-failback snapshot, Veeam Backup & Replication needs to transfer fewer changes and therefore puts less load on the network when replication activities are resumed.

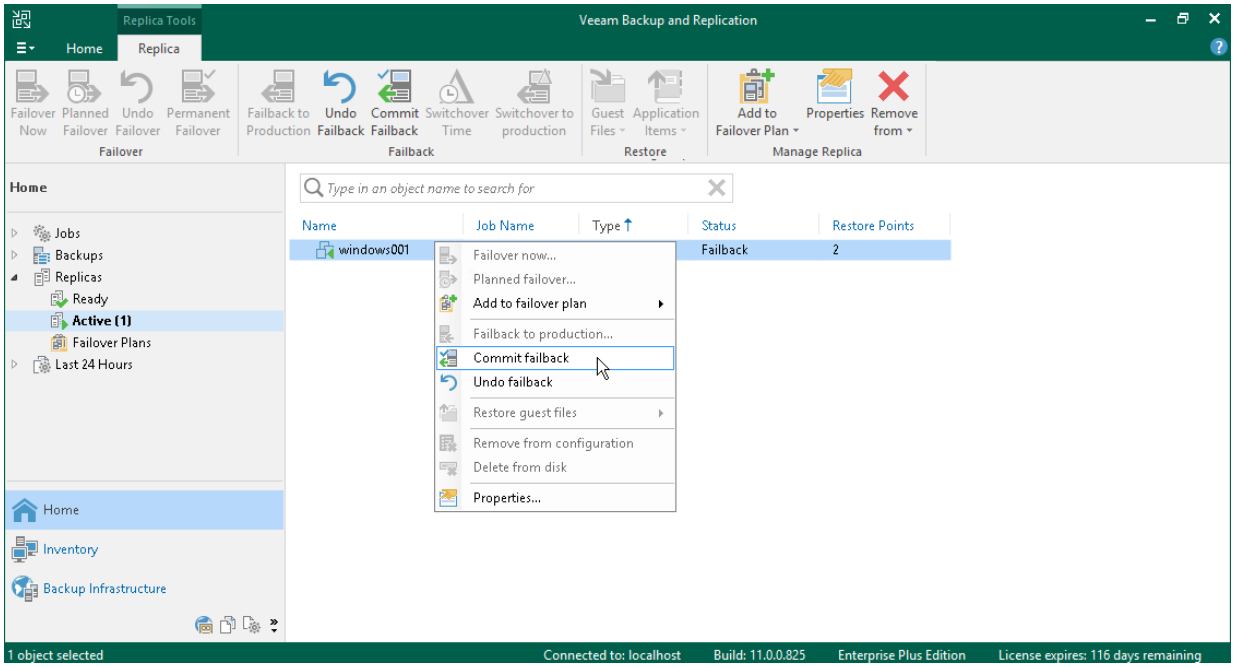


Committing Failback

To commit failback:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Replicas** > **Active** node.

3. In the working area, select the necessary replica and click **Commit Failback** on the ribbon. As an alternative, you can right-click the replica and select **Commit failback**.

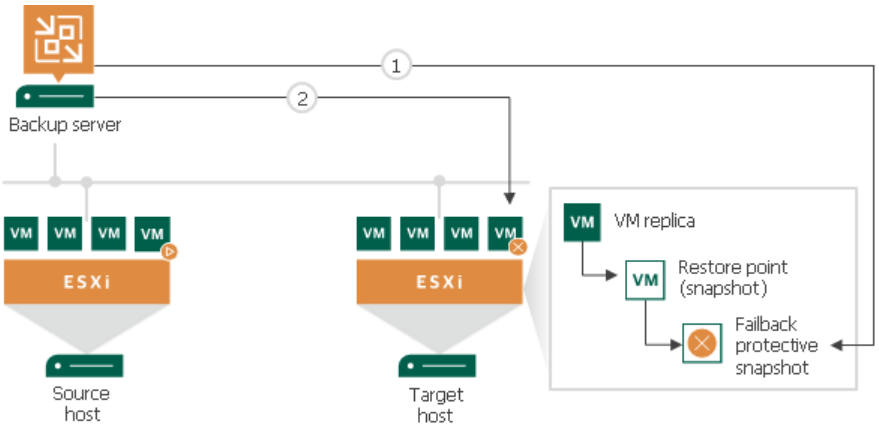


Failback Undo

Failback undo is one of the ways to finalize failback. When you undo failback, you confirm that the VM to which you failed back (the production VM) works in a wrong way and you want to get back to the replica.

The undo failback operation is performed in the following way:

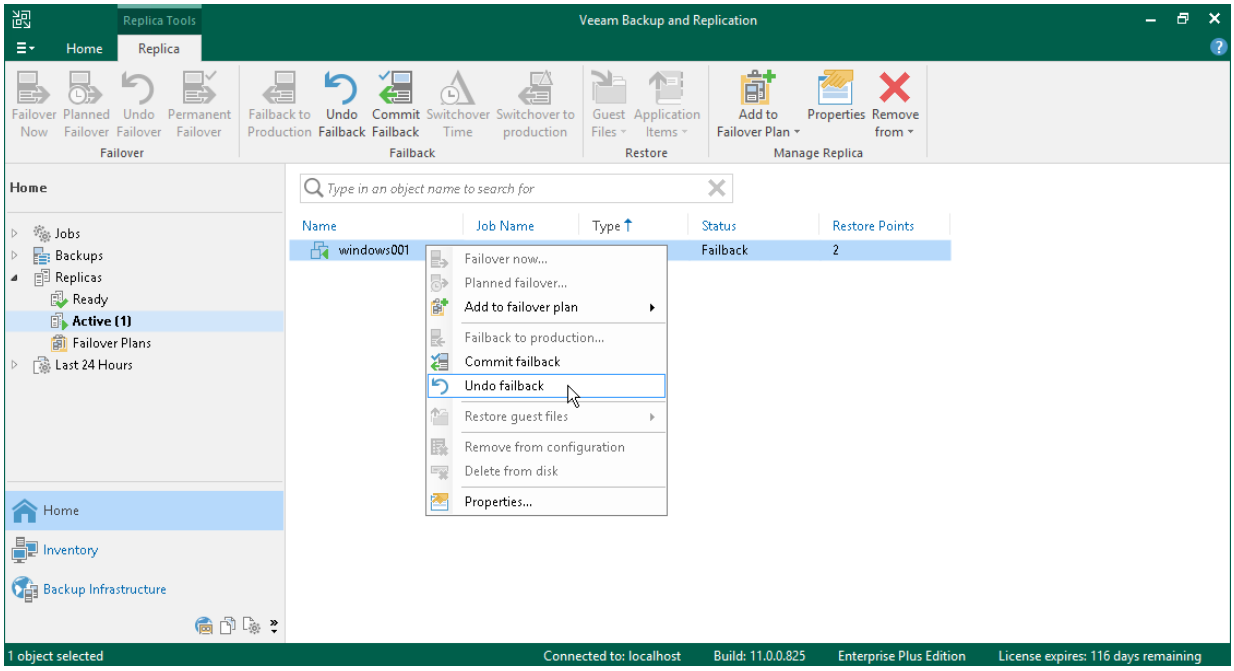
1. Veeam Backup & Replication deletes the protective failback snapshot on the VM replica.
2. Veeam Backup & Replication powers on the VM replica and changes the VM replica state from *Failback* to *Failover*.



Undoing Failback

To undo failback:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Replicas** > **Active** node.
3. In the working area, select the necessary replica and click **Undo Failback** on the ribbon. As an alternative, you can right-click the replica and select **Undo failback**.



Continuous Data Protection (CDP)

Continuous data protection (CDP) is a technology that helps you protect mission-critical VMware virtual machines when data loss for seconds or minutes is unacceptable. CDP also provides minimum recovery time objective (RTO) in case a disaster strikes because VM replicas are in a ready-to-start state.

Data Replication

First, CDP creates VM replicas and, then, keeps these replicas up to date.

CDP constantly replicates I/O operations performed on VMs. To read and process I/O operations in transit between the protected VMs and their underlining datastore, CDP uses vSphere APIs for I/O filtering (VAIO) that gives an option not to create snapshots. Because CDP is always on and does not create snapshots, it allows reaching a lower recovery point objective (RPO) compared to the snapshot-based replication — near-zero RPO which means almost no data loss.

Data of I/O operations is stored on the target datastore and relates to short-term restore points. The short-term restore points allow you to recover a VM to a state back to seconds or minutes ago (depending on the RPO that you specify) in case a disaster strikes. Information about short-term restore points is maintained in a special journal. This journal stores records about short-term restore points for maximum 24 hours. If you want to recover a VM to an older state, Veeam Backup & Replication allows you to create additional restore points that contain a VM state back to hours or days ago. Such restore points are called long-term restore points.

To facilitate replication over slow connections, Veeam Backup & Replication optimizes traffic transmission. It filters out unnecessary data blocks such as duplicate data blocks, zero data blocks and compresses replica traffic.

To replicate a VM, you need to configure required backup infrastructure components and create a CDP policy.

Data Recovery

To recover a VM to a short-term or long-term restore point, you need to fail over to its replica.

When you fail over to a replica, the replica takes over the role of the original VM. After your original VM is repaired, you can fail back to it and transfer all changes that occurred to replica to the original VM. If your original VM cannot be repaired, you can perform permanent failover, that is, permanently switch from the original VM to the VM replica and use this replica as the original VM. For more information, see Failover and Failback for CDP.

Requirements and Limitations

CDP has a set of requirements and limitations.

Requirements

The following requirements apply to CDP:

- CDP is included in the Veeam Universal License. When using a legacy socket-based license, the Enterprise Plus edition is required.
- VMware vSphere edition must be VMware vSphere Standard Edition or higher. VMware vSphere Essentials Kits editions are not supported.
- For supported platforms, see Platform Support.
- For supported source and target datastores, see Veeam CDP Source and Target.
- Cisco HyperFlex 4.5 (2a) and later can be used as a source or target only with with VMware vSphere 7.0 U2 and later. With the previous versions of VMware vSphere, Cisco HyperFlex is not supported.
- All hosts in a cluster must be of the same major version: 7.x or 6.x (6.5, or 6.7, or a combination of 6.5 and 6.7 is supported). In turn, all clusters managed by the same vCenter Server must also be of the same major version.
- The target cluster must support hardware versions of VMs on the source cluster (VMs that you plan to protect).
- VMs on the same cluster must be protected with CDP only by one backup server.
- The backup server must have at least 16 GB RAM.
- VMs that you plan to protect must not have snapshots at the moment when the CDP policy starts for the first time. Further, you can add these VMs to backup jobs that will create snapshots. However, you must not create snapshots by other means, for example, manually or by another software. This may disrupt further replication operations or cause loss of important data.
- During the initial synchronization, Veeam Backup & Replication needs double space allocation for VMs with thick-provisioned disks on the target host. After the initial synchronization finishes, Veeam Backup & Replication frees half of the allocated space for such VMs. This is a part of mechanism required to support protection of disks newly added to the original VMs.
- Network between infrastructure components required for CDP must be minimum 1 Gbps. We recommend to use 10 Gbps or faster and MTU 9000.
- If you enable VM encryption on VMware site, make sure that the **Allow I/O filters before encryption** parameter is set to *True* in the **Default encryption properties** storage policy component.
- For VMware CDP proxy requirements, see VMware CDP Proxy.

Limitations

When planning to use CDP, mind the following limitations:

- CDP works only for powered on VMs. However, for VMs with IDE hard drives you must power off VMs before the initial synchronization. This is required to apply a new storage policy to VMs. For more information, see How CDP Works.

- One VM can be processed only by one CDP policy.

- CDP is not supported for VMs to which VM storage policies with multiple datastore specific rule sets apply. If you need to define datastores to be used for placement of the VMs, you can add a tag rule for vSAN instead of a tag based placement datastore specific rule. For more information on how to create tag rules for vSAN, see VMware Docs.

  - VM replicas can be powered on only using the failover operation; manual power on is disabled.

  - Testing CDP replicas with Virtual Labs is currently not supported.

  - On the target host, Veeam Backup & Replication does not allow you to migrate VM replicas using VMware vSphere Storage vMotion. Note that host vMotion is supported.

  - Shared disks, physical RDM and SCSI bus sharing are not supported. Note that vRDM disks are supported.

  - The maximum supported number of disks for one VM is 50. The maximum number of disks for one ESXi host is 500.

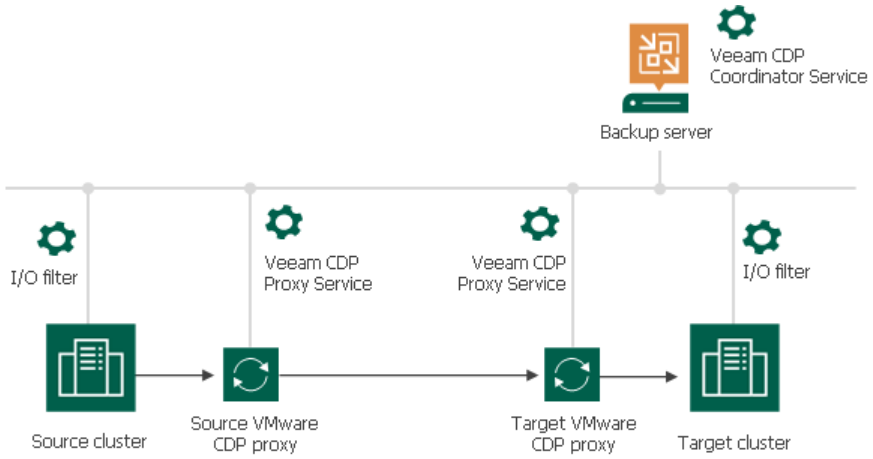
  - The maximum number of long-term restore points for one disk is 95.

  - Data encryption rules do not apply to traffic between the ESXi host and VMware CDP proxy. For more information on the rules, see Enabling Data Encryption.

Backup Infrastructure for CDP

The following backup infrastructure components are required for CDP:

- Backup server
- Source and target hosts
- VMware CDP proxies



Backup Server

The backup server is the configuration, administration and management core of the backup infrastructure. The backup server runs the Veeam CDP Coordinator Service. This service coordinates replication and data transfer tasks, and controls resource allocation. We recommend you to place the backup server in the target site or as a separate unit.

For more information on the backup server, see Backup Server.

Source and Target Hosts

The source and the target hosts are two terminal points between which replicated VM data is moved. The source and target hosts must be a part of the same cluster or two different clusters. In turn, clusters must be managed by the same vCenter Server or two different vCenter Servers connected to the same backup server. For more information on requirements to the hosts and how to add vCenter Servers to the backup infrastructure, see the Requirements and Limitations and Adding VMware vSphere Servers sections.

The source and target hosts perform the following tasks:

- The source host reads VM disk data, reads and processess I/O operations and sends data to source proxies. The data is sent uncompressed.
- The target host receives data from target proxies and saves this data to VM replicas on the datastore. Also, the target host manages VM replicas: creates replicas, retains restore points and so on.

I/O Filter on Hosts

To be able to use hosts for CDP, you must install the I/O filter on each cluster where hosts reside. After you install the I/O filter on the clusters, Veeam Backup & Replication automatically installs the filter on all hosts added to the clusters. For more information on how to install the filter, see Installing I/O Filter.

It is the I/O filter that reads and processes I/O operations in transit between the protected VMs and their underlining datastore and that sends/receives data to/from VMware CDP proxies. Also, the filter communicates with the Veeam CDP Coordinator Service on the backup server and notifies the service that the backup infrastructure must be reconfigured if any proxy becomes unavailable. This I/O filter is built on the basis of vSphere API for I/O filtering (VAIO).

VMware CDP Proxies

A VMware CDP proxy is a component that operates as a data mover and transfers data between the source and target hosts. We recommend you to configure at least two VMware CDP proxies: one (source proxy) in the production site and one (target proxy) in the disaster recovery site.

The source and target VMware CDP proxies perform the following tasks:

- The source proxy prepares data for short-term restore points from data received from the source host, compresses and encrypts the data (if encryption is enabled in the network traffic rules). Then sends it to the target proxy.
- The target proxy receives the data, decompresses and decrypts it, and then sends to the target host.

For more information on VMware CDP proxies, their requirements, limitations and deployment, see VMware CDP Proxy.

How CDP Works

This section describes how CDP works during replication. To learn how CDP works during data recovery, see Failover and Failback for CDP.

CDP workflow during replication is divided into two parts: backup infrastructure component configuration and data transfer.

During the configuration, Veeam Backup & Replication configures the required backup infrastructure components. Veeam Backup & Replication also reconfigures the components if something changes in the infrastructure. During data transfer, Veeam Backup & Replication creates short-term and long-term restore points by sending disk data blocks and changes made to them. For more information on restore points, see Replication Chain.

Backup infrastructure component configuration and data transfer are constant processes.

Component Configuration Algorithm

The following steps apply during the initial configuration, that is, when a CDP policy starts for the first time or after the policy is enabled:

1. Veeam CDP Coordinator Service reads policy settings from the configuration database and creates a list of VM tasks to process. For every VM added to the CDP policy, Veeam Backup & Replication creates a new task.
2. Veeam CDP Coordinator Service checks that required backup infrastructure components are available.
3. Veeam CDP Coordinator Service queries information about VMs added to the CDP policy and virtualization hosts from the vCenter Server.
4. Veeam CDP Coordinator Service requests vCenter Server to create on the target host empty VM replicas with the same configuration as original VMs, but with empty virtual disks. vCenter Server registers the created VM replicas.
5. Veeam CDP Coordinator Service requests vCenter to apply the Veeam CDP Replication storage policy to virtual disks of VMs on the source ESXi hosts. This storage policy adds the component that is required for CDP and that gets data of all I/O operations. For more information on storage policies, see VMware Docs.

   The Veeam CDP Replication storage policy itself is created on the vCenter Server when you install the I/O filter. For more information on how to install the filter, see Installing I/O Filter.
6. Veeam CDP Coordinator Service selects which VMware CDP proxies will be used for data transfer and sets a number of rules for data transfer, such as network traffic throttling rules and so on.

   If you select automatic proxy selection when configuring the CDP policy, Veeam Backup & Replication analyzes the current workload on CDP proxies and selects a VMware CDP proxy according to the following priority rules (starting from the most preferable one):

   - VMware CDP proxy on a physical machine
   - VMware CDP proxy on a VM located on the same ESXi host — that is, the source proxy on the host where original VMs are located, the target proxy on the host where VM replicas are located
   - Other VMware CDP proxies

   For more information on how to specify the proxy selection mode, see Specify Data Transfer and Replica Settings.

7. Veeam CDP Coordinator Service sends to the backup infrastructure components configurations required for CDP. This configuration includes such information as RPO, short-term and long-term retention settings.

After the initial configuration finishes, Veeam Backup & Replication starts monitoring the backup infrastructure. If something changes in the infrastructure or CDP policy settings, Veeam Backup & Replication reconfigures the components. Consider the following examples:

- If a VMware CDP proxy becomes unavailable, Veeam CDP Coordinator Service on the backup service gets a notification that this proxy is no longer available. Then, the service selects another proxy.

  > **NOTE**
  >
  > If you add new proxies to the backup infrastructure, Veeam Backup & Replication does not use these proxies for the already created CDP policies — that is, Veeam Backup & Replication does not reconfigure the infrastructure. If you have selected automatic proxy selection for a CDP policy and want to use the newly added proxies, disable and then enable the CDP policy. If you have selected the proxies manually, edit the CDP policy settings and add the required proxies.

- If virtual disks are added to original VMs, Veeam CDP Coordinator Service requests vCenter to create the disks on the target host, applies the storage policy and selects VMware CDP proxies to transfer disk data.
- If new VMs are added to the CDP policy, Veeam CDP Coordinator Service requests vCenter to create VM replicas with empty disks, applies the storage policy to the VM disks and selects which VMware CDP proxies to use for data transfer.
- If the original VMs were migrated to other host or datastore using VMware vSphere vMotion, Veeam CDP Coordinator Service analyzies how data will be transferred after the migration and selects VMware CDP proxies to optimize data transfer.

Component reconfiguration requires Veeam CDP Coordinator Service to be working. If the coordinator goes down, existing CDP policies still work, create and remove short-term restore points. Long-term restore points are not created and removed because it is the coordinator who manages them. However, if any of the components goes out of service, for example, VMware CDP proxy goes offline or VMware vSphere vMotion changes the infrastructure, CDP policies start failing until Veeam CDP Coordinator Service is repaired.

Data Transfer Algorithm

Data transfer starts right after Veeam Backup & Replication configures the backup infrastructure components (performs the initial configuration). Data transfer differs during the initial synchronization and during the incremental synchronization.

As a rule, the initial synchronization is performed when disk data is sent to the target host for the first time. During the initial synchronization, Veeam Backup & Replication sends data for full copies of virtual disks and creates the very first restore points.

During the incremental synchronization, Veeam Backup & Replication mainly sends data for incremental changes made to virtual disks and creates short-term and long-term restore points. For more information on restore points and files created for VM replicas, see Replication Chain.

Data Transfer Algorithm During Initial Synchronization

The following steps apply to data transfer during the initial synchronization:

1. On the source host, I/O filter reads all data from VM disks and sends it to the source VMware CDP proxies.

   As the original VMs are still running, data for the already transferred data blocks can change. I/O filter intercepts these changes and sends them to the proxies. Sending changes instead of whole changed data blocks helps minimize traffic sent over the network.

2. The Veeam CDP Proxy Service on the source proxies compresses, encrypts the received data and sends it to the target proxy.

3. The Veeam CDP Proxy Service on the target proxies decompresses and decrypts the received data. Then sends data to the target host.

4. The I/O filter on the target host saves the received data to virtual disks. The saved data relates to the very first long-term restore point. This restore point is crash-consistent.

   > **NOTE**
   >
   > Veeam Backup & Replication starts creating long-term and short-term restore points only after the initial synchronization finishes.

After the initial synchronization finishes, Veeam Backup & Replication starts the incremental synchronization.

Data Transfer Algorithm During Incremental Synchronization

During the incremental synchronization, Veeam Backup & Replication creates short-term and long-term restore points. To create short-term restore points, Veeam Backup & Replication intercepts changes made by transactions on VM disks and sends these changes to the target datastore. Changes are constantly transferred and are saved to transaction logs on the target datastore. To create long-term restore points, Veeam Backup & Replication uses data of short-term restore points and saves the restore points to delta disks. Long-term restore points are created by schedule.

The following steps apply when Veeam Backup & Replication transfers data for short-term restore points:

1. On the source host, I/O filter intercepts data of all I/O operations and sends this data to the source VMware CDP proxies.

2. Once in the RPO, the Veeam CDP Proxy Service on the source proxies prepares data required for a short-term restore point. For this, Veeam CDP Proxy Service gets the latest state of the data that the source VMware CDP proxies has accumulated.

3. The source Veeam CDP Proxy Service compresses, encrypts data and sends it to the target proxy.

4. The target Veeam CDP Proxy Service decompresses and decrypts the received data. Then sends data to the target host.

5. The I/O filter on the target host saves the received data to transaction logs.

The following steps apply when long-term restore point creation is scheduled:

1. [If application-aware image processing is enabled for the CDP policy] Veeam Backup & Replication connects to VM guest OSes, deploys non-persistent runtime components or connects/deploys persistent agent components on VM guest OSes and performs in-guest processing tasks such as quiescing applications on the VM and creating a consistent view of application data. For more information on guest agent and non-persistent components, see Non-Persistent Runtime Components and Persistent Agent Components.

2. On the target datastore, the I/O filter forms a long-term restore point using data of short-term restore points created since the creation of the previous long-term restore point. Data for the new long-term restore point is saved to a delta disk.

CDP Replication Chain

A replication chain is a sequence of files that allows you to roll back a VM replica to a specific point in time during failover. Veeam Backup & Replication creates replication chains for each VM added to a CDP policy.

The replication chain contains short-term and long-term restore points. Short-term restore points allow you to roll back VM replica data to a state created seconds or minutes ago, while long-term restore points — to a state created hours or days ago. Short-term restore points are always crash-consistent, long-term restore points can be crash-consistent or application-consistent depending on how long-term restore points are configured in CDP policy settings.

The replication chain is stored on the target datastore in the *<replica_VM_name>* folder. The replication chain consists of the files of the following types (only the key file types are listed):

- VMX — the configuration file of the VM replica.

  The replication chain contains one .vmx file, other files from the list are created per virtual disk.

- VMDK — virtual disk files that store contents of replica hard disk drives.

  On the datastore, you can see files under the following names:

  - *<disk_name>*.vmdk — files that store full copies of virtual disks, that is, store base disk data. These files are created during the initial synchronization and relate to the very first long-term restore point. This restore point is crash-consistent.
  - *<disk_name>*-*<index>*.vmdk — files that store incremental changes made to virtual disks, that is, store delta disk data. Delta disk files relate to the following:
    - Long-term restore points. In this case, files are created according to the long-term schedule configured in CDP policy settings. These files remain unchanged till they are retained by long-term retention.
    - Technical points. These technical points are required during short-term retention. The technical points creation is connected to the transaction log creation. The technical point is created when the transaction log file reaches its maximum size or when transaction log is created after the short-term retention is reached. For more information on transaction log files, see the TLOG description.
  - [For VSAN and VVOL starting from Veeam Backup & Replication 11a (build 11.0.1.1261)] *<disk_name>*-*<index>*.tlog.vmdk — see the TLOG description.
  - *<disk_name>*-interim.vmdk — files for protective virtual disks. Changes made to virtual disks will be written in these files when you perform failover.

  > **NOTE**
  >
  > Although VMDK files look like VMware snapshot files, they are not real snapshots. These files are created by I/O filter installed on the target host.

- VMFD — files that contains metadata for disks.
- TLOG — transaction log files that store incremental changes made to virtual disks during RPO set in CDP policy settings. These files relate to short-term restore points that are created once in RPO set in CDP policy settings. One transaction log file stores data for multiple short-term restore points.

  [For VSAN and VVOL starting from Veeam Backup & Replication 11a (build 11.0.1.1261)]
  Veeam Backup & Replication saves transaction log files as virtual disk files (*<disk_name>*-*<index>*.tlog.vmdk) to avoid the VSAN limit of 255 GB and the VVOL limit of 4 GB for a file.

New transaction log files are created in the following cases:

- [For Veeam Backup & Replication 11a (build 11.0.1.1261) and later] When a transaction log file reaches a specific size: 2 TB on VMFS datastores, 512 GB on VSAN and VVOL.
- [For Veeam Backup & Replication prior to 11a (build 11.0.1.1261)] When a .tlog file reaches a specific size: 2 TB on VMFS datastores, 255 GB on VSAN or 100 MB on VVOL.
- When a long-term restore point is created.
- [For Veeam Backup & Replication 11a (build 11.0.1.1261) and later] When the short-term retention is reached. For example, if you set short-term retention to 1 hour, a new transaction log file will be created every hour.
- [For Veeam Backup & Replication prior to 11a (build 11.0.1.1261)] When one quarter of the short-term retention is reached. For example, if you set short-term retention to 1 hour, a new transaction log file will be created every 15 minutes.
- After failback commit.

- META — files that contain metadata for transaction log files.

To roll back a VM replica to a specific point in time, the chain of files created for the replica must contain files with data for the base disk (*<disk_name>*-flat.vmdk) and a set of files that contain incremental changes for disks (*<disk-name>*-*<index>*.vmdk + .tlog or .tlog.vmdk for VSAN and VVOL). If any file in the replication chain is missing, you will not be able to roll back to the necessary state. For this reason, you must not delete files from the datastore manually. Instead, you must specify retention policy settings that will let you maintain the desired number of files. For more information on retention policies and how to configure them, see Retention Policies.

Retention Policies

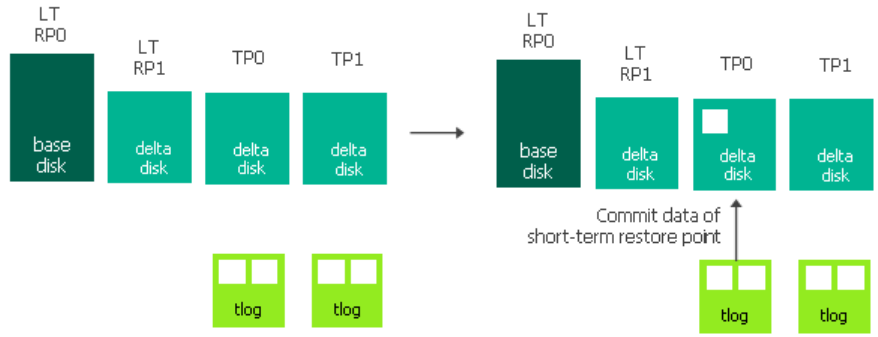
A retention policy defines for how long Veeam Backup & Replication must store restore points for VM replicas. Veeam Backup & Replication offers two retention policy schemes:

- Short-term retention
- Long-term retention

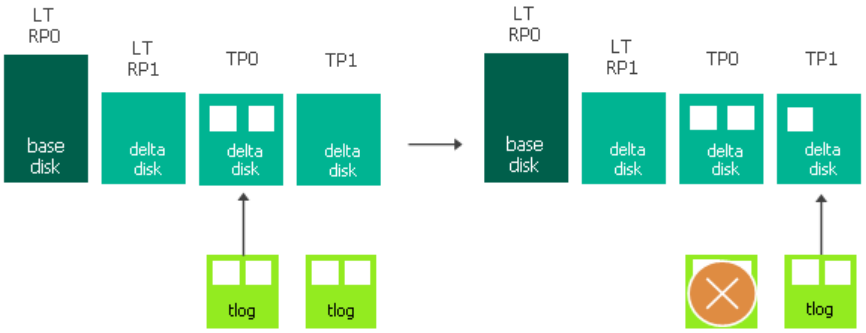
Short-term Retention

Veeam Backup & Replication retains short-term restore points for the number of hours specified in CDP policy settings. When the retention period is exceeded, Veeam Backup & Replication transforms the replication chain in the following way. The example shows how short-term retention works for a VM replica with one virtual disk.

1. Veeam Backup & Replication checks whether the replication chain contains outdated short-term restore points.
2. If an outdated restore point exists, Veeam Backup & Replication commits data for the short-term restore point from the transaction log file into the nearest technical point (TP). For more information on technical points, short-term restore points, long-term restore points and when they are created, see CDP Replication Chain.

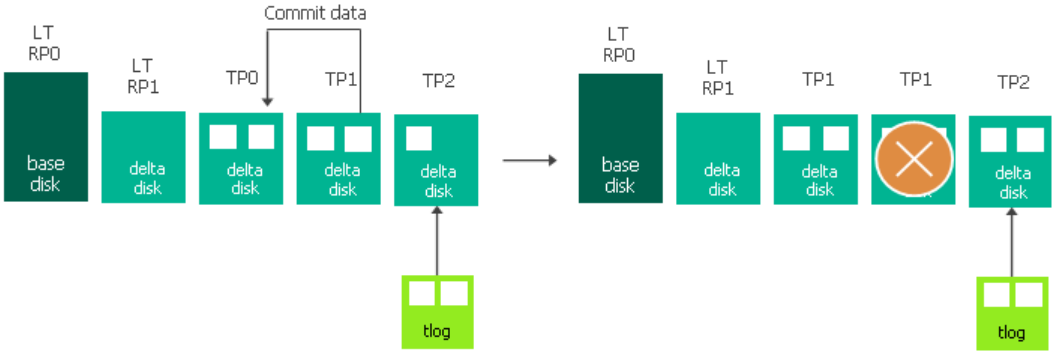


3. If the transaction log file does not contain data for further short-term restore points, Veeam Backup & Replication deletes the transaction log file as redundant — its data has already been committed into the technical delta disk file.



4. After a technical point remains without the related transaction log file, Veeam Backup & Replication considers this technical point outdated and commits data of a newer technical point or a long-term restore point into the outdated technical point. Mind the following:
   - The newer technical restore point must not have the related transaction log.

- The data of the long-term restore point can be committed into a technical point but not vice versa. Long-term restore points remain unchanged till they are retained by the long-term retention.



NOTE

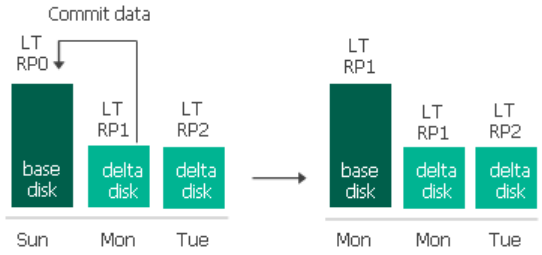
Veeam Backup & Replication can store short-term restore points for a longer period than specified in the short-term retention policy. This period is maximum 25% longer.

Long-term Retention

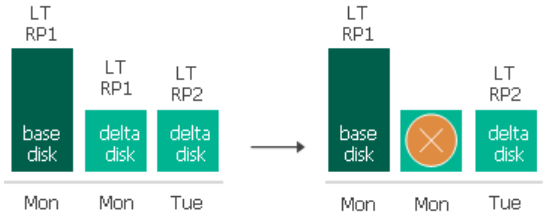
Veeam Backup & Replication retains long-term restore points for the number of days specified in CDP policy settings. When the retention period is exceeded, Veeam Backup & Replication transforms the replication chain in the following way. The example shows how long-term retention works for a VM replica with one virtual disk.

1. Veeam Backup & Replication checks whether the replication chain contains outdated long-term restore points.
2. If an outdated restore point exists, Veeam Backup & Replication rebuilds the file of an outdated long-term restore point (LT RP) to include data of a newer long-term restore point. To do that, Veeam Backup & Replication commits into the base disk file data from the earliest delta disk file that relates to a long-term restore point. This way, the base disk file 'moves' forward in the replication chain.

   For more information on technical points, short-term restore points, long-term restore points and when they are created, see CDP Replication Chain.



3. Veeam Backup & Replication removes the earliest delta disk file from the chain as redundant — this data has already been committed into the base disk file.



Guaranteed Delivery

In CDP, data delivery between any two backup infrastructure components is guaranteed by means of the TCP protocol. However, there can be situations when Veeam Backup & Replication is not able to send all data changes generated during CDP in time. For example, if RPO is small, a VM generates a lot of changes, but the infrastructure performance is not enough; or a VMware CDP proxy gets out of work because of a power outage, and all data changes stored on it are lost. This makes data inconsistent and leads to the loss of restore points.

CDP has a special mechanism and tools that return data into the consistent state. The I/O filter on the source host has a mechanism that tracks data blocks that have changed — change tracking (CT). The I/O filter deletes data block addresses from the list of changed blocks only after a confirmation message that data blocks were successfully saved to the VM replica is received from the target host. In addition, VMware CDP proxies also store data changes until a confirmation message is received from the target host.

How Veeam Backup & Replication uses these tools depend on whether issues occur on the source or target site:

- **Source site**. If the source host does not send the data changes because of the high system load, Veeam Backup & Replication waits till the load decreases. Then, it gets the list of changed data blocks, reads data blocks from the disk and sends the data to the target host. The target host saves the received data to the datastore. Until this moment, the data on the target host remains inconsistent. As a result, there are "gaps" in the restore point journal and you are not able to restore a VM to some points in time. However, after the target host saves the data, the CDP policy resumes the creation of consistent restore points.

  If the source proxy gets out of work, Veeam Backup & Replication selects another VMware CDP proxy and then behaves as described above except for waiting for the load decrease.

- **Target site**. If the target proxy gets out of work, Veeam Backup & Replication selects another VMware CDP proxy and requests data changes from the source VMware CDP proxy. The proxy sends the changes to the target host once again, and the data on the target host becomes consistent almost immediately, so that you are able to restore a VM to any point in time.

  If the connection between the target VMware CDP proxy and the target host is lost, Veeam Backup & Replication checks the host state. If the host is in the maintenance mode or has disappeared from the cluster, Veeam Backup & Replication starts writing data changes to VM replicas using another target ESXi host (provided that the host exists and is connected to the datastore where VM replicas are stored). If the host is not in the maintenance mode or has not disappeared, Veeam Backup & Replication considers that these are temporary problems with the network and sends data changes again after some time.

Replica Seeding and Mapping

Replica seeding and mapping are technologies that help reduce the amount of traffic sent over a network. With these technologies, Veeam Backup & Replication does not have to transfer all of VM data from the source host to the target host across the sites during the initial synchronization. For more information on the initial synchronization, see How CDP Works.

You can use seeding and mapping in the following scenarios:

- **Seeding**

  Configure replica seeding if, in a backup repository located in the disaster recovery (DR) site, you have backups of VMs that you plan to replicate. During replication, Veeam Backup & Replication will restore VMs from these backups and will synchronize the state of the restored VMs with the latest state of the original VMs. Then Veeam Backup & Replication will use these restored VMs as replicas.

  For more information on how to create backups that can be used as "seeds" for replica, see Creating Replica Seeds.

- **Mapping**

  Configure replica mapping if, on the host in the DR site, you have ready-to-use copies of the original VMs. These can be restored VMs or replicas created by other CDP policies. Veeam Backup & Replication will synchronize the state of these ready-to-use VMs with the latest state of the original VMs and will use these VMs as replicas.

You can also configure both replica seeding and replica mapping in the same CDP policy. For example, if a policy includes 2 VMs, you can use seeding for one VM and map the other VM to an existing VM.

> **IMPORTANT**
>
> If seeding or mapping is enabled in a policy, all VMs in the policy must be covered with seeding or mapping. If a VM neither has a seed, nor is mapped to an existing VM, it will be skipped from processing.

Algorithm for Seeding

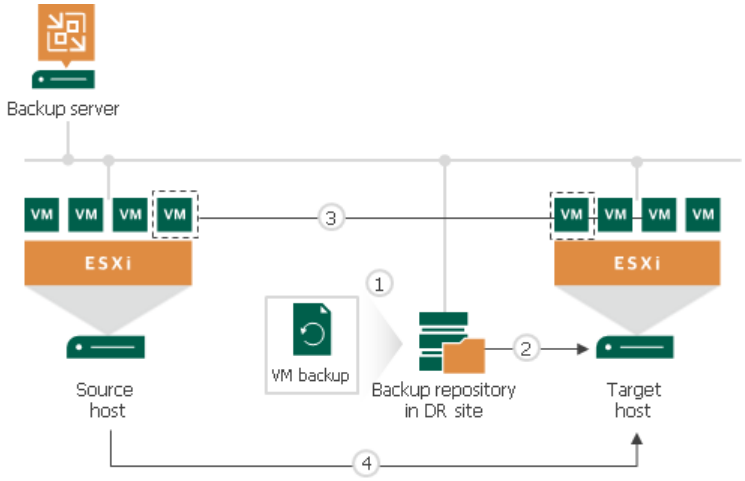
Replica seeding includes the following steps:

1. As a preparatory step for replica seeding, you need to create a backup of a VM that you plan to replicate. For more information on how to create a backup that will be used as a "seed" for replica, see Creating Replica Seeds.
2. When you create a CDP policy, you should point it to a backup repository in the DR site. During the initial synchronization, Veeam Backup & Replication accesses the backup repository where the replica seed is located, and restores the VM from the backup. The restored VM is registered on the target host in the DR site. Files of the restored VM are placed to the location you specify as the replica destination datastore.

   Virtual disks of a replica restored from the backup preserve their format (that is, if the original VM used thin provisioned disks, virtual disks of the VM replica are restored as thin provisioned).
3. Veeam Backup & Replication synchronizes the restored VM with the latest state of the original VM.

   After successful synchronization, in the **Home** view in the Veeam Backup & Replication console, under **Replicas** node you will see a VM replica with two restore points. One point will contain the state of the VM from the backup file; the other point will contain the latest state of the original VM you want to replicate.

4. During incremental synchronization, Veeam Backup & Replication transfers only incremental changes in a regular manner.



Replica seeding dramatically reduces traffic sent over WAN or slow connections because Veeam Backup & Replication does not send the full contents of the VM image. Instead, it transmits only differential data blocks.

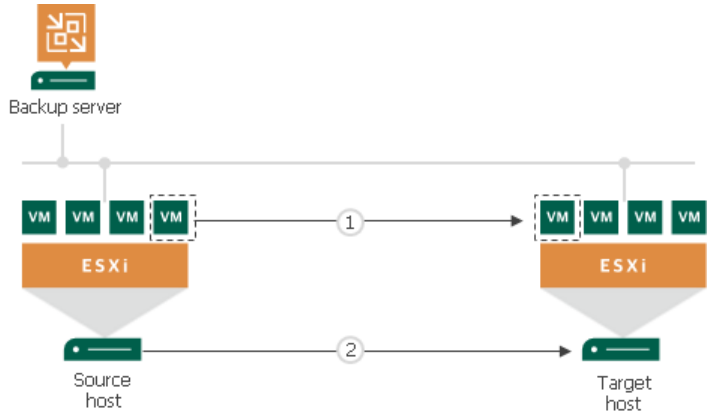
Algorithm for Mapping

Replication to a mapped VM is performed in the following way:

1. The first step differs depending on which VM you have selected for mapping:
   - If you have selected a regular VM, Veeam Backup & Replication calculates the differences between the original and mapped VM.
   - If you have selected a snapshot VM replica, Veeam Backup & Replication deletes all restore points and delta disks and then calculates the differences between the original and mapped VM.
   - If you have selected a CDP VM replica, Veeam Backup & Replication imports all restore points of this replica and then calculates the differences between the original and mapped VM. Note that if disk sizes of the original and mapped VM differ, Veeam Backup & Replication will delete all restore points of the mapped VM.
2. To synchronize the state of the mapped VM with the state of the original VM, Veeam Backup & Replication sends the calculated changes to the mapped VM.

   The first and second steps take place during the initial synchronization.
3. During the incremental synchronization, Veeam Backup & Replication transfers only incremental changes in a regular manner.

After the successful initial synchronization, in the **Home** view of Veeam Backup & Replication, under **Replicas** node you will see a VM replica with restore points. If you have selected for mapping a regular VM or snapshot VM replica, you will see two restore points: one restore point will contain the latest state of the mapped VM, the other will contain the state of the original VM. If you have selected a CDP VM replica, you will see all restore points of the mapped VM plus one restore point that will contain the state of the original VM.



Installing I/O Filter

To be able to protect VMs with CDP, you must install I/O filter on each cluster where the VMs that you plan to protect reside and where VM replicas will reside. For more information on the filter, see Source and Target Hosts.

To install the filter on a specific cluster, open the **Inventory** view. In the inventory pane, navigate to the **Virtual Infrastructure** > **VMware vSphere** > **vCenter Servers** > *<vCenter Server Name>* > *<Cluster Name>* node and right-click it. In the menu, select **Install I/O filter**.

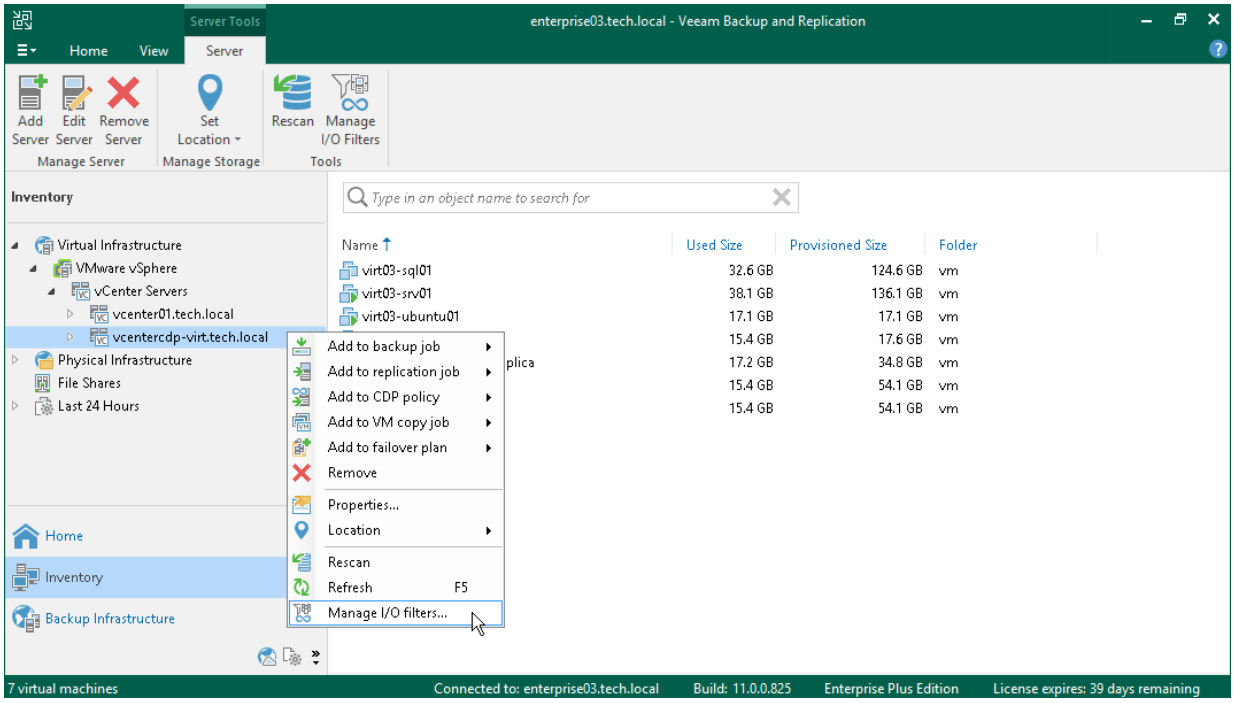
> **IMPORTANT**
>
> If vSphere LifeCycle Management is enabled for your cluster, you need first follow the instructions from this section and then follow the instructions from this Veeam KB article.

To install I/O filter on multiple clusters in a vCenter Server, use the **I/O Filter Management** wizard.

Step 1. Launch I/O Filter Management Wizard

To launch the **I/O Filter Management** wizard, do one of the following:

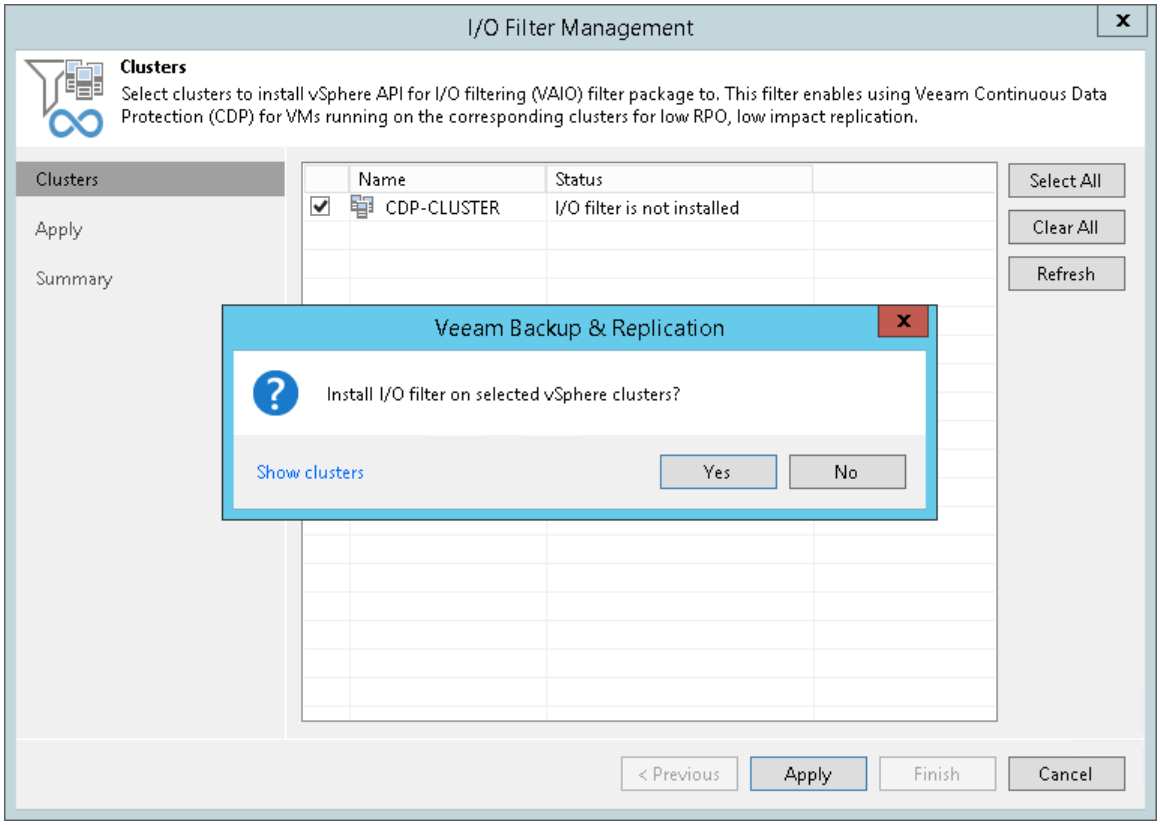
- Open the **Backup Infrastructure** view. In the inventory pane, navigate to the **Managed Servers** > **VMware vSphere** > **vCenter Servers** > *<vCenter Server Name>* node and right-click it. Select **Manage I/O filters**.
- Open the **Backup Infrastructure** view. In the inventory pane, navigate to the **Managed Servers** > **VMware vSphere** > **vCenter Servers** > *<vCenter Server Name>* node. On the ribbon, click **Manage I/O Filters**.
- Open the **Inventory** view. In the inventory pane, navigate to the **Virtual Infrastructure** > **VMware vSphere** > **vCenter Servers** > *<vCenter Server Name>* node and right-click it. Select **Manage I/O filters**.
- Open the **Inventory** view. In the inventory pane, navigate to the **Virtual Infrastructure** > **VMware vSphere** > **vCenter Servers** > *<vCenter Server Name>* node. On the ribbon, click **Manage I/O Filters**.



Step 2. Select Clusters

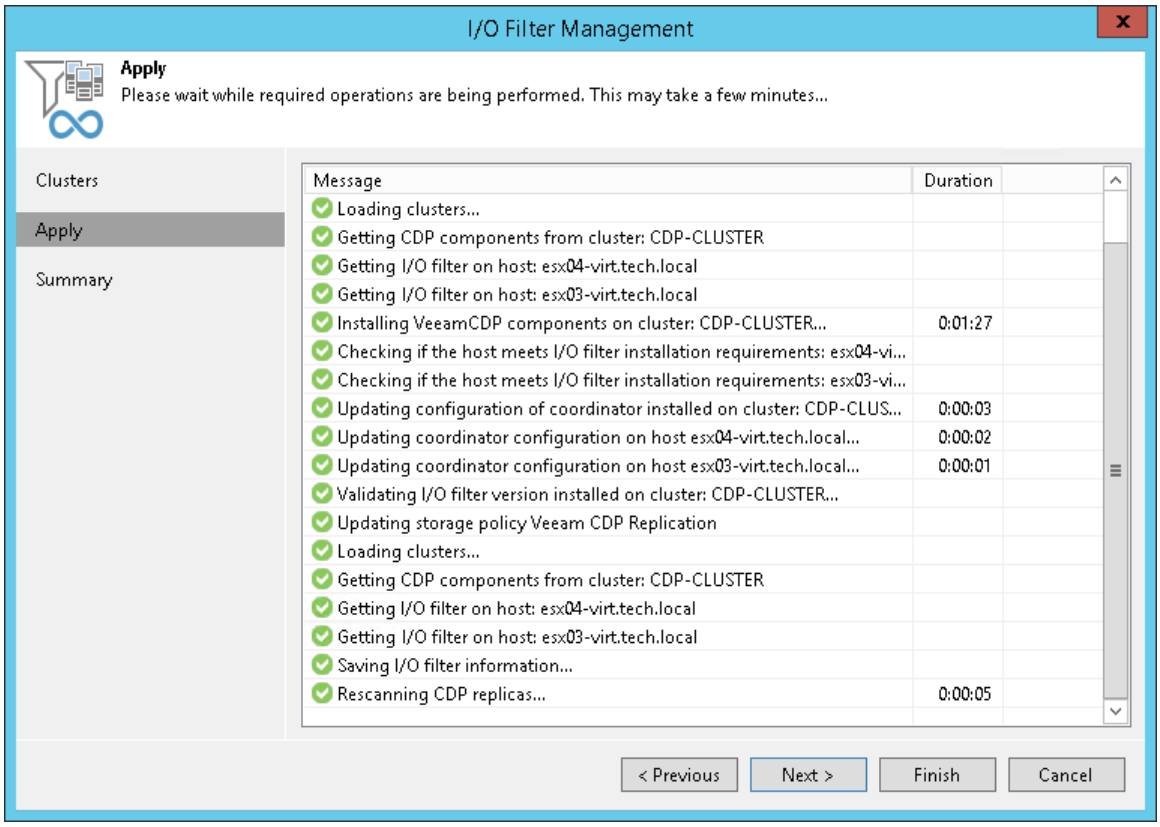
At the **Clusters** step of the wizard, select check boxes near clusters on which you want to install I/O filter.

If you select check boxes near clusters where the filters are already installed, Veeam Backup & Replication will update the filters. If you clear check boxes, Veeam Backup & Replication will delete the I/O filter from these clusters.



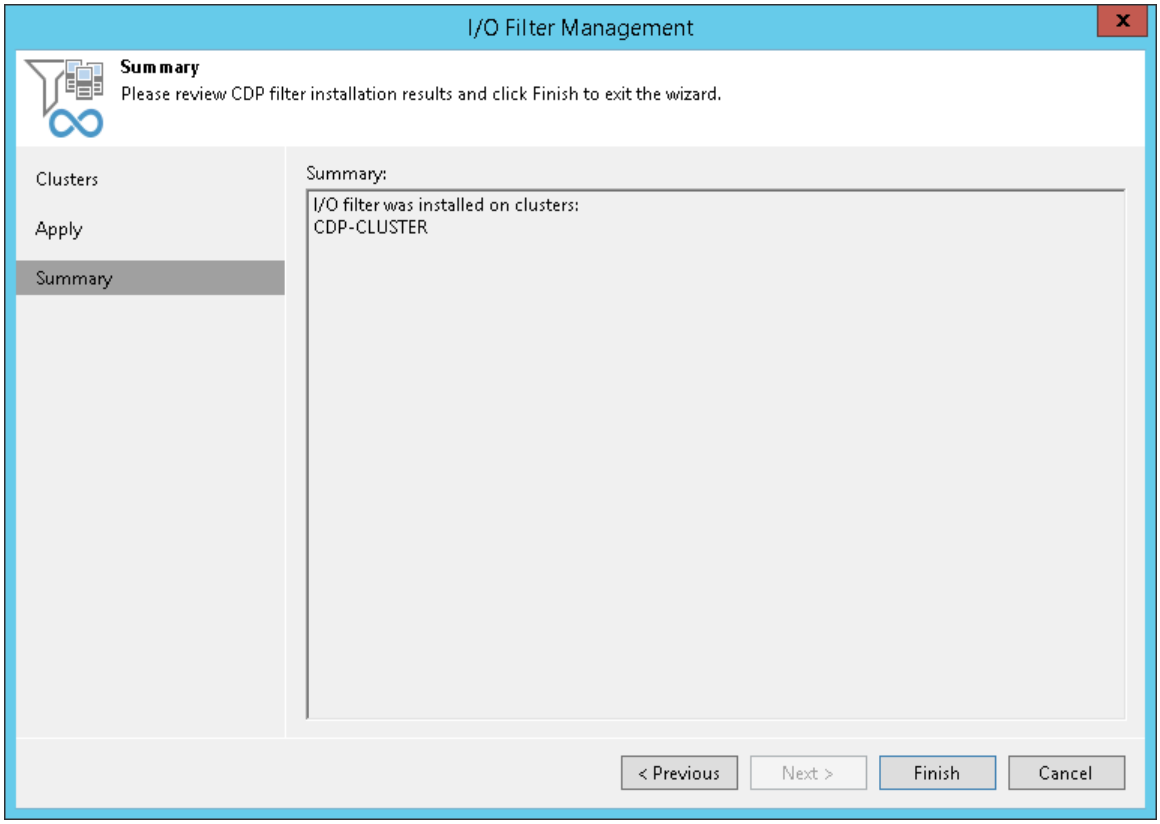
Step 3. Apply Filter Settings

At the **Apply** step of the wizard, wait till Veeam Backup & Replication installs I/O filter. Click **Next**.



Step 4. Finish Working with Wizard

At the **Summary** step of the wizard, review on which clusters I/O filter is installed and click **Finish** to exit the wizard.



Upgrading and Uninstalling I/O Filter

Veeam Backup & Replication allows you to upgrade or uninstall I/O filter from clusters using the Veeam Backup & Replication console.

To upgrade or uninstall the I/O filter from clusters:

1. Make sure that you have disabled or deleted all CDP policies as described in the Disabling and Deleting Policies section.
2. [When uninstalling the filter] If you have manually assigned the Veeam CDP Replication storage policy to VMs that are parts of the clusters, or replicas are still present in the Veeam Backup & Replication configuration database, you must change the storage policy for these VMs.

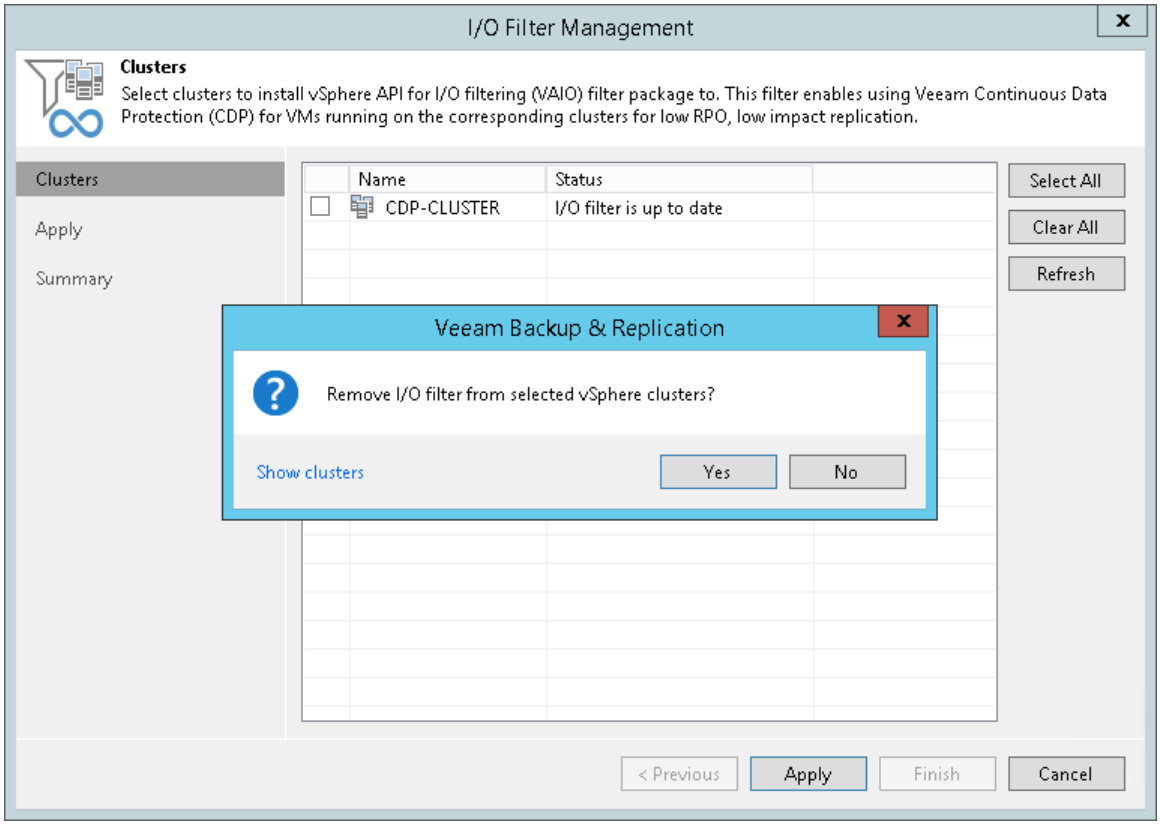
   For more information on how to change storage policies, see VMware Docs. To see the list of replicas, open the **Home** view. In the inventory pane, click the **Replicas** node.
3. If VMware vSphere Distributed Resource Schedule (DRS) is disabled for the clusters, place the hosts of the clusters in the maintenance mode.

   You can simultaneously place all hosts in the maintenance mode or place them one by one for each cluster. Veeam Backup & Replication will uninstall/upgrade the filter only on the hosts in the maintenance mode. That is why you must repeat the procedure (place a host in the maintenance mode and launch installation/upgrade) for each host in a cluster. When Veeam Backup & Replication uninstalls/upgrades the filter on the last host in the cluster, it also uninstalls/upgrades the filter on the cluster.

   If DRS is enabled and the hosts use the shared storage, you do not have to put hosts in the maintenance mode manually. VMware vSphere will automatically place hosts in the maintenance mode one by one and migrate VMs from one host to another.
4. Follow the instructions provided in the Installing Filters section. At the **Clusters** step of the wizard, do the following:
   - To upgrade the filter, make sure that check boxes are selected near the necessary clusters.
   - To uninstall the filter, clear the check boxes near the necessary clusters.

> **NOTE**
>
> If the upgrade attempt has failed, fix the problems listed in the wizard and launch the upgrade again. During the next upgrade operation, you may see the following error: *"The specified key, name or identifier 'vibUrl' already exists. The VIB contains the same filter as the one to be upgraded"*. If this is the only error, it can be ignored. The error is shown because during the previous upgrade filters were already upgraded on some components. The remaining components are upgraded on the current run. To make sure that upgrade completed successfully, launch the upgrade wizard once again.

As an alternative, you can uninstall the I/O filter from a specific cluster. To do this, open the **Inventory** view. In the inventory pane, navigate to the **Virtual Infrastructure** > **VMware vSphere** > **vCenter Servers** > *<vCenter Server Name>* > *<Cluster Name>* node and right-click it. Select **Uninstall I/O filter**.



Creating CDP Policies

To protect VMs with CDP, you must configure a CDP policy. The CDP policy defines which VMs to protect, where to store VM replicas, how often create short-term and long-term restore points, and so on. One CDP policy can process one or multiple VMs.

To create a CDP policy use the **New CDP Policy** wizard.

Before You Begin

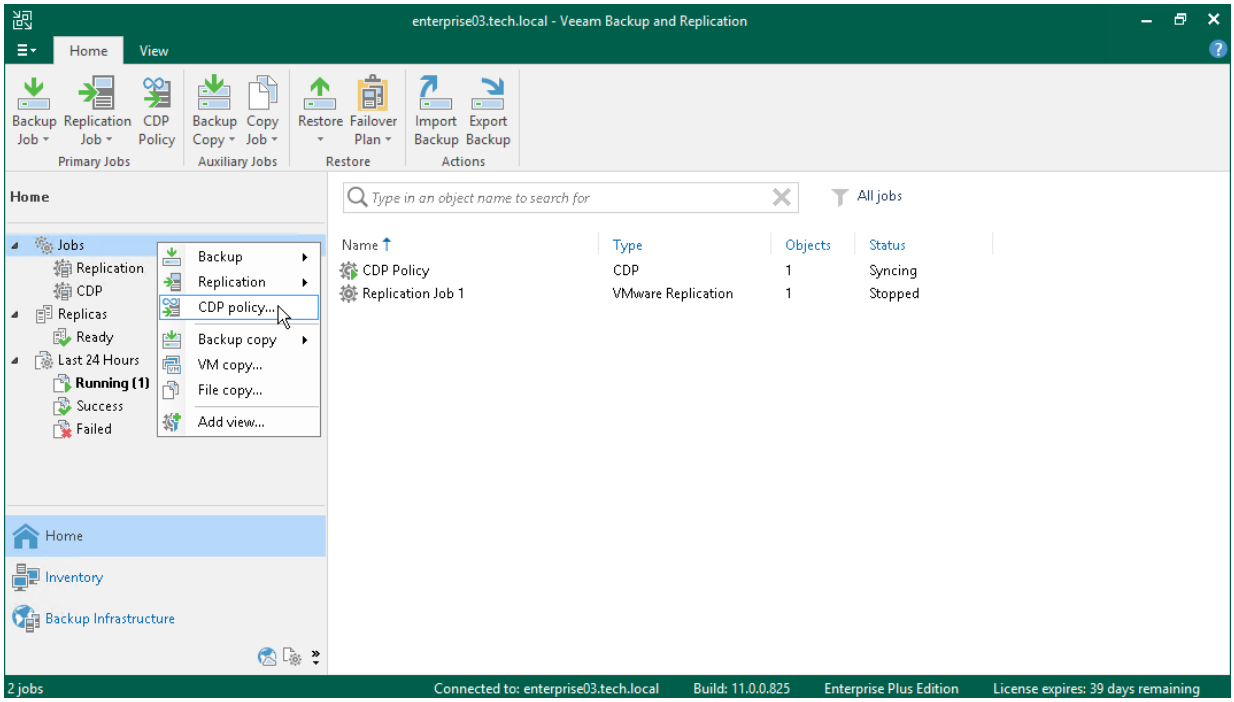
Before you create a CDP policy, check the following prerequisites:

- Check that all the required components are added to the backup infrastructure. For more information on the required components, see Backup Infrastructure.
- I/O filter must be installed on each cluster where VMs that you plan to protect reside. For more information on how to install the filter, see Installing I/O Filter.
- If you plan to use replica seeding, you must create a seed as described in the Creating Replica Seeds section.

Step 1. Launch New CDP Policy Wizard

To launch the **New CDP Policy** wizard, do one of the following:

- Open the **Home** view. On the ribbon, click **CDP Policy**.
- Open the **Home** view. In the inventory pane, right-click **Jobs** and select **CDP Policy**.
- Open the **Inventory** view. In the working area, right-click VMs that you want to replicate. Select **Add to CDP policy** > **New job** if you want to create a new CDP policy, or **Add to CDP policy** > *<Policy Name>* if you want to add VMs to an existing CDP policy.



Step 2. Specify Policy Name and Advanced Settings

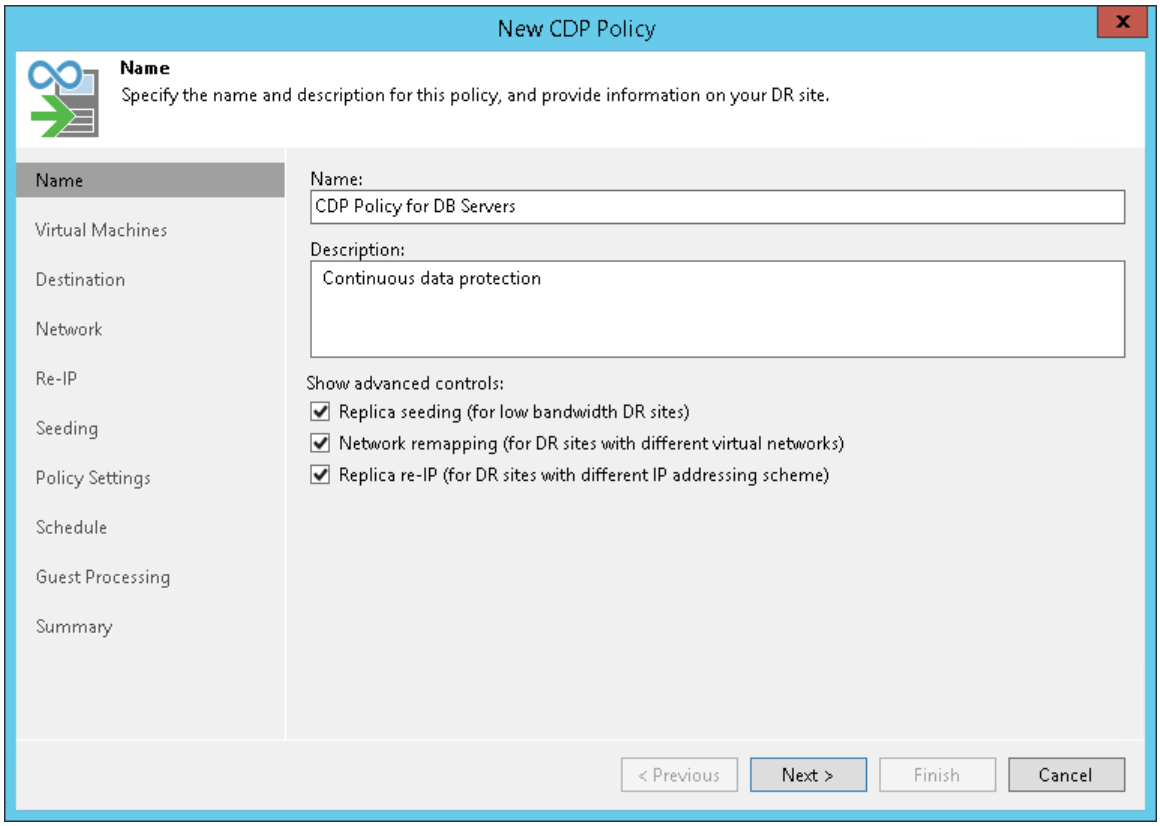
At the **Name** step of the wizard, specify a name and description for the CDP policy, and choose whether you want to use replica seeding, network mapping or re-IP:

1. In the **Name** field, enter a name for the CDP policy.
2. In the **Description** field, provide a description for future reference.
3. If a network between your production and disaster recovery (DR) sites has low bandwidth, and you want to reduce the amount of traffic sent during the initial synchronization of the CDP policy, select the **Replica seeding (for low bandwidth DR sites)** check box.

   When selected, this check box enables the **Seeding** step where you will have to configure replica seeding and mapping. For more information on seeding and mapping, see Replica Seeding and Mapping.
4. If your DR site networks do not match your production site networks, select the **Network remapping (for DR sites with different virtual networks)** check box.

   When selected, this check box enables the **Network** step where you will have to configure a network mapping table.
5. If the IP addressing scheme in your production site differs from the scheme in the DR site, select the **Replica re-IP (for DR sites with different IP addressing scheme)** check box.

   When selected, this check box enables the **Re-IP** step where you will have to configure replica re-IP rules.



Step 3. Select VMs to Replicate

At the **Virtual Machines** step of the wizard, select VMs or VM containers that you want to replicate:

1. Click **Add**.
2. In the **Add Object** window, select the necessary VMs or VM containers and click **Add**.

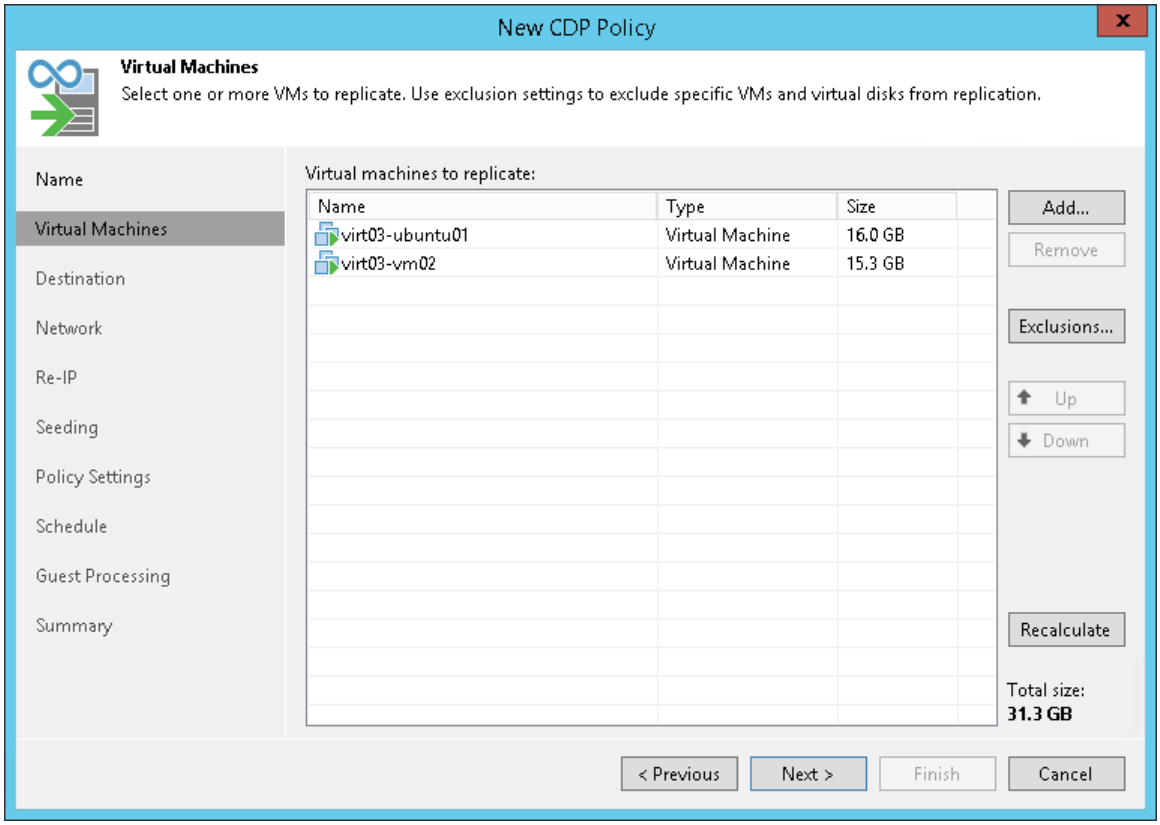
   If you select VM containers (hosts, clusters, folders, resource pools, VirtualApps or datastores) and add new VMs to this container in future, Veeam Backup & Replication will update CDP policy settings automatically to include these VMs.

   > **NOTE**
   >
   > Even if you add VM containers to a CDP policy, only VM are replicated. VM templates, Vm logs, folders and so on are not replicated.

   You can use the toolbar at the top right corner of the window to switch between views. Depending on the view you select, some objects may not be available. For example, if you select **Tags combination** view, no resource pools, hosts or clusters will be displayed in the tree. In the **Tags combination** view, you can select multiple tags and only those VMs that have all the selected tags will be processed by the policy.

> **IMPORTANT**
>
> Mind the following:
>
> - You can replicate only VMs that are turned on, the turned off VMs will be skipped from processing.
> - You cannot add to a CDP policy VMs that were already added to other CDP policies created on the same backup server.
> - We recommend to add not more than 50 VMs or 200 virtual disks in one CDP policy.

To quickly find the necessary VMs, you can use the search field at the bottom of the **Add Object** window. If you want to switch between types of VMs you want to search through, use the button to the left of the search field.



Step 4. Exclude Objects

After you have added VMs or VM containers to the CDP policy, you can specify which objects you want to exclude from being replicated. You can exclude the following types of objects:

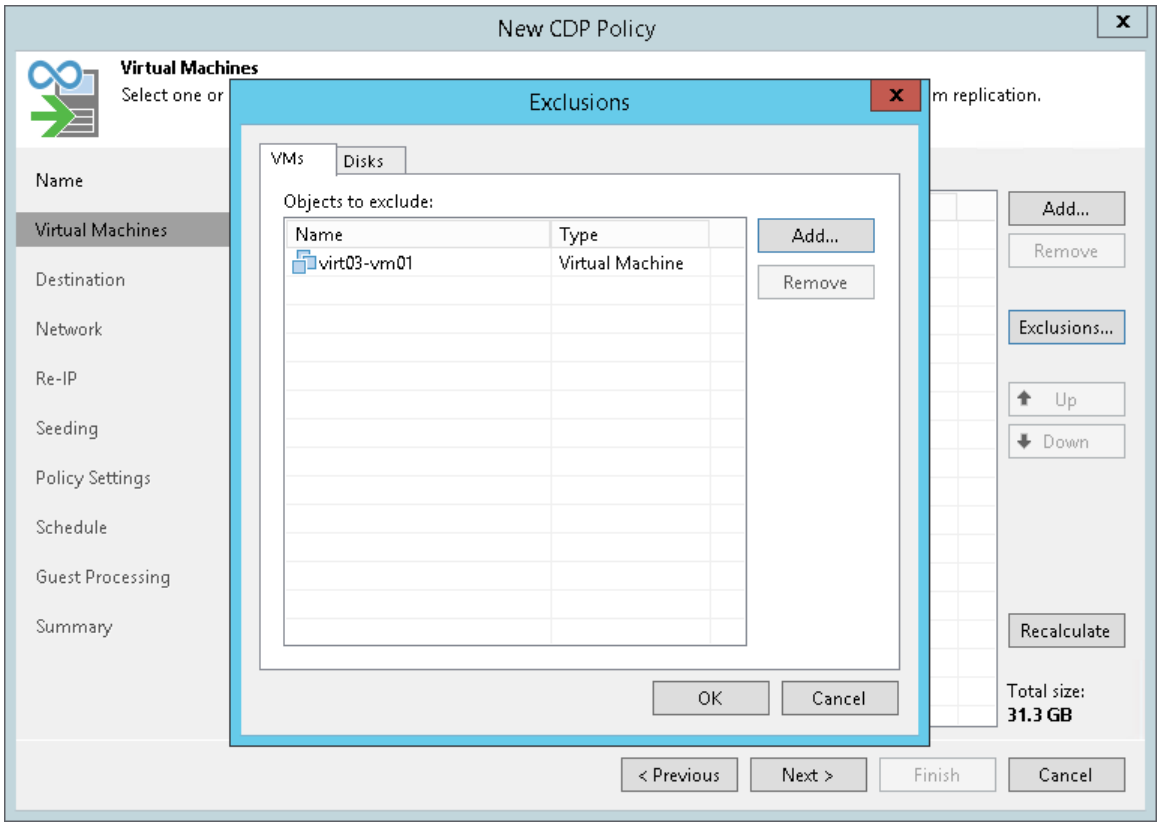
- VMs or VM containers
- Disks

Excluding VMs or VM Containers

To exclude VMs or VM containers:

1. At the **Virtual Machines** step of the wizard, click **Exclusions**.
2. In the **Exclusions** window, check that the **VMs** tab is selected and click **Add**.
3. In the **Add Objects** window, select VMs or VM containers that you want to exclude from being replicated and click **Add**.

   You can use the toolbar at the top right corner of the window to switch between views. Depending on the view you select, some objects may not be available. For example, if you select **Tags combination** view, no resource pools, hosts or clusters will be displayed in the tree. In the **Tags combination** view, you can select multiple tags and only those VMs that have all the selected tags will be excluded from the policy.
4. In the **Exclusions** window, click **OK**.

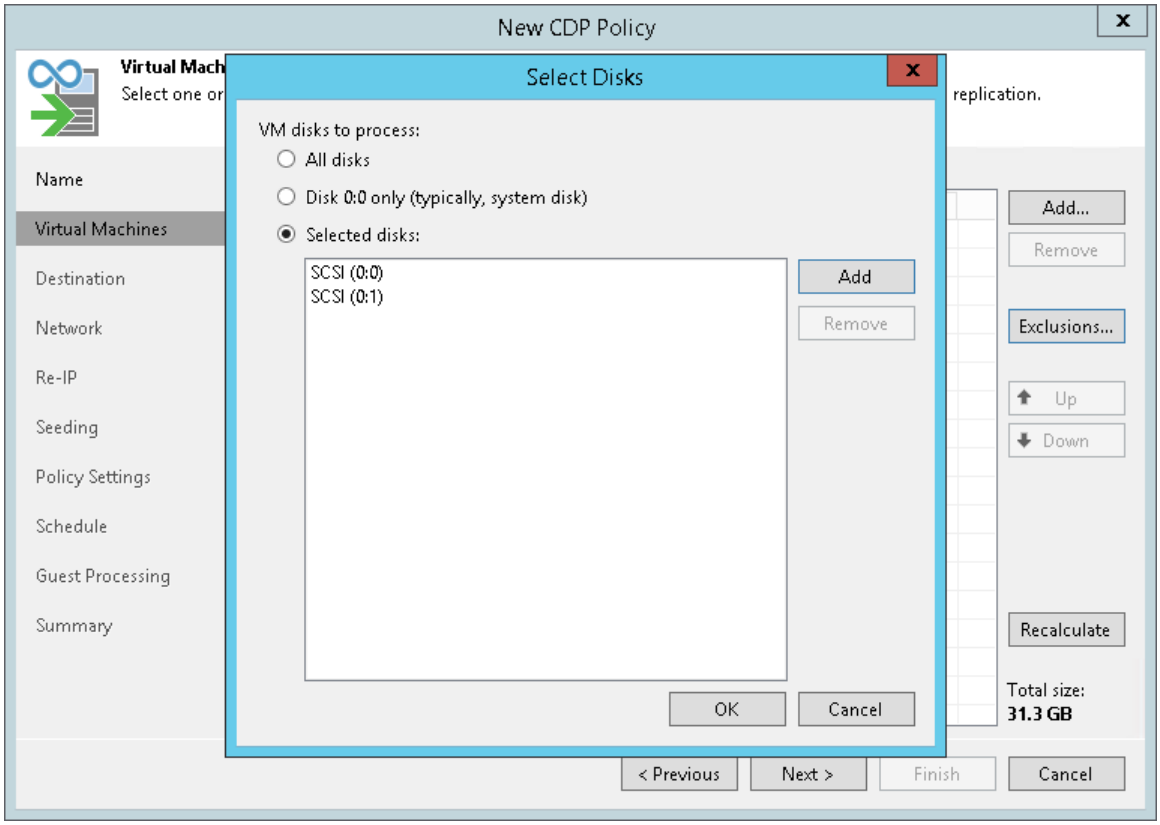


Excluding Disks

To exclude disks:

1. At the **Virtual Machines** step of the wizard, click **Exclusions**.
2. In the **Exclusions** window, do the following:
   a. Switch to the **Disks** tab.
   b. If you want to exclude disks of VMs that are added as a part of containers, click **Add**. In the **Add Objects** window, select the necessary VMs and click **Add**. Veeam Backup & Replication will include these VMs in the list as standalone objects.
   c. In the **Disks to process** list, select the necessary VMs or VM containers.
   d. Click **Edit**.
3. In the **Select Disks** window, select disks that you want to replicate: all disks, 0:0 disks (as a rule, system disks) or specific IDE, SCSI, SATA or NVMe disks. Disks that you do not select will be excluded from processing. Click **OK**.
4. In the **Exclusions** window, click **OK**.

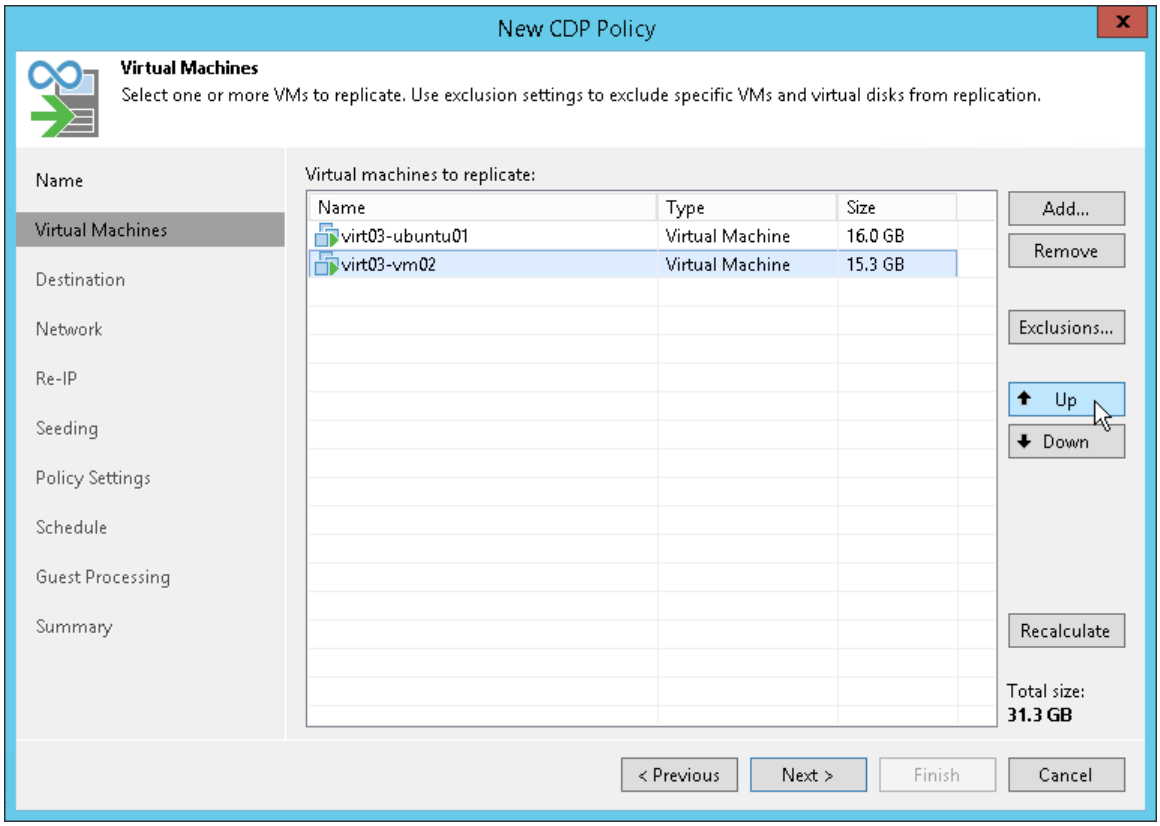
> **NOTE**
>
> If you exclude disks from a backup and enable application-aware processing, Veeam Backup & Replication will still perform application-aware processing for the excluded disks. This means that VSS will process disk data.



Step 5. Define VM Processing Order

At the **Virtual Machines** step of the wizard, click **Up** and **Down** to change the processing order. VMs at the top of the list have a higher priority and will be processed first.

> **NOTE**
>
> Mind the following:
>
> - VMs inside a VM container are processed at random. To ensure that VMs are processed in the defined order, you must add them as standalone VMs, not as a part of containers.
> - The processing order may differ from the order that you have defined. For example, if resources of a VM that is higher in the priority are not available, and resources of a VM that is lower in the priority are available, Veeam Backup & Replication will process the VM with the lower priority first.



Step 6. Select Replica Destination

At the **Destination** step of the wizard, select a target host or cluster, resource pool, folder and datastore for replicas, and types of replica disks:

1. Next to the **Host or cluster** field, click **Choose** and select a host or cluster where replicas must be registered. If you select a cluster or vCenter Server, the replication process will become more sustainable — the replication process will not fail if there is at least one available host in the cluster.

   If you select a cluster as a destination, Veeam Backup & Replication will request VMware to send the list of available hosts, and will select the first host in this list as the destination for the replicas. These replicas will be stored on the datastore with the most free disk space.

2. Next to the **Resource pool** field, click **Choose** and select a resource pool to which replicas will be added.

   If you have selected to replicate multiple VMs and want to add individual replicas to other resource pools:

   a. Click the **Pick resource pool for selected replicas** link.

   b. In the **Choose Resource Pool** window, click **Add VM**.

   c. In the **Add Objects** window, select the necessary VMs and click **Add**.

   d. In the **Choose Resource Pool** window, select the necessary VMs in the **Replica VM resource pool** list. At the bottom of the window, click **Resource Pool**.

   e. In the **Select Resource Pool** window, select the necessary resource pool and click **OK**.

3. Next to the **VM folder** field, click **Choose** and select a folder where all VM files will be stored. Note that the **VM folder** section is disabled if you have selected a standalone ESXi host as the target for replicas.

   If you have selected to replicate multiple VMs and want to place individual replicas to other folders:

   a. Click the **Pick VM folder for selected replicas** link.

   b. In the **Choose Folder** window, click **Add VM**.

   c. In the **Add Objects** window, select the necessary VMs and click **Add**.

   d. In the **Choose Folder** window, select the necessary VMs in the **Replica VM folder** list. At the bottom of the window, click **VM Folder**.

   e. In the **Select Folder** window, select the necessary folder.

4. Next to the **Datastore** field, click **Choose** and select a datastore where replica files will be stored. Note that if you have chosen to replicate VMs to a cluster, Veeam Backup & Replication displays only shared datastores.

   If you have selected to replicate multiple VMs and want to place individual replicas to other datastores:

   a. Click the **Pick datastore for selected virtual disks** link.

   b. In the **Choose VM Files Location** window, click **Add VM**.

   c. In the **Add Objects** window, select the necessary VMs and click **Add**.

   d. In the **Choose VM Files Location** window, select the necessary VMs in the **Files location** list. At the bottom of the window, click **Datastore**.

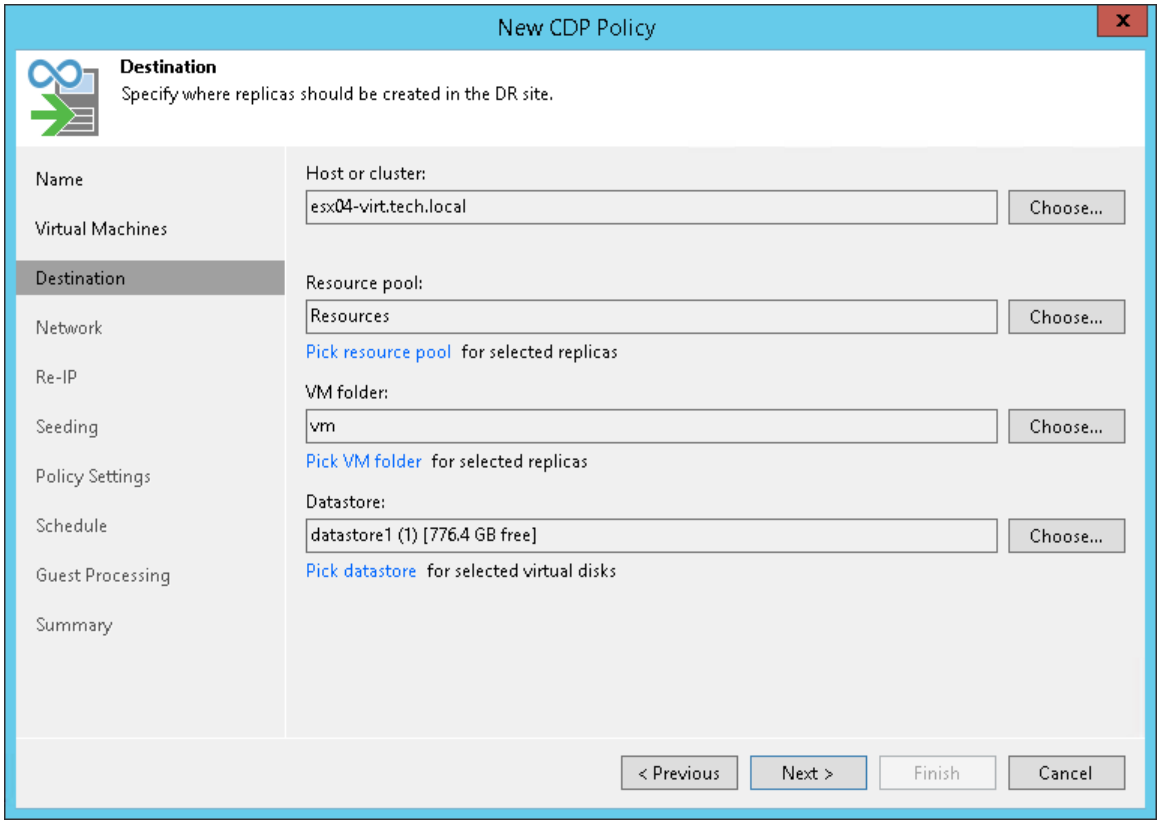
   e. In the **Select Datastore** window, select the necessary datastore.

5. If you want to store replica configuration files and disk files in different datastores:
   a. Click the **Pick datastore for selected virtual disks** link.
   b. In the **Choose VM Files Location** window, click **Add VM**.
   c. In the **Add Objects** window, select the necessary VMs and click **Add**.
   d. In the **Choose VM Files Location** window, expand the necessary VMs in the **Files location** list, and select the necessary files. At the bottom of the window, click **Datastore**.
   e. In the **Select Datastore** window, select the destination for the selected type of files.
6. You can change types of replica disks. By default, Veeam Backup & Replication saves disks in the thin type.

   To change replica disk types:
   a. Click the **Pick datastore for selected virtual disks** link.
   b. In the **Choose VM Files Location** window, click **Add VM**.
   c. In the **Add Objects** window, select VMs whose disk type you want to change and click **Add**.
   d. In the **Choose VM Files Location** window, select the necessary VMs in the **Files location** list. At the bottom of the window, click **Disk type**.
   e. In the **Disk Type Settings** window, select a type that will be used to restore replica disk files: same as source, thin, thick lazy zeroed or thick eager zeroed.

      For more information about disk types, see VMware Docs.

> **NOTE**
>
> Disk type change is available only for VMs that use virtual hardware version 7 or later.



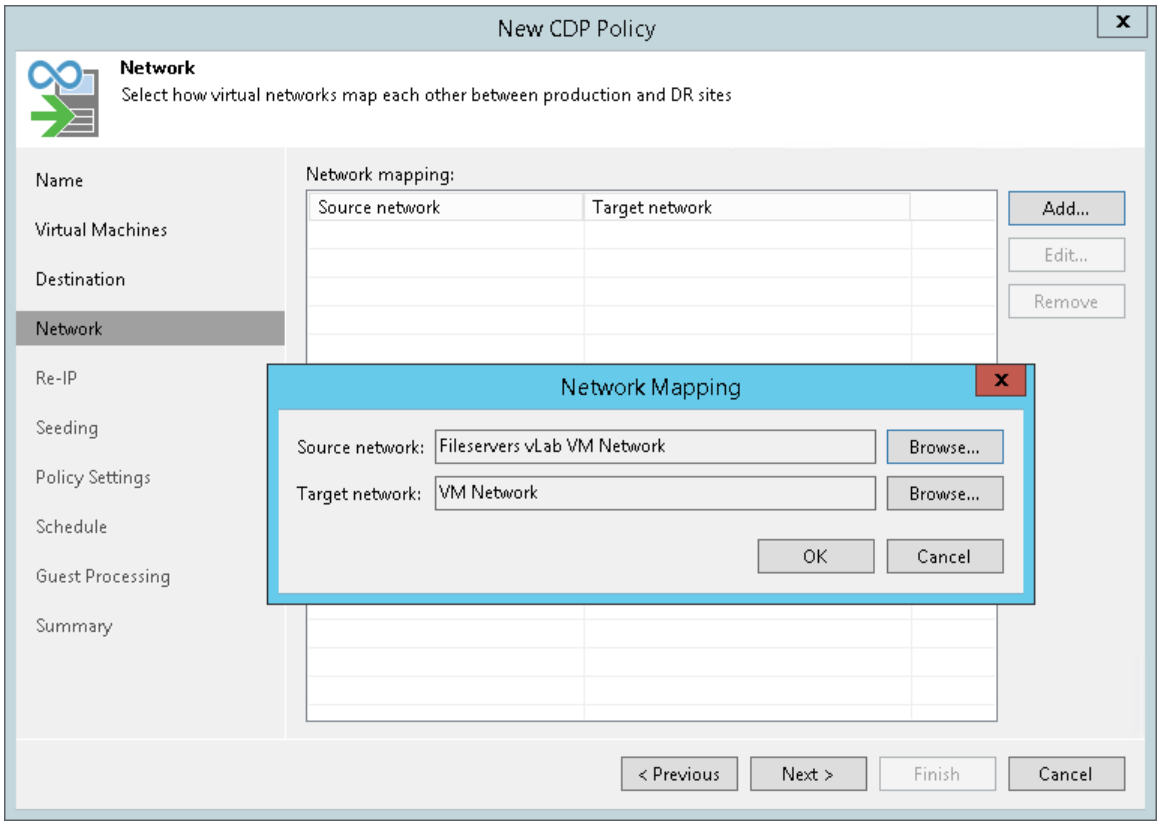
Step 7. Configure Network Mapping

The **Network** step is available if you have selected the **Network remapping** check box at the Name step of the wizard.

At the **Network** step of the wizard, configure a network mapping table. This table maps networks in the production site to networks in the disaster recovery (DR) site. When the policy starts, Veeam Backup & Replication will check the network mapping table. Then Veeam Backup & Replication will update replica configuration files to replace the production networks with the specified networks in the DR site. As a result, you will not have to re-configure network settings manually.

To add a row to a network mapping table:

1. Click **Add**.
2. In the **Network Mapping** window, click **Browse** next to the **Source network** field.
3. In the **Select Network** window, select the production network to which the original VMs are connected and click **OK**.
4. In the **Network Mapping** window, click **Browse** next to the **Target network** field.
5. In the **Select Network** window, select a network in the DR site to which replicas will be connected and click **OK**.
6. In the **Network Mapping** window, click **OK**.



Step 8. Configure Re-IP Rules

The **Re-IP** step is available if you have selected the **Replica re-IP** check box at the Name step of the wizard. This step applies only to VMs with Microsoft Windows OSes.

At the **Re-IP** step of the wizard, configure re-IP rules. These rules map IPs in the production site to IPs in the disaster recovery (DR) site. When you perform failover, Veeam Backup & Replication will check the configured re-IP rules and will change replica IPs if the rules apply. VM replicas will get new IP addresses according to the network masks specified in the rules, so that you will be able to reach replicas in the DR site.

To configure a re-IP rule:

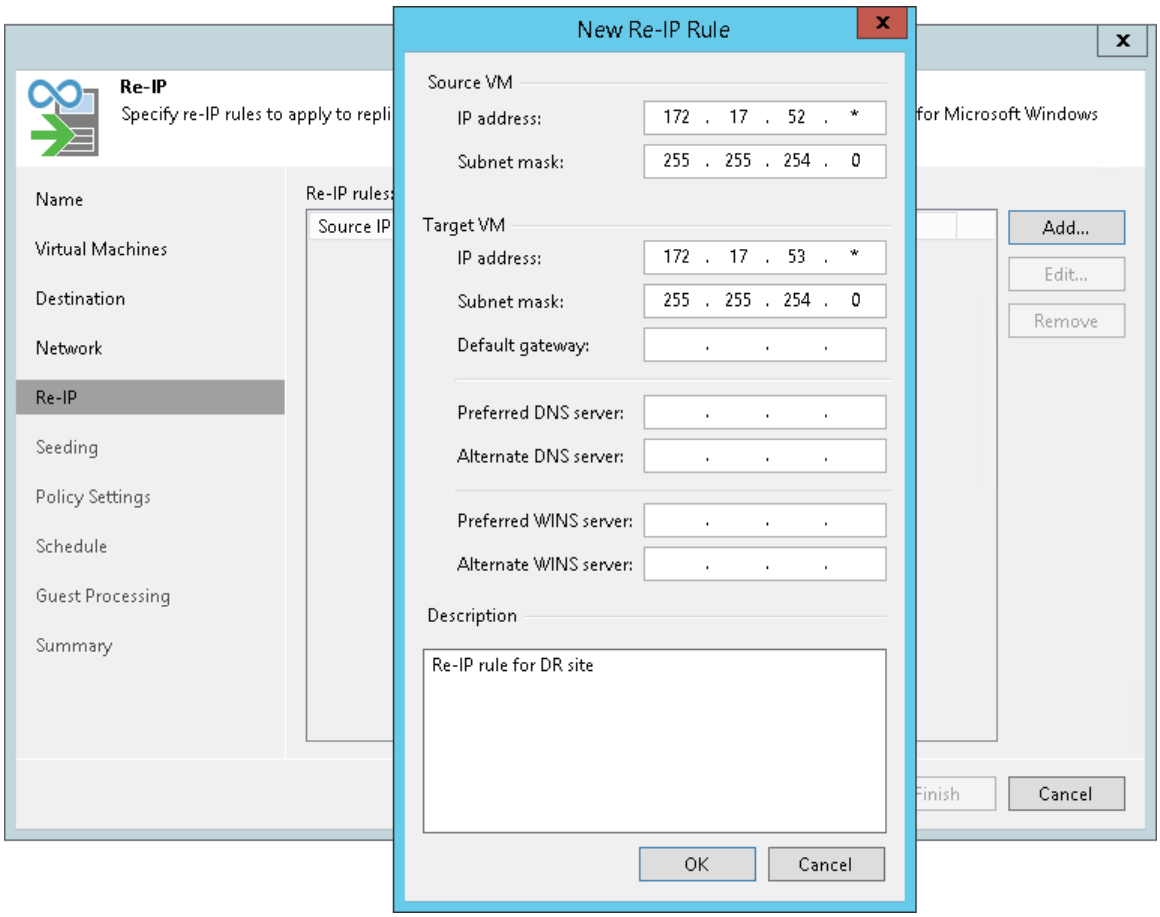
1. Click **Add**.
2. In the **Source VM** section, specify an IP numbering scheme used in the production site.

   To facilitate the configuration, Veeam Backup & Replication detects an IP address and subnet mask for the backup server and pre-populates values in the **Source VM** section.
3. In the **Target VM** section, specify an IP address, subnet mask and default gateway that will be used for replicas in the DR site. If required, specify the DNS and WINS server addresses.
4. In the **Description** field, provide a description the rule.
5. Click **OK**.

NOTE

Mind the following:

- Do not use 0 to specify IP address ranges. In Veeam Backup & Replication, value 172.16.17.0 means a regular IP address 172.16.17.0, not an IP address range. To specify a range, use the asterisk character (*).
- Replica re-IP works only if you perform replica failover using Veeam Backup & Replication. If you power on a VM replica in some other way, for example, manually using vSphere Client, re-IP rules will not be applied to it.
- The backup server OS must support mounting of the system disks of VMs that will be replicated.



Step 9. Configure Seeding and Mapping

The **Seeding** step is available if you have selected the **Replica seeding** check box at the Name step of the wizard.

At the **Seeding** step of the wizard, configure replica seeding and mapping. Seeding and mapping help reduce the amount of traffic sent during the initial replica synchronization. For more information on when to use seeding and mapping, see Replica Seeding and Mapping.

> **IMPORTANT**
>
> If the **Replica seeding** check box is enabled in a policy, all VMs in the policy must be covered with seeding or mapping. If a VM is neither has a seed, nor is mapped to an existing VM, it will be skipped from processing.

Configuring Replica Seeding

To configure replica seeding:

1. Make sure that you have backups of replicated VMs in a backup repository in the DR site. If you do not have the backups, create them as described in the Creating Replica Seeds section.

   > **IMPORTANT**
   >
   > Mind the following:
   >
   > - Backups must be created by Veeam Backup & Replication.
   > - Backups must not reside in a scale-out backup repository.

2. Select the **Get seed from the following backup repository** check box.
3. From the list of available backup repositories, select the repository where your replica seeds are stored.

> **NOTE**
>
> If a VM has a seed and is mapped to an existing replica, replication will be performed using replica mapping because mapping has a higher priority.

Configuring Replica Mapping

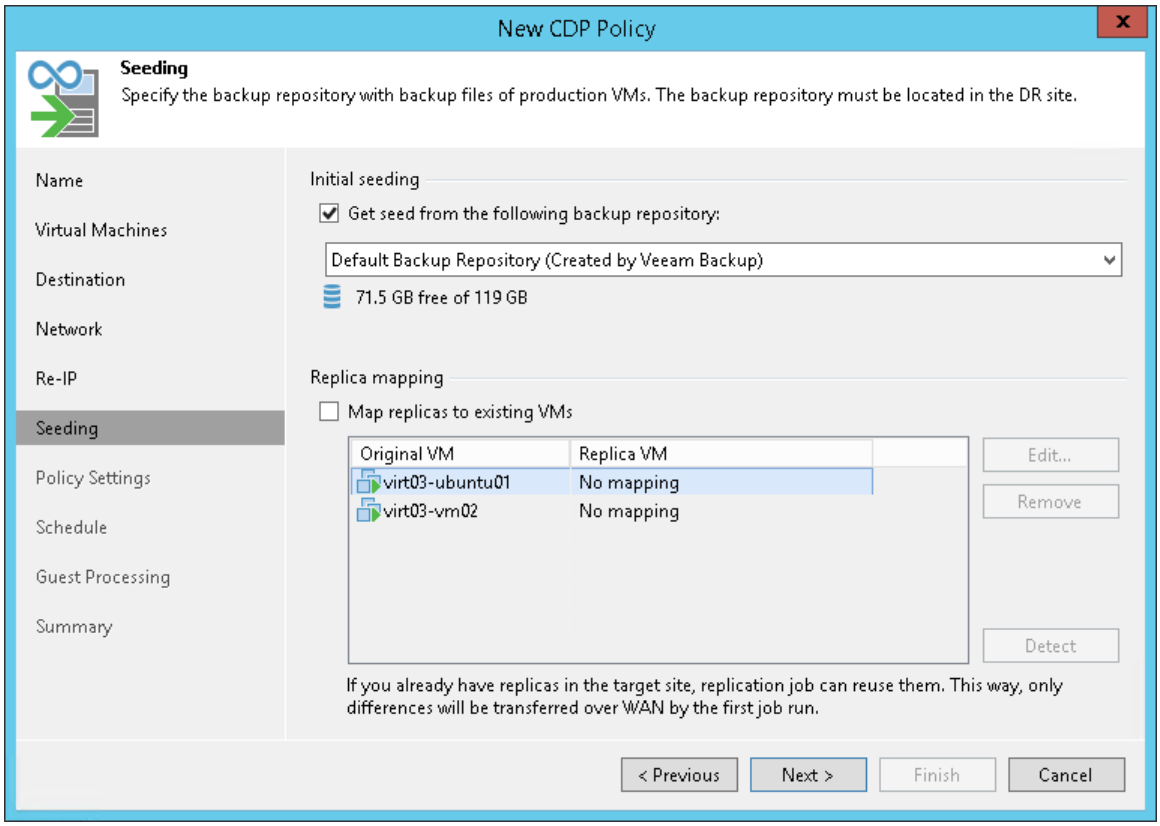
To configure replica mapping:

1. Select the **Map replicas to existing VMs** check box.
2. If you want Veeam Backup & Replication to scan the DR site to detect existing copies of VMs that you plan to replicate, click **Detect**.

   If any matches are found, Veeam Backup & Replication will populate the mapping table. If Veeam Backup & Replication does not find a match, you can map a VM to its copy manually.

3. If you want to map a VM manually, select an original VM from the list, click **Edit** and select the copy of this VM on the target host in the DR site.

To remove a mapping association, select a VM in the list and click **Remove**.



Step 10. Specify Data Transfer and Replica Settings

At the **Policy Settings** step of the wizard, select VMware CDP proxies that must be used for the CDP policy and which suffix to add to replica names:

1. Specify which VMware CDP proxies you want to use:
   - If you want Veeam Backup & Replication to select proxies automatically, leave **Automatic selection** in the **Source proxy** and **Target proxy** fields.

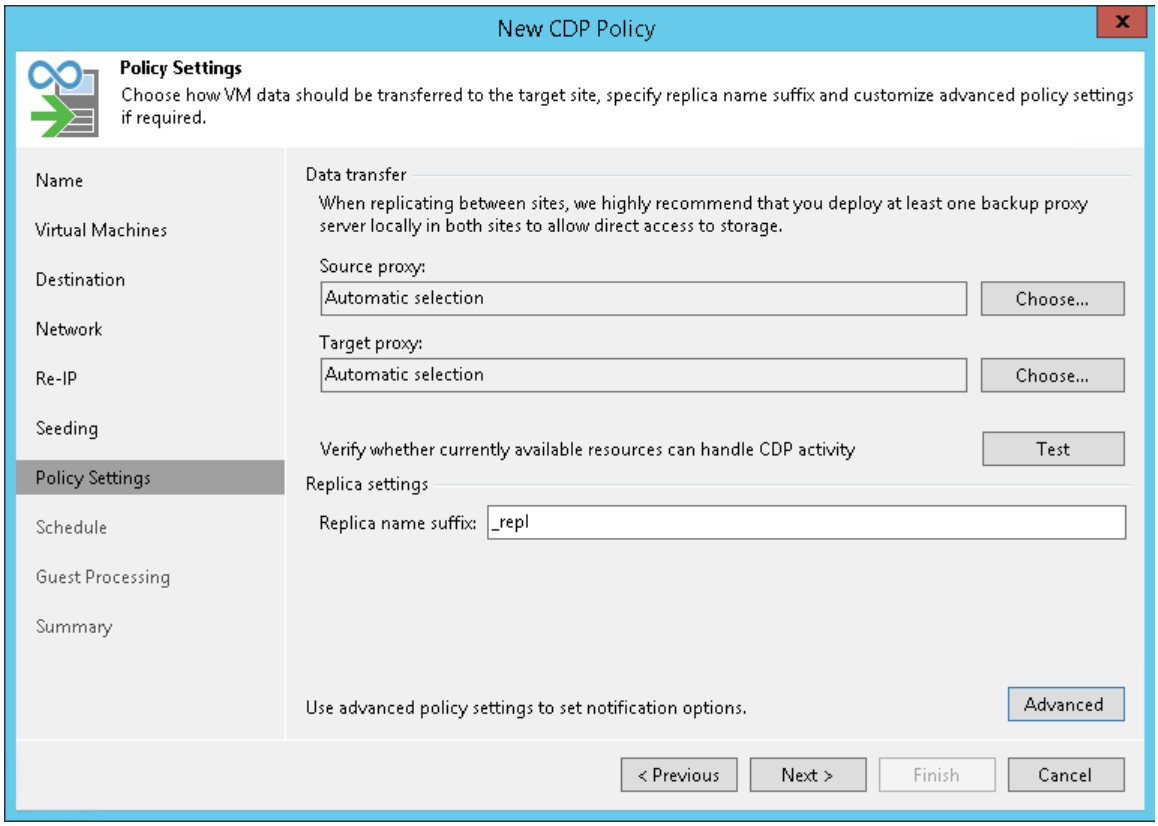
     Veeam Backup & Replication will assign VMware CDP proxies for VM processing one by one. Before processing a new VM from the list, Veeam Backup & Replication will check available VMware CDP proxies. For more information on how Veeam Backup & Replication selects proxies, see How CDP Works.
   - If you want to select VMware CDP proxies manually, do the following:
     i. Click **Choose** next to the **Source proxy** field if you want to select VMware CDP proxies in the production site, or next to the **Target proxy** field if you want to select VMware CDP proxies in the disaster recovery site.
     ii. In the **Backup Proxy** window, click **Use the selected backup proxy servers only**. Select proxies that you want to use and click **OK**.

   > **NOTE**
   >
   > We recommend that you deploy at least two VMware CDP proxies: one CDP proxy in the production site and one CDP proxy in the disaster recovery site.

2. To test whether VMware CDP proxies available in the backup infrastructure can handle replication, click **Test**.

   Veeam Backup & Replication will analyze available CPU on all source and all target VMware CDP proxies, the maximum VM disk write speed during the last hour, and will calculate approximate requirements for VMware CDP proxies. In the **CDP Infrastructure Assessment** window, you will see the calculated values:
   - The **CPU** rows show CPU cores available on all proxies (source or target).
   - The **Proxy RAM** rows show RAM required for CDP and, in parenthesis, RAM available on all proxies (source or target). If values in the parentheses and near them are the same, you need to upgrade proxies for which values coincide to provide more resources. RAM on proxies must be doubled up because CDP uses 50% of RAM by default. You can change this value with a registry key. For more information, contact Veeam Customer Support.
   - The **Proxy Bandwidth** rows show the maximum disk write speed during the last hour and, in parenthesis, available bandwidth based on available cores of source or target proxies.

3. In the **Replica name suffix** field, specify a suffix that will be added to names of replicas.



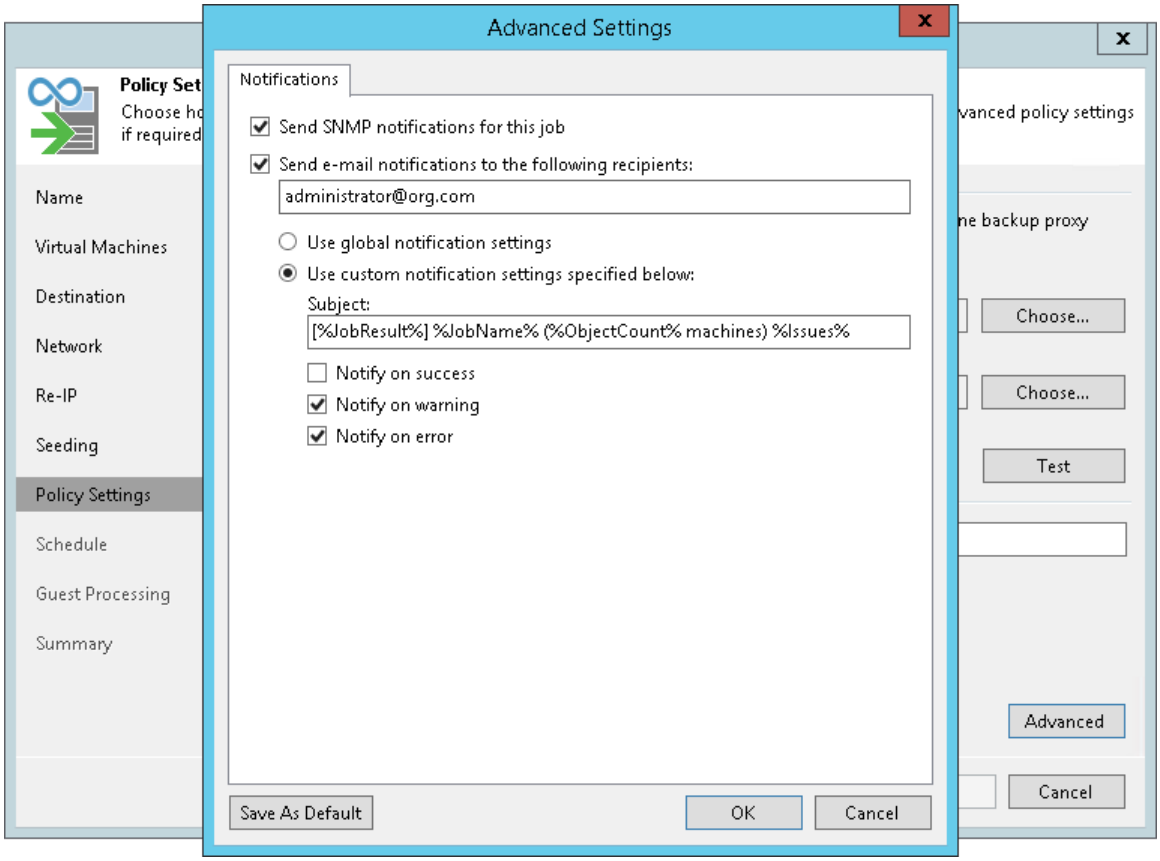
Step 11. Specify Notification Settings

At the **Policy Settings** wizard, specify notification settings:

1. At the lower right corner, click **Advanced**.
2. To receive SNMP traps on the CDP policy, select the **Send SNMP notifications for this job** check box.

   SNMP traps will be sent if you configure global SNMP settings in Veeam Backup & Replication and configure software on recipient machine to receive SNMP traps. For more information, see Specifying SNMP Settings.
3. To receive notifications by email in case of policy failure, success or warning, select the **Send email notifications to the following recipients** check box. Then configure notification settings:
   a. Check that you have configured global email notification settings as described in the Configuring Global Email Notification Settings section.
   b. In the text field, specify a recipient email address. If you want to specify multiple addresses, separate them by a semicolon.
   c. To use global notification settings, select **Use global notification settings**.
   d. To specify a custom notification subject and redefine at which time notifications must be sent, select **Use custom notification settings specified below.** Then specify the following settings:
      i. In the **Subject** field, specify a notification subject. You can use the following variables in the subject: *%JobResult%*, *%JobName%*, *%ObjectCount%* (number of VMs in the policy) and *%Issues%* (number of VMs in the policy that have been processed with the *Warning* or *Failed* status).
      ii. Select the **Notify on success**, **Notify on error** or **Notify on warning** check boxes to receive email notification if the policy gets the *Warning*, *Success* of *Error* status.

> **NOTE**
>
> A CDP policy will get the *Warning* or *Error* status according to the reporting settings configured at the Schedule step of the wizard. A policy will get the *Success* status after the initial configuration succeeds and every day at 8 A.M. if no error or warning occurs.



Step 12. Specify Replica Schedule

At the **Schedule** step of the wizard, configure the schedule and retention policies:

1. Specify scheduling options:

   a. In the **Recovery point objective** field, specify the necessary RPO in seconds or minutes, that is, how often to create short-term restore points.

      The minimum RPO is 2 seconds, however it can be not optimal if your CDP policy contains many VMs with high workload. The optimal RPO is not less than 15 seconds. The maximum RPO is 60 minutes.

      During every specified period, Veeam Backup & Replication will prepare data for short-term restore points for VM replicas and send this data to the target destination. Note that short-term restore points are crash-consistent. For more information on how restore points are created, see How CDP Works.

   b. If you want to prohibit the policy to run at specific time intervals, click **Schedule**. In the schedule box, select the necessary time area and click **Denied**.

2. To instruct the CDP policy to display a warning or error if a newly created restore points are not transferred to the target within the set RPO, click **Reporting**. Then specify when the policy must display error and warning.

   If you have configured email notification settings, Veeam Backup & Replication will mark the policy with the *Warning* or *Error* status and will also send email notifications.

3. In the **Short-term retention** section, configure the short-term retention policy, that is, specify for how long to store short-term restore points.

   For more information on short-term retention, see Short-term Retention Policy.

4. In the **Long-term retention** section, specify when to create long-term restore points and for how long to store them:

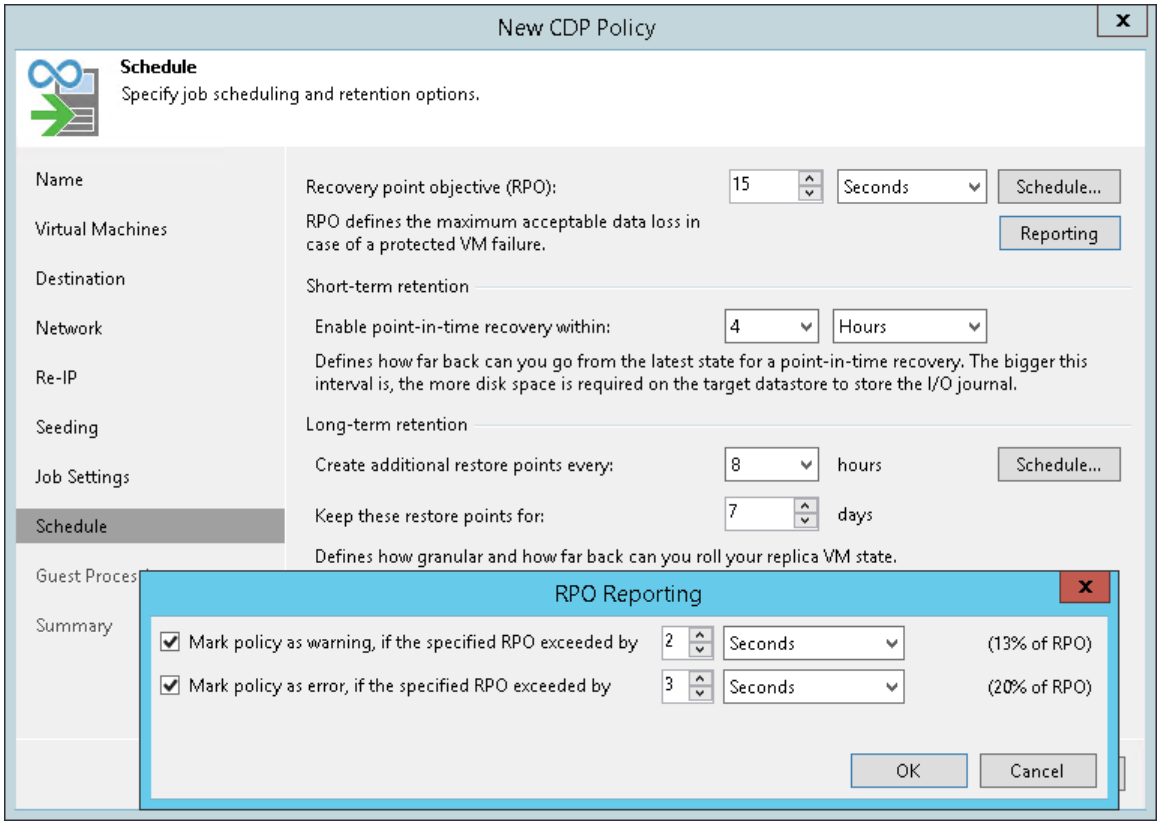
   a. In the **Create additional restore points every** field, specify how often you want to create long-term restore points.

   b. In the **Keep restore points for** field, specify for how long to retain these long-term restore points.

      For more information on what happens after the retention period ends, see Long-term Retention Policy.

c. To specify time periods when Veeam Backup & Replication must create application-consistent and when crash-consistent long-term restore points, click **Schedule**. In the schedule box, select the necessary time area and click **Crash-consistent** or **Application-consistent**. By default, Veeam Backup & Replication creates application-consistent backups if you enable application-aware processing. If you do not enable application-aware processing, Veeam Backup & Replication will create crash-consistent long-term restore points.

   If you want to shift the schedule, specify the offset in the **Start time within an hour** field. For example, you schedule creation of crash-consistent restore points from 00:00 to 01:00, and set the offset value to 25. The schedule will be shifted forward, and the crash-consistent restore points will be created from 0:25 and to 01:25.



Step 13. Specify Guest Processing Settings

Settings configured at this step apply only to long-term restore points.

At the **Guest Processing** step of the wizard, enable guest OS processing that allows you to create transactionally consistent replicas and use pre-freeze and post-thaw scripts. For more information on guest OS processing, see the Guest Processing section.

To enable guest OS processing:

1. Select **Enable application-aware processing**.
2. If you have added Microsoft Windows VMs to the policy, specify which guest interaction proxy Veeam Backup & Replication can use to perform application-aware processing and log processing:
   - If you want Veeam Backup & Replication to select the guest interaction proxy automatically, leave **Automatic selection** on the **Guest interaction proxy** field.
   - If you want to explicitly specify which servers will perform the guest interaction proxy role, click **Choose**. In the **Guest Interaction Proxy** window, click **Prefer the following guest interaction proxy server**, and select the necessary proxies.

   For more information on the guest interaction proxy, see Guest Interaction Proxy.
3. From the **Guest OS credentials** list, select a user account that has enough permissions. For more information on the permissions and requirements for the user account, see Required Permissions for Guest Processing.

   Veeam Backup & Replication will use the account to connect to guest OSes and deploy the non-persistent runtime components or connect to persistent agent. For more information on guest agent and non-persistent components, see Non-Persistent Runtime Components and Persistent Agent Components.

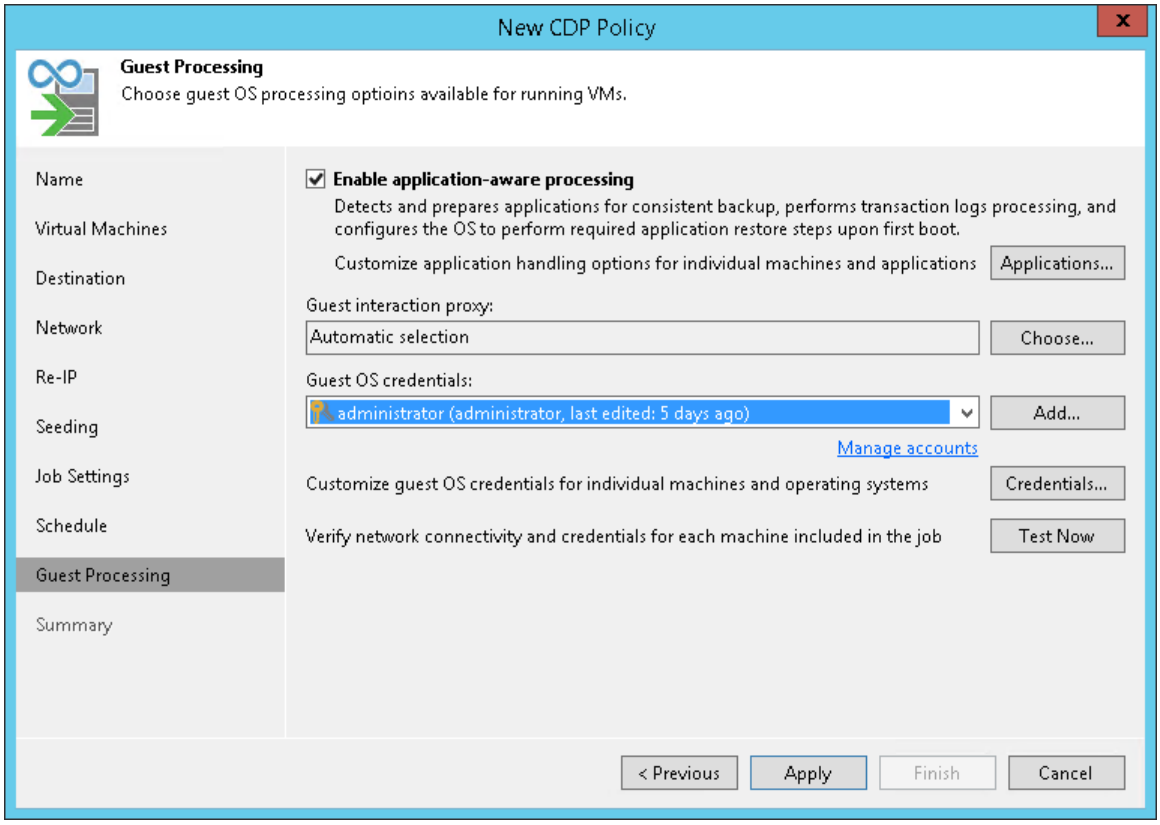
   If you have not set up credentials beforehand, click the **Manage accounts** link or click the **Add** button to add credentials. For more information on adding credentials, see the Credentials Manager section.

   > **NOTE**
   >
   > If you plan to use Kerberos authentication, check limitations and requirements listed in the Guest Processing section.
4. To specify other credentials for individual VMs, click **Credentials**. Then select the necessary VM and set user credentials for it.
5. To check whether Veeam Backup & Replication can connect to VMs using the specified guest OS credentials and can deploy the non-persistent runtime components or connect to persistent agent components on the guest OSes, click **Test Now**.

After you have enabled guest OS processing, configure the following guest processing options:

- Application-aware processing
- Microsoft SQL Server transaction log settings
- Oracle archived log settings

- Pre-freeze and post-thaw scripts



Application-Aware Processing

If VMs run Microsoft Active Directory, Microsoft SQL Server, Microsoft SharePoint, Microsoft Exchange or Oracle, you can enable application-aware processing to create transactionally consistent replicas. The transactionally consistent replicas guarantee proper recovery of applications without data loss.

To configure application-aware processing:

1. Check that you have configured guest OS processing.
2. At the **Guest Processing** step of the wizard, click **Applications**.
3. In the **Application-Aware Processing Options** list, select a VM for which you want to configure application-aware processing, and click **Edit**.

   To define custom settings for a VM added as a part of a VM container, you must include the VM in the list as a standalone object. To do this, click **Add** and choose the necessary VM. Then select the VM in the list and define the necessary settings.
4. On the **General** tab, in the **Applications** section, specify the behavior scenario for application-aware processing:
   - Select **Require successful processing** if you want Veeam Backup & Replication to stop the replication process if any error occurs during application-aware processing.
   - Select **Try application processing, but ignore failures** if you want to continue the replication process even if an error occurs during application-aware processing. This option guarantees that the CDP policy will continue working. However, the resulting replica will be crash consistent, not transactionally consistent.
   - Select **Disable application processing** if you want to disable application-aware processing for the VM.

5. [For Microsoft Exchange, Microsoft SQL and Oracle] In the **Transaction logs** section, specify if Veeam Backup & Replication must process transaction logs or create copy-only replicas:

   a. Select **Process transaction logs with this job** if you want Veeam Backup & Replication to process transaction logs.

      [For Microsoft Exchange VMs] Transaction logs will be truncated after the CDP policy creates a long-term restore point. If the creation fails, the logs will remain untouched until the next start of the long-term restore point creation.

      [For Microsoft SQL Server VMs and Oracle VMs] You will need to configure how to process transaction logs. For more information on how to configure log handling, see Microsoft SQL Server Transaction Log Settings and Oracle Archived Log Settings.

   b. Select **Perform copy only** if you use another replication tool to perform guest level replication, and this tool maintains consistency of the database state. Veeam Backup & Replication will create a copy-only replica for the selected VM. The copy only replica preserves the chain of full and differential files and transaction logs on the VM. For more information, see Microsoft Docs.

6. In the **Persistent guest agent** section, select the **Use persistent guest agent** check box to use for application-aware processing persistent guest agents on each protected VM.

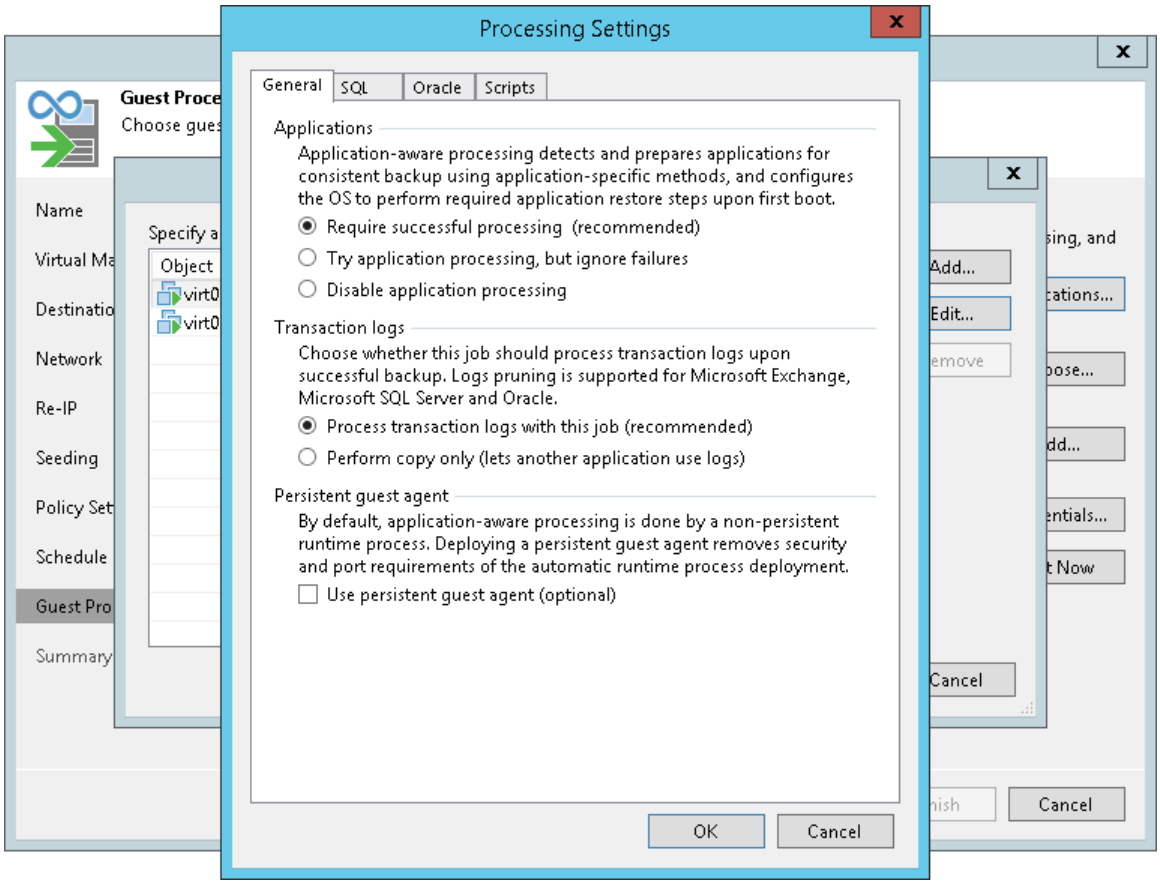
   By default, Veeam Backup & Replication uses non-persistent runtime components.
   Veeam Backup & Replication deploys runtime components on each protected VM before the CDP policy creates an application-consistent long-term restore point, and removes the runtime components as soon as the restore point is created.

   For more information on guest agent and non-persistent components, see Non-Persistent Runtime Components and Persistent Agent Components.

To configure application-aware processing for multiple VMs at once, do the following:

1. At the **Guest Processing** step of the wizard, click **Applications.**
2. In the **Application-Aware Processing Options** list, select multiple VMs.
3. Click **Edit** > **General**.

4. Configure options described in the previous list.



Microsoft SQL Server Transaction Log Settings

The **SQL** tab applies to VMs that run Microsoft SQL Server and if you have selected **Process transaction logs with this job** when configuring application-aware processing.

To configure how Veeam Backup & Replication must process transaction logs of an Microsoft SQL Server:

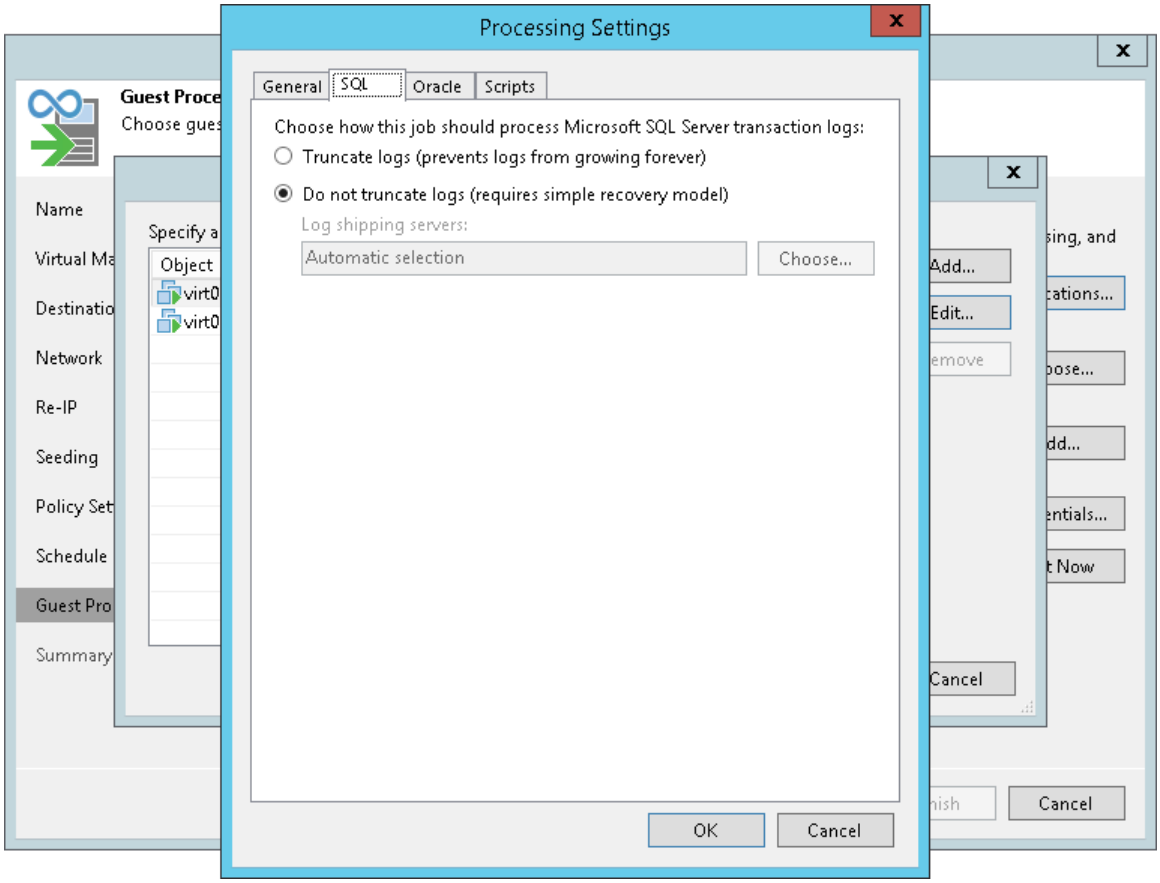
1. At the **Guest Processing** step of the wizard, click **Applications**.
2. In the **Application-Aware Processing Options** list, select a VM that runs the Microsoft SQL Server.
3. Click **Edit**.
4. In the **Processing Settings** window, switch to the **SQL** tab and specify how transaction logs must be processed:
   - If you want Veeam Backup & Replication to trigger truncation of transaction logs after the CDP policy creates a long-term restore point, select **Truncate logs**.

     In this case, transaction logs will be truncated after the CDP policy creates a long-term restore point. If the creation fails, the logs will remain untouched until the next start of the long-term restore point creation.
   - If you do not want Veeam Backup & Replication to truncate logs at all, select **Do not truncate logs**.

     This option is recommended if you use another tool to perform VM guest-level replication, and this tool maintains consistency of the database state.

To configure transaction log handling for multiple VMs at once, do the following:

1. At the **Guest Processing** step of the wizard, click **Applications.**
2. In the **Application-Aware Processing Options** list, select multiple VMs.
3. Click **Edit** > **Microsoft SQL Server**.
4. Configure options described in the previous list.



Oracle Archived Log Settings

The **Oracle** tab applies to VMs that run Oracle and if you have selected **Process transaction logs with this job** when configuring application-aware processing.

To configure how Veeam Backup & Replication must process archive logs of an Oracle server:

1. At the **Guest Processing** step of the wizard, click **Applications.**
2. In the **Application-Aware Processing Options** list, select a VM that runs Oracle.
3. Click **Edit.**
4. In the **Processing Settings** window, switch to the **Oracle** tab.

5. From the **Specify Oracle account with SYSDBA privileges** drop-down list, select a user account that Veeam Backup & Replication will use to connect to the Oracle databases. The account must have SYSDBA rights on the Oracle database.

   You can select **Use guest credentials** from the list of user accounts. In this case, Veeam Backup & Replication will use the account specified at the **Guest Processing** step of the wizard to access the VM guest OS and connect to the Oracle databases. Note that this account must be a member of specific groups or granted specific privileges. For more information, see Required Permissions.

6. In the **Archived logs** section, specify how to process archived logs:

   - If you want to preserve archived logs on the VM guest OS, select **Do not delete archived logs**.

     It is recommended that you select this option for databases where the ARCHIVELOG mode is turned off. If the ARCHIVELOG mode is turned on, archived logs may grow large and consume all disk space.

   - If you want to delete archived logs older than <N> hours, select **Delete logs older than <N> hours** and specify the number of hours.

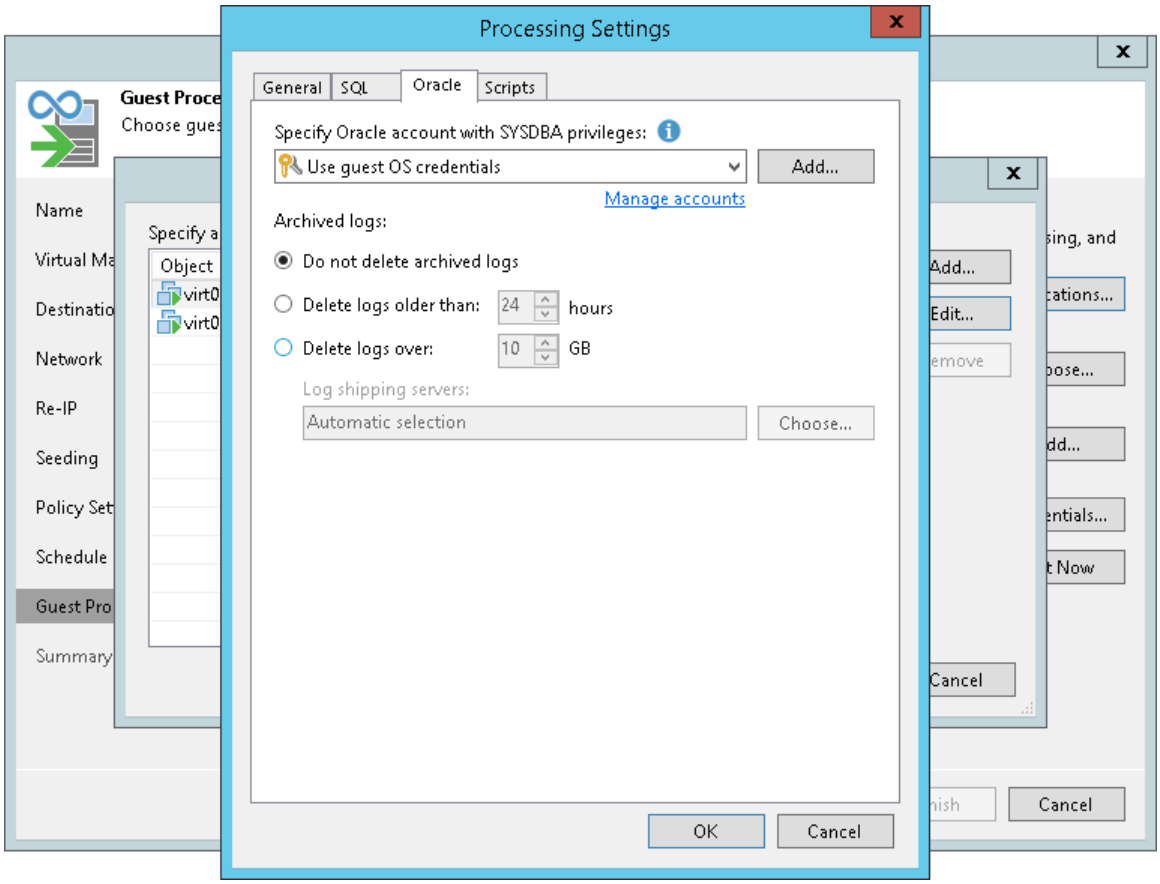
   - If you want to delete archived logs larger than <N> GB, select **Delete logs over <N> GB** and specify the size. The specified size refers to the log size of each database, not all databases on the selected Oracle server.

   Transaction logs will be deleted using Oracle Call Interface after the CDP policy creates a long-term restore point. If the creation fails, the logs will remain untouched until the next start of the long-term restore point creation.

To configure archive log handling for multiple VMs at once, do the following:

1. At the **Guest Processing** step of the wizard, click **Applications.**
2. In the **Application-Aware Processing Options** list, select multiple VMs.
3. Click **Edit** > **Oracle**.

4. Configure options described in the previous list.



Script Settings

You can instruct Veeam Backup & Replication to run custom scripts before the CDP policy starts the creation of a long-term restore point and after the policy finishes the creation. For example, these can be pre-freeze and post-thaw scripts for a VM that does not support VSS. The scripts will quiesce the VM file system and application data to bring the VM to a consistent state before the creation of a restore point, and bring the VM and applications to their initial state after the creation finishes.

To specify scripts:

1. Check that you have enabled guest OS processing.
2. At the **Guest Processing** step of the wizard, click **Applications**.
3. In the **Application-Aware Processing Options** list, select a VM for which you want to configure scripts, and click **Edit**.

   To define custom settings for a VM added as a part of a VM container, you must include the VM in the list as a standalone object. To do this, click **Add** and choose the necessary VM. Then select the VM in the list and define the necessary settings.
4. Click the **Scripts** tab.
5. In the **Script processing mode** section, select a scenario for script execution:
   - Select **Require successful script execution** if you want Veeam Backup & Replication to stop the replication process if scripts fail.

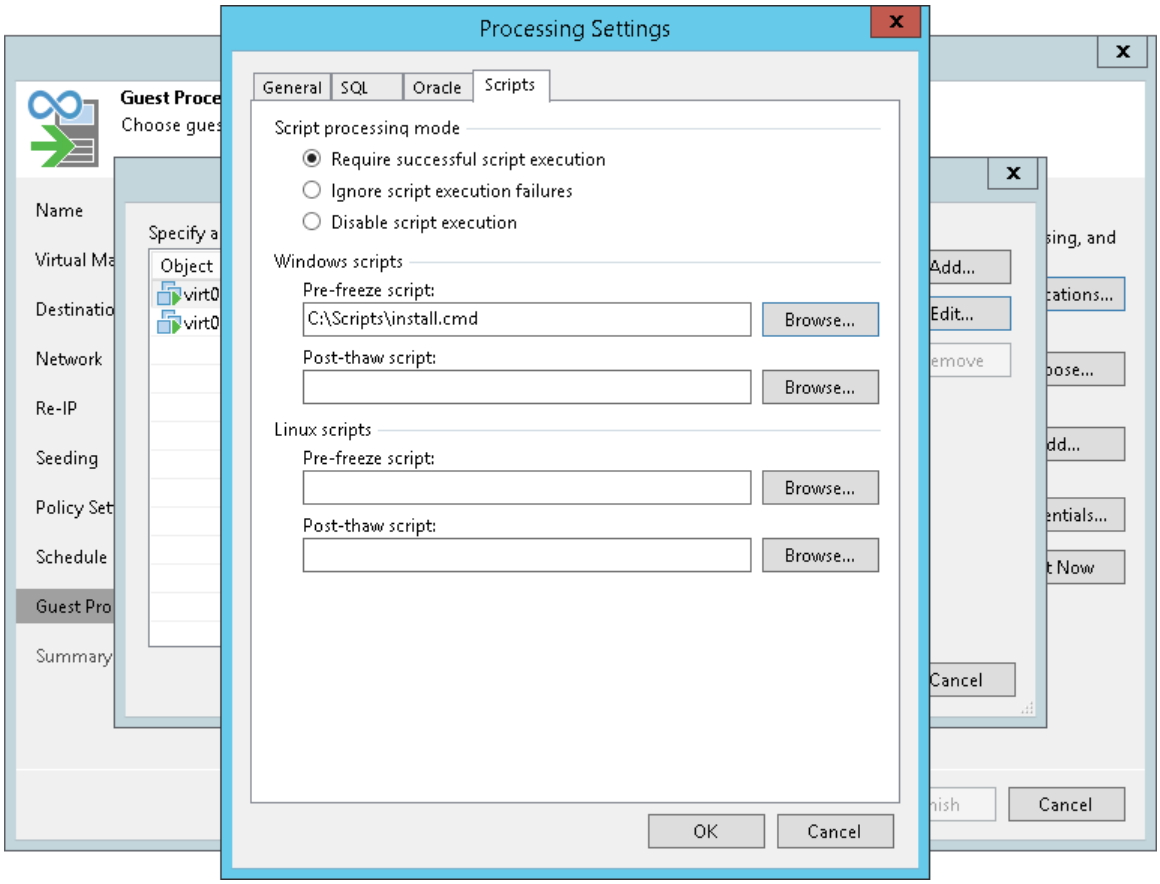
- Select **Ignore script execution failures** if you want to continue the replication process even if script errors occur.
- Select **Disable script execution** if you do not want to run scripts for the VM.

5. In the **Windows scripts** section, specify paths to scripts for Microsoft Windows VMs. For the list of supported script formats, see Pre-Freeze and Post-Thaw Scripts.
6. In the **Linux scripts** section, specify paths to scripts for Linux VMs. For the list of supported script formats, see Pre-Freeze and Post-Thaw Scripts.

   If you have added to the policy a VM container with Microsoft Windows and Linux VMs, you can select to execute both Microsoft Windows and Linux scripts. When the CDP policy starts, Veeam Backup & Replication will automatically determine which OS type is installed on the VM and use the required scripts for this VM.

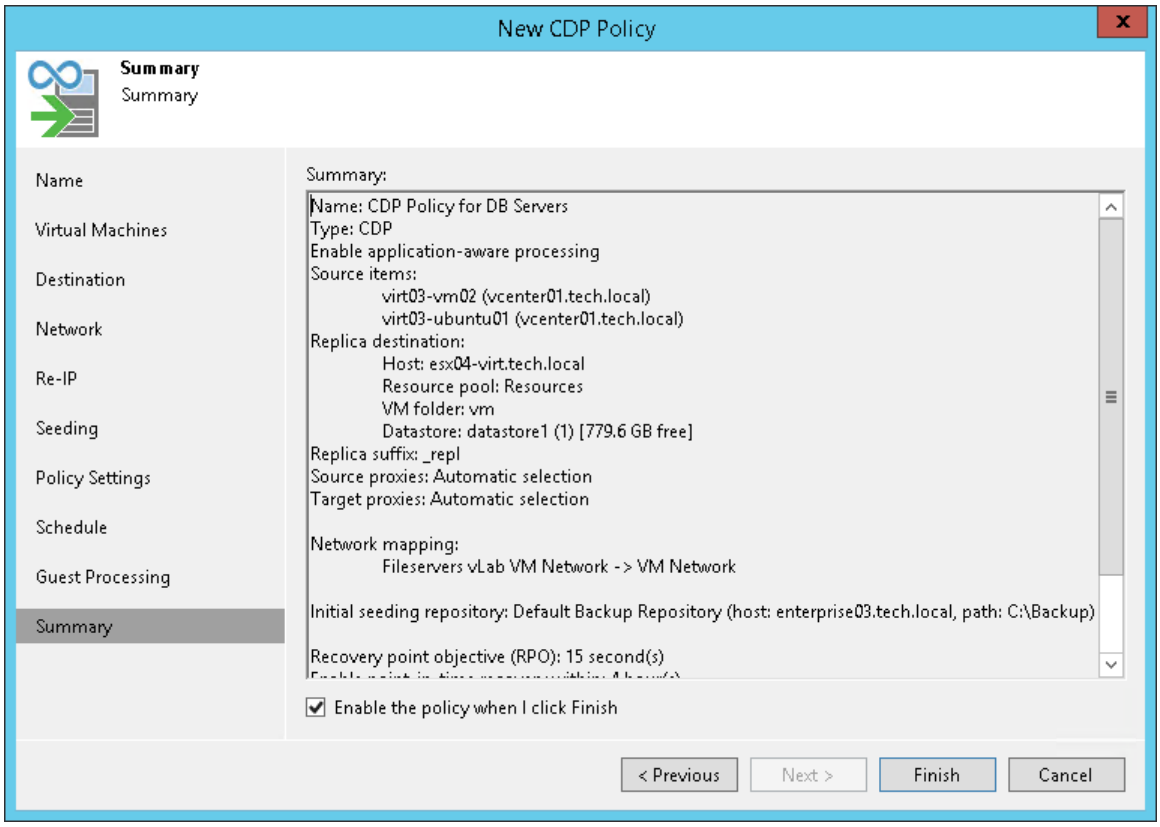
To configure scripts for multiple VMs at once, do the following:

1. At the **Guest Processing** step of the wizard, click **Applications.**
2. In the **Application-Aware Processing Options** list, select multiple VMs.
3. Click **Edit** > **Scripts**.
4. Configure options described in the previous list.



Step 14. Finish Working with Wizard

At the **Summary** step of the wizard, review details of the CDP policy. If you want to start the policy right after you close the wizard, leave the **Enable the policy when I click Finish** check box selected, otherwise clear the check box. Then click **Finish** to close the wizard.



Creating Replica Seeds for CDP

To use replica seeding in a CDP policy, you must have backups of replicated VMs in a backup repository in the disaster recovery (DR) site. These backups are known as replica seeds. For more information on seeding and when to use it, see Replica Seeding and Mapping.

If you do not have replica seeds in the DR site, do the following:

1. Create a backup of VMs that you plan to replicate as described in the Creating Backup Jobs section. As the target repository for this job, select a backup repository in the production site. Then run the job.

   If you already have backups containing the necessary VMs, there is no need to configure and run a new backup job. For seeding, you can use any existing backups created with Veeam Backup & Replication. The backup must include VBK and VBM files. If you have a full backup and a chain of forward increments, you can use VIB files together with the VBK and VBM files. In this case, Veeam Backup & Replication will restore VMs from the seed to the latest available restore point.

2. Copy the backup from the backup repository in the production site to a backup repository in the DR site.

   You can move the backup using a file copy job or any other appropriate method, for example, copy the backup to a removable storage device, ship the device to the DR site and copy backups to the backup repository in the DR site.

   If you do not have a backup repository in the DR site, you need to create the repository as described in Backup Repository.

   > **IMPORTANT**
   >
   > You cannot copy backups to a scale-out backup repository in the DR site.

3. After the backup is copied to the backup repository in the DR site, perform rescan of this backup repository as described in the Rescanning Backup Repositories section. Otherwise, Veeam Backup & Replication will not be able to detect the copied backup.

Managing CDP Policies

After you create CDP policies, you can edit, disable and delete them.

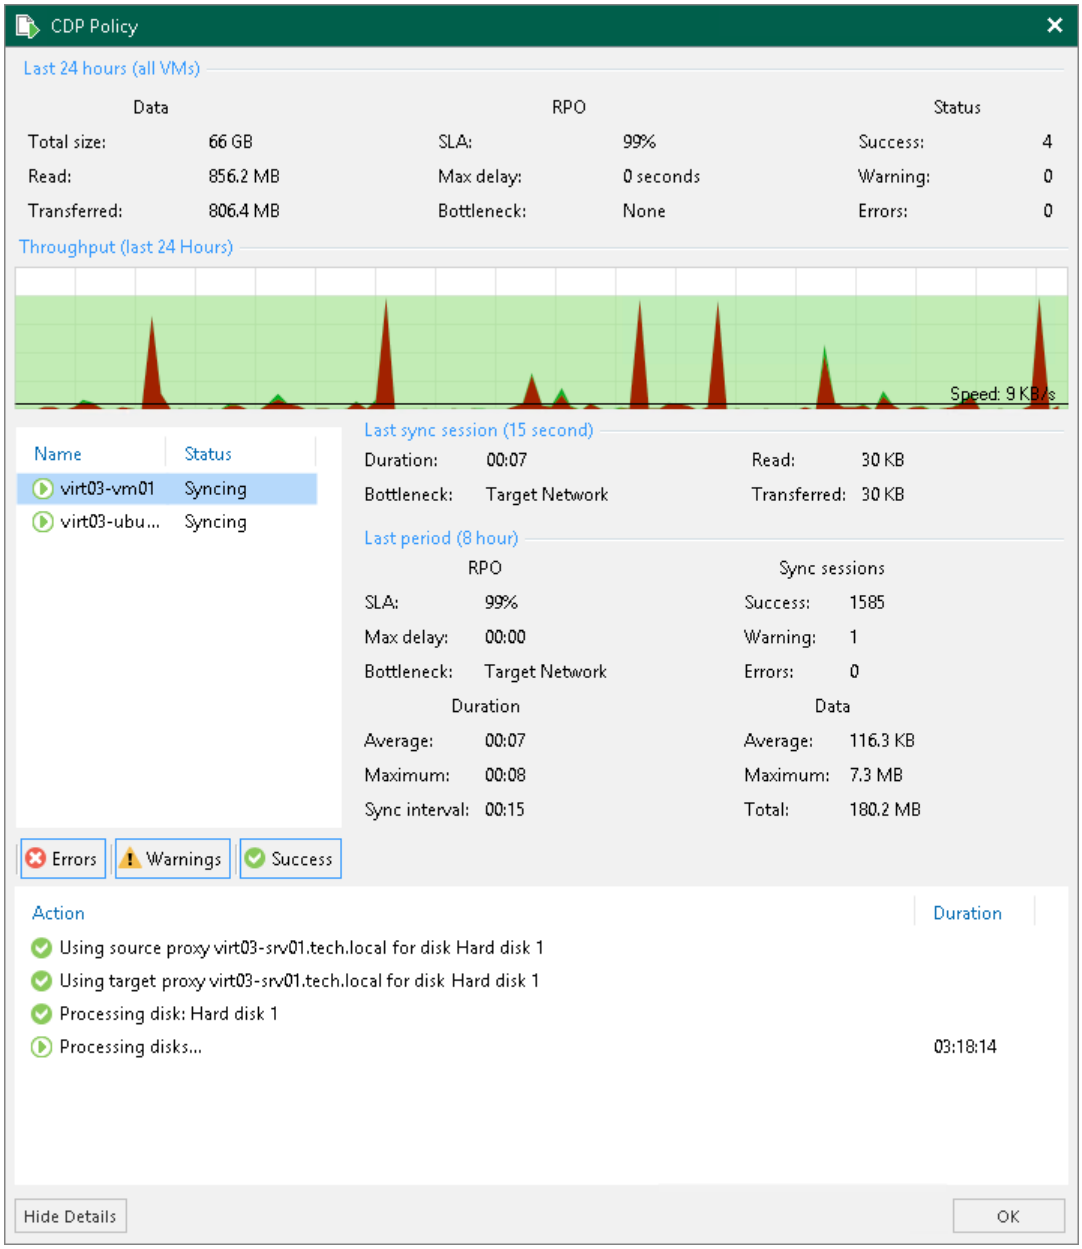
To view all created CDP policies, open the **Home** view and navigate to the **Jobs** > **CDP** node. The working area displays the full list of the created policies. Here, you can manage the policies.

Viewing Session Statistics and Results

Veeam Backup & Replication allows you to view real-time statistics and session results for a VM added to a CDP policy:

- To view real-time statistics for the policy, open the **Home** view. In the inventory pane select **Jobs** > **CDP**. In the working area, double-click the necessary policy. Alternatively, right-click the policy and select **Statistics**. In the opened window, you will be able to switch between statistics of individual VMs.
- To view real-time statistics for an individual VM, open the **Home** view. In the inventory pane select **Jobs, Last 24 hours or Running**. In the working area, double-click the necessary VM. Alternatively, right-click the VM and select **Statistics**.
- To view statistics on the finished policy sessions, open the **History** view. In the inventory pane select **CDP**. In the working area, double-click the necessary policy session.

The statistics provides detailed data on policy sessions: duration, performance bottlenecks, amount of processed data, read and transferred data.



Statistics Counters

Veeam Backup & Replication displays policy statistics for the following counters:

- The **Last 24 hours (all VMs)** section shows the general information for all VMs for which you have opened the statistics. This information is collected during 24 hours.
  - The **Data** box shows information about processed VM data:
    - **Total size** – total size of the processed data.
    - **Read** – amount of data read from the datastore prior to applying compression and deduplication. The value of this counter is typically lower than the value of the **Total size** counter. Veeam Backup & Replication reads only data blocks that have changed since the last policy session, processes and copies these data blocks to the target.
    - **Transferred** – amount of data transferred from the source VMware CDP proxy to the target VMware CDP proxy. The data is transferred after compression and deduplication.
  - The **RPO** box shows information related to RPO:
    - **SLA** – percentage of sessions completed within the desired RPO.
    - **Max delay** – difference between the configured RPO and time required to transfer and save data.
    - **Bottleneck** – bottleneck in the data transmission process.
  - The **Status box** shows information about task session results. This box informs how many task sessions have completed with the *Success*, *Warning* and *Error* statuses during the 24-hour session. One task session lasts for the period between the creation of two long-term restore points.
- The **Last sync session** section shows general information collected during the last synchronization session of the 24-hour session. One synchronization session lasts for the period between the creation of the two short-term restore points.
  - **Duration** – time period during which data was collected and sent to the target host.
  - **Bottleneck** – bottleneck in the data transmission process.
  - **Read** – amount of data read from the datastore prior to applying compression and deduplication.
  - **Transferred** – amount of data transferred from the source VMware CDP proxy to the target VMware CDP proxy. The data is sent after compression and deduplication.
- The **Last period** section shows general information collected during the last task session of the 24-hour session. One task session lasts for the period between the creation of two long-term restore points.
  - The **RPO** box shows information related to RPO:
    - **SLA** – percentage of sessions completed within the desired RPO.
    - **Max delay** – difference between the configured RPO and time required to transfer and save data.
    - **Bottleneck** – bottleneck in the data transmission process.
  - The **Sync session** box shows information about synchronization session results. This box informs how many synchronization sessions have completed with the *Success*, *Warning* and *Error* statuses during the last task session. One synchronization session lasts for the period between the creation of the two short-term restore points.

- The **Duration** box shows information about duration of synchronization sessions:
  - **Average** — average duration of a synchronization session.
  - **Maximum** — maximum duration of a synchronization session.
  - **Sync interval** — duration of a synchronization session configured in the policy, that is, the specified RPO.
- The **Data** box shows information about processed VM data:
  - **Average** — average amount of data processed within one synchronization session.
  - **Maximum** — maximum amount of data processed within one synchronization session.
  - **Total** — total size of data sent during the task session.

- The pane at the left shows the list of VMs included into the session statistics.
- The pane at the bottom shows the list of operations performed during the session. If you open statistics for a policy, you can see the list of operations for the whole policy or an individual VM. To see the list of operations for an individual VM, click the VM in the pane on the left. To see the list of operations for the whole policy, click anywhere on the blank area in the left pane.

Colored Graph

To visualize the data transfer process, Veeam Backup & Replication displays a colored graph in the statistics window:

- The green color defines the amount of data read from the source.
- The brown color defines the amount of data transferred to the target.
- The horizontal line defines the current data processing speed.

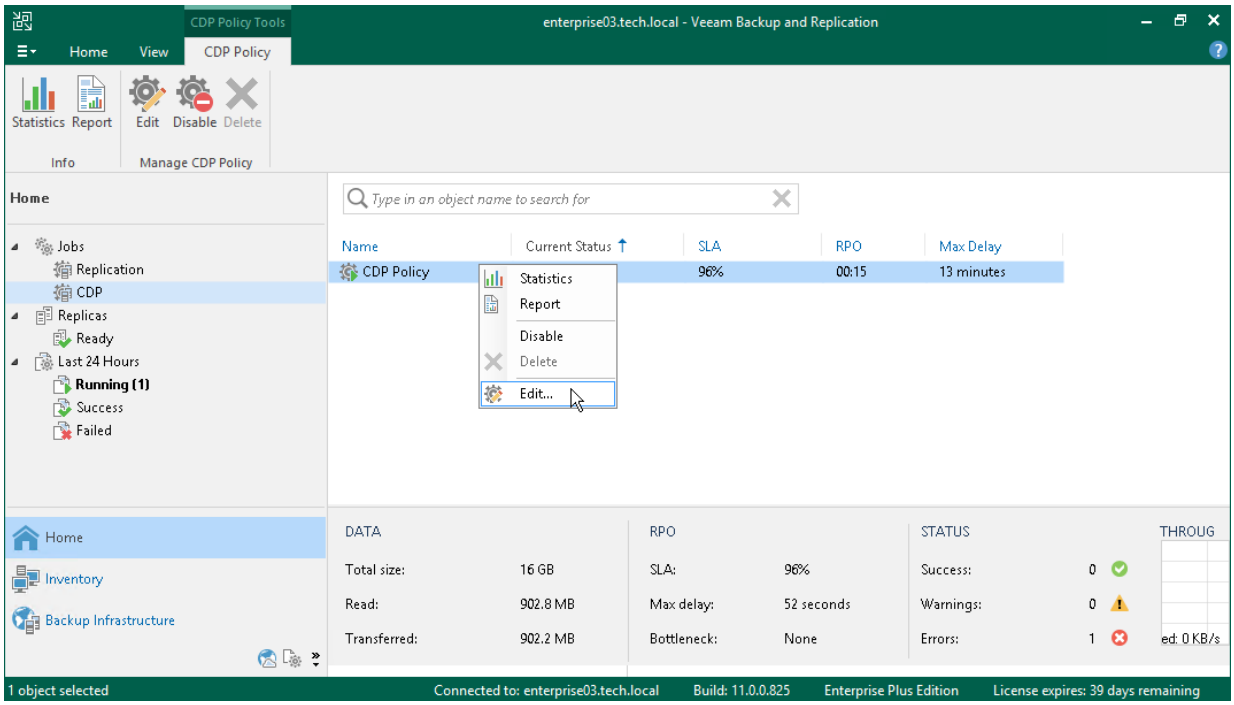
If the policy session is still being performed, you can click the graph to view data rate for the last 5 minutes or the last 24 hours. If the policy session has already ended, the graph will display information for the last 24 hours only.

The colored graph is displayed only for the currently running policy session or the latest finished session. If you open statistics for past sessions other than the latest one, the colored graph will not be displayed.

Editing Policies

To edit a CDP policy:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Jobs** > **CDP** node.
3. In the working area, select the necessary policy and select **Edit** on the ribbon. As an alternative, right-click the necessary policy and select **Edit**.
4. Follow the instructions provided in Creating CDP Policies.



Disabling and Deleting Policies

Veeam Backup & Replication allows you to temporarily disable or permanently delete created CDP policies.

Disabling CDP Policies

To disable a CDP policy:

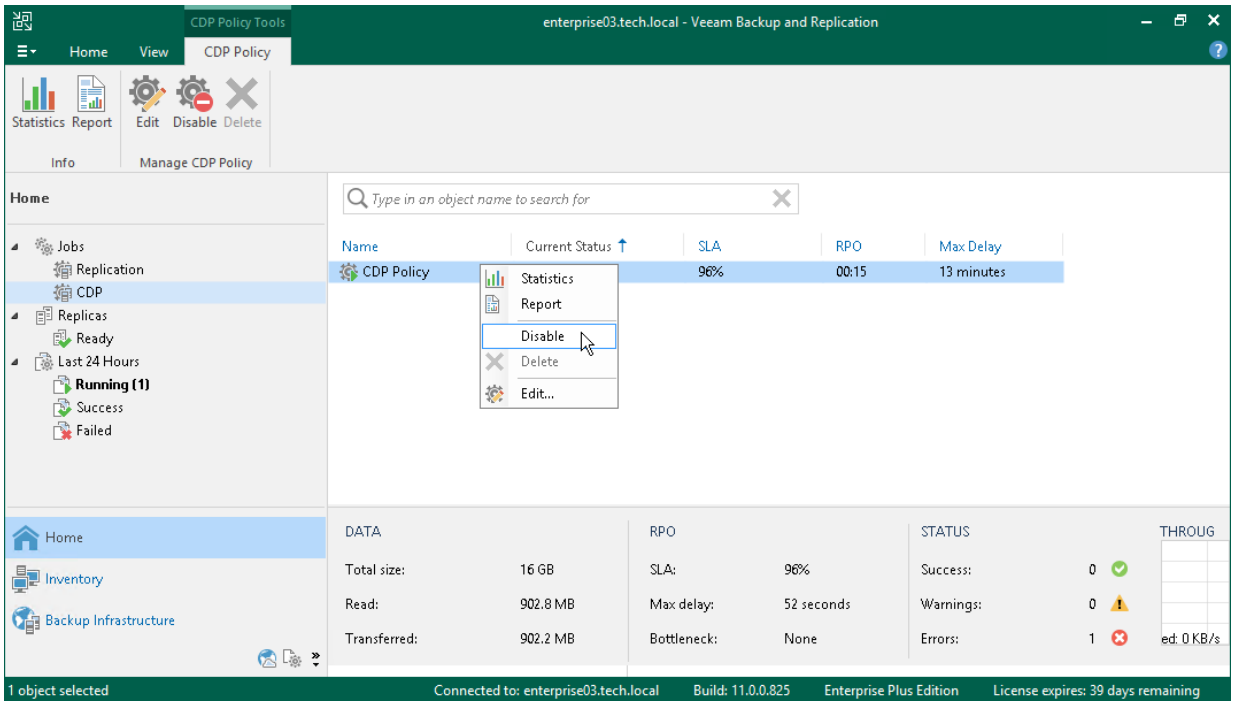
1. Open the **Home** view.
2. In the inventory pane, navigate to the **Jobs** > **CDP** node.
3. In the working area, select the necessary policy and select **Disable** on the ribbon. As an alternative, right-click the necessary policy and select **Disable**.

> **TIP**
>
> To enable a disabled policy, select it and click **Disable** once again.

Deleting CDP Policies

Veeam Backup & Replication allows you to delete only disabled policies. To delete a CDP policy:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Jobs** > **CDP** node.
3. In the working area, select the necessary policy and select **Delete** on the ribbon. As an alternative, right-click the necessary policy and select **Delete**.



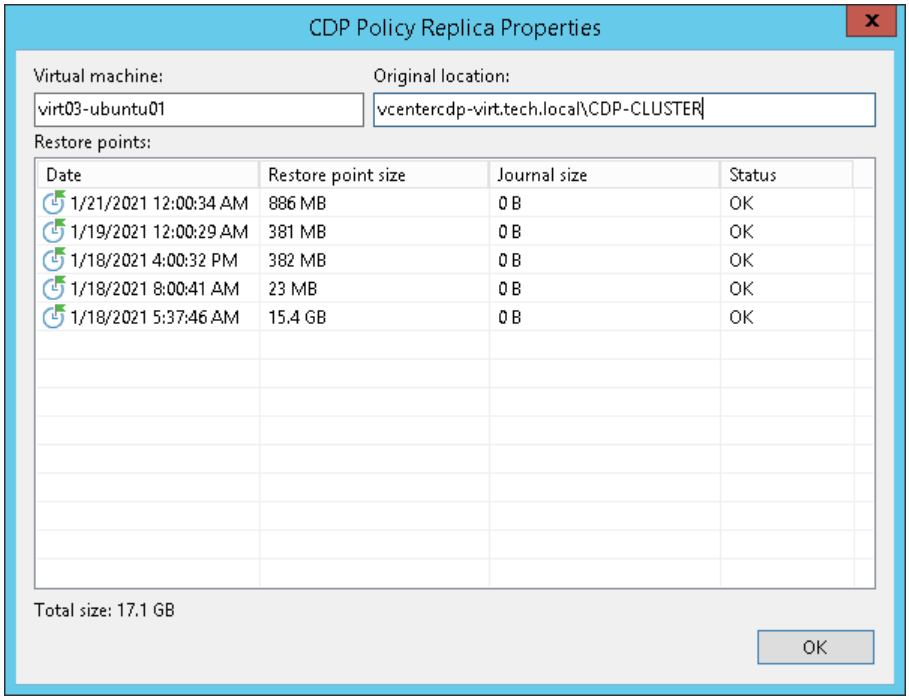
Managing CDP Replicas

To view all created CDP replicas, open the **Home** view and navigate to the **Replicas** node. The working area displays the full list of the created replicas. Here, you can view replica properties and delete replicas from the configuration database or disk.

Viewing Replica Properties

To view replica properties:

1. Open the **Home** view.
2. In the inventory pane, click the **Replicas** node.
3. In the working area, select the necessary replica and select **Properties** on the ribbon. As an alternative, right-click the necessary replica and select **Properties**.



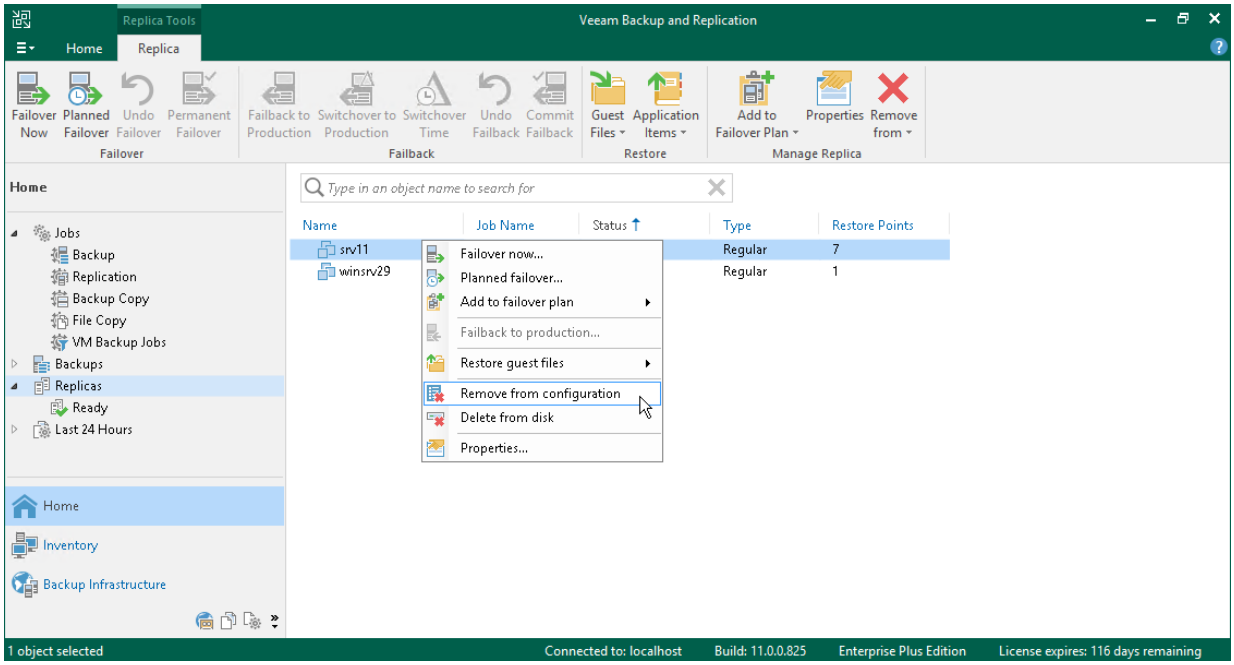
Removing Replicas from Configuration

When you remove replicas from the configuration, Veeam Backup & Replication removes the records about the replicas from the configuration database, stops showing the replicas in Veeam Backup & Replication console and also stops synchronizing their state with the state of the original VMs. However, the actual replicas remain on hosts, so that you can start the replicas manually in the VMware vSphere client.

> **NOTE**
>
> Mind the following:
>
> - You can remove records only about replicas that are in the *Ready* state.
> - When you remove from the configuration a VM that is replicated as a standalone object, Veeam Backup & Replication removes this VM from the initial replication job. When you remove from the configuration a VM that is replicated as part of a VM container, Veeam Backup & Replication adds this VM to the list of exclusions in the CDP policy.

To remove records about VM replicas from the Veeam Backup & Replication console and configuration database:

1. Open the **Home** view.
2. In the inventory pane, click the **Replicas** node.
3. In the working area, select the necessary replica and click **Remove from** > **Configuration** on the ribbon. As an alternative, right-click the necessary replica and select **Remove from configuration**.



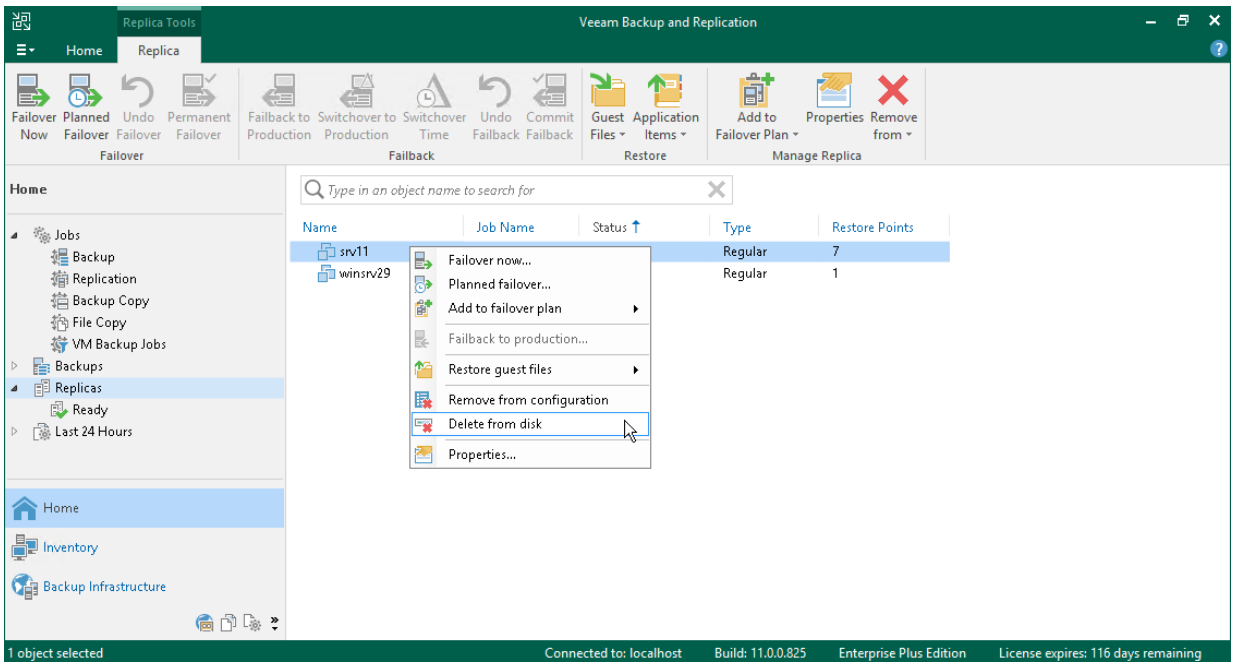
Deleting Replicas from Disk

When you delete replicas from disks, Veeam Backup & Replication removes the replicas not only from the Veeam Backup & Replication console and configuration database, but also from host storage.

> **NOTE**
>
> Mind the following:
>
> - You can delete records only about replicas that are in the *Ready* state.
> - Do not delete replica files from host storage manually, use the **Delete from disk** option instead. If you delete replica files manually, subsequent replication sessions will fail.

To delete replica files from disks:

1. Open the **Home** view.
2. In the inventory pane, click the **Replicas** node.
3. In the working area, select the necessary replica and click **Remove from** > **Disk** on the ribbon. As an alternative, right-click the replica and select **Delete from disk**.



Failover and Failback for CDP

Failover and failback operations help you ensure that your business will function even if a disaster strikes your production site. Failover is a process of switching from the original VM on the source host to its VM replica on a host in the disaster recovery site. Failback is a process of returning from the VM replica to the original VM.

Veeam Backup & Replication provides the following failover and failback operations:

- **Perform failover**

  When you perform failover, you shift all processes from the original VM in the production site to the VM replica in the disaster recovery site. Failover is an intermediate step that needs to be finalized: you can undo failover, perform permanent failover or perform failback.

  For more information on how failover is performed, see Failover.

- **Create failover plan**

  When you create a failover plan, you define the order in which Veeam Backup & Replication must perform failover for VMs, and an interval of time for which Veeam Backup & Replication must wait before starting the failover operation for the next VM in the list.

  For more information on failover plans, see Failover Plan.

- **Perform permanent failover**

  When you perform permanent failover, you permanently switch from the original VM to a VM replica and use this replica as the original VM. You can use this scenario if the original VM and VM replica are located in the same site and are nearly equal in terms of resources. Otherwise, perform failback.

  For more information on how permanent failover is performed, see Permanent Failover.

- **Undo failover**

  When you undo failover, you switch back to the original VM and discard all changes made to the VM replica while it was running. You can use the undo failover scenario if you have failed over to the VM replica for testing and troubleshooting purposes, and you do not need to synchronize the original VM state with the current state of the replica.

  For more information on how failover undo is performed, see Failover Undo.

- **Perform failback**

  When you perform failback, you switch back to the original VM and send to the original VM all changes that took place while the VM replica was running. If the source host is not available, you can recover a VM with the same configuration as the original VM and switch to it. For more information on how failback is performed, see Failback.

  When you perform failback, changes are only sent to the original/recovered VM but not published. You must test whether the original/recovered VM works with these changes. Depending on the test results, you can do the following:

  - **Commit failback**. When you commit failback, you confirm that the original/recovered VM works as expected and you want to get back to it.

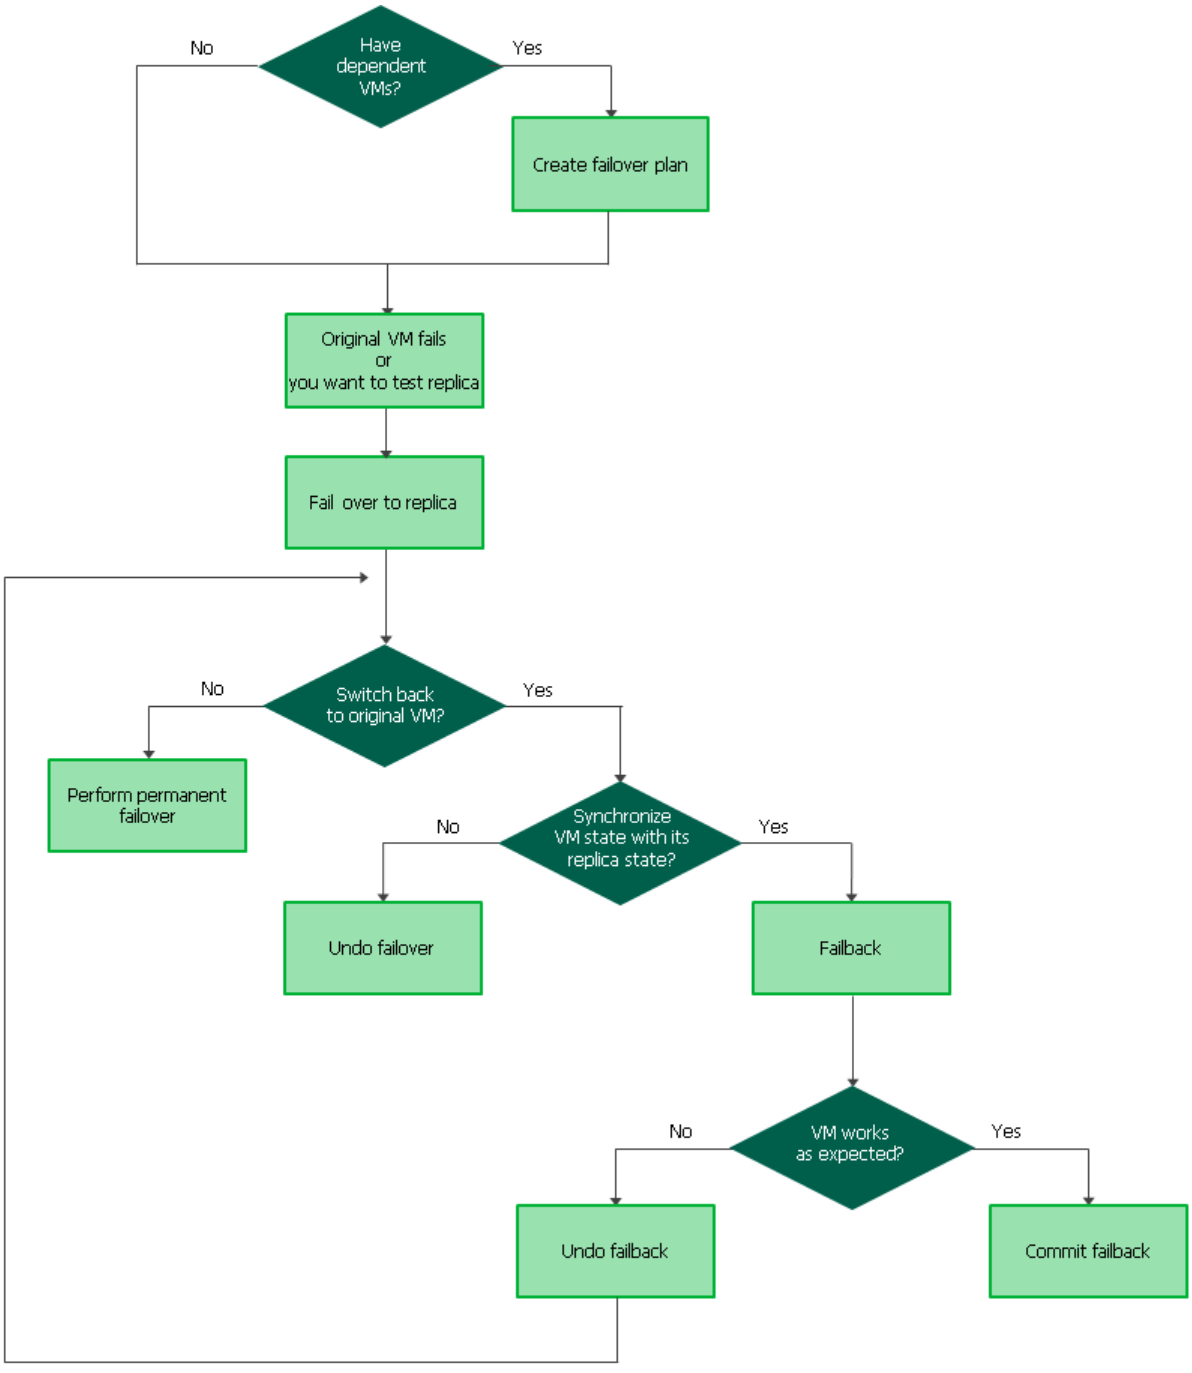
    For more information on how failback commit is performed, see Failback Commit.

  - **Undo failback**. When you undo failback, you confirm that the original/recovered VM is not working as expected and you want to get back to the VM replica.

    For more information on how failback undo is performed, see Failback Undo.

Veeam Backup & Replication supports failover and failback operations for one VM and for several VMs. In case one or several hosts fail, you can use batch processing to restore operations with minimum downtime.

The following scheme can help you decide which steps are preferable when you fail over to a replica.



Related Topics

- Performing Failover
- Performing Permanent Failover
- Failover Plans
- Undoing Failover

- Performing Failback
- Committing Failback
- Undoing Failback

Failover

Failover is a process of switching from the original VM in the production site to its VM replica in the disaster recovery site. During failover, Veeam Backup & Replication recovers the VM replica to the required restore point and shifts all I/O processes from the original VM to its replica. As a result, you have a fully functional VM within a couple of seconds, and your users can access services and applications with minimum disruption.

You can fail over to replicas not only when a disaster strikes the production site, but also to test replicas for recoverability. If the original VMs and VM replicas are located in the same network, consider temporary disconnecting the original VMs from the network to avoid IP address or machine name conflicts.

The failover operation is performed in the following way:

1. Veeam Backup & Replication puts all replication activities on hold.
2. The state of the VM replica is changed from *Ready* to *Failover*.
3. Veeam Backup & Replication recovers a VM replica to the required restore point.
4. Veeam Backup & Replication powers on the VM replica.

   If you perform failover for testing purposes, and the original VM still exists and is running, the original VM remains powered on.
5. All changes made to the VM replica disks while the replica is running in the *Failover* state are written to the protective virtual disks (*<disk_name>-interim.vmdk* files).

Failover is an intermediate step that needs to be finalized. You can perform the following operations:

- Undo failover
- Perform permanent failover
- Perform failback.

Performing Failover

To perform failover, use the **VMware Failover** wizard.

Before You Begin

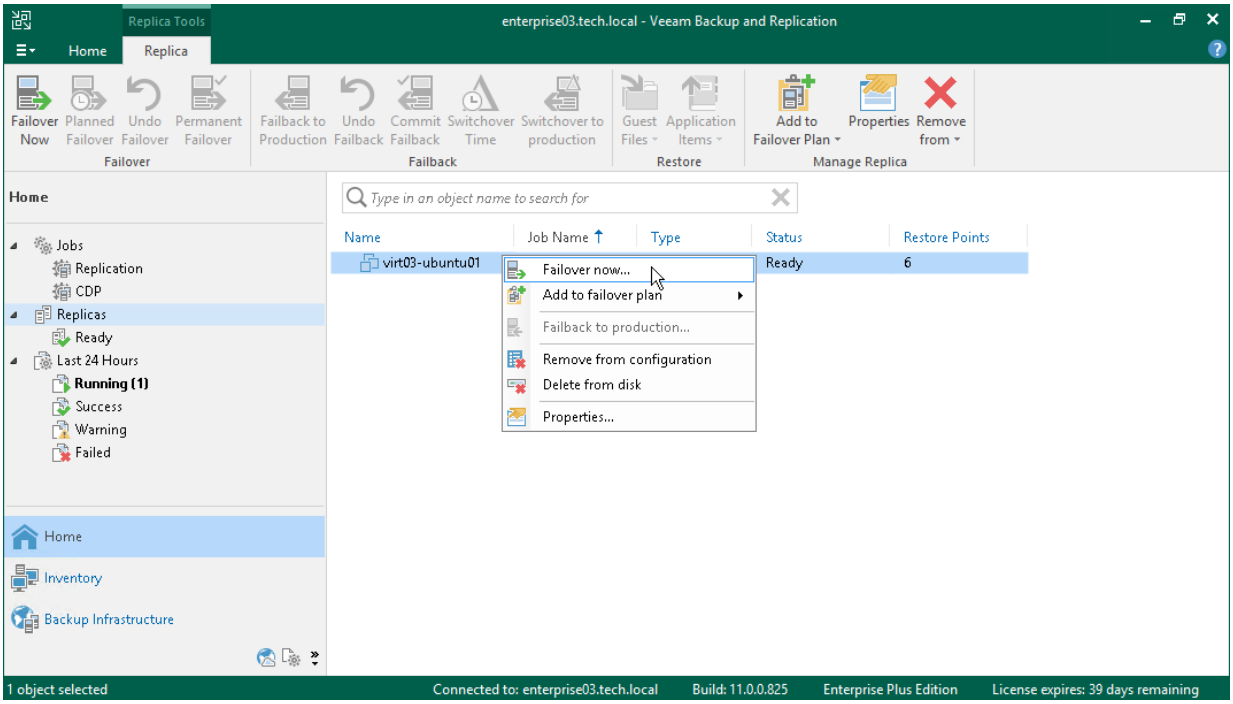
Before you fail over to a CDP replica, check the following prerequisites:

- You can perform failover for VMs that have been successfully replicated at least once.
- CDP replicas must be in the *Ready* state.

Step 1. Launch VMware Failover Wizard

To launch the **VMware Failover** wizard, do one of the following:

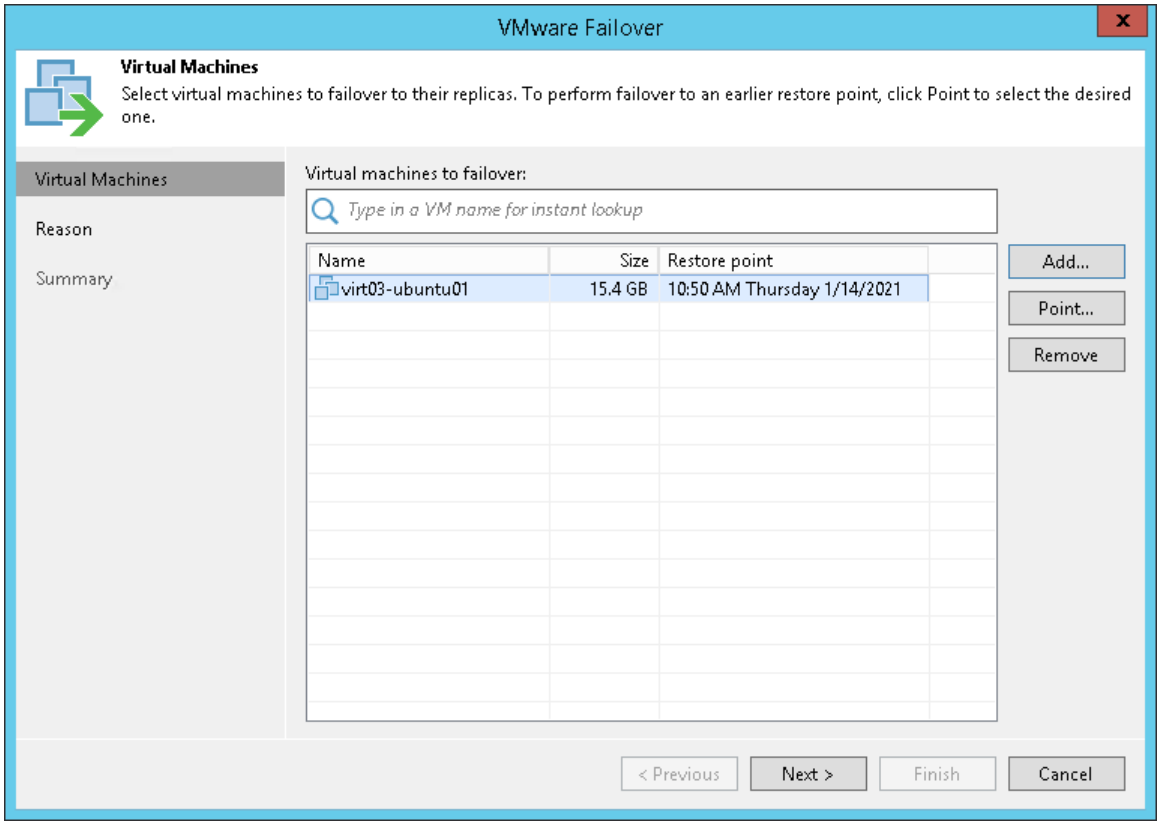
- On the **Home** tab, click **Restore** > **VMware vSphere** > **Restore from replica** > **Entire replica** > **Failover to a replica**.
- Open the **Home** view. In the inventory pane select the **Replicas** > **Ready** node. In the working area, select the necessary replica and click **Failover Now** on the ribbon.
- Open the **Home** view. In the inventory pane select the **Replicas** > **Ready** node. In the working area, right-click the necessary replica and select **Failover now**.



Step 2. Select VMs

At the **Virtual Machines** step of the wizard, you can modify a list of VMs from which you fail over. To add VMs or VM containers, click **Add** > **From infrastructure** if you want to add VMs from the virtual infrastructure, or **Add** > **From replicas** if you want to add VMs from existing replicas. Then select the necessary VMs or VM containers. If you select VM containers, Veeam Backup & Replication will expand them to a plain VM list.

To quickly find the necessary VMs within existing replicas, you can use the search field at the top of the window.

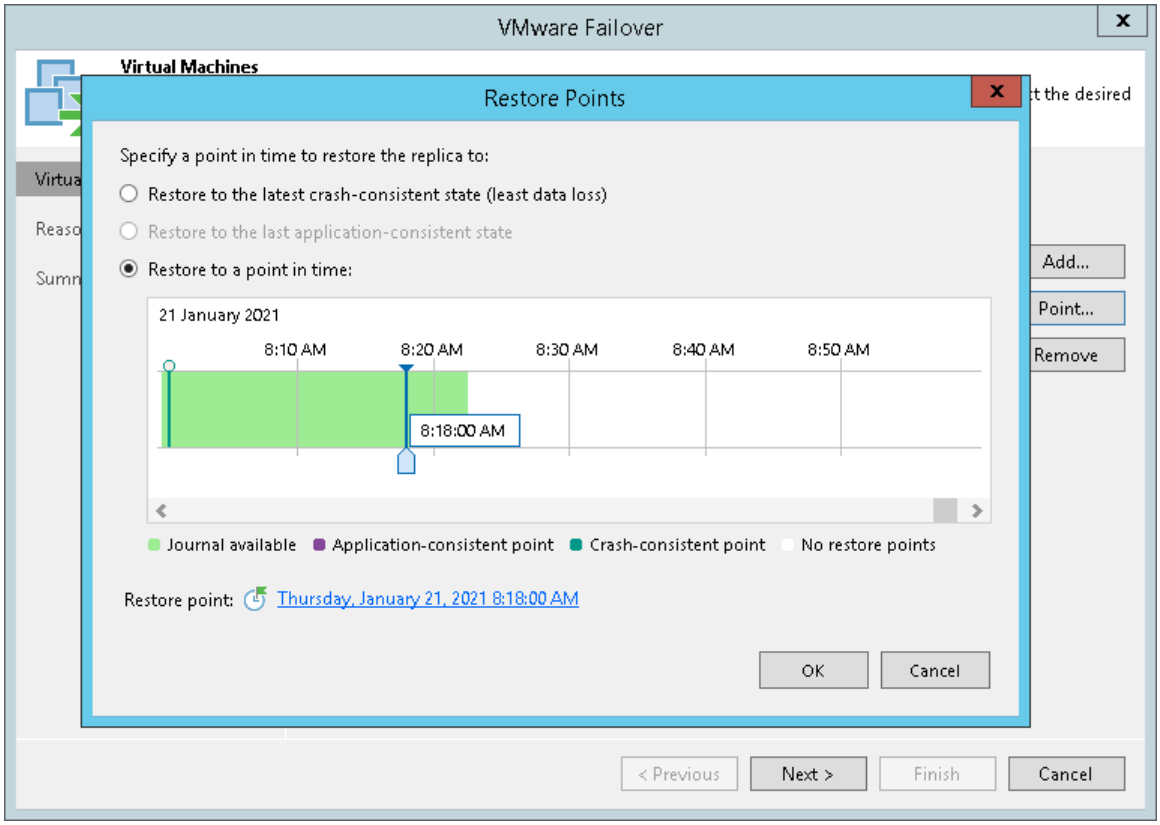


Step 3. Select Restore Points

At the **Virtual Machines** step of the wizard select to which state of VM replicas you want to fail over:

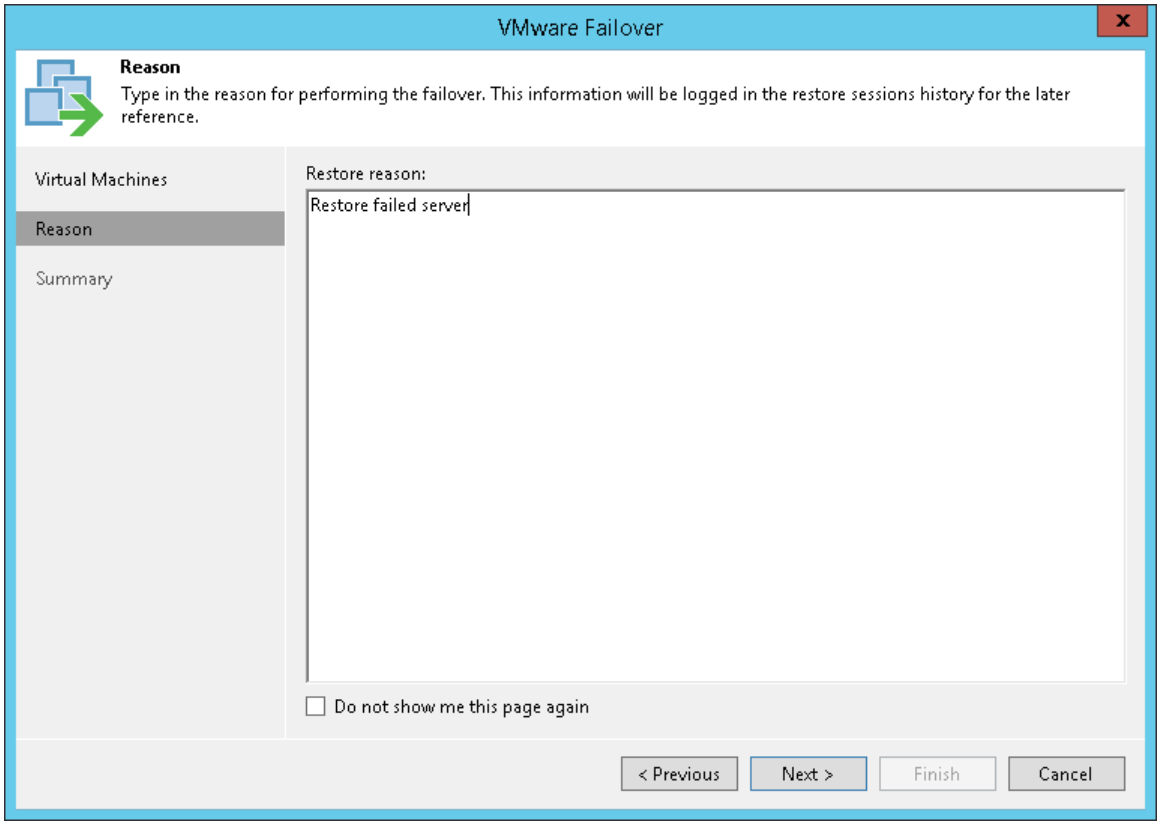
1. In the **Virtual machines to failover** list, select the necessary VM and click **Point**.
2. In the **Restore Points** window, select whether you want to fail over to the latest available crash-consistent restore point, to the latest long-term application-consistent restore point or to a specific point in time.

   If you fail over to a specific point in time, use the right and left arrows on the keyboard to select the required restore point. To quickly find a long-term restore point, click a link that shows a date. In the opened window, you will see a calendar where you can select the necessary day. In the **Timestamp** section, you will see long-term restore points created during the selected day.



Step 4. Specify Failover Reason

At the **Reason** step of the wizard, enter a reason for failing over to the replicas. The information you provide will be saved in the session history and you can reference it later.



Step 5. Finish Working with Wizard

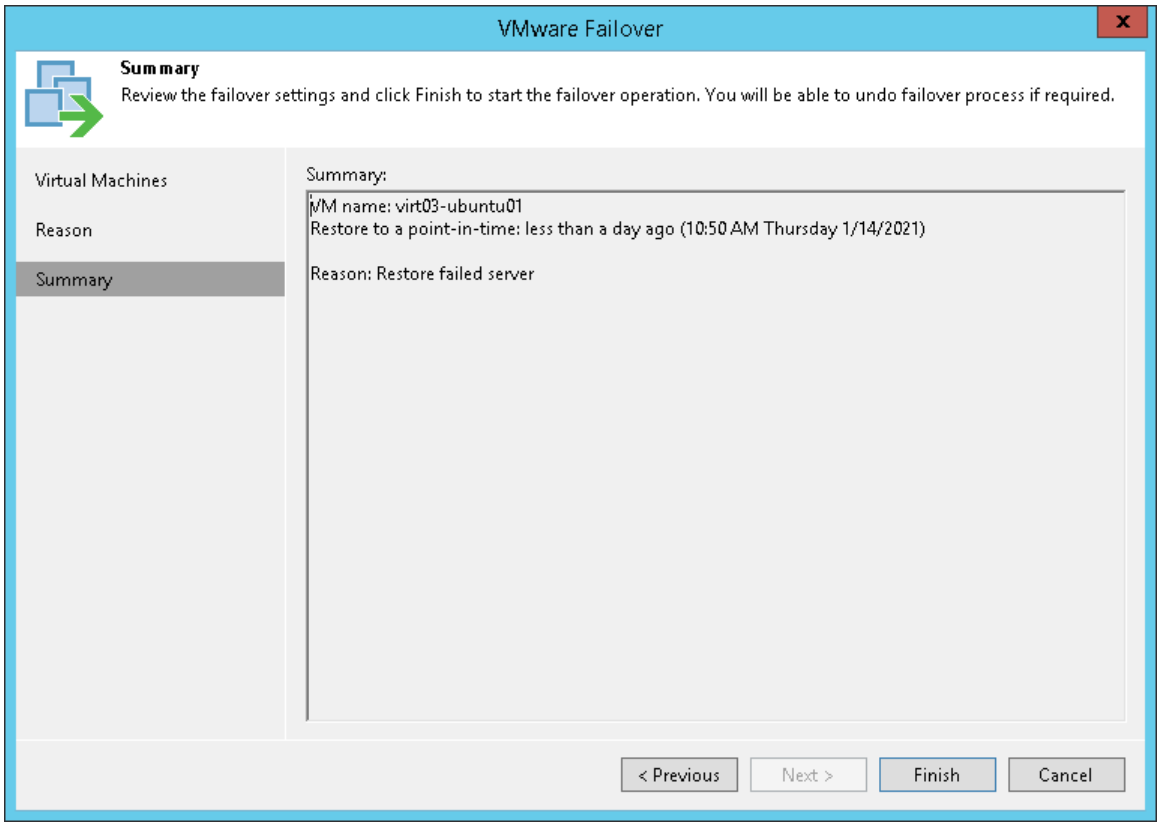
At the **Summary** step of the wizard, review details of the failover task and click **Finish** to exit the wizard.

During the failover process, Veeam Backup & Replication will request the target host to to power on the CDP replicas.

What You Do Next

Failover is an intermediate step that needs to be finalized. You can finalize failover in the following ways:

- Perform permanent failover
- Undo failover
- Perform failback



Failover Plans

A failover plan helps you perform failover for dependent VMs one by one. In the failover plan, you define the order in which VMs must be processed and an interval of time for which Veeam Backup & Replication must wait before starting the failover operation for the next VM in the list. For more information on the failover plan, see Failover Plan.

You can add a VM to an existing failover plan or to a new plan as described in Creating Failover Plans. For more information on how to manage failover plans, see Running Failover Plans and Undoing Failover by Failover Plans.

Permanent Failover

Permanent failover is one of the ways to finalize failover. When you perform permanent failover, you permanently switch from the original VM to it replica. As a result of permanent failover, the VM replica stops acting as a replica and starts acting as the production VM.

> **NOTE**
>
> We recommend you to perform permanent failover only if the original VM and its replica are located in the same site and are nearly equal in terms of resources. In this case, users will not experience any latency in ongoing operations. Otherwise, perform failback.

The permanent failover operation is performed in the following way:

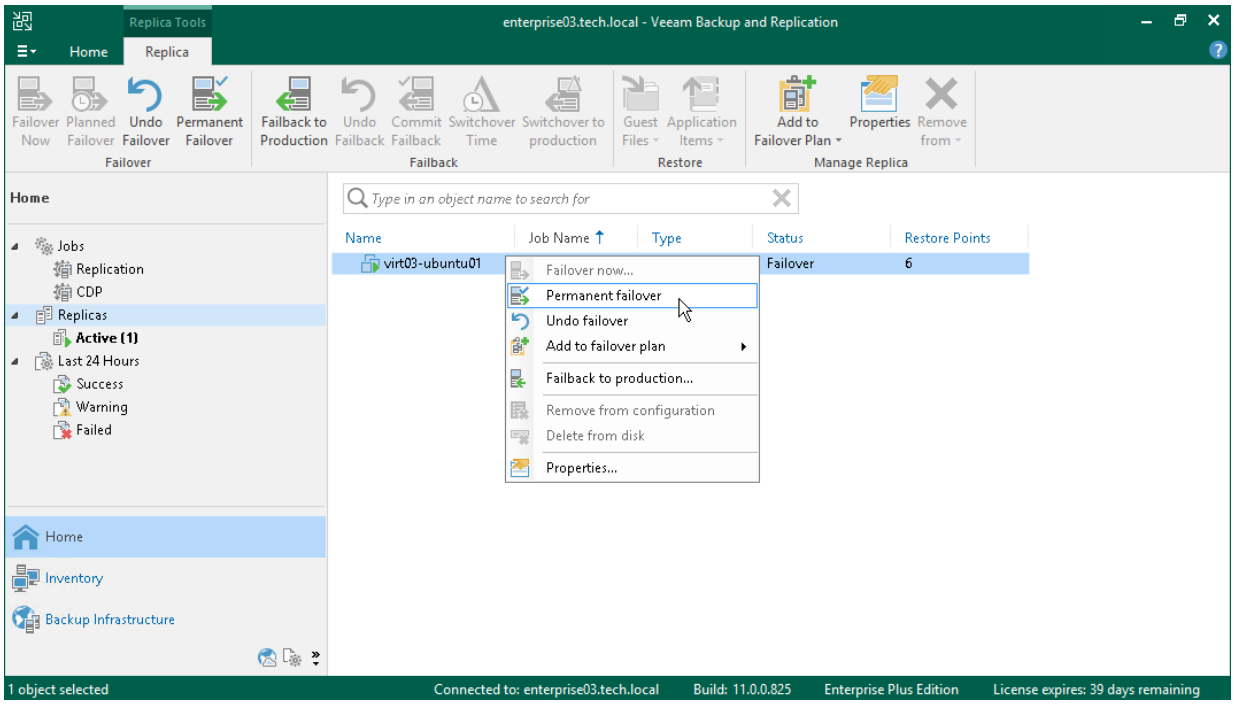
1. Veeam Backup & Replication powers off the VM replica.
2. Veeam Backup & Replication removes short-term and long-term restore points of the VM replica from the replication chain and deletes associated files from the datastore. Changes that were written to the protective virtual disks (*<disk_name>-interim.vmdk*) are committed to the VM replica to bring the VM replica to the most recent state.
3. Veeam Backup & Replication removes the VM replica from the list of replicas in the Veeam Backup & Replication console.
4. To protect the VM replica from corruption after permanent failover is complete, Veeam Backup & Replication reconfigures the CDP policy: adds the VM replica in the list of processed VMs and adds the original VM to the list of exclusions. When the CDP policy starts, the original VM is skipped from processing. As a result, no data is written to the working VM replica.

Performing Permanent Failover

To perform permanent failover, do one of the following:

- Open the **Home** view. In the inventory pane navigate to the **Replicas** > **Active** node. In the working area, select the necessary replica and click **Permanent Failover** on the ribbon.

- Open the **Home** view. In the inventory pane navigate to the **Replicas** > **Active** node. In the working area, right-click the necessary replica and select **Permanent failover**.



Failover Undo

Failover undo is one of the ways to finalize failover. When you undo failover, you switch back from a VM replica to the original VM. Veeam Backup & Replication discards all changes made to the VM replica while it was in the *Failover* state.

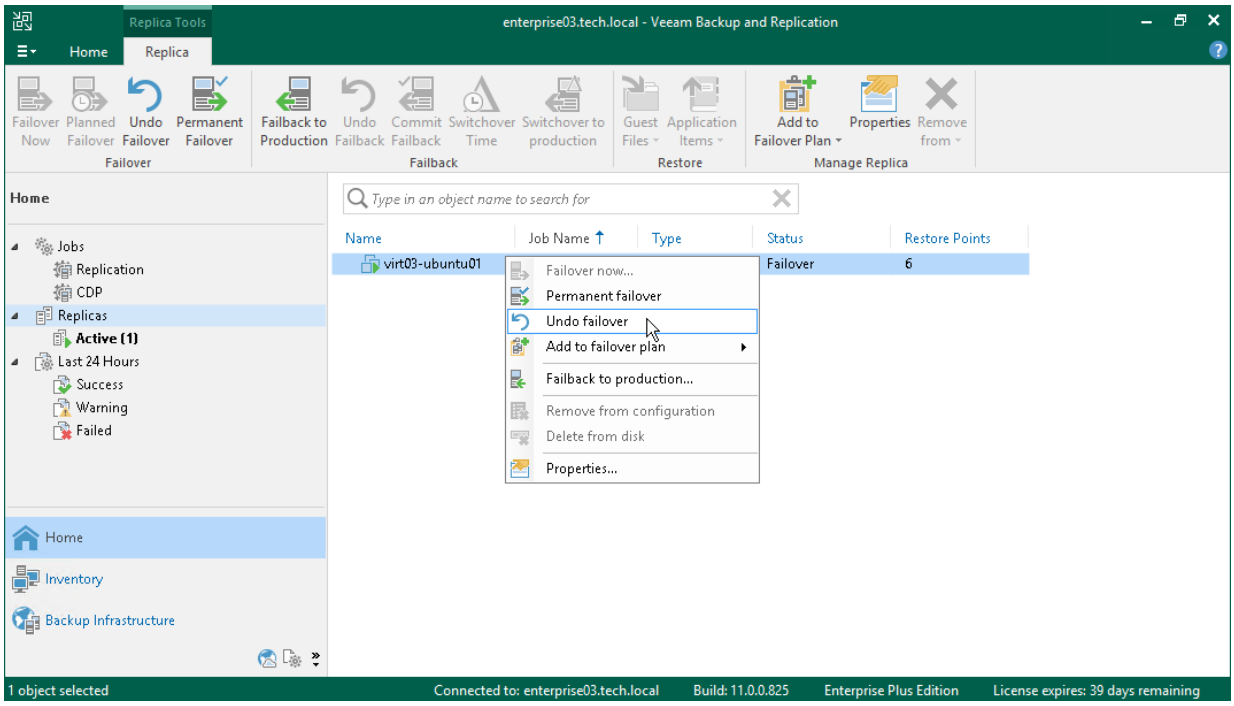
The undo failover operation is performed in the following way:

1. Veeam Backup & Replication reverts the VM replica to its pre-failover state. To do this, Veeam Backup & Replication powers off the VM replica and gets it back to the latest restore point in the replication chain. Changes that were written to the protective virtual disks (*<disk_name>-interim.vmdk*) while the VM replica was in the *Failover* state are discarded.
2. The state of the VM replica gets back to *Ready*, and Veeam Backup & Replication resumes replication activities for the original VM on the source host.

Undoing Failover

To undo failover:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Replicas** > **Active** node.
3. In the working area, select the necessary replica and click **Undo Failover** on the ribbon. As an alternative, you can right-click the replica and select **Undo failover**.



Failback

Failback is one of the ways to finalize failover. When you perform failback, you switch back to the production VM from a VM replica, shift I/O processes from the disaster recovery site to the production site.

Veeam Backup & Replication provides you the following options to perform failback:

- You can fail back to the original VM in the original location.
- You can fail back to a VM already recovered to a new location. This VM must be recovered before you perform failback. For example, you can recover the VM from a backup.
- You can fail back to a VM recovered from a replica to a new location, or to any location but with different settings. The VM will be recovered from the replica during the failback process.

The first two options help you decrease recovery time and the use of the network traffic because Veeam Backup & Replication needs to transfer only differences between the original/recovered VM and VM replica. For the third option, Veeam Backup & Replication needs to transfer the whole VM data, including its configuration and virtual disk content. Use the third option if there is no way to use the original VM or restore it from a backup.

Veeam Backup & Replication performs failback in two phases:

- **First phase**: Veeam Backup & Replication synchronizes the state of the production VM (the original VM, an already recovered VM or a VM that will be recovered from the replica) with the current state of its replica. This phase may take a lot of time especially if VM is large. While Veeam Backup & Replication performs the first phase of failback, VM replicas are still up and running, users can access these VMs and perform daily routine tasks as normal.
- **Second phase**: Veeam Backup & Replication switches all processes from the VM replica to the production VM, turns off the replica and also sends to the production VM changes made to the VM replica since the end of the first phase.

  The time when the second phase starts depends on how you want to switch from the replica to the production VM. You can switch to the production VM automatically, at the scheduled time or manually. If you select to switch automatically, the second phase will start right after the first phase finishes. If you select to switch at the scheduled time or manually, the second phase will start at the time you want.

The process of failing back to the original VM or an already recovered VM differs from the process of failing back to a VM recovered from a replica:

- How failback to the original VM and already recovered VM works
- How failback to a VM recovered from a replica works

How Failback to Original VM or Already Recovered VM Works

When you fail back to the original VM or an already recovered VM, Veeam Backup & Replication performs the following operations during the first phase:

1. If the original VM is running, Veeam Backup & Replication powers it off.

2. Veeam Backup & Replication calculates the difference between disks of the production VM and disks of the VM replica in the *Failover* state. Difference calculation helps Veeam Backup & Replication understand what data needs to be transferred to the production VM to synchronize its state with the state of the VM replica.

   If you fail back to the original VM in the original location and you have enabled the **Quick rollback** option, difference calculation can be performed much faster than without this option enabled. For more information on quick rollback, see Quick Rollback.

3. Veeam Backup & Replication transfers the data that was detected at the previous step to the production VM. The transferred data is written to the production VM.

4. Veeam Backup & Replication changes the state of the VM replica from *Failover* to *Ready to switch*.

During the second phase, Veeam Backup & Replication performs the following operations:

1. The guest OS of the VM replica is shut down or the VM replica is powered off.

   If VMware Tools are installed on the VM replica, Veeam Backup & Replication tries to shut down the replica guest OS. If nothing happens after 15 minutes, Veeam Backup & Replication powers off the VM replica. If VMware Tools are not installed on the VM or the VM is suspended, Veeam Backup & Replication powers off the VM. The VM replica remains powered off until you commit failback or undo failback.

2. Veeam Backup & Replication calculates the difference between disks of the production VM and disks of the VM replica. Difference calculation helps Veeam Backup & Replication understand what data was changed while the VM replica was in the *Ready to switch* state.

3. Sends data changed on the VM replica while it was in the *Ready to switch* state to the production VM.

4. The state of the VM replica is changed from *Ready to switch* to *Failback*.

5. [If you fail back to a recovered VM] Veeam Backup & Replication updates the ID of the original VM in the Veeam Backup & Replication configuration database. The ID of the original VM is replaced with the ID of the recovered VM.

6. If you have selected to power on the production VM after failback, Veeam Backup & Replication powers on the production VM on the host.

How Failback to VM Recovered from Replica Works

When you fail back to a VM recovered from a replica, Veeam Backup & Replication performs the following operations during the first phase:

1. Veeam Backup & Replication requests vCenter Server to create on the target host an empty VM with the same configuration as the VM replica. vCenter Server registers the created production VM.

2. Veeam Backup & Replication transfers data of the VM replica to the production VM to update the production VM state to the VM replica state.

3. Veeam Backup & Replication changes the state of the VM replica from *Failover* to the *Ready to switch*.

During the second phase, Veeam Backup & Replication performs the same operations as described in How Failback to Original VM or Already Recovered VM Works except for the step 2. When you fail back to a VM recovered from a replica, Veeam Backup & Replication does not calculate the difference between disks.

Failback is an intermediate step that needs to be finalized. If the production VM works as expected and you want to get back to it, commit failback. If the VM does not work as expected, undo failback.

Quick Rollback

Quick rollback helps you significantly reduce the failback time. You can use quick rollback if you fail back from a VM replica to the original VM in the original location.

During failback, Veeam Backup & Replication calculates differences between disks of the original VM and disks of the VM replica. With the quick rollback option enabled, Veeam Backup & Replication compares only those disk sectors that have changed during the VM replica was in the *Failover* state instead of comparing entire disks. To get information about the changed disk sectors, Veeam Backup & Replication uses VMware vSphere Changed Block Tracking (CBT).

As a result of enabling quick rollback, difference calculation becomes much faster. After the differences are calculated, Veeam Backup & Replication performs failback in a regular way: transport changed blocks to the original VM, powers off the VM replica and synchronizes the original VM with the VM replica once again.

Requirements for Quick Rollback

To perform quick rollback, make sure that the following requirements are met:

- You fail back to the original VM in the original location.
- Do not use quick rollback if the problem occurred at the VM hardware level, storage level or due to a power loss.

  Use quick rollback if you fail back to the original VM that had a problem at the guest OS level — for example, there was an application error or a user accidentally deleted a file on the original VM guest OS.
- CBT must be enabled for the original VM.

Limitations for Quick Rollback

During the first replication job session after failback with quick rollback, CBT on the original VM is reset. Due to that Veeam Backup & Replication will read data of the entire VM.

Performing Failback

To perform failback, use the **Failback** wizard.

Before You Begin

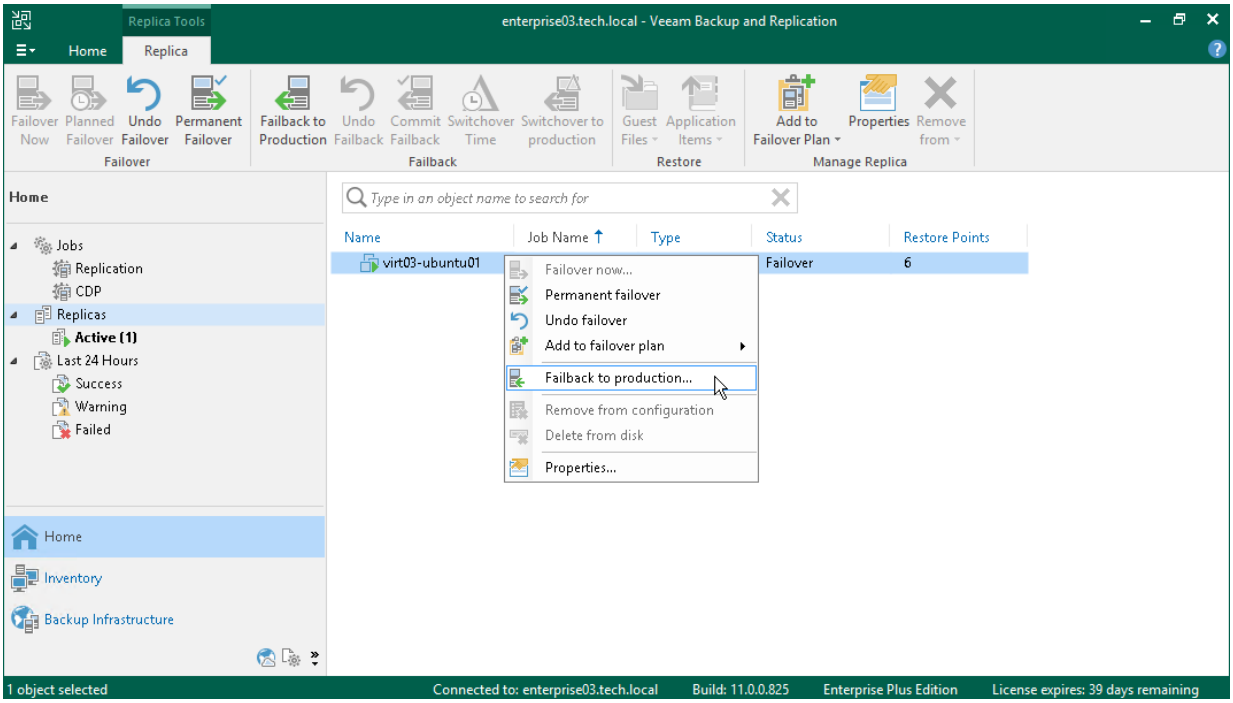
Before you perform failback, check the following prerequisites:

- VMs for which you plan to perform failback must be successfully replicated at least once.
- CDP replicas from which you want to fail back must be in the *Failover* state.

Step 1. Launch Failback Wizard

To launch the **Failback** wizard, do one of the following:

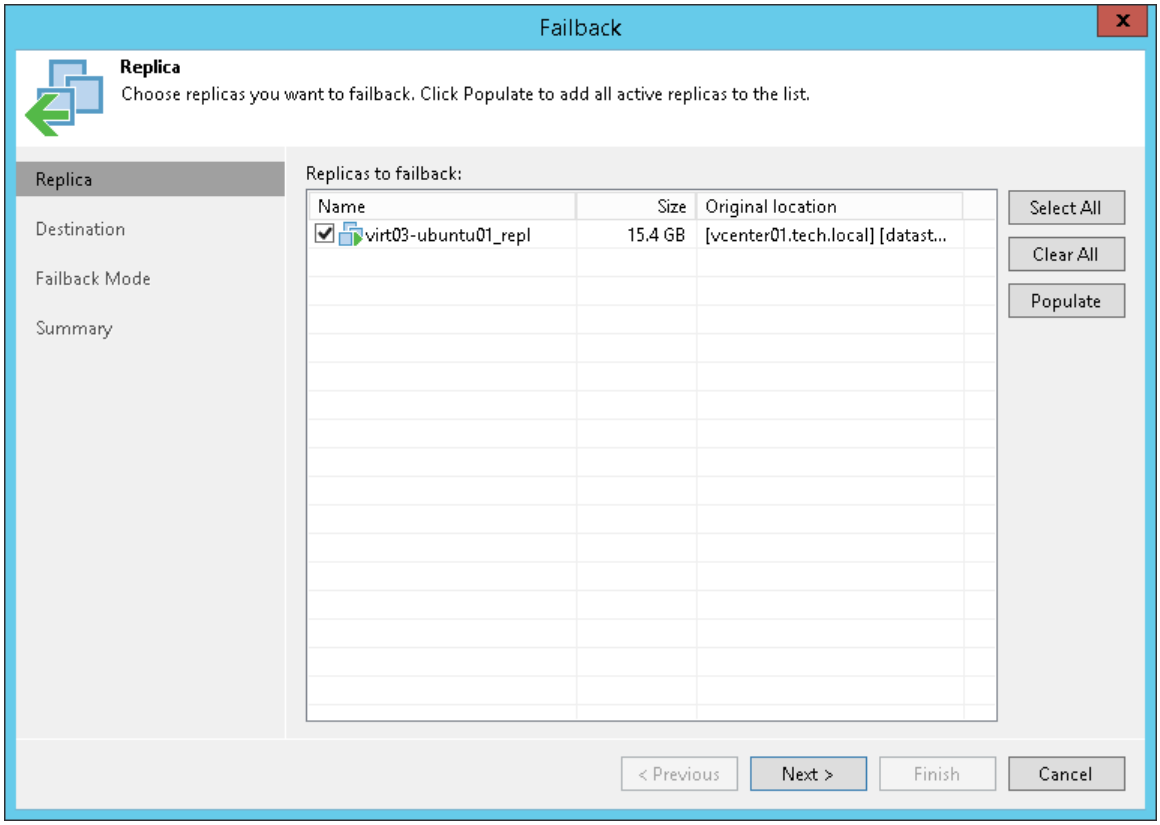
- On the **Home** tab, click **Restore** > **VMware vSphere** > **Restore from replica** > **Entire replica** > **Failback to production**.
- Open the **Home** view. In the inventory pane navigate to the **Replicas** > **Active** node. In the working area, select the necessary replica and click **Failback to Production** on the ribbon.
- Open the **Home** view. In the inventory pane navigate to the **Replicas** > **Active** node. In the working area, right-click the necessary replica and select **Failback to production**.



Step 2. Select Replicas

At the **Replica** step of the wizard, select CDP replicas from which you want to fail back.

To update the list of replicas that are ready for failback (replicas in the *Failover* state), click **Populate**.



Step 3. Select Destination

At the **Destination** step of the wizard, select a failback destination and backup proxies for data transfer during failback:

1. Select a destination for failback. Veeam Backup & Replication supports the following options:
   - **Failback to the original VM** — select this option if you want to fail back to the original VMs that reside on the source hosts. Veeam Backup & Replication will synchronize the state of the original VMs with the current state of their replicas to apply any changes that occurred to the VM replicas while running in the DR site.

     If you select this option, you will proceed to the Summary step of the wizard.
   - **Failback to the original VM restored in a different location** — select this option if the original VMs have already been recovered to a new location, and you want to switch to the recovered VMs from their replicas. Veeam Backup & Replication will synchronize the state of the recovered VMs with the current state of the VM replicas to apply any changes that occurred to the replicas while running in the DR site.

     If you select this option, you will proceed to the Target VM step of the wizard.
   - **Failback to the specified location** — select this option if you want to recover VMs from replicas. You can recover VMs to a new location, or to any location but with different settings (such as network settings, virtual disk type, configuration file path and so on). Select this option if there is no way to fail back to the original VM or an already recovered VM.

     If you select this option, the wizard will include additional steps.

   If you select one of the first two options, Veeam Backup & Replication will send to the original/recovered VMs only differences between the existing virtual disks. Veeam Backup & Replication will not send replica configuration changes such as different IP address or network settings (if replica Re-IP and network mapping were applied), new hardware or virtual disks added while the replicas were in the *Failover* state.

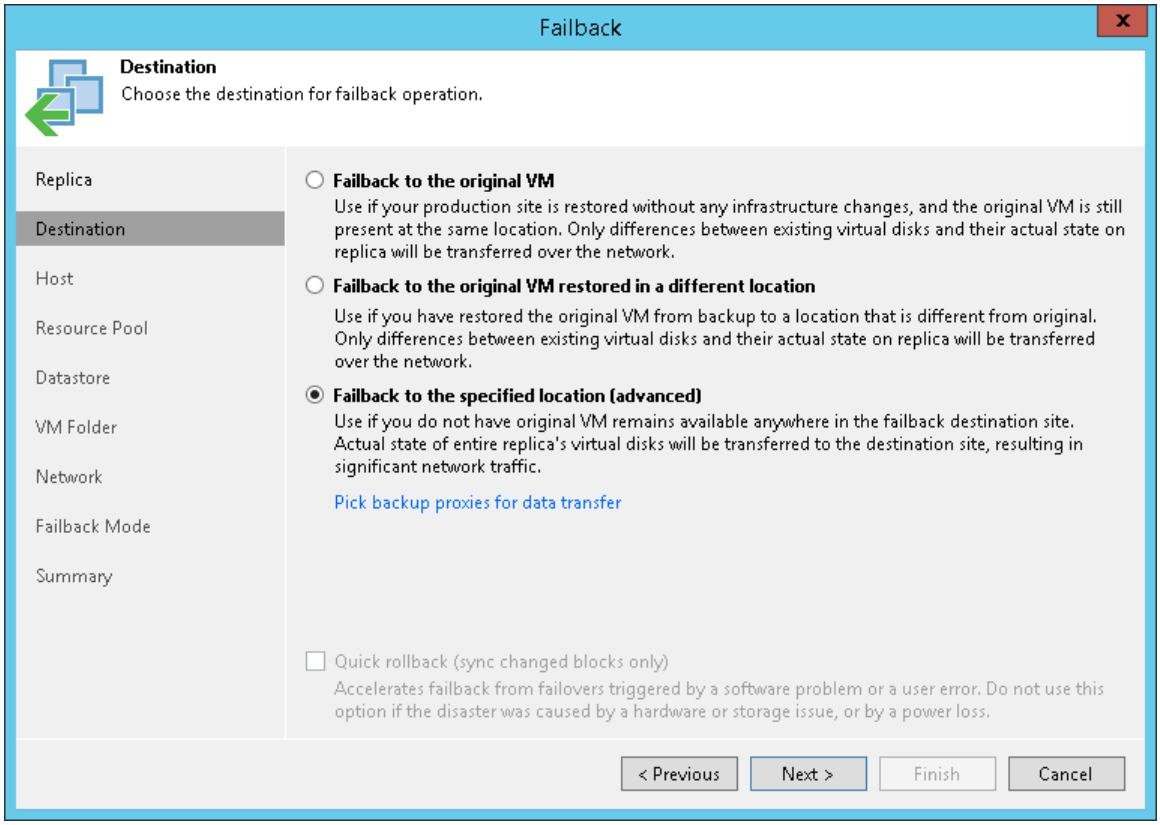
   If you select **Failback to the specified location**, Veeam Backup & Replication will send to the specified location whole VM data, including VM configurations and virtual disk content.

2. To select which backup proxies will be used for data transfer, click **Pick backup proxies for data transfer**. By default, Veeam Backup & Replication selects proxies automatically.

   If VMs and their replicas reside in different sites, select at least one backup proxy in the production site and one proxy in the disaster recovery site. If VMs and replicas reside in the same site, you can use the same backup proxy as the source and target one.

   We recommend you to select at least two backup proxies in each site to ensure that failback will be performed in case one proxy fails or looses the network connection.

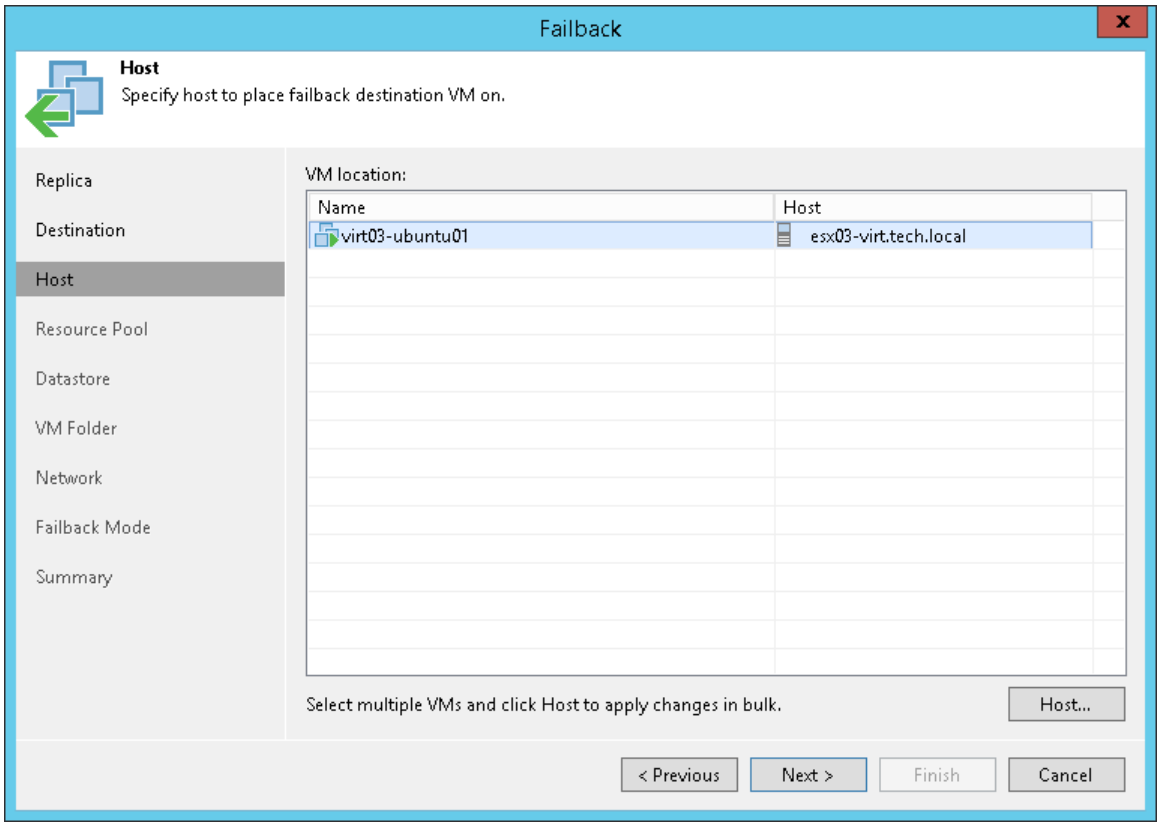
3. [For failback to the original VMs] If you want to fasten failback, and the original VMs had problems at the guest OS level, select the **Quick rollback** check box. For more information on quick rollback, its requirements and limitations, see Quick Rollback.



Step 4. Select Hosts

The **Host** step is available if you have selected the **Failback to the specified location** option at the Destination step.

At the **Host** step of the wizard, specify hosts or clusters where the recovered VMs will be registered. To do this, select the necessary VMs, click **Host** and select a host or cluster.

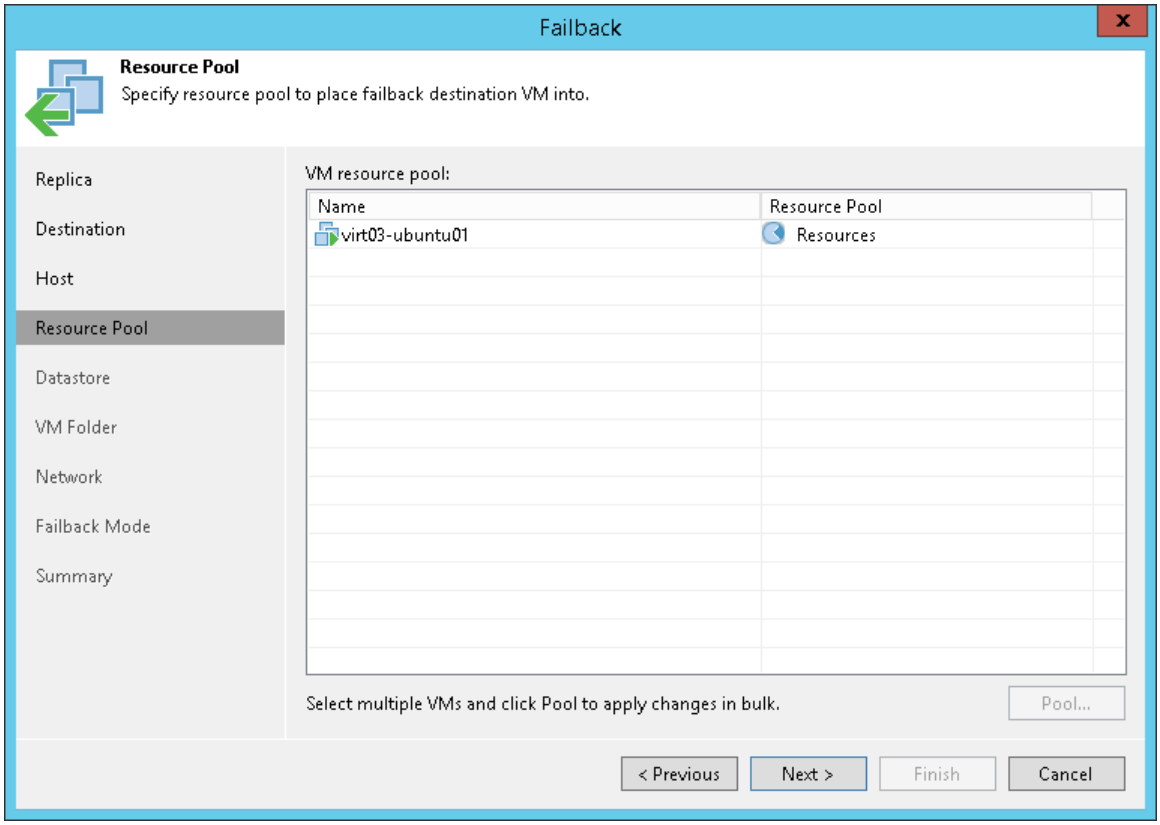


Step 5. Select Resource Pools

The **Resource Pool** step is available if you have selected the **Failback to the specified location** option at the Destination step.

At the **Resource Pool** step of the wizard, select resource pools to which the recovered VMs will be added. To do this, select VMs that you want to add to the same resource pool, click **Pool** and select the necessary resource pool in the **Select Resource Pool** window.

As an alternative, you can select a vAPP to which the restored VM will be included. To find the necessary vAPP, at the left bottom corner of the **Select Resource Pool** window, click the resource pool icon ( ) and select *VirtualApp*.



Step 6. Select Datastores

The **Datastore** step is available if you have selected the **Failback to the specified location** option at the Destination step.

At the **Datastore** step of the wizard, specify datastores where you want to store configuration files and disk files of VMs that will be recovered. Also, you can change disk types.

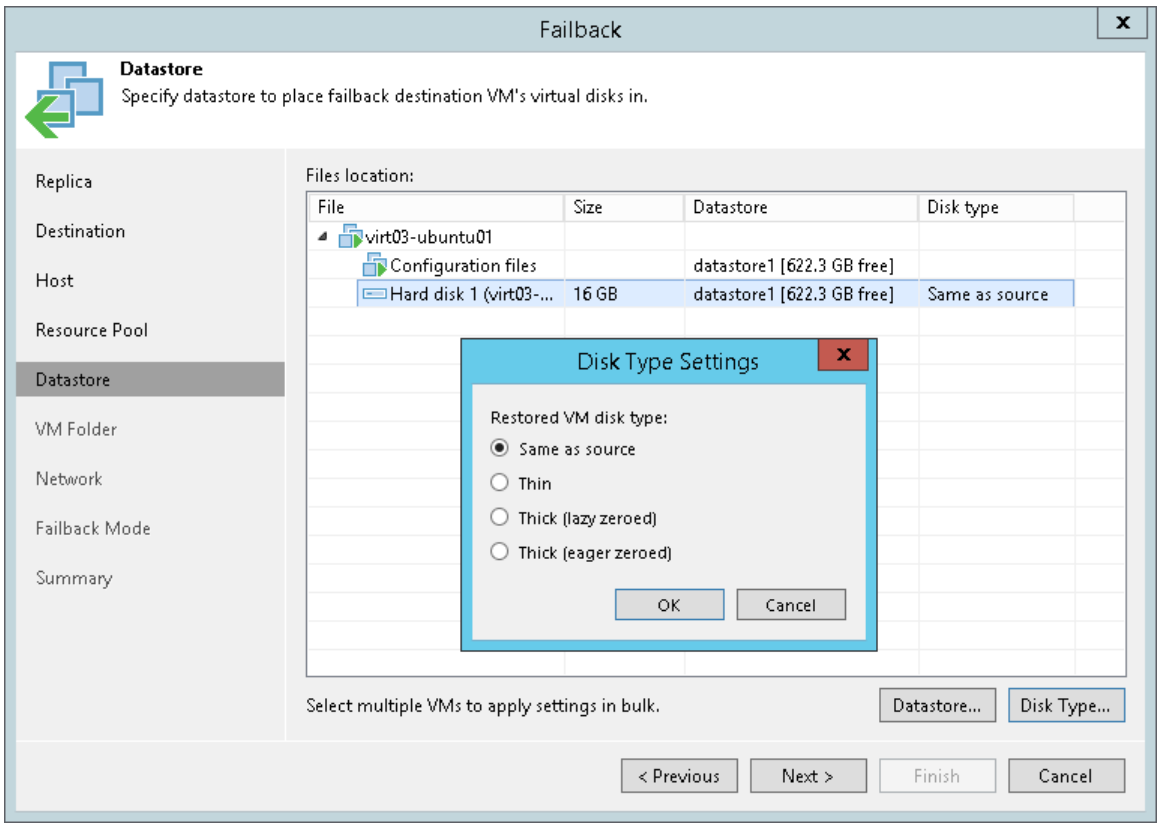
1. To change a datastore where VM files will be stored, select the necessary VMs and click **Datastore**. From the list of available datastores, select the necessary datastore.

   If configuration and disk files of VMs must be placed to different datastores, select files of the necessary type, click **Datastore** and select the necessary datastore.

2. To change disk type settings, select the necessary disk files and click **Disk Type**. In the **Disk Type Settings** window, select the necessary disk type.

   By default, Veeam Backup & Replication preserves disk types of the original VMs.

   > **NOTE**
   >
   > You can change disk types only for VMs with Virtual Hardware version 7 or later.



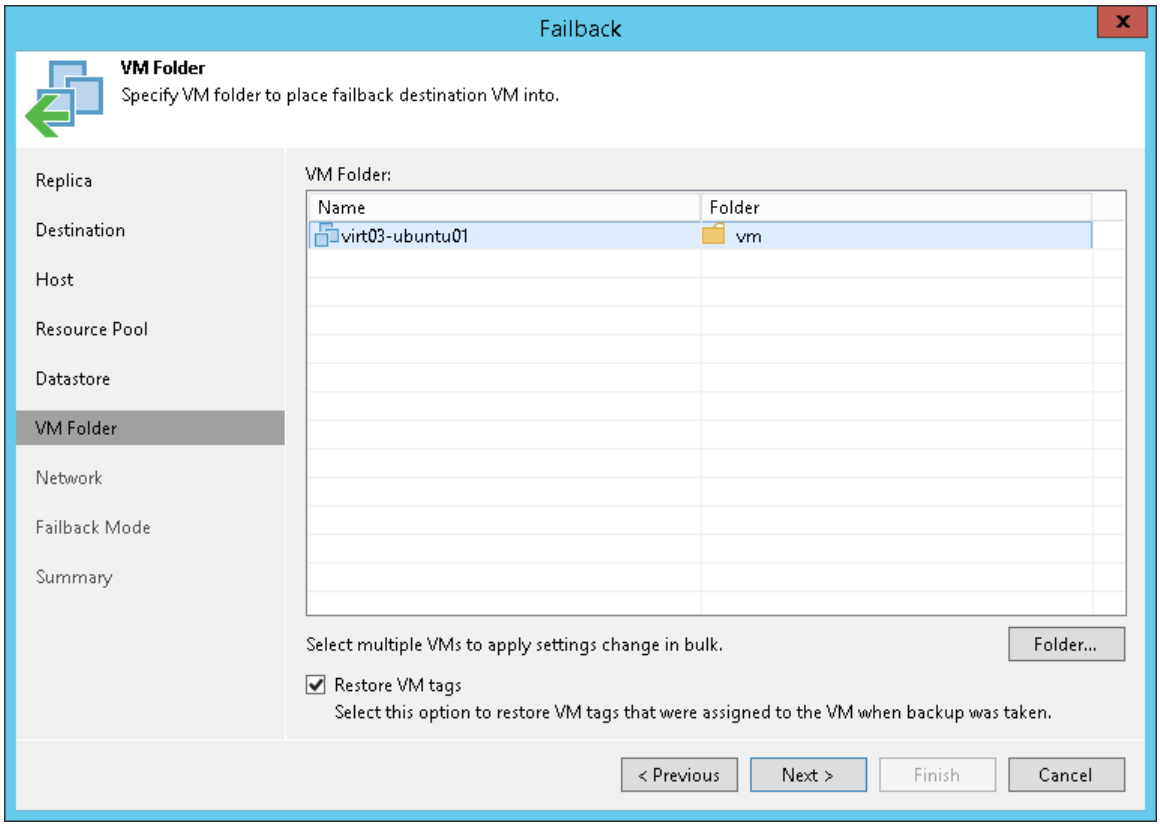
Step 7. Select Folders

The **VM Folder** step is available if you have selected the **Failback to the specified location** option at the Destination step.

At the **VM Folder** step of the wizard, specify folders in the target datastores where all files of the recovered VMs will be stored.

If you want the recovered VMs to have the same tags as the source VMs, select the **Restore VM Tags** check box.

> **NOTE**
>
> Mind the following:
>
> - You can select destination folders only if you recover VMs to destinations other than standalone hosts.
> - You can recover VM tags only if you recover VMs to their original locations, and the original VM tags are still available on the source vCenter Server.



Step 8. Configure Network Mapping

The **Network** step is available if you have selected the **Failback to the specified location** option at the Destination step. This step applies if you fail back to VMs recovered to new locations, and if networks in those locations differ from networks in the disaster recovery (DR) site.

At the **Network** step of the wizard create a network mapping table. This table maps networks in the DR site to networks in the site where the recovered VMs will reside. Veeam Backup & Replication will use the network mapping table to update configuration files of VMs on the fly, during the failback process.

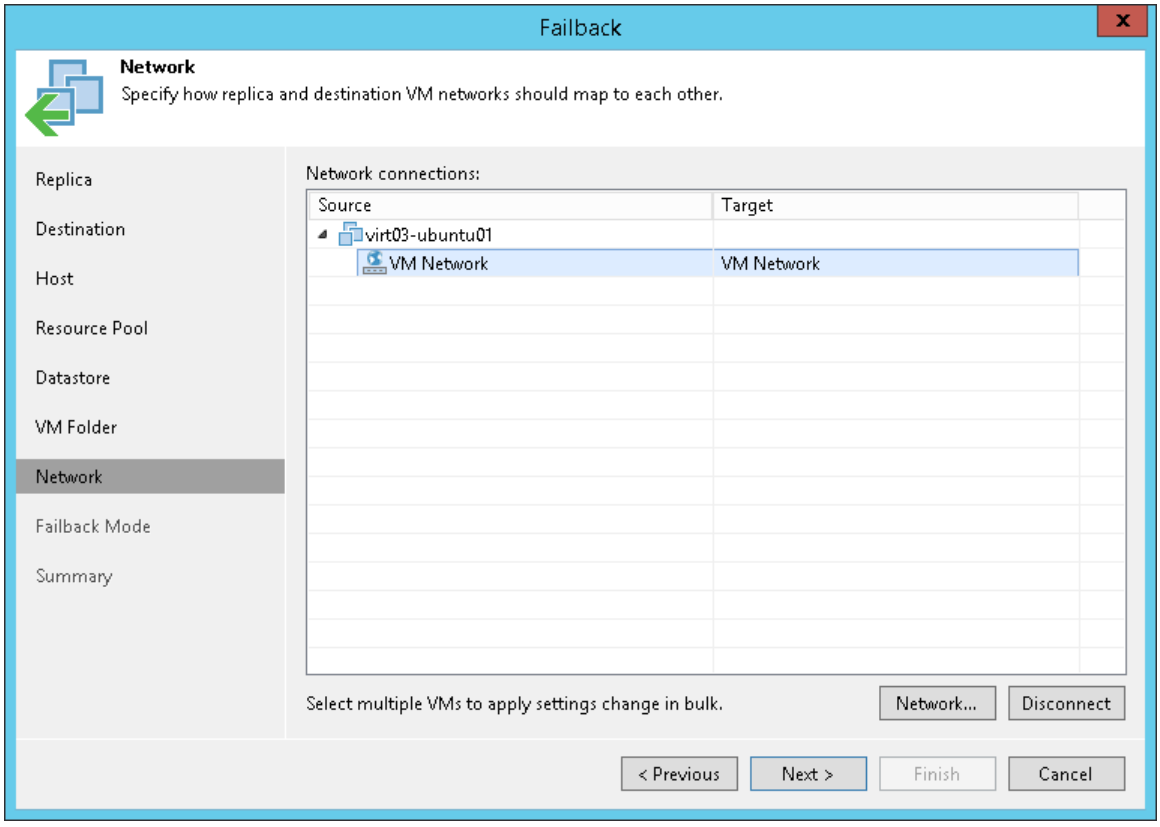
To change networks to which the restored VMs will be connected:

1. In the **Network connections** list, select the necessary VMs and click **Network**.

   If VMs are connected to multiple networks, select networks which you want to map.

2. In the list of available networks, select a network to which the recovered VMs will be connected.

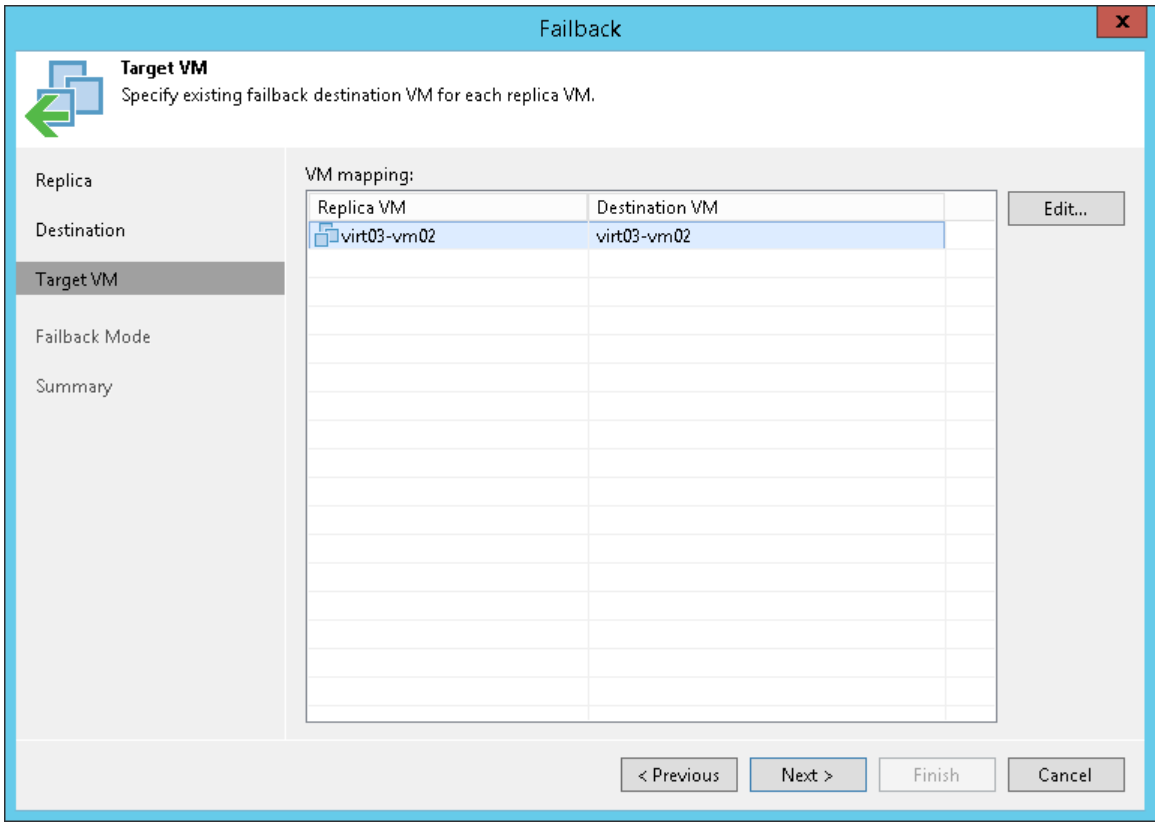
If you do not want to connect the recovered VMs to any virtual network, select the necessary VMs and click **Disconnect**.



Step 9. Map Replicas to Restored VMs

The **Target VM** step is available if you have selected the **Failback to the original VM restored in a different location** option at the Destination step.

At the **Target VM** step of the wizard, specify to which VMs you want to fail back from replicas. These VMs must be already restored from backups in the required location.

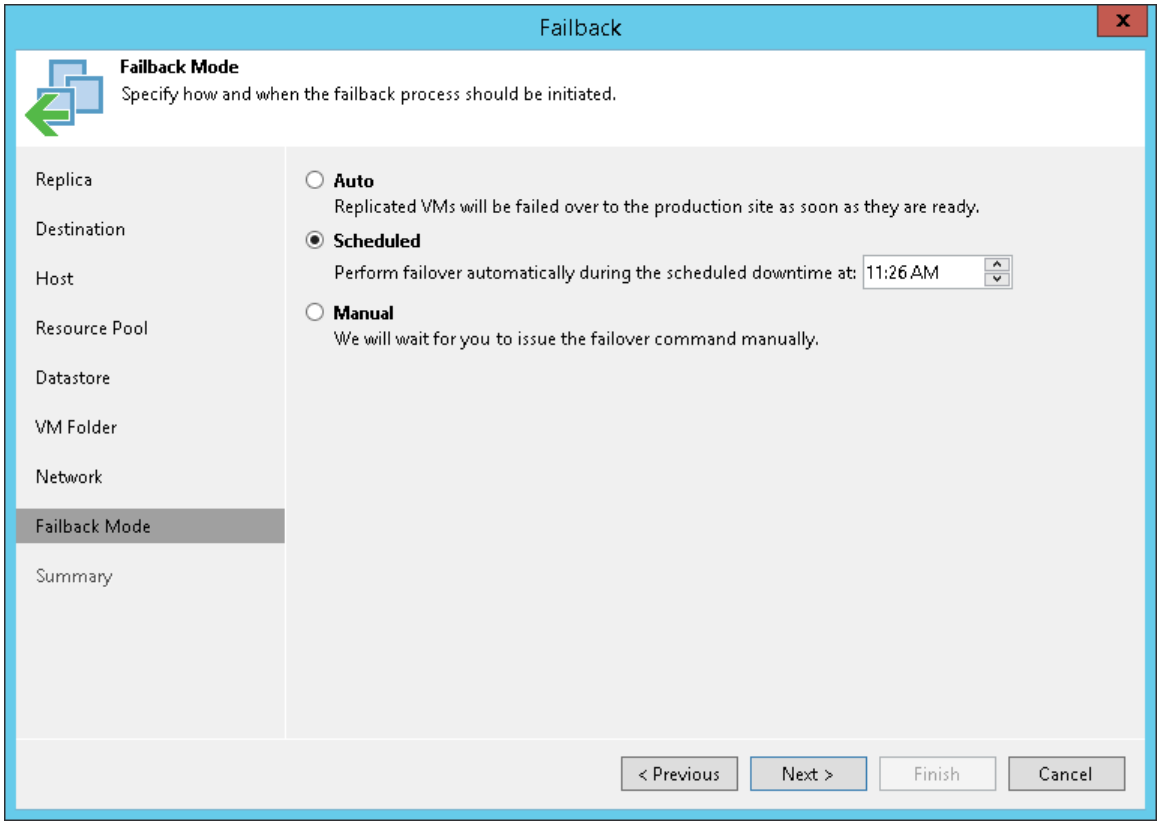


Step 10. Schedule Switch to Production VMs

At the **Failback Mode** step of the wizard, specify when switch from replicas to production VMs must be performed:

- Select **Auto** if you want Veeam Backup & Replication to perform the switch automatically right after the state of the production VMs is synchronized with the state of their replicas.
- Select **Scheduled** if you want Veeam Backup & Replication to perform the switch at a specific time.
- Select **Manual** if you want to perform the switch manually.

If you select the **Scheduled** or **Manual** option, you can further reset/set the scheduled time or switch to the production VM manually. For more information, see Changing Switching Time and Switching Replicas to Production VMs Manually.



Step 11. Finish Working with Wizard

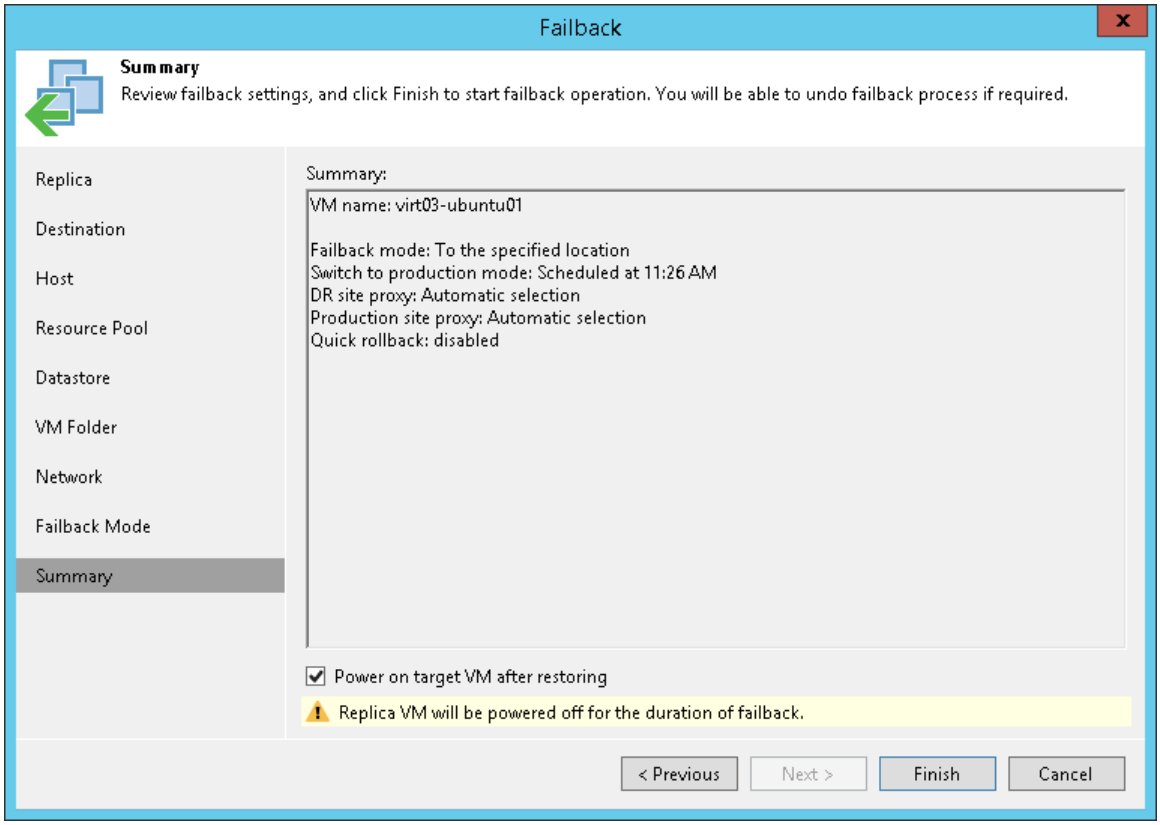
At the **Summary** step of the wizard, review the configured failback settings and click **Finish**.

If you want to power on the production VMs right after the switch to production operation is performed, select the **Power on target VM after restoring** check box.

What You Do Next

Failback is an intermediate step that needs to be finalized. You can finalize failback in the following ways:

- Commit failback
- Undo failback



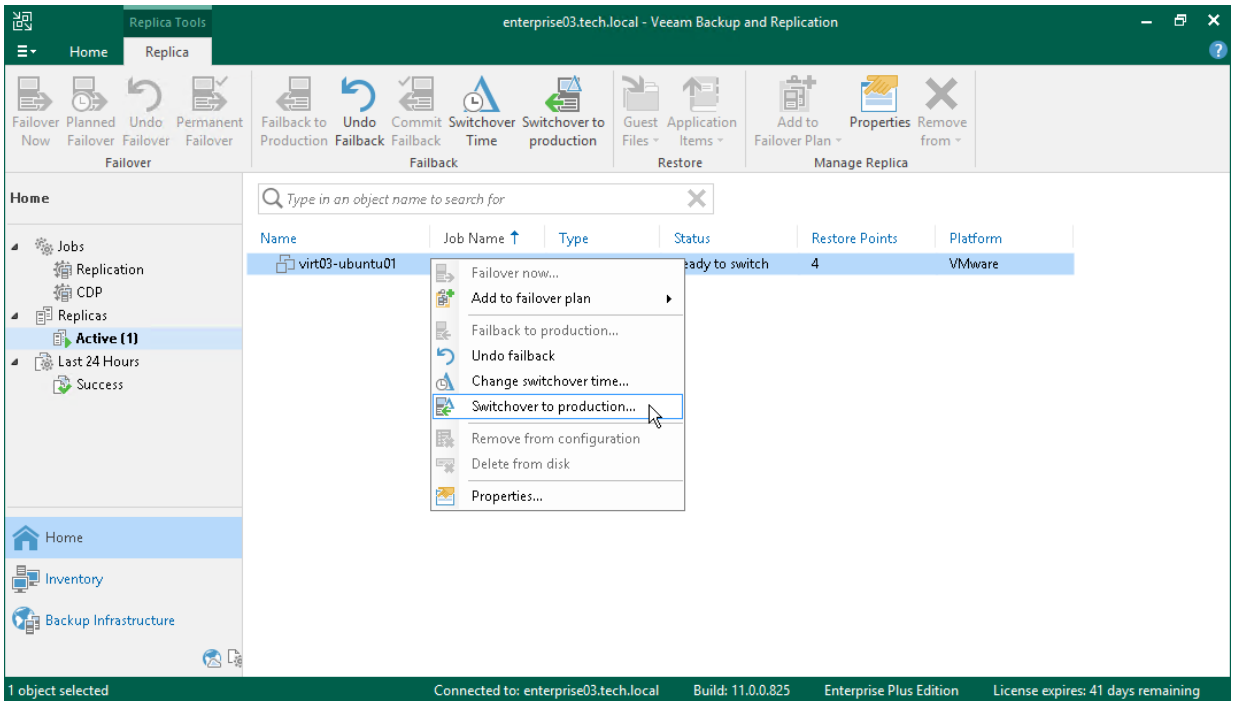
Switching to Production VMs Manually

The following instructions apply if you have selected to switch from replicas to production VMs manually or at the scheduled time at the **Failback Mode** step of the **Failback** wizard.

To switch to a production VM from its replica, do the following:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Replicas** > **Active** node.

3. Right-click a replica in the *Ready to switch* state and select **Switch to production**.



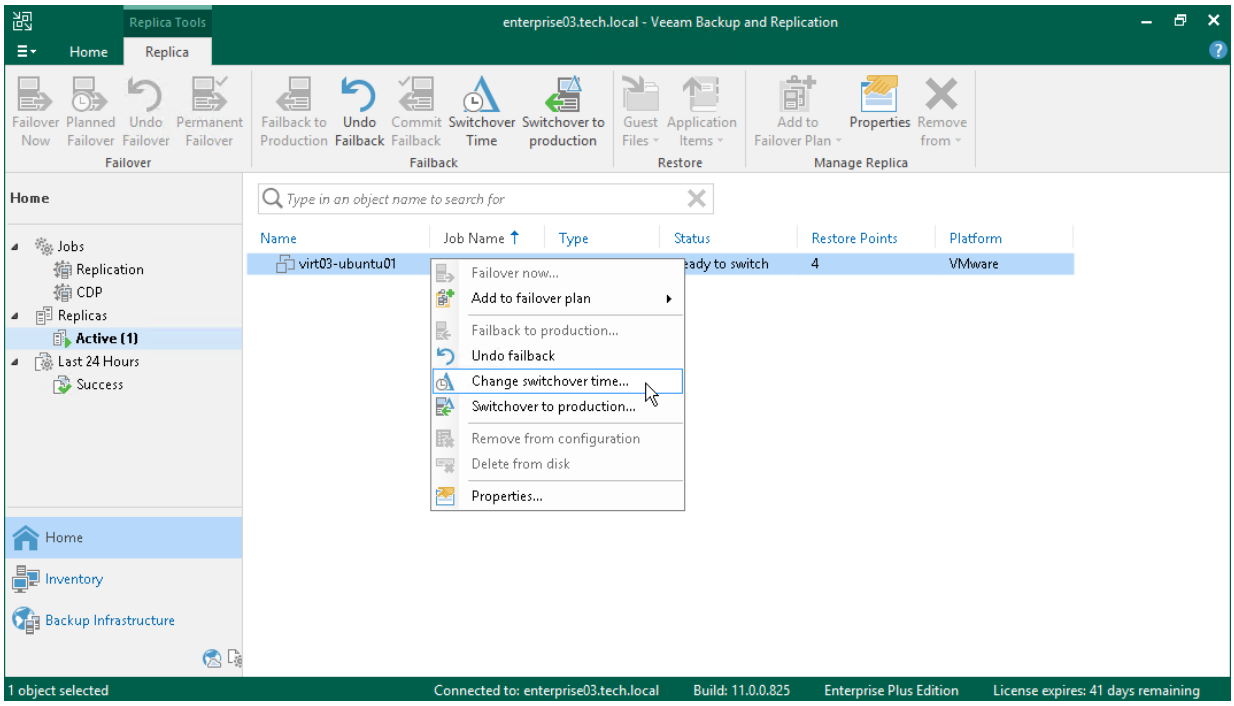
Changing Switching Time

The following instructions apply if you have selected to switch from replicas to production VMs manually or at the scheduled time at the **Failback Mode** step of the **Failback** wizard.

To change the time when Veeam Backup & Replication will switch from replicas to production VMs:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Replicas** > **Active** node.

3. Right-click a replica in the *Ready to switch* state and select **Change switching time.**



Failback Commit

Failback commit is one of the ways to finalize failback. When you commit failback, you confirm that the VM to which you failed back (the production VM) works as expected. After the commit operation, Veeam Backup & Replication resumes replication activities for the production VM.

The commit failback operation is performed in the following way:

1. Veeam Backup & Replication changes the state of the replica from *Failback* to *Ready*.
2. Further operations depend on whether you have failed back to the original VM or recovered VM:
   - If you have failed back to a VM recovered from a backup or replica, Veeam Backup & Replication reconfigures all existing jobs where the original VM is present and adds the original VM to the list of exclusions. The recovered VM takes the role of the original VM and is included into all jobs instead of the excluded VM. When the CDP policy starts, Veeam Backup & Replication processes the recovered VM instead of the former original VM.
   - If you have failed back to the original VM, the CDP policy is not reconfigured. When the CDP policy starts, Veeam Backup & Replication still processes the original VM.

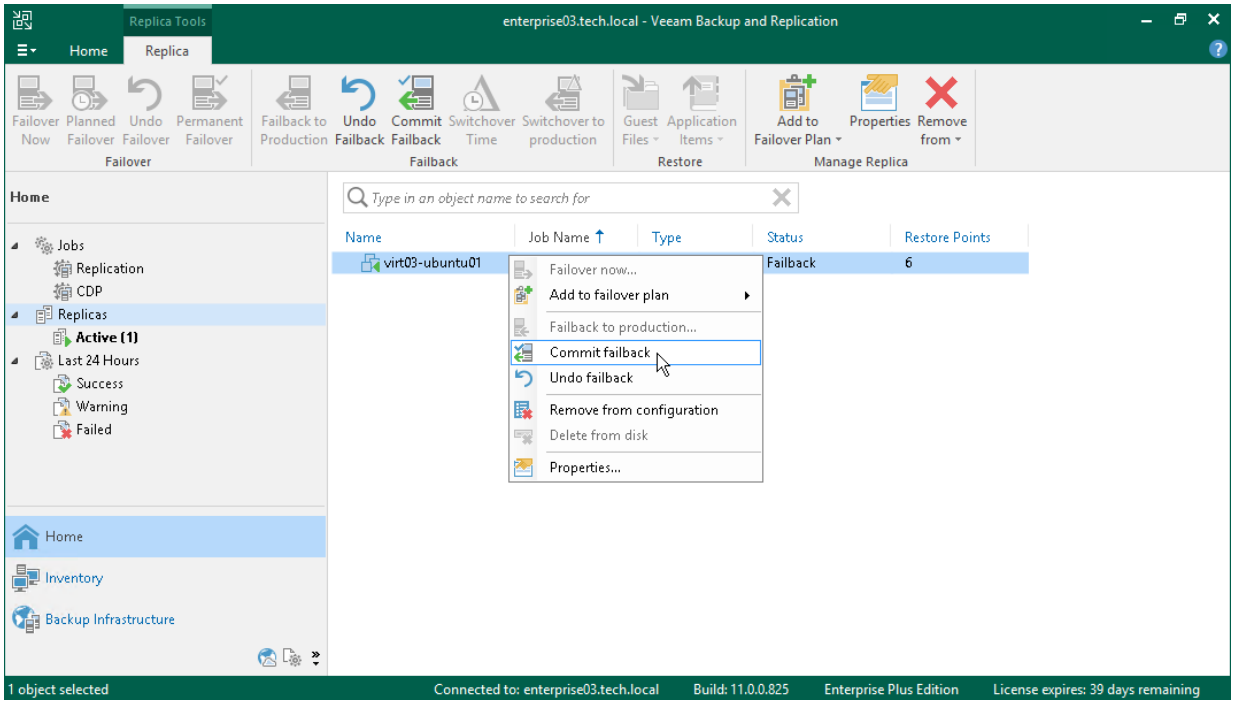
Committing Failback

To commit failback:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Replicas** > **Active** node.

> **IMPORTANT**
>
> If you have failed back to a VM with IDE disks, you must manually power off this VM before committing the failback.

3. In the working area, select the necessary replica and click **Commit Failback** on the ribbon. As an alternative, you can right-click the replica and select **Commit failback**.



Failback Undo

Failback undo is one of the ways to finalize failback. When you undo failback, you confirm that the VM to which you failed back (the production VM) works in a wrong way and you want to get back to the replica.

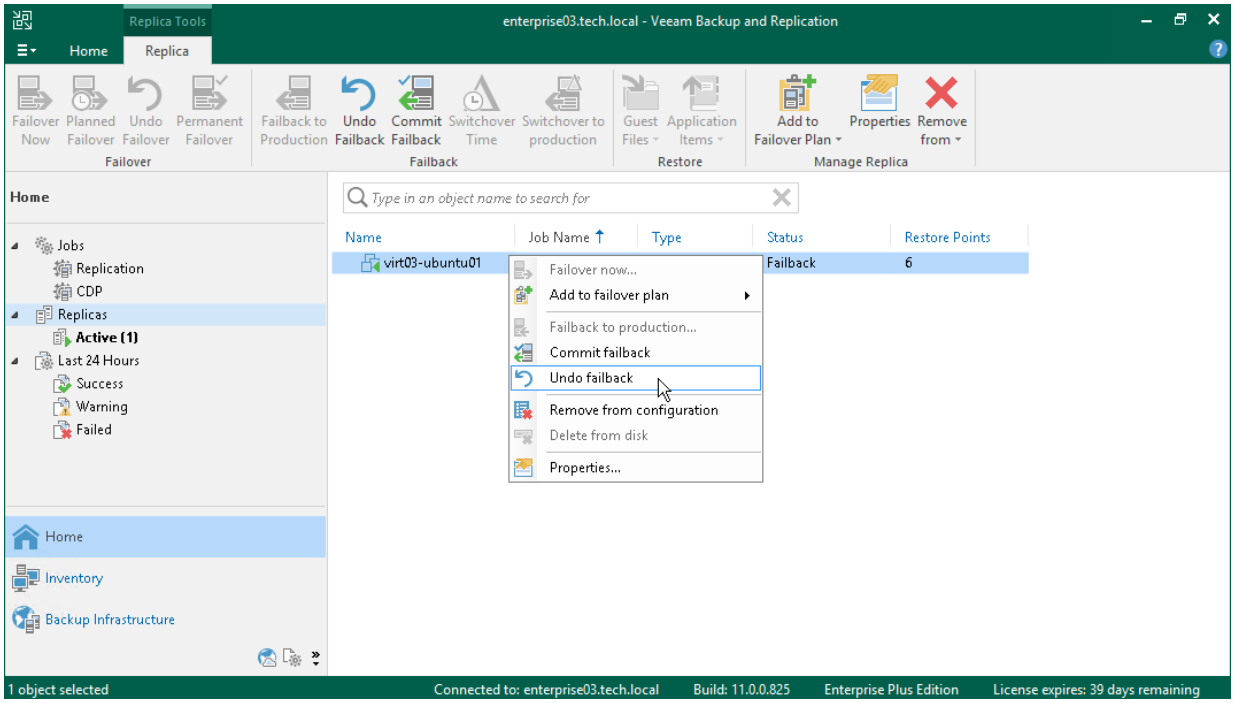
The undo failback operation is performed in the following way:

1. Veeam Backup & Replication powers off the production VM.
2. Veeam Backup & Replication reverts the VM replica to its pre-failback state.
3. Veeam Backup & Replication powers on the VM replica and changes the VM replica state from *Failback* to *Failover*.

Undoing Failback

To undo failback:

1. Open the **Home** view.
2. In the inventory pane, navigate to the **Replicas** > **Active** node.
3. In the working area, select the necessary replica and click **Undo Failback** on the ribbon. As an alternative, you can right-click the replica and select **Undo failback**.



VeeamZIP

With Veeam Backup & Replication, you can quickly perform backup of one or several VMs with VeeamZIP.

VeeamZIP is similar to a full VM backup. The VeeamZIP job always produces a full backup file (VBK) that acts as an independent restore point. You can store the backup file in a backup repository, in a local folder on the backup server or in a network share.

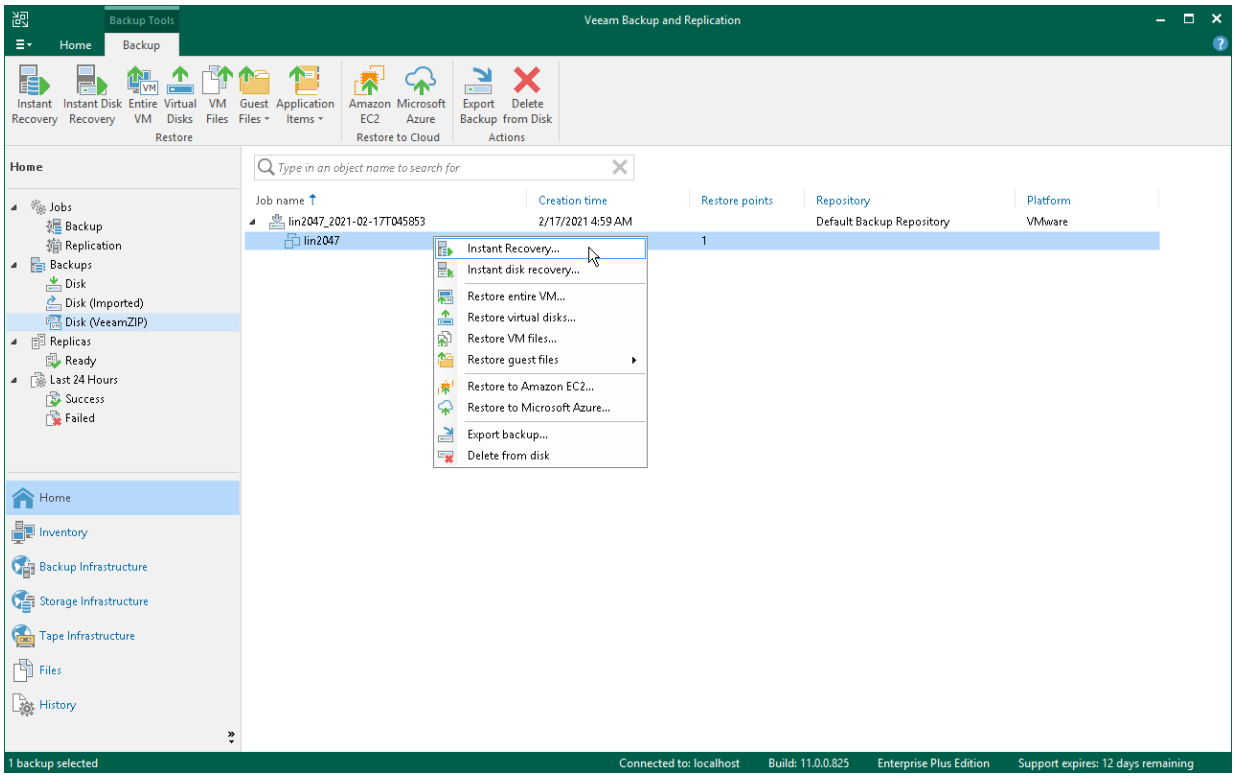
> **IMPORTANT**
>
> Consider the following:
>
> - Veeam Backup & Replication does not enforce backup repository throttling rules during VeeamZIP jobs.
> - You cannot use a Veeam Cloud Connect repository as a target for VeeamZIP jobs.

Backup files produced with VeeamZIP jobs are displayed in the **Home** view, under the **Backups** > **Disk (VeeamZIP)** node.

When you perform backup with VeeamZIP, you do not have to configure a backup job and schedule it. Instead, you can start the backup process for selected VMs immediately. This type of a backup requires minimum settings — you should only select the backup destination, choose the necessary compression level and enable or disable encryption and application-aware processing if necessary. For more information, see Creating VeeamZIP Backups.

To view the progress or results of the VeeamZIP job session, you can use the **History** view. For more information, see Viewing Real-Time Statistics.

To restore VM data from VeeamZIP backups, you can right-click it in the **Home** view and select the necessary restore option. You can also double-click the necessary VeeamZIP backup file on the machine where Veeam Backup & Replication is installed.



Creating VeeamZIP Backups

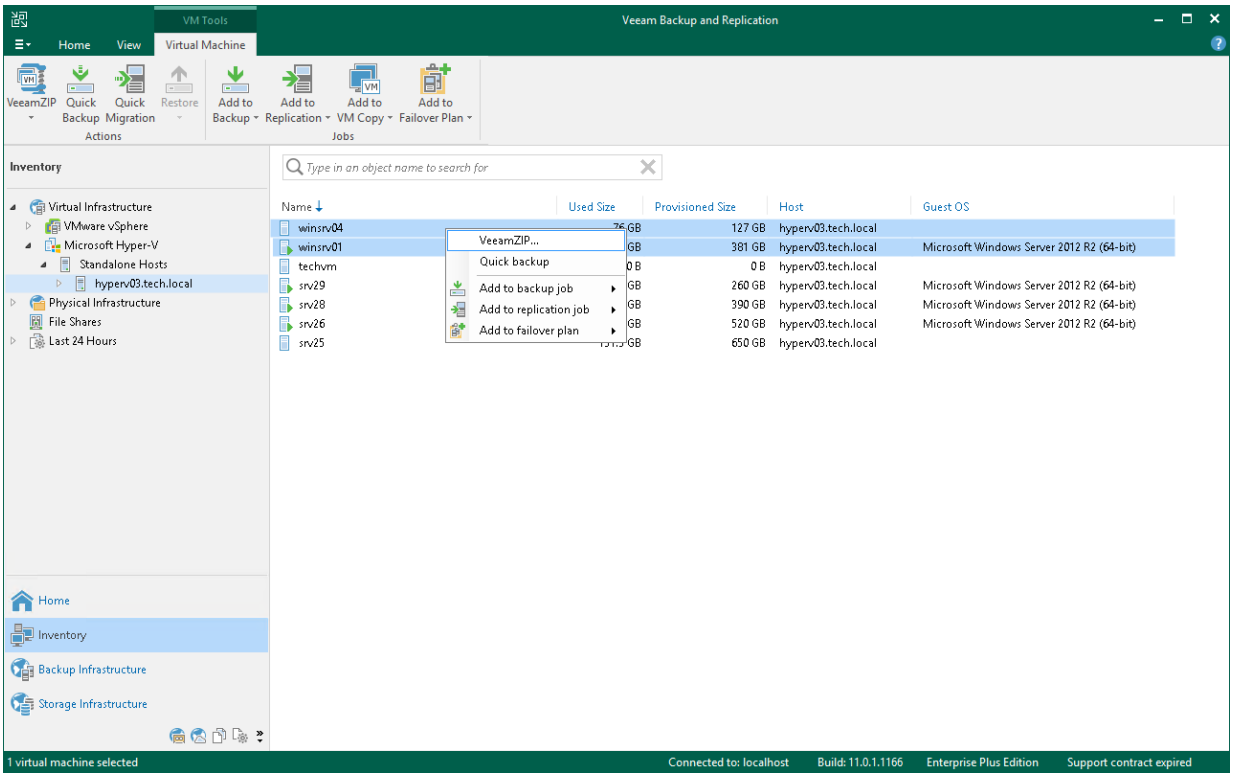
You can quickly back up running and powered off VMs with VeeamZIP. VeeamZIP can be helpful if you want to create an ad-hoc backup for VMs, archive VMs before decommissioning and so on. You can create VeeamZIP backups for one or more VMs.

> **TIP**
>
> Veeam Backup & Replication keeps settings of the latest VeeamZIP task. To quickly create VeeamZIP backups with the same settings and store backups in the same location, right-click the necessary VM and select **VeeamZIP to**.

To create VeeamZIP backups:

1. Open the **Inventory** view. In the infrastructure tree, select a host or VM container (host, cluster, folder, resource pool, VirtualApp, datastore or tag) in which the VMs that you want to back up reside.
2. In the working area, select the VMs and click **VeeamZIP > VeeamZIP** on the ribbon or right-click the VMs and select **VeeamZIP**.

   To quickly find the necessary VMs, type the VM name or a part of it in the search field at the top of the working area and click the **Start search** button on the right or press [Enter] on the keyboard.



3. In the opened **VeeamZIP <N> VM** window:
   a. In the **Destination** section, specify a location in which you want to store VeeamZIP backups.
      - To store VeeamZIP backups in a backup repository, select **Backup repository** and choose the necessary backup repository from the list. In this case, VeeamZIP backups will be saved to the `VeeamZIP` subfolder of the folder where the backup repository stores backups. You can check this folder at the Configure Backup Repository Settings step of the backup repository wizard.

- To store VeeamZIP backups in a local folder on the backup server, select **Local or shared folder**, click **Browse** on the right and select a folder in which VeeamZIP backups must be stored.
- To store VeeamZIP backups in a shared folder, select **Local or shared folder** and type in the UNC name of the shared folder in the field below. Keep in mind that the UNC name always starts with two back slashes (\\).

  If the shared folder requires authentication, select the necessary credentials from the **Credentials** list. If you have not set up credentials beforehand, click the **Manage accounts** link or click **Add** on the right to add necessary credentials. For more information, see Managing Credentials.

b. If you want to specify retention settings for the created VeeamZIP backups, select the **Delete this backup automatically** check box. From the drop-down list, select the retention period. If you do not want to delete VeeamZIP backups, leave the **Delete this backup automatically** check box unselected.

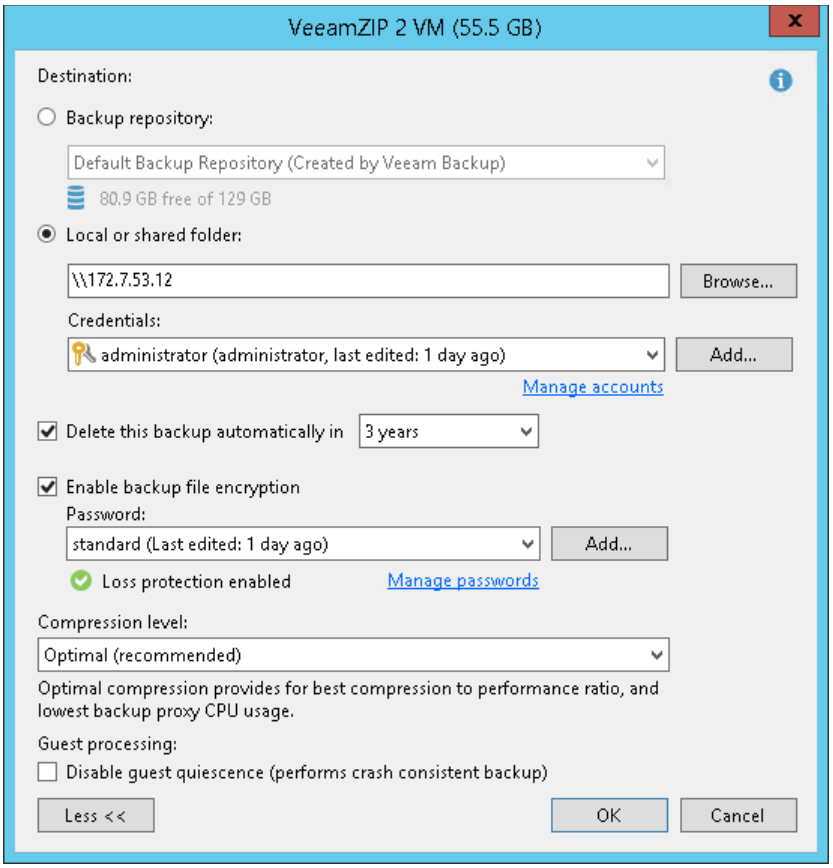
   [For hardened repository] Veeam Backup & Replication sets an immutability period for backup files with retention period as equal to the longest of these periods. For more information, see Limitations and Considerations.

c. To encrypt VeeamZIP backups, select the **Enable backup file encryption** check box. From the **Password** list, select a password that you want to use for encryption. If you have not created a password beforehand, click **Add** or use the **Manage passwords** link to specify a new password. For more information, see Managing Passwords for Data Encryption.

d. From the **Compression level** list, select a compression level for created backups.

e. By default, Veeam Backup & Replication uses VMware Tools quiescence to create a transactionally consistent image of VMs. You can disable VM quiescence. To do this, select the **Disable guest quiescence** check box. In this case, Veeam Backup & Replication will create a crash-consistent VM backup.

f. Click **OK**. The VeeamZIP task will start immediately. Veeam Backup & Replication will create a full backup file (VBK) and store it in the specified location. The VM name, date and time of the file creation are appended to the file name so you can easily find the necessary backups afterwards.



4. As the job runs, you can track the job performance in the real-time mode. To see the job results once it completes, open the **History** view, expand the **Jobs** node and click **Backup**. Then double-click the job session in the list.

Backup Copy

The main backup purpose is to protect your data against disasters and virtual or physical machine failures. However, having just one backup does not provide the necessary level of safety. The primary backup may get destroyed together with production data, and you will have no backups from which you can restore data.

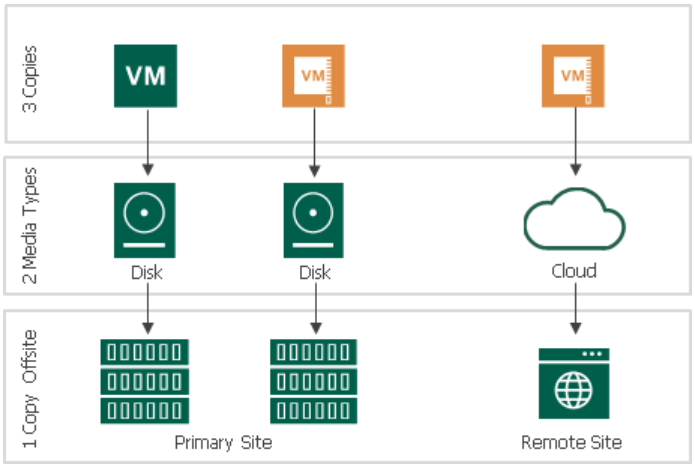
To build a successful data protection and disaster recovery plan, it is recommended that you follow the 3-2-1 rule:

- 3: You must have at least three copies of your data: the original production data and two backups.
- 2: You must use at least two different types of media to store the copies of your data, for example, local disk and cloud.
- 1: You must keep at least one backup off-site, for example, in the cloud or in a remote site.

Thus, you must have at least two backups and they must be in different locations. If a disaster takes out your production data and local backup, you can still recover from your off-site backup.

To help you adopt the 3-2-1 rule, Veeam Backup & Replication offers backup copy capabilities. Backup copy allows you to create several instances of the same backup data in different locations, whether on-site or off-site. Backup copies have the same format as those created by backup jobs and you can recover your data from them when you need it.

Backup copy is a job-driven process. Veeam Backup & Replication fully automates the backup copy process and lets you specify retention settings to maintain the desired number of restore points, as well as full backups for archival purposes.



About Backup Copy

With backup copy, you can create several instances of the same backup file and copy them to secondary (target) backup repositories for long-term storage. Target backup repositories can be located in the same site as the source backup repository or can be deployed off-site. The backup copy file has the same format as the primary backup, so you can restore necessary data directly from it in case of a disaster.

Veeam Backup & Replication supports backup copy for the following types of backups:

- Backups of VMware vSphere or vCloud Director virtual machines created with Veeam Backup & Replication
- Backups of Microsoft Hyper-V virtual machines created with Veeam Backup & Replication
- Backups of virtual and physical machines created with Veeam Agent for Microsoft Windows, Veeam Agent for Linux, Veeam Agent for Mac, Veeam Agent for Oracle Solaris or Veeam Agent for IBM AIX
- Backups of Nutanix AHV virtual machines created with Veeam Backup for Nutanix AHV
- Backups of Oracle and SAP HANA databases created with Veeam Plug-ins for Enterprise Applications
- Backups stored in an HPE StoreOnce backup repository
- File share backups created with Veeam Backup & Replication
- Backups of Amazon EC2 instances created with Veeam Backup for AWS
- Backups of Microsoft Azure virtual machines created with Veeam Backup for Microsoft Azure
- Backups of VM instances created with Veeam Backup for Google Cloud Platform
- Backups created with Veeam Backup for RHV
- [Starting from Veeam Backup & Replication cumulative patch P20211211 for v11a] Backups exported by K10 policies

> **IMPORTANT**
>
> Mind the following for copying backups created with Veeam Backup for AWS, Veeam Backup for Microsoft Azure or Veeam Backup for Google Cloud Platform:
>
> - You can copy such backups only if you have paid types of license. Veeam Backup Community Edition is not supported.
> - You can copy such backups from external repositories but not to them.

When the backup copying process starts, Veeam Backup & Replication accesses backup files in the source backup repository, retrieves data blocks for a specific machine from the backup file, copies them to the target backup repository, and composes copied blocks into a backup file in the target backup repository. The backup copying process does not affect virtual and physical infrastructure resources, does not require creation of additional VM snapshots or VSS snapshots and does not produce load on machines whose backups are copied.

Backup copy is a job-driven process. To copy backups, you need to configure backup copy jobs. The backup copy job defines when, what, how and where to copy. For more information on how to create backup copy jobs, see Creating Backup Copy Jobs for VMs and Physical Machines. Note that to copy file share backups, you need to configure a file share backup job, not the backup copy job. For more information, see Creating File Share Backup Jobs.

How Backup Copy Works

Veeam Backup & Replication performs backup copy in the following way:

1. [For VM backup copy jobs only] Veeam Backup & Replication connects to vCenter Servers and ESXi hosts to gather information about VMs whose restore points you want to copy.
2. For backup copying process, Veeam Backup & Replication starts two Veeam Data Movers — source Veeam Data Mover and target Veeam Data Mover. Veeam Data Movers location depends on the backup repository type and data transport path. For more information, see Backup Copy Architecture.
3. The first backup copy run always produces a full backup file. Veeam Backup & Replication copies data blocks that are necessary to build a full backup of a machine as of the most recent state.

   Veeam Backup & Replication can copy data blocks from one or more backup files in the backup chain in the source backup repository.

   - If the backup chain is created in the reverse incremental backup method, Veeam Backup & Replication copies data blocks of the latest full backup.
   - If the backup chain is created in the forward or forever forward incremental backup method, Veeam Backup & Replication copies data blocks from the first full backup and a set of incremental backups.

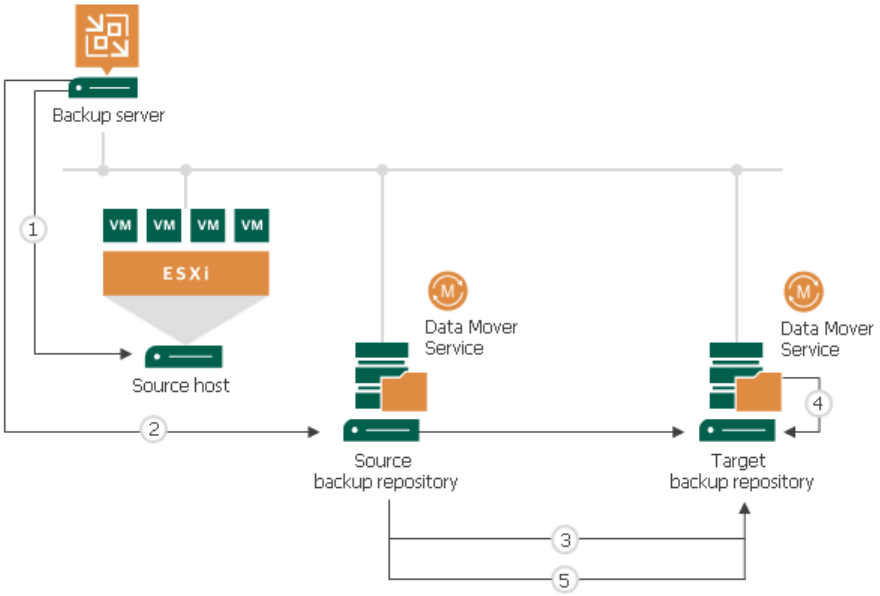
   To minimize the amount of traffic going over the network, Veeam Backup & Replication uses the data compression and deduplication technologies.
4. Veeam Backup & Replication transfers copied data to the target backup repository and writes all copied data blocks to the full backup file.

   - If you do not enable the **Use per-machine backup files** option for the target backup repository, Veeam Backup & Replication creates one backup file in the target backup repository and stores to it data for all machines processed by the job.
   - If you enable the **Use per-machine backup files** option, data of every machine in the job is stored to separate backup files in the target backup repository.

   You can check whether the Use per-machine option is enabled at the Repository step of the repository wizard.
5. During every next backup copy run, when a new restore point appears in the source backup repository, Veeam Backup & Replication copies incremental changes from this most recent restore point and transfers them to the target backup repository. Veeam Backup & Replication writes the copied data blocks to the incremental backup file in the target backup repository, that is, Veeam Backup & Replication creates a new restore point in the forever forward incremental backup chain.

   To retain the desired number of restore points, Veeam Backup & Replication uses a retention policy. For more information, see Short-Term Retention Policy.

> **NOTE**
>
> If you want to store some restore points for longer periods (for weeks, months or years) and enable long-term retention policy (GFS retention policy), Veeam Backup & Replication creates a forward incremental backup chain. For more information on GFS retention policy, see Long-Term Retention Policy (GFS).



In some cases, the source backup job and backup copy job may overlap. This situation can occur if the source backup job needs to transform the source backup chain.

If a specific task in the backup copy job locks the source backup chain to read data from it, and the source backup job that needs to write data to this backup chain starts at this moment (for example, for reverse incremental backup), the task in the backup copy job is put on hold. The backup copy job can continue processing other tasks that use other sources (for example, backup files created by other backup jobs). After the source backup job releases the backup chain, the backup copy job resumes processing machines in this backup chain.

Backup Copy Architecture

To transport data from the source backup repository to the target backup repository, the backup copy job uses one of the following paths:

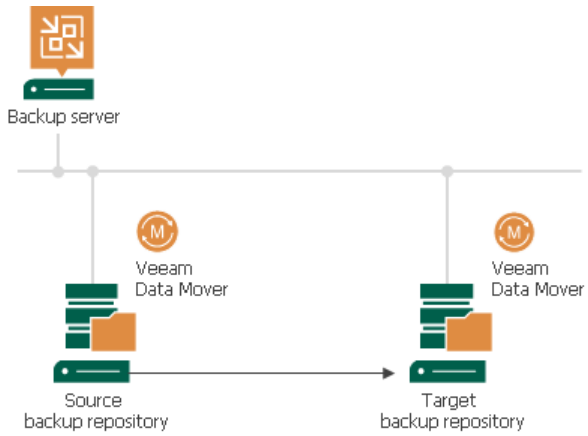
- Direct transport path
- Transport path over WAN accelerators

Direct Transport Path

Veeam Backup & Replication transports data directly from the source backup repository to the target backup repository. This type of data transport is recommended for copying backups to on-site backup repositories or off-site backup repositories over fast connections.

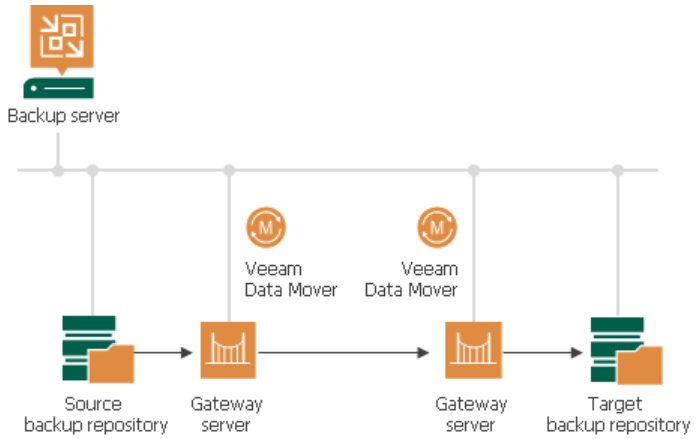
When Veeam Backup & Replication transports data over the direct data path, it uses Veeam Data Movers on the following backup infrastructure components:

- **Microsoft Windows and Linux repositories**. Veeam Backup & Replication uses the source Veeam Data Mover on the source backup repository and target Veeam Data Mover on the target backup repository.



- **Shared folder backup repository**. If you have instructed Veeam Backup & Replication to automatically select the gateway server, Veeam Backup & Replication will use Veeam Data Movers deployed on mount servers associated with backup repositories. In case mount servers cannot be used for some reason, Veeam Backup & Replication will fail over to the backup server.

  If you have explicitly defined the gateway server, Veeam Backup & Replication will use the source Veeam Data Mover on the gateway server in the source site and target Veeam Data Mover on the gateway server on the target site.



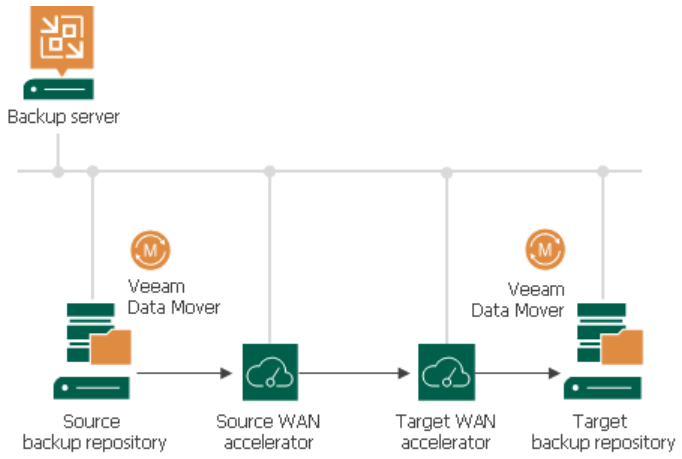
Transport Path over WAN Accelerators

Veeam Backup & Replication transports data through a pair of WAN accelerators: one deployed on the source side and the other one deployed on the target side. WAN accelerators remove redundant blocks before transferring data and thus significantly reduce the amount of traffic going over the network. This type of data transport is recommended for copying backups off-site over slow connections or WAN.

> **IMPORTANT**
>
> The WAN acceleration technology is included in the Veeam Universal License. When using a legacy socket-based license, the Enterprise Plus edition is required. For more information, see WAN Acceleration.

When Veeam Backup & Replication transports data using WAN accelerators, it uses Veeam Data Movers on the following backup infrastructure components:

- **Microsoft Windows and Linux repositories**. Veeam Backup & Replication uses the source Veeam Data Mover on the source backup repository and target Veeam Data Mover on the target backup repository.
- **Shared folder backup repository**. If you have instructed Veeam Backup & Replication to automatically select the gateway server, Veeam Backup & Replication will use Veeam Data Mover deployed on the source and/or target WAN accelerator. If you have explicitly defined the gateway server, Veeam Backup & Replication will use the source Veeam Data Mover on the gateway server in the source site and target Veeam Data Mover on the gateway server on the target site.



Backup Copy Modes

Veeam Backup & Replication offers two backup copy modes:

- **Immediate copy**

  In the immediate copy mode, Veeam Backup & Replication copies restore points as soon as they appear in a source backup repository. Veeam Backup & Replication copies only restore points created by source backup jobs (backup jobs that you select when configuring a backup copy job).
  Veeam Backup & Replication can also copy transaction log backups if you enable this capability in job settings.

  The immediate copy mode is supported for the following backup types:

  - Backups of VMware vSphere and Microsoft Hyper-V VMs created with Veeam Backup & Replication
  - Backups created with Veeam Agent for Microsoft Windows, Veeam Agent for Linux or Veeam Agent for Mac operating in the managed mode

    Only backups created by backup jobs managed by the backup server are supported.
  - Backups created with Veeam Plug-ins for Enterprise Applications (Oracle RMAN, SAP HANA, SAP on Oracle)

- **Periodic copy**

  In the periodic copy mode, Veeam Backup & Replication copies restore points once in an interval specified in backup copy job settings. This interval is also known as backup copy interval. For more information, see Backup Copy Intervals.

  Veeam Backup & Replication can copy restore points created by backup jobs or restore points of individual machines. If you select backup jobs when configuring copy job settings, Veeam Backup & Replication copies only restore points that selected backup jobs create. If you select machines,
  Veeam Backup & Replication copies all restore points of the selected machines — even restore points created by different backup jobs.

  The periodic copy mode is supported for the following backup types:

  - Backups of VMware vSphere and Microsoft Hyper-V VMs created with Veeam Backup & Replication
  - Backups created with Veeam Agent for Microsoft Windows, Veeam Agent for Linux or Veeam Agent for Mac operating in the standalone or managed mode

    For the managed mode, backups created by both backup jobs managed by the backup server and jobs managed by Veeam Agent are supported.
  - Backups created with Veeam Agent for Oracle Solaris
  - Backups created with Veeam Backup for Nutanix AHV
  - Backups created with Veeam Backup for AWS
  - Backups created with Veeam Backup for Microsoft Azure
  - [Starting from Veeam Backup & Replication 11a (build 11.0.1.1261)] Backups created with Google Cloud Platform
  - [Starting from Veeam Backup & Replication cumulative patch P20211211 for v11a] Backups exported with K10 policies

> **NOTE**
>
> Veeam Backup & Replication does not allow you to change the selected backup copy mode because the operating logic and created data are very different in these modes. If you want to change the mode, you need to create a new backup copy job in the required mode and create a seed for it. Note that the seed must also be created in the same mode as the new backup copy job. For more information, see Creating Seed for Backup Copy Job.

Backup Copy Intervals

A backup copy interval is a time span in which a backup copy job must copy a restore point from the source backup repository to the target backup repository.

> **NOTE**
>
> Veeam Backup & Replication uses backup copy intervals only in the periodic copy mode.

The backup copy interval affects the restore point selection process. For more information, see Restore Point Selection.

At the beginning of a new interval, Veeam Backup & Replication checks if a new restore point is available in the source backup repository:

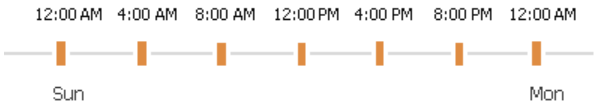
- If a new restore point is found, the backup copy job starts the synchronization process and copies the latest restore point to the target backup repository.
- If a new restore point is not found or is locked by the source backup job, the backup copy job enters the *Idle* state.

By default, the backup copy interval is set to 1 day. You can change this interval when configuring a backup copy job and set the interval in minutes or hours. Note that if you specify a too short backup copy interval or change the interval, some issues can occur. For details, see Issues with Backup Copy Intervals.

Minutely and Hourly Backup Copy Intervals

The first minutely and hourly intervals start when the backup copy job runs for the first time. Each subsequent backup copy interval starts after the period that you specified in the backup copy job settings.

For example, if you set the backup copy interval to 4 hours and start the backup copy job at 12:00 AM, Veeam Backup & Replication will start new backup copy intervals at 12:00 AM, 4:00 AM, 8:00 AM and so on.



Daily Backup Copy Intervals

The daily backup intervals start every day at the time specified in the backup copy job settings.

For example, if you set the backup copy interval to 1 day and instruct Veeam Backup & Replication to start new intervals at 12:00 PM, Veeam Backup & Replication will start new backup copy intervals at 12:00 PM daily.



In some cases, the start time of the backup copy job and backup copy interval may differ.

For example, you configure a backup copy job and set the start time of the backup copy interval to 12:00 PM. You also specify that the job must start right after the wizard is closed, and you close the wizard at 12:00 AM. In this case, the first backup copy interval will start immediately and will run for a shorter period of time — for 12 hours instead of one day. All subsequent backup copy intervals will start as defined by backup copy job schedule.

Issues with Backup Copy Intervals

Being a scheduled activity, the backup copy job may fail to run as expected. Veeam Backup & Replication automatically handles some issues that can occur with the backup copy job.

Short Backup Copy Intervals

In some cases, Veeam Backup & Replication may fail to transport the restore point within the backup copy interval of the backup copy job. This can happen, for example, if the backup copy interval is too short and is not sufficient for the amount of data to be copied.

Veeam Backup & Replication handles this situation differently for the first and subsequent backup copy intervals.

- The first backup copy interval always produces a full backup file — the starting point in the backup chain. If Veeam Backup & Replication fails to copy data for the full backup file during the first backup copy interval, it marks the job session as finished with the *Warning* status. During the next backup copy interval, Veeam Backup & Replication attempts to copy data for the full backup file in the following manner:

  a. When a new backup copy interval begins, the restore point that was previously copied no longer corresponds to the restore point selection rules. That is, the time of the restore point creation falls out of the search scope. For this reason, Veeam Backup & Replication waits for a new restore point to appear in the source backup repository.

  b. When a new restore point appears in the source backup repository, Veeam Backup & Replication detects what data blocks still need to be copied to make up a full backup file in the target backup repository, and copies these data blocks.

     This process continues until there is a full backup file in the target backup repository.

- At subsequent backup copy intervals, Veeam Backup & Replication copies incremental restore points. If Veeam Backup & Replication fails to transport an incremental restore point, it marks the synchronization task as failed. Veeam Backup & Replication waits for the expiration of the backup copy interval; after that, Veeam Backup & Replication marks the job session as finished with the *Error* status.

Veeam Backup & Replication does not mark the backup copy job session with the *Error* status in the following cases:

- The source backup job has not started during the backup copy interval of the backup copy job (that is, the backup copy job has nothing to copy to the target backup repository).
- A task in the backup copy job processes a VM template, and the source backup job is set to exclude the VM template during incremental backup jobs sessions.

Change of the Backup Copy Interval Start Time

If you have selected to run a backup copy job with a daily backup copy interval, you must define the start time of the backup copy interval. However, you may want to change the start time afterwards. After the start time change, Veeam Backup & Replication behaves in the following manner:

1. Veeam Backup & Replication finishes the current backup copy interval running according to the 'old' start time value as usual.
2. After the current backup copy interval is over, Veeam Backup & Replication immediately starts the backup copy interval, not waiting for the 'new' start time point to come. At that, Veeam Backup & Replication "stretches" the started interval: the interval lasts for the time remaining till the new start time plus the time of the backup copy interval itself.

3. All subsequent backup copy intervals are created and started in a regular manner by the new schedule.

For example, when you first created a backup copy job, you set a daily backup copy interval with the start time at 8 AM. After that, you changed the start time to 10 AM. In this case, Veeam Backup & Replication will first finish the backup copy interval that is currently running — that is, the backup copy interval that was started at 8 AM — as usual. After that, it will immediately start a new backup copy interval. This interval will run for 26 hours — from 8 AM of the current day until 10 AM of the next day. All subsequent backup copy intervals will be started at 10 AM every day.

The first backup copy interval that is run after the start time change is typically longer than a regular one. This happens because of the backup copy interval "stretch" mentioned above. To start the synchronization process right away, you can use the **Sync Now** option after you change the start time value. In this case, Veeam Backup & Replication will behave in the following manner:

1. When you start the synchronization process manually, Veeam Backup & Replication forcibly finishes the current backup copy interval and begins a new backup copy interval according to the new start time value. This backup copy interval lasts until a new backup copy interval by the new schedule must be started.
2. All subsequent backup copy intervals are created and started in a regular manner.

As a result, the first backup copy interval after the start time change will begin immediately.

For example, when you first created a backup copy job, you set a daily backup copy interval with the start time at 8 AM. After that, you changed the start time to 10 AM. On the start time change, you started the manual synchronization process at 1 PM. In this case, Veeam Backup & Replication will finish the current backup copy interval — that is, the backup copy interval that was started at 8 AM — immediately at 1 PM. After that, it will start a new backup copy interval. This interval will run for 21 hours — from 1 PM of the current day until 10 AM of the next day. All subsequent backup copy intervals will be started at 10 AM every day.

Restore Point Selection

Veeam Backup & Replication always copies the most recent restore points, even if a backup copy job runs for the first time and source backup repositories already contain chains of restore points.

In the immediate copy mode, the most recent restore point is the latest complete restore point created by a source backup job.

In the periodic copy mode, Veeam Backup & Replication uses the backup copy interval to identify a restore point as the most recent one — the time of restore point creation must be greater or equal to the difference between the current time and the backup copy interval.

```
Time of restore point creation >= current time – backup copy interval
```

Consider the following example. You create a backup copy job and set the backup copy interval to 24 hours. The source backup job will create a restore point on 6/30/2020 at 1:00 PM. The backup copy job will start on 7/1/2020, 12:00 PM. According to the rule, Veeam Backup & Replication will copy the created restore point:

```
6/30/2020, 1:00 PM >= 7/1/2020, 12:00 PM – 24 hours
```

If there are no restore points considered as recent, Veeam Backup & Replication does not copy data from source backup repositories. Instead, it waits for new restore points to appear. Only after that, Veeam Backup & Replication copies the most recent data blocks to the target repository.

In the periodic copy mode, you can also specify the search scope for restore points. For more information, see Select Machines to Process.

Limitations for Restore Points Selection

The following limitations apply when Veeam Backup & Replication selects restore points that must be copied to the target repository:

- Veeam Backup & Replication does not copy restore points from the target backup repository.
- Veeam Backup & Replication does not copy restore points from imported backups.
- Veeam Backup & Replication does not copy restore points that have already been copied by the same backup copy job to the target backup repository.
- Veeam Backup & Replication does not copy incomplete restore points.
- Veeam Backup & Replication does not copy restore points that are locked by the backup transformation process (merge, transform).
- A backup copy job does not copy a restore point if its data block size differs from the data block size of restore points that the job has already copied to the target backup repository. To copy restore points with the changed block size, you need to create active full backups. For details, see Change Storage Optimization Settings for Backup Copy Job.

  For example, if you have changed the block size for restore points in the source backup job (the **Storage optimization** option in the Storage Settings), Veeam Backup & Replication will not copy newly created restore points and will display the *Restore point is located in backup file with different block size* message.

Transformation Processes

Veeam Backup & Replication can perform additional transformations in the target backup repository after the backup copying task or at the end of the backup copy interval. Transformation processes are the following:

- **Backup chain transformation**

  When a new restore point is copied to the target backup repository, Veeam Backup & Replication checks the retention policy settings for the backup copy job. If the limit in restore points is exceeded, Veeam Backup & Replication does the following:

  - If only short-term retention policy is enabled, Veeam Backup & Replication transforms the backup chain to make room for a new restore point. For more information, see Short-Term Retention Policy.
  - If long-term retention policy is enabled, Veeam Backup & Replication removes unnecessary restore points. Veeam Backup & Replication removes restore points in a way similar to the one described in Forward Incremental Backup Retention Policy.

  For more information on retention policies, see Retention Policy for Backup Copy Jobs.

  After the transformation process, Veeam Backup & Replication can perform additional operations: remove data of deleted workloads from the backup chain and compact a full backup file.

- **Removal of deleted items**

  In the backup copy job settings, you can specify after which period you want to delete data of deleted workloads from backups created by backup copy jobs. After the period ends, Veeam Backup & Replication checks the list of workloads included in the job and removes data of deleted workloads from the backup chain in the target backup repository. For more information on how data is deleted and which limitations apply, see Deleted Items Retention. For more information on how to configure the deleted items retention, see Specifying Advanced Settings.

- **Full backup file compact**

  In the backup copy job settings, you can select to periodically compact a full backup file to reduce its size and increase the speed of read and write operations. For more information, see Compact of Full Backup File.

Backup Copy Window

A backup copy window is a period of time when a backup copy job is allowed to transport data between source and target repositories.

By default, the backup copy window is configured to allow data transfer at any time. If you do not want the backup copy job to overlap the production hours, you can reduce the backup copy window and specify "prohibited" hours.

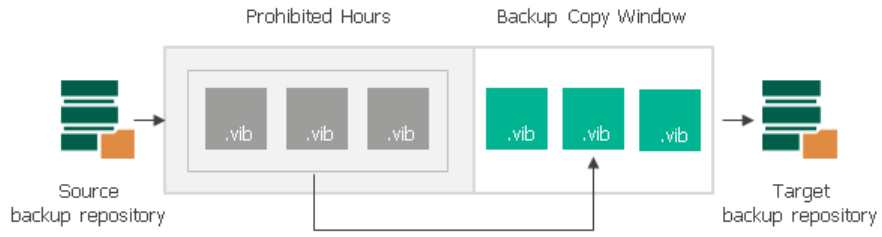
During the prohibited hours Veeam Backup & Replication cannot transfer data between source and target repositories. Other aspects of how Veeam Backup & Replication behaves during the prohibited hours and backup copy window depend on the selected backup copy mode. For more information, see Backup Copy Window and Prohibited Hours in Immediate Copy Mode and Backup Copy Window and Prohibited Hours in Periodic Copy Mode.

> **NOTE**
>
> During the prohibited hours, Veeam Backup & Replication stops only data transferring operations. Transformation processes in the target repository are still performed. For more information, see Transformation Processes.

Backup Copy Window and Prohibited Hours in Immediate Copy Mode

When the prohibited hours start, Veeam Backup & Replication stops backup copy job sessions during which data transfer was not finished and assigns *Failed* to the session statuses. During the prohibited hours, Veeam Backup & Replication does not start new backup copy job sessions.

When the backup copy window starts, Veeam Backup & Replication transfers all restore points that were not transferred and creates as many incremental backup files as were omitted.



Backup Copy Window and Prohibited Hours in Periodic Copy Mode

How Veeam Backup & Replication behaves during the prohibited hours and backup copy window depend on the length of the backup copy interval:

- If the backup copy interval is longer than the duration of the prohibited hours, Veeam Backup & Replication puts backup copy job sessions to the *Idle* state and postpones data transfer operations. When the backup copy window starts, Veeam Backup & Replication renews data transfer.

- If the backup copy interval is shorter than the duration of the prohibited hours, Veeam Backup & Replication finishes all backup copy job sessions and assigns *Failed* to the session statuses. During the first backup copy interval on the backup copy window, Veeam Backup & Replication aggregates all data changes between the latest restore point on the target backup repository and latest restore point in the source backup repository. Then, Veeam Backup & Replication transfers these changes to the target repository and writes them into one restore point.



Retention Policy for Backup Copy Jobs

The retention policy of a backup copy job defines for how long Veeam Backup & Replication must retain copied restore points in the target backup repository. The retention policy of a backup copy job does not depend on retention policy of the source backup job. The backup copy job has its own retention policy settings.

Veeam Backup & Replication offers two retention policy schemes for backup copy jobs:

- Short-Term Retention Policy
- GFS Retention Policy (Weekly, Monthly, Yearly)

Also, there is a separate retention policy for machines that has been removed from the infrastructure. For details, see Deleted Items Retention.

Short-Term Retention Policy

The short-term retention policy allows to retain restore points created by backup copy jobs for a specified number of days or until the number of restore points reaches the specified number in the retention settings.

During the first backup copy session, Veeam Backup & Replication creates the first restore point — a full backup. The next backup copy sessions add incremental backups to the backup chain. As a result, the regular backup cycle produces a chain of a full backup and set of incremental backups in the target backup repository. When the retention policy is exceeded, Veeam Backup & Replication removes the earliest restore points from backup chains in the target backup repositories.

Since Veeam Backup & Replication creates forever forward incremental backup chains while backup copy jobs run, Veeam Backup & Replication applies the forever forward incremental retention policy to remove restore points and maintain the desired number of restore points.

When configuring short-term retention policy settings for a backup copy job, you have two options:

- Specify the number of restore points.

  Veeam Backup & Replication keeps the last *N* restore points, where *N* is the number of restore points that you specify in the settings. The minimum number that you can specify is 2.

- Specify the number of days.

  Veeam Backup & Replication keeps restore points created during the last *N* days, where *N* is the number of days that you specify in the settings.

  Mind the following for the daily retention policy:

  - The minimum number of retained restore points is 3. This number does not depend on the number of days set in the retention policy. For example, the retention policy is set to 5 days. You launch the job after it was stopped for 10 days. Normally, Veeam Backup & Replication would delete all previous restore points. However, due to the minimum number of retained restore points, you will still have at least 3 restore points: the newly created restore point and the two previous ones.

    You can change the minimum number of retained restore points with a registry key. For more information, contact Veeam Customer Support.

  - If the backup job starts at the end of the day and finishes the next day, Veeam Backup & Replication assumes that the restore point is created at the moment when the backup job started. However, Veeam Backup & Replication starts counting retention policy days only after the backup job finishes processing workloads.

- When determining whether the number of allowed days is exceeded, Veeam Backup & Replication ignores the day when the daily retention policy runs. In fact, Veeam Backup & Replication keeps restore points for the *N + 1* days, where *N* is the number of days that you specify in the settings.
- When determining whether the number of allowed days is exceeded, Veeam Backup & Replication also counts days when the backup job did not create any backups.

NOTE

If you want to create full backups periodically (weekly, monthly, yearly), enable the GFS retention policy. If you do not enable the GFS retention, the regular backup copy cycles will create only incremental backups.

IMPORTANT

This section describes how the short-term retention policy functions when the GFS retention policy is disabled. If you enable the GFS retention policy, mind the following:

- The backup copy chain will contain more restore points than you have specified in the short-term retention policy.
- With enabled GFS retention policy, Veeam Backup & Replication applies the forward-incremental retention policy to the backup copy chain.

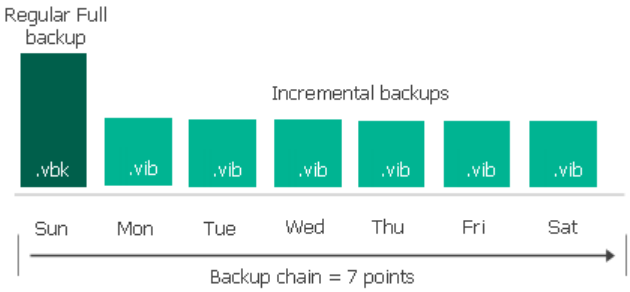
For details, see GFS Retention Policy.

Example

The regular backup cycle is based on the short-term retention policy scheme. When you specify retention policy settings, you define how many restore points you want to retain in the backup chain in the target backup repository.

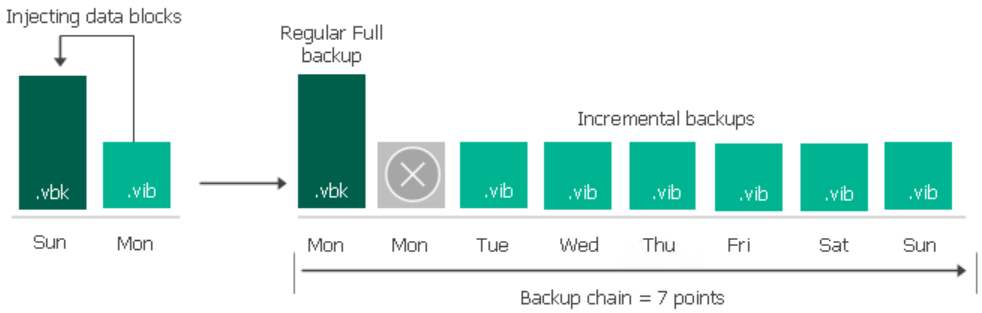
For example, you have selected to retain 7 restore points. The backup copy runs once a day and starts on Sunday.

1. Veeam Backup & Replication creates a full backup on Sunday and add 6 incremental backups Monday through Saturday.



2. On Sunday, Veeam Backup & Replication creates another increment. As a result, there will be 8 restore points, which exceeds the retention policy. Thus, the oldest increment is merged to the full backup.

   After the oldest increment is merged to the full backup, Veeam Backup & Replication removes the increment as it is no longer needed.



Related Topics

Long-Term Retention Policy (GFS)

Long-Term Retention Policy (GFS)

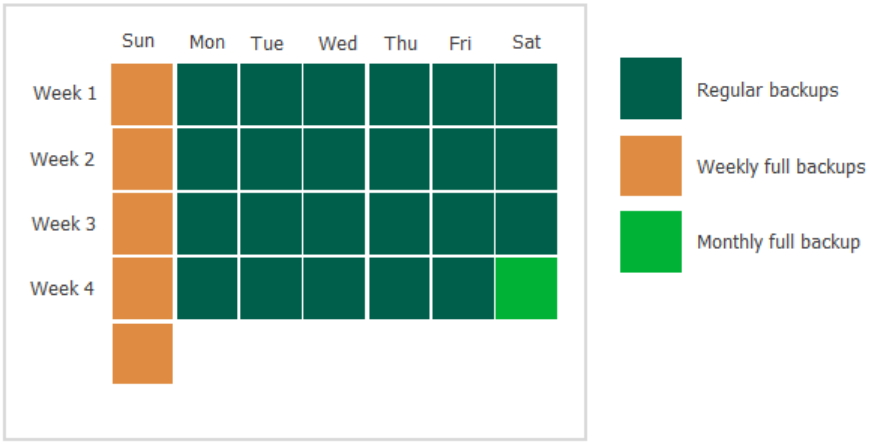
The long-term or Grandfather-Father-Son (GFS) retention policy allows you to store VM backups for long periods of time — for weeks, months and years. For this purpose, Veeam Backup & Replication creates synthetic or active full backup files and marks them with GFS flags. These GFS flags can be of three types: weekly, monthly or yearly. Depending on which flag is assigned to the full backup, it will be stored for specified number of weeks, months or years.

The GFS retention also helps you to mitigate risks that the short-term retention policy has, such as large number of subsequent incremental backups. Large number of subsequent incremental backups can increase recovery time, because Veeam Backup & Replication has to read data through the whole backup chain. Also, one corrupted increment can make the whole chain useless. When you configure the GFS retention, Veeam Backup & Replication creates weekly/monthly/yearly full backups, so instead of one backup chain consisting of one full backup and incremental backups, you will have several backup chains.

GFS backups are always full backup files that contain data of the whole machine image as of a specific date. GFS is a tiered retention policy and it uses a number of cycles to retain backups for different periods of time:

- Weekly backup cycle
- Monthly backup cycle
- Yearly backup cycle

In the GFS retention policy, weekly backups are known as 'sons', monthly backups are known as 'fathers' and yearly backups are known as 'grandfathers'. Weekly, monthly and yearly backups are also called archive backups.



NOTE

GFS retention policy functions in combination with short-term retention policy. After you enable the GFS retention, the backup chain switches from the forever-forward incremental policy to forward incremental policy. Thus, the increments are no longer merged to the full backup file.

If you enable only yearly full backups without monthly and weekly backups, this can result in a large number of increments in a backup chain. To avoid this, it is recommended to enable an additional weekly GFS cycle. Weekly GFS cycle will update the backup chain every week which will allow to remove excessive increment files.

Related Topics

- How GFS Retention Works
- Limitations and Considerations for GFS Cycles

How GFS Retention Works

To understand how GFS retention works, see the following sections:

- Backup Copy GFS Methods
- Backup Copy GFS Cycles
- Backup Chain for GFS Backups
- Limitations and Considerations for GFS Retention Policy

Backup Copy GFS Methods

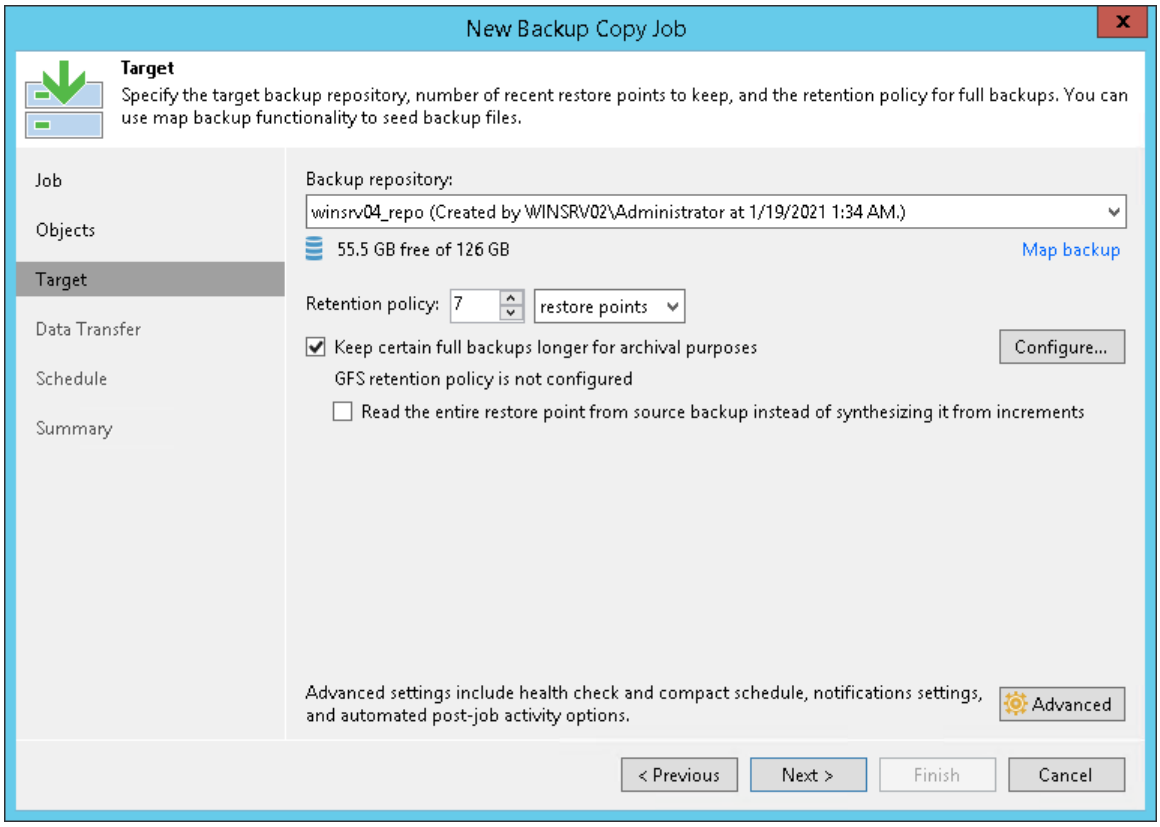
You can instruct Veeam Backup & Replication to create archive full backups with the following methods:

- Synthetic full method — Veeam Backup & Replication synthesizes archive full backups using restore points in the target backup repository.
- Active full method — Veeam Backup & Replication copies data for archive full backups from the source backup repository. This method is available only for the periodic copy mode.

Synthetic Full Method for Archive Backups

The synthetic full backup is the default method to create archive full backups. Veeam Backup & Replication does not copy data for archive full backups from the source backup repository. It synthesizes archive full backups from backup files that are already stored in the target backup repository. This approach helps reduce load on the network and production environment.

> **NOTE**
>
> The synthetic full method is not recommended if you use a deduplication storage appliance as a target backup repository. Performing a synthetic full backup in such repositories requires additional time and resources to download and decompress backup data blocks.
>
> This recommendation does not apply to HPE StoreOnce, Dell EMC Data Domain and ExaGrid:
>
> - HPE StoreOnce and Dell EMC Data Domain use virtual synthetics. Veeam Backup & Replication creates archive full backups by virtually synthesizing data blocks from existing backup files.
> - ExaGrid uses adaptive deduplication. Veeam Backup & Replication creates archive full backups from existing backup files that are stored in complete form in ExaGrid high-speed cache.



Active Full Method for Archive Backups

You can instruct Veeam Backup & Replication to create archive full backups (backups retained by the GFS scheme) with the active full backup method. The active full backup method is recommended if you use a deduplicating storage appliance as the target backup repository. Active full backup helps improve the backup copy job performance and reduce the load on the target backup repository.

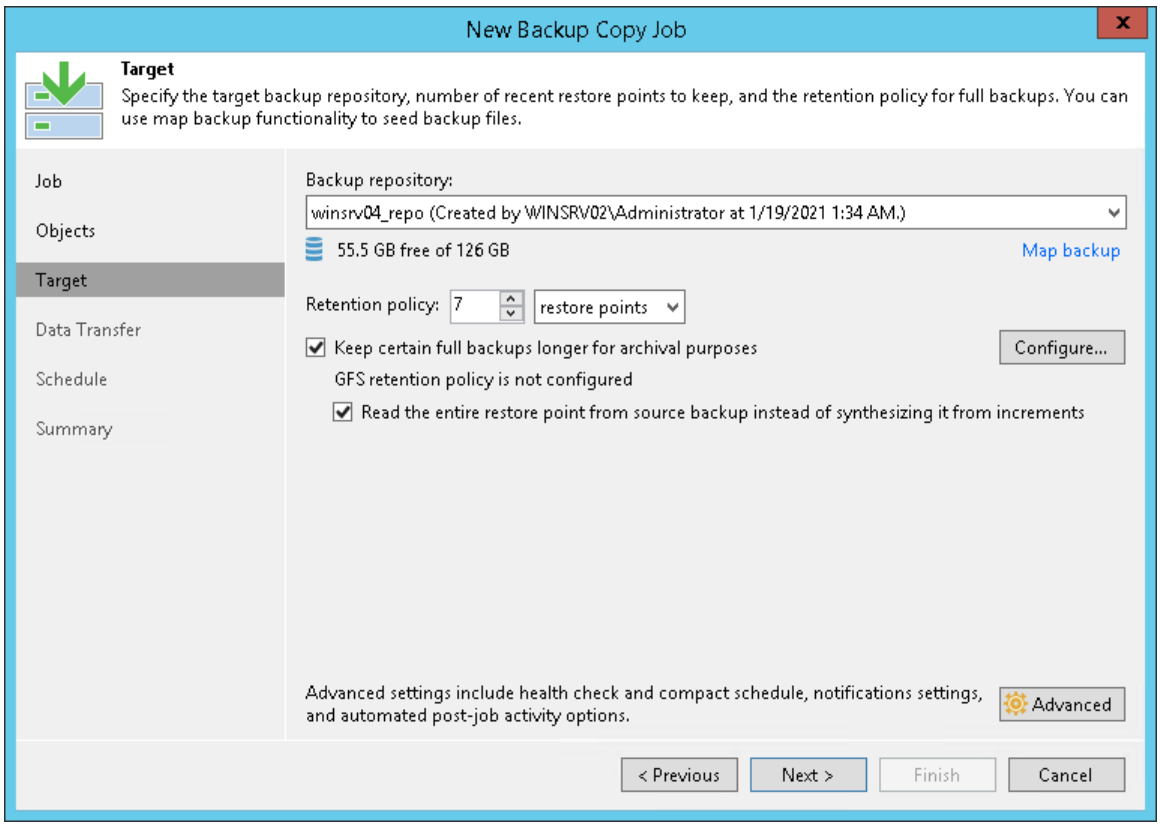
By default, Veeam Backup & Replication uses the synthetic backup method to create archive full backups. However, synthesizing archive full backups can cause problems with storage performance on deduplicating storage appliances. Deduplicating storage appliances are optimized for sequential data access. The synthetic backup creation, however, takes random I/O operations — Veeam Backup & Replication reads data from existing backup files and writes data to the synthesized archive full backup file. As a result, the storage performance can degrade.

In addition, backups reside in the target backup repository in the deduplicated and compressed state. Before creating synthetic full backups, Veeam Backup & Replication needs to download and decompress data blocks of backups, which requires additional time and resources.

NOTE

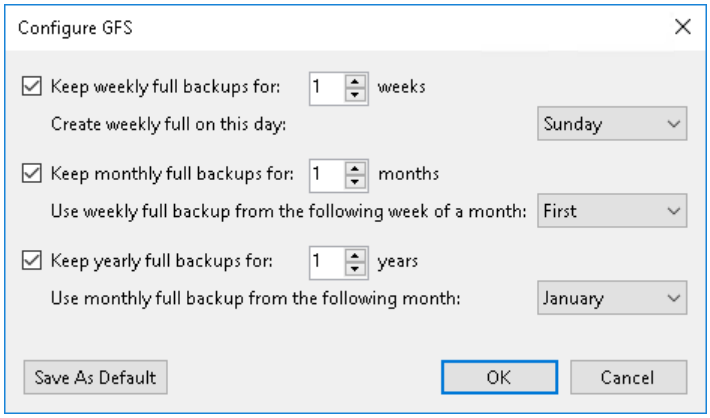
Mind the following:

- The active full backup method does not always copy the most recent restore point from the source backup repository. If the recent restore point is not created by the time the GFS task must start, Veeam Backup & Replication copies the latest available restore point from the source backup repository.
- If Veeam Backup & Replication does not manage to transfer the restore point during the backup interval, the interval is extended to finalize the transfer.



Backup Copy GFS Cycles

When you configure the GFS schedule in the backup copy job settings, you specify for how long archive backups must be stored and on which day certain GFS backup must be created. You can configure the GFS schedule in the settings of a backup copy job. For details, see Backup Copy Job: Define Backup Copy Target.



There are separate schedules for weekly, monthly and yearly full cycles. For details on settings of the GFS retention, see the following table.

GFS Backup Option	Description
Weekly GFS cycle	If you want to create weekly full backups, select the **Keep weekly full backups for** check box. Then, specify the number of weeks during which the weekly backup must be stored on the target repository. During this period the weekly backup cannot be deleted or modified. Veeam Backup & Replication creates a weekly full backup on the specified day of the week. On this day, creation of a weekly full backup starts as soon as the backup copy interval starts.

GFS Backup Option	Description
Monthly GFS cycle	To create monthly restore points, you can select the **Keep monthly full backups for** check box. Then, specify the number of months during which the monthly backup must be stored on the target repository. During this period the monthly backup cannot be deleted or modified. Veeam Backup & Replication creates monthly full backups according to a schedule that depends on whether the weekly cycle is enabled or disabled: • If weekly backups are enabled, Veeam Backup & Replication uses the weekly backup schedule and adds a monthly flag to the weekly backup. • If weekly backups are disabled, Veeam Backup & Replication creates monthly full backups on the first day of the selected week. If you select **First**, monthly backups are created on the first day of each month. If you select **Last**, monthly backups are created on different dates depending on the number of days in a month: o For 31 days: 25th of the month o For 30 days: 24th of the month o For 29 days: 23rd of the month o For 28 days: 22nd of the month If the first day of the current week has already passed and you select the first week for the monthly full backup cycle, Veeam Backup & Replication creates a monthly full backup even if it is not the first day of the current week. If you enable the monthly GFS cycle and select the first week that has already passed, Veeam Backup & Replication will create a monthly full backup only in the next month. Monthly full backup for the current month will not be created.

GFS Backup Option	Description
Yearly GFS cycle	If you want to create yearly restore points, select the **Keep yearly full backups for** check box. Then, specify the number of years during which the yearly backup must be stored on the target repository. During this period the yearly backup cannot be deleted or modified. Veeam Backup & Replication creates yearly full backups according to a schedule that depends on whether the monthly cycle is enabled or disabled: • If monthly backups are enabled, Veeam Backup & Replication uses the monthly backup schedule and adds a yearly flag to the monthly backup. • If monthly backups are disabled, Veeam Backup & Replication creates yearly full backups on the first day of the selected month. If the first day of the current month has already passed and you select the current month for the yearly full backup cycle, Veeam Backup & Replication creates a yearly full backup even if it is not the first day of the current month. If you enable the yearly GFS cycle and select the month that has already passed, Veeam Backup & Replication will create a yearly full backup only in the next year. Yearly full backup for the current year will not be created.

Backup Chain for GFS Backups

GFS retention creates yearly, monthly and weekly full backups and functions in a combination with the short-term retention. Short-term retention policy cannot delete or merge these GFS full backups.
Veeam Backup & Replication removes GFS backups only after the specified retention period for yearly/monthly/weekly backup is exceeded. Thus, the backup chain may contain more restore points than specified in the short-term retention policy.

When you enable the GFS retention, Veeam Backup & Replication no longer merges increments to full backups because GFS full backups cannot be modified. Thus, the short-term retention policy counts retention points only in the active backup chain not in the whole combination of backup chains.

> **NOTE**
>
> Veeam Backup & Replication removes GFS backup files only during running backup copy job sessions. This means that if the backup copy job does not run on the expected retention date,
> Veeam Backup & Replication will remove the GFS backup file later during the next job session.

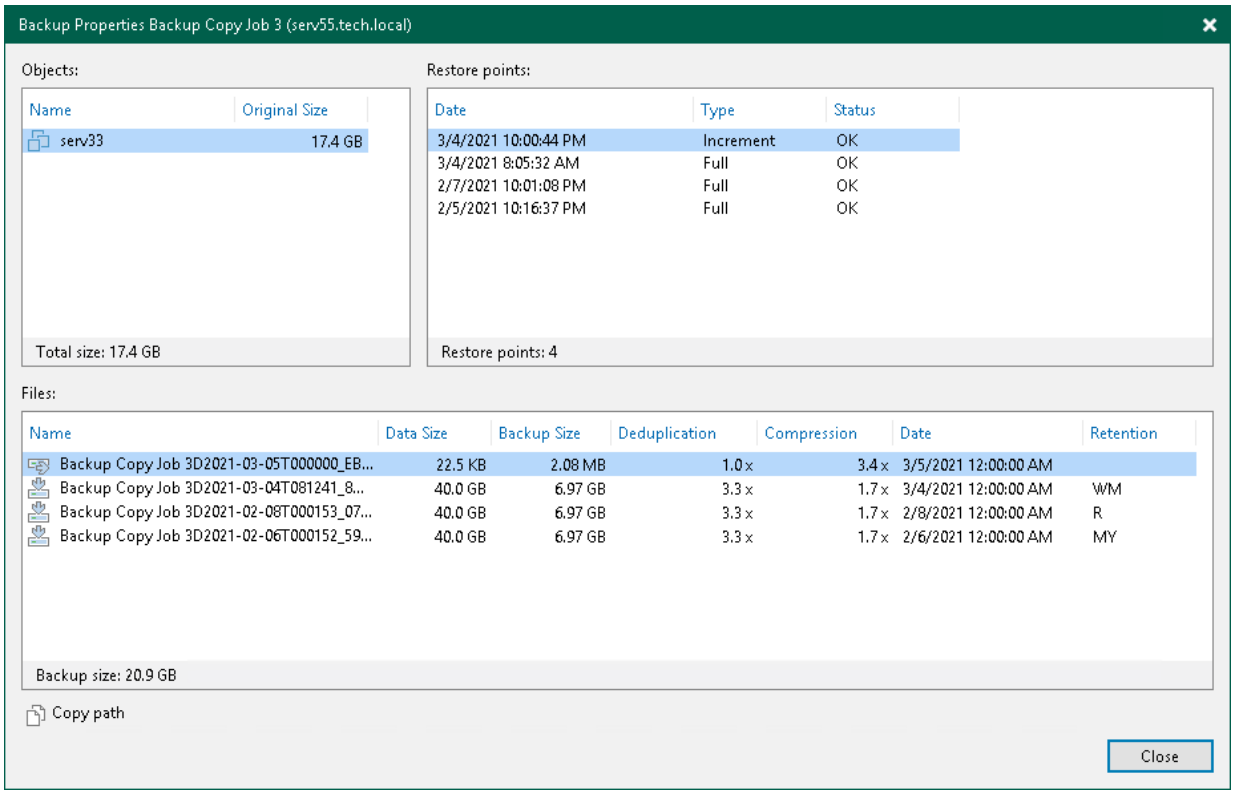
Multiple GFS Flags

If you schedule a monthly and/or yearly full backup on the same day when the weekly full backup is scheduled, Veeam Backup & Replication creates only one archive full backup. The created backup will be marked at the same time as weekly, monthly and yearly GFS backup. In the Veeam Backup & Replication console, you will see all GFS flags assigned to the backup.

The full backup can be marked as weekly, monthly and/or yearly. When transforming weekly, monthly and yearly backup chains, Veeam Backup & Replication checks flags set for the full backup file. If the full backup file belongs to some other retention policy tier and must be retained in the target backup repository, such backup file will not be removed.

Checking Which Restore Point Has GFS Flag

To check whether a restore point has a GFS flag, you can open the backup properties in the Veeam Backup & Replication console. Weekly, monthly, yearly backups have "*W*", "*M*" and/or "*Y*" flag in the *Retention* column. For instructions, see Viewing Backup Properties.



Limitations and Considerations for GFS Cycles

Before configuring the GFS retention policy for a backup copy job, mind the following limitations and considerations:

- General Settings
- Periodic Copy Mode
- Changes in GFS Retention After Upgrading from Veeam Backup & Replication 10 to version 11

General Settings

Mind the following for general settings of a backup copy job:

- You cannot enable GFS retention settings if you use a backup repository with rotated drives as the target backup repository.
- [For yearly GFS cycle] If you enable only the yearly GFS cycle, you can encounter the case when there is one full backup and a large number of increments for the whole year. To avoid this case, it is recommended to enable an additional weekly GFS cycle. Weekly GFS cycle will update the backup chain every week which will allow removing excessive increment files.
- If it is the day when the GFS full backup must be created and there were no new backup files since the last run of the backup copy, the GFS full backup will be created from the latest available backup chain.

- GFS full backups cannot be merged or deleted by short-term retention. However, regular (R) full backups can be merged and removed by short-term retention.

Periodic Copy Mode

If you want to use the periodic copy mode, mind the following:

- If the backup copy interval is not enough to get the active full backup, the interval will be prolonged until the active full process is finished.
- Veeam Backup & Replication creates a GFS full backup even if the GFS full backup creation is scheduled when the backup copy interval is not finished. On the day when the GFS full backup must be created, Veeam Backup & Replication shows a warning that the current backup copy interval will be completed, completes the interval, and starts a new one. The way how Veeam Backup & Replication behaves further depends on the selected backup copy GFS method:
  - In case of the synthetic full method, Veeam Backup & Replication first copies data for an incremental backup from the source backup repository and then, on the target backup repository, synthesizes the GFS full backup using this data and data of the already stored backup files.
  - In case of the active full method, Veeam Backup & Replication copies data for the GFS full backup from the source backup repository and creates the GFS full backup on the target backup repository.

Changes in GFS Retention After Upgrading from Veeam Backup & Replication 10 to version 11

In Veeam Backup & Replication version 11, the backup copy GFS retention settings are different from the settings in version 10. If you upgrade to Veeam Backup & Replication 11, all GFS retention settings of existing backup copy jobs are automatically switched to the new format with minimal changes:

- **Weekly GFS retention**: If in version 10 the GFS setting was to keep *5* weekly backups, in version 11 the setting is changed to keep weekly backups for *5* weeks.
- **Monthly GFS retention**: If in version 10 the monthly GFS schedule was set to a period between the 1st Monday and 2nd Sunday, in version 11 the monthly GFS settings is changed to *First week*.

  If the monthly GFS schedule was set to a period between the 3rd Monday and last Sunday, in version 11 the monthly GFS settings is changed to *Last week*.
- **Yearly GFS retention**: If the yearly GFS schedule was set to a certain day of the month, in version 11 the schedule is set to the first day of the specified month.

  If in version 10 the yearly GFS schedule was set to the first-fourth Monday-Sunday of the year, in version 11 the schedule is changed to the first day of January.

  If the yearly GFS schedule was set to the last Monday-Sunday of the year, in version 11 the schedule is set to the first day of December.
- **Quarterly GFS retention**: Since Veeam Backup & Replication version 11, quarterly GFS retention policy option is deprecated.

  If in version 10 the quarterly GFS policy was enabled, in version 11 three additional months are added to the monthly GFS policy to compensate the quarterly full backups.

  If in version 10, the GFS policy was set to X monthly backups and Y quarterly backups. Then, in version 11, the retention policy is switched to store monthly backups for (X + 3Y) months.

IMPORTANT

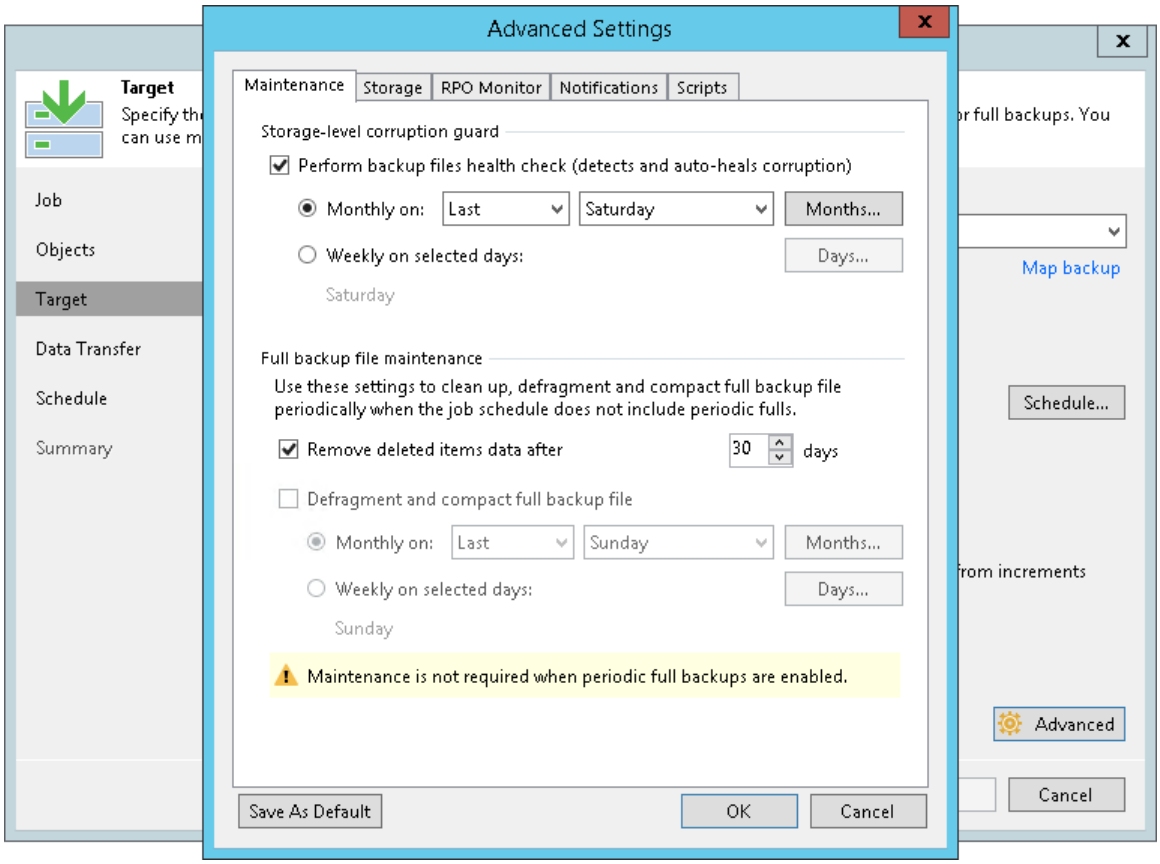
[For synthetic method] Before the upgrade to Veeam Backup & Replication 11 or later, make sure that all GFS candidates (incremental restore points created on days when GFS was scheduled and that are expected to be transformed into full GFS restore points) are already transformed into GFS restore points. To force the backup copy job to transform all GFS candidates, you can temporarily decrease the short-term retention to a value less than the number of restore points between the latest restore point and the most recent GFS candidate and then wait till all the candidates are transformed.

Before Veeam Backup & Replication version 11, Veeam Backup & Replication created GFS candidates on days when GFS was scheduled and only then transformed them into full GFS restore points according to the short-term retention. For more information on how restore points were transformed, see Synthetic Weekly Full Backups. Starting from Veeam Backup & Replication version 11, Veeam Backup & Replication creates GFS restore points according to a new schedule and creates them right on the scheduled days. After the upgrade, Veeam Backup & Replication no longer transforms previous GFS candidates into full GFS restore points. This means, that all GFS candidates lose their GFS status, they become regular incremental restore points and are deleted according to the short-term retention policy.

Deleted Items Retention

After you configure a backup copy job, you may want to change something in the virtual infrastructure. For example, you may remove some virtual or physical machines or move VMs to another location. You may also exclude VMs from the backup copy job that has already run for some time.

By default, when you remove a machine protected by Veeam Backup & Replication from the virtual infrastructure, exclude a machine from the backup copy job or stop protecting a machine with Veeam Agent, the copied data still remains in backup files in the target backup repository. To avoid keeping redundant data on disk, you can enable the **Remove deleted items data after** option in the backup copy job settings. With this option enabled, at the end of every synchronization cycle Veeam Backup & Replication will remove data for deleted machines from backup files in the target backup repository.



Veeam Backup & Replication removes data for deleted machine only if two conditions are met:

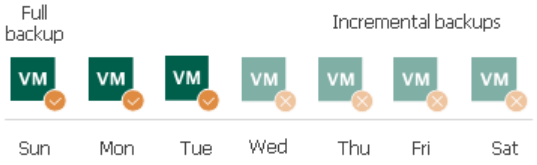
1. Veeam Backup & Replication has not created a valid restore points for the deleted machine for the number of days specified in the **Remove deleted items data after** field.
2. The backup chain in the target backup repository does not contain any successful incremental restore points for the deleted machine.

This approach helps ensure that data for deleted machines can be saved by the GFS retention.

For example:

- The retention for the backup copy job is set to 7.
- The retention period for deleted machine is set to 3 days.

The backup copy job has created 3 successful restore points — a full backup and two incremental backups. During the next 4 days, no successful restore points were created. At the next synchronization cycle, Veeam Backup & Replication will not remove data for the deleted machine from the target backup repository as the backup chain contains successful incremental restore points for this machine.



IMPORTANT

Mind the following:

- The deleted items retention applies only to regular backup chains. Veeam Backup & Replication does not remove data for deleted machines from weekly, monthly and yearly backups.
- When Veeam Backup & Replication removes data for deleted machines from regular backup chains, it does not free up space in the backup repository. It marks the space as available to be overwritten, and this space is overwritten during subsequent job sessions or the backup file compact operation.
- When Veeam Backup & Replication removes data for deleted machines from per-machine backup chains, it does not mark the space as available but deletes backup files since they contain data for 1 machine only.

Veeam Backup & Replication does not analyze the reason for which the machine has not been processed during the backup copy session. For example, a VM may be regarded as deleted if Veeam Backup & Replication has failed to obtain data for the VM from the virtual infrastructure, the VM has failed to be processed in time during the backup copy session and so on.

For this reason, you must be careful when specifying the retention period for deleted machines. If the retention period is too short, Veeam Backup & Replication may remove from the backup chain restore points that you still require.

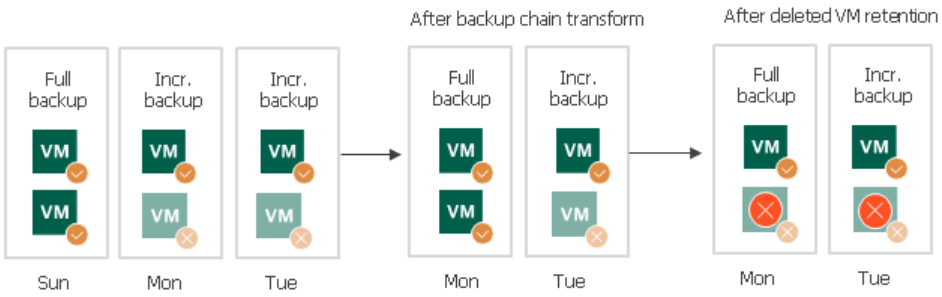
For example, a backup copy job is configured to process 2 VMs and has the following settings:

- The backup copy starts once a day.
- The retention for the backup copy job is set to 2.
- The retention period for deleted VMs is set to 1 day.

The backup copy job runs in the following way:

1. On Sunday, the backup copy job creates a full backup for 2 VMs – VM1 and VM2.
2. On Monday, the backup copy job creates an incremental backup for VM1. The backup copy job does not manage to process VM2 in time.
3. On Tuesday, the backup copy job creates an incremental backup for VM1. The backup copy job does not manage to process VM2 in time.
4. At the end of the backup copy job session on Tuesday, Veeam Backup & Replication transforms the backup chain and detects deleted VMs. Veeam Backup & Replication regards VM2 as a deleted VM — the deleted VMs retention is set to 1 day, and after transformation, there are no valid restore points for this VM in the backup chain.

As a result, after the backup copy session on Tuesday backup files in the target backup repository will not contain data for VM2.



Health Check for Backup Files

You can instruct Veeam Backup & Replication to periodically perform a health check for the latest restore point or backup file in the backup chain. The health check helps Veeam Backup & Replication make sure that further restore will be possible.

The health check starts as soon as a backup copy job starts. The health check verifies restore points (full backup files or related full and incremental backup files). Only the latest restore points are verified. On the day when the health check is scheduled, Veeam Backup & Replication performs the health check only once, even if the backup copy job runs several times.

Note that the health check process differs for backup files stored in the HPE StoreOnce repository. For details, see Health Check for Backup Files Stored on HPE StoreOnce.

> **NOTE**
>
> If you perform the health check for the encrypted backup files, Veeam Backup & Replication will pass encryption keys to the regular backup repository or cloud repository. For more information on encryption, see Data Encryption.

How Health Check Works

Veeam Backup & Replication performs the health check in the following way:

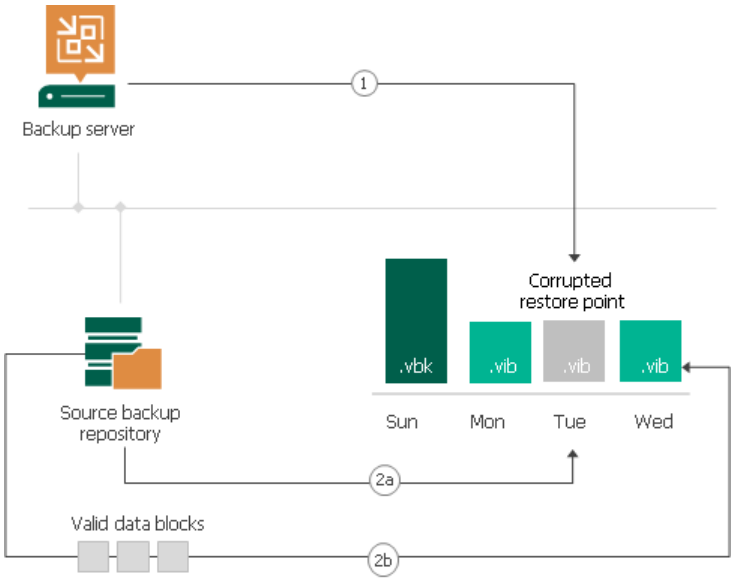
1. As soon as a restore point is saved to the backup repository, Veeam Backup & Replication calculates CRC values for backup metadata and hash values for data blocks of a disk in the backup file and saves these values in the metadata of the backup file, together with copied data.
2. On the day when the health check is scheduled, Veeam Backup & Replication performs the following actions:
    a. At the beginning of the backup copy session, Veeam Backup & Replication performs the health check for the latest restore point in the backup chain. If the latest restore point in the backup chain is incomplete, Veeam Backup & Replication checks the restore point preceding the latest one.

       Veeam Backup & Replication calculates CRC values for backup metadata and hash values for disks data blocks in the backup file and compares them with the CRC and hash values that are already stored in the backup file.

    b. If the health check detects corrupted data blocks, together with data blocks for the new restore point, Veeam Backup & Replication transports valid data blocks for the corrupted restore point. The valid data blocks are stored to the new incremental restore point created with this backup copy session. As a result, the backup chain gets "fixed", and you get a possibility to restore data from restore points following the corrupted restore point.

NOTE

If the backup copy job uses WAN accelerators, Veeam Backup & Replication attempts to find data blocks in the global cache not to transfer data over the network. For more information, see WAN Acceleration.



Health Check for Backup Files Stored on HPE StoreOnce

The health check starts as soon as a backup copy job transfers backup files to the target repository. The health check verifies backup files, not restore points, and only those backup files transferred during the current job session.

How Health Check Works

Veeam Backup & Replication performs the health check in the following way:

1. As soon as a backup file is transferred to the target repository, Veeam Backup & Replication calls the HPE StoreOnce internal method that calculates checksums for data blocks of this backup file.

   The checksum calculation requires the rehydration of the verified data. This may lower the performance on the target repository especially during the first backup copy session. During this session, the health check verifies all backup files in source backup chains.

2. Veeam Backup & Replication compares the calculated checksums and the checksums already stored in the backup file.

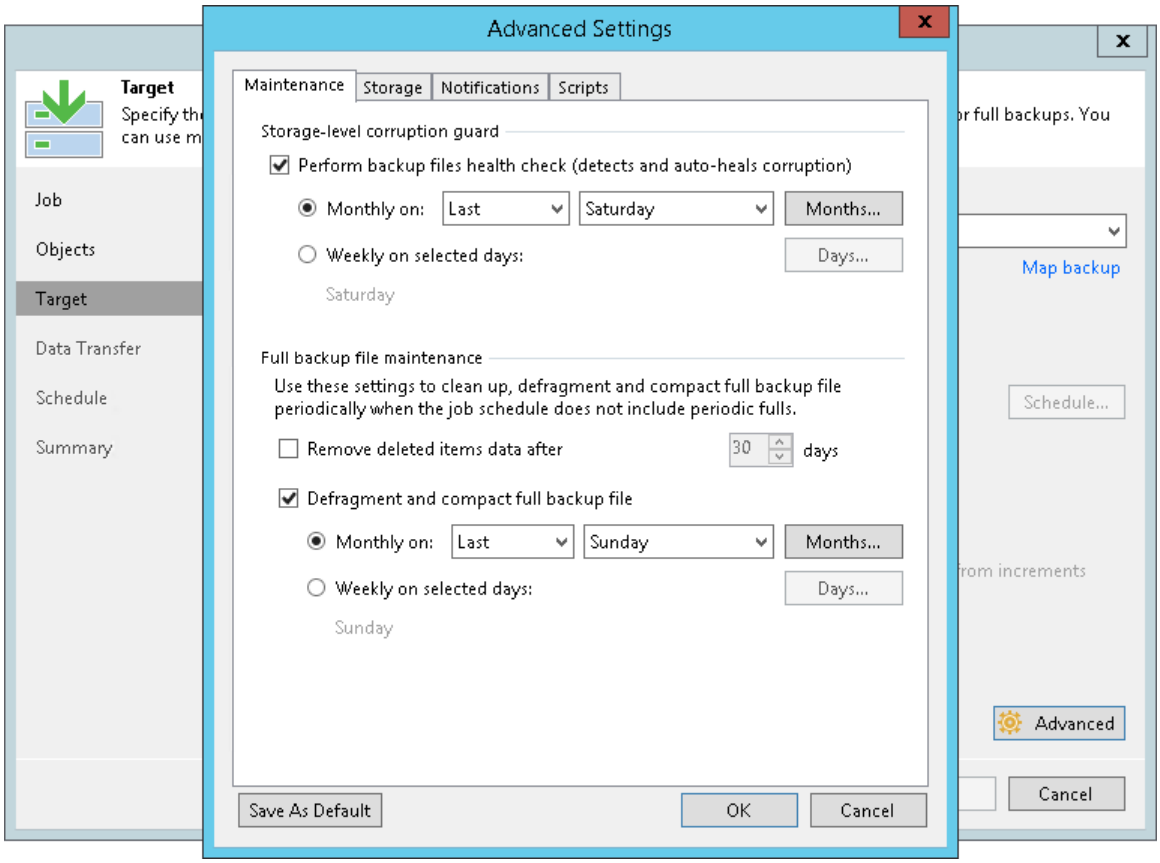
3. If the health check detects corrupted data blocks, Veeam Backup & Replication deletes the backup file from the target repository. On the next backup copy job session, Veeam Backup & Replication transfers and performs the health check for this file again.

Compact of Full Backup File

The backup copy job constantly transforms the full backup file in the backup chain to meet retention policy settings. The transformation process, however, has a side effect. In the long run, the full backup file grows large and gets fragmented. The file data occurs to be written to non-contiguous clusters on disk, and operations of reading and writing data from and to the backup file slow down.

To resolve the fragmentation problem, you can instruct Veeam Backup & Replication to compact the full backup file periodically. During the file compact operation, Veeam Backup & Replication creates a new full backup file in the target repository: it copies existing data blocks from the old backup file, rearranges and stores them close to each other. As a result, the full backup file gets defragmented, its size reduces and the speed of reading and writing from and to the file increases.

To compact the full backup file periodically, you must enable the **Defragment and compact full backup file** option in the backup copy job settings and define the compact operation schedule. By default, the compact operation is performed on the last Sunday of every month. You can change the compact operation schedule and instruct Veeam Backup & Replication to perform it weekly or monthly on specific days.



Backup Copy Job Mapping

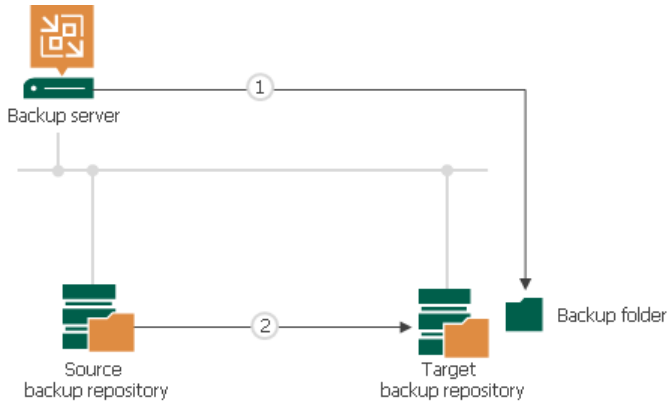
Backup copy job mapping helps you reduce the amount of data transferred over network and decrease the load on WAN accelerators or slow connections.

If you use the target backup repository also as a target for other backup copy or backup jobs, you can already have a backup of machines that you want to copy. In this case, you can map the backup copy job to this backup.

A backup copy job mapped to a backup is performed in the following way:

1. Veeam Backup & Replication accesses a backup to which you map the backup copy job. The backup may have any number of restore points in the chain. This backup chain will be used as a seed for the further backup copying process.
2. During subsequent backup copy sessions, Veeam Backup & Replication copies restore points in a regular manner. It copies only incremental changes and stores them as new restore points next to the seed backup chain.

A mapped backup copy job does not store copied restore points in a dedicated folder in the target backup repository. Instead, it stores restore points to the same folder where the "seed" backup chain resides.



Creating Seed for Backup Copy Job

Backup copy jobs have limitations for backups that can be used as seeds. The limitations are listed in the Map Backup File section.

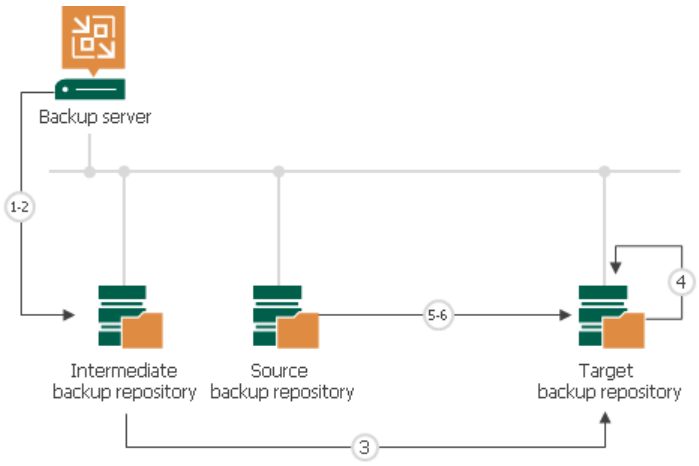
To overcome these limitations, you can create a seed using a secondary backup copy job. To create a seed for the primary backup copy job, do the following:

1. Create a backup copy job. Add machines or backups whose restore points you want to copy to this backup copy job. Target the backup copy job to some backup repository on the source side. This backup repository will be used as an intermediate one.
2. Run the backup copy job to create a full backup file (VBK) in the intermediate backup repository.
3. Transfer the created VBK file and VBM file from the intermediate backup repository to the target backup repository.
4. Perform repository rescan to populate the target backup repository.

   If the initial backup file was encrypted, you must enter a password to unlock the full backup file. Otherwise, Veeam Backup & Replication will not display the full backup file in the list of backups in the backup repository. For more information, see Importing Encrypted Backups.

5. Remap the backup copy job to the full backup file that you have created and transferred to the target backup repository.
6. Click **Sync Now** to start the backup copy.

As a result, Veeam Backup & Replication will use the full backup file as a seed. When a new restore point for the machine is available in the source backup repository, Veeam Backup & Replication will copy the restore point to the target backup repository and store it next to the full backup seed.



Specifying Backup Copy Interval for Periodic Copy Mode

When you configure a backup copy job, make sure that its backup copy interval covers the whole chain of restore points in the backup repository from which you plan to copy backups. The length of the backup copy interval has an impact on the algorithm of restore point selection. Veeam Backup & Replication copies only restore points that match the following criterion:

```
Time of restore point creation >= current time – backup copy interval
```

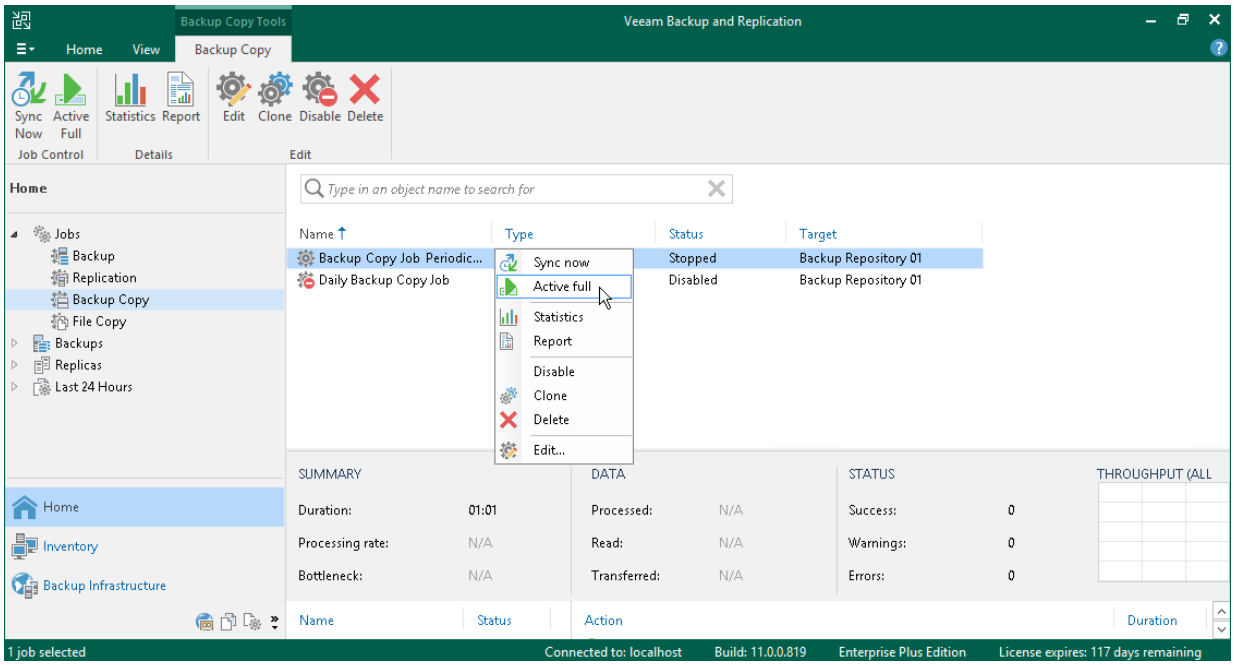
That is, if you have a backup chain whose earliest restore point is 1 week old, you need to set the backup copy interval to 1 week. If you set the backup copy interval to a smaller time interval, for example, 1 day, all restore points that are older than 1 day will fall out of the search scope, and Veeam Backup & Replication will not transfer such restore points. For more information, see Restore Point Selection.

Active Full Backup Copies

You can manually create an ad-hoc full backup for the backup copy job — active full backup copy, and add it to the backup chain in the target backup repository. To do this, you can use the **Active Full** button on the ribbon or the **Active Full** command from the shortcut menu.

Active full backup copy can be helpful if you want to change backup copy job settings, for example, enable or disable encryption. Veeam Backup & Replication will apply new settings starting from this full backup.

Veeam Backup & Replication treats archive full backups created with the active full backup method as regular backups and applies regular retention policy rules to maintain the necessary number of restore points.



Retention Policy for Active Full Backups

If you create active full backups for backup copy jobs, Veeam Backup & Replication applies to the backup chain retention rules of the forward incremental backup method. Veeam Backup & Replication waits until the number of restore points in the new backup chain is equal to the retention policy setting, and then removes the previous backup chain on the whole. For more information, see Retention for Forward Incremental Backup.

Automatic Job Retries

Veeam Backup & Replication automatically retries several operations that are performed within a backup copy job session.

Job Tasks Retry

By default, Veeam Backup & Replication automatically retries a failed backup copy task 5 times within one backup copy job session. A new task is started immediately after the previous one, without any interval.

The backup copy task is retried only if the previous task has failed and a restore point has not been copied to the target backup repository. Veeam Backup & Replication does not perform a retry if a task has finished with the *Success* or the *Warning* status.

The backup copy task is retried during the same backup copy session only. If a restore point fails to be copied during all retries in the current backup copy session, Veeam Backup & Replication marks the current task as failed. In case of the periodic copy mode, Veeam Backup & Replication also waits for the expiration of the backup copy interval. After that, Veeam Backup & Replication performs the necessary transformation processes and starts a new backup copy session.

A backup copy job can process several machines. If only some machines are successfully processed by the backup copy task, Veeam Backup & Replication creates a restore point holding data for these machines in the target backup repository. Veeam Backup & Replication will attempt to process restore points for all machines during the next backup copy session.

> **NOTE**
>
> Some errors from WAN accelerators can block backup copy job retries. For example, if there is no space in the global cache on the target WAN accelerator, Veeam Backup & Replication puts backup copy operations on hold and waits for the expiration of the backup copy session.

Transformation Retry

After the backup copying task, Veeam Backup & Replication may perform a number of additional transformation processes in the target backup repository. These processes include the backup chain transformation, removing of deleted machines from restore points and compacting a full backup file. For more information, see Transformation Processes.

Veeam Backup & Replication may fail to perform transformation: for example, if the backup file in the target backup repository is locked by the file-level restore session. By default, Veeam Backup & Replication automatically retries transformation processes for 5 times. The first interval between retries is 1 minute; the interval doubles with every new attempt. If all retries of transformation processes fail, Veeam Backup & Replication does the following:

- [For the immediate copy mode] Stops the job with the *Fail* status and waits for the new job session.
- [For the periodic copy mode] Puts the job to the idle state and waits for the new backup copy interval to begin.

Virtual Infrastructure Access Retry

At the beginning of every backup copy session, Veeam Backup & Replication accesses the virtual infrastructure to make up a list of machines processed by the job.

Veeam Backup & Replication may fail to access the virtual infrastructure for some reason: for example, in case the vCenter Server is not responding. By default, Veeam Backup & Replication automatically retries access operations for 5 times with a 5 minute interval.

Creating Backup Copy Jobs for VMs and Physical Machines

To copy backups to a secondary location, you must configure a backup copy job. The backup copy job defines how, where and when to copy backups. One backup copy job can be used to process one or multiple machines. Machines included in the job are processed in parallel. If a machine included in the backup copy job has multiple disks, disks are processed sequentially, one after another.

> **NOTE**
>
> If you want to copy backups between HPE StoreOnce repositories, follow the instructions listed in Creating Backup Copy Jobs for HPE StoreOnce Repositories.
>
> If you want to copy file share backups, follow the instructions listed in Creating File Share Backup Jobs.

Before you create a job, check prerequisites. Then use the **New Backup Copy Job** wizard to configure the backup copy job.

Before You Begin

Before you create a backup copy job, check the following prerequisites:

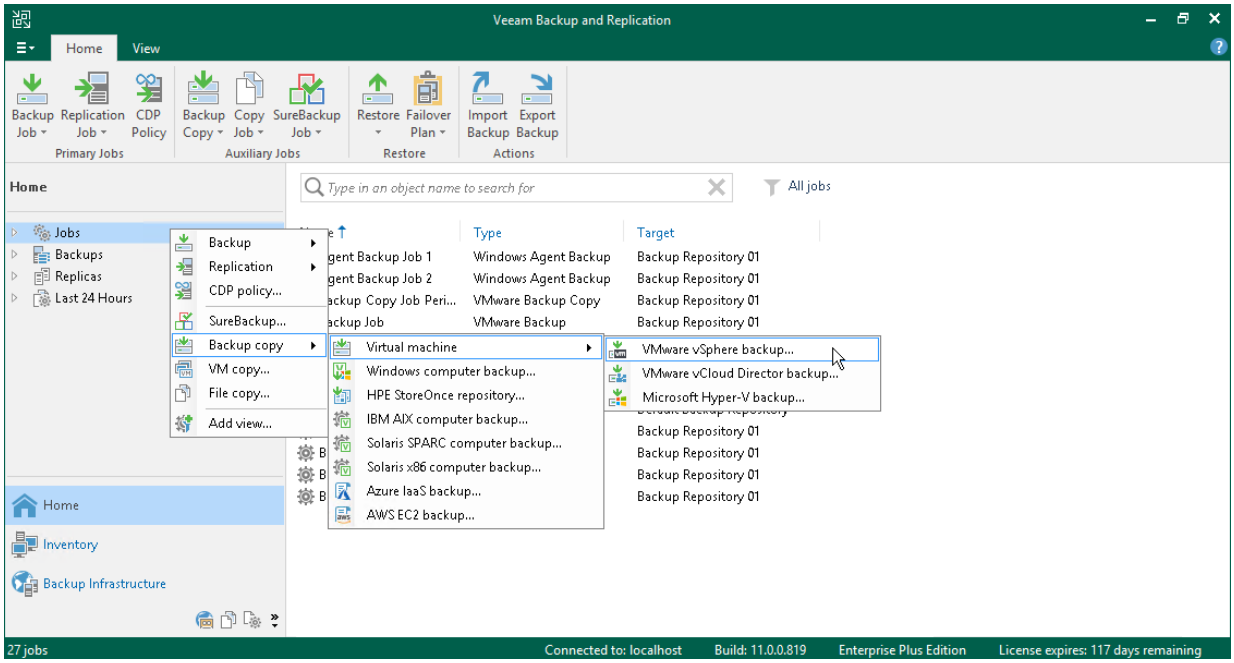
- Backup infrastructure components that will take part in the backup copy process must be added to the backup infrastructure. This include target backup repository to which backups must be copied. For more information on adding components, see Backup Infrastructure Components.
- If you plan to use pre-job and/or post-job scripts, you must create scripts before you configure the backup copy job.
- If you plan to copy backups to an HPE StoreOnce repository, check limitations and requirements for it.
- If you plan to use WAN accelerators, check that you use the Enterprise Plus edition of Veeam Backup & Replication and that target and source WAN accelerators are added to the backup infrastructure. For more information, see Adding WAN Accelerators.

Step 1. Launch New Backup Copy Job Wizard

To run the **New Backup Copy Job** wizard, do one of the following:

- On the **Home** tab, click **Backup Copy** and select a platform:
  - Virtual Machine > VMware vSphere backup
  - Virtual Machine > VMware vCloud Director backup
  - Windows computer backup
  - Linux computer backup
  - Veeam Agent for Mac backup
  - Veeam Agent for IBM AIX backup
  - Solaris x86 backup
  - Solaris SPARC backup
  - AWS EC2 backup
  - Azure IaaS backup
  - GCE backup
  - Nutanix backup
- Open the **Home** view, in the inventory pane right-click **Jobs** or right-click anywhere in the working area, and select the necessary option:
  - **Backup copy** > **Virtual machine** > **VMware vSphere backup** – if you want to create a copy of a VM backup.
  - **Backup copy** > **Virtual machine** > **VMware vCloud director backup** – if you want to create a copy of a vCD backup.
  - **Backup copy** > **Windows computer backup** – if you want to create a copy of a Veeam Agent backup created for Microsoft Windows machines.
  - **Backup copy** > **Linux computer backup** – if you want to create a copy of a Veeam Agent backup created for Linux machines.
  - **Backup copy** > **Mac computer backup** – if you want to create a copy of a Veeam Agent backup created for Mac machines.
  - **Backup copy** > **IBM AIX computer backup** – if you want to create a copy of a Veeam Agent backup created for IBM AIX machines.
  - **Backup copy** > **Solaris SPARC computer backup** – if you want to create a copy of a Veeam Agent backup created for Oracle Solaris machines based on SPARC architecture.
  - **Backup copy** > **Solaris x86 computer backup** – if you want to create a copy of a Veeam Agent backup created for Oracle Solaris machines based on Intel x86 architecture.
  - **Backup copy** > **AWS EC2 backup** – if you want to create a copy of an EC2 instance backup created with Veeam Backup for AWS.

- **Backup copy** > **Azure IaaS backup** – if you want to create a copy of an Microsoft Azure VM backup created with Veeam Backup for Microsoft Azure.
- **Backup copy** > **GCE backup** – if you want to create a copy of a VM instance backup created with Veeam Backup for Google Cloud Platform.
- **Backup copy** > **Nutanix backup** – if you want to create a copy of a backup created with Veeam Backup for Nutanix AHV.

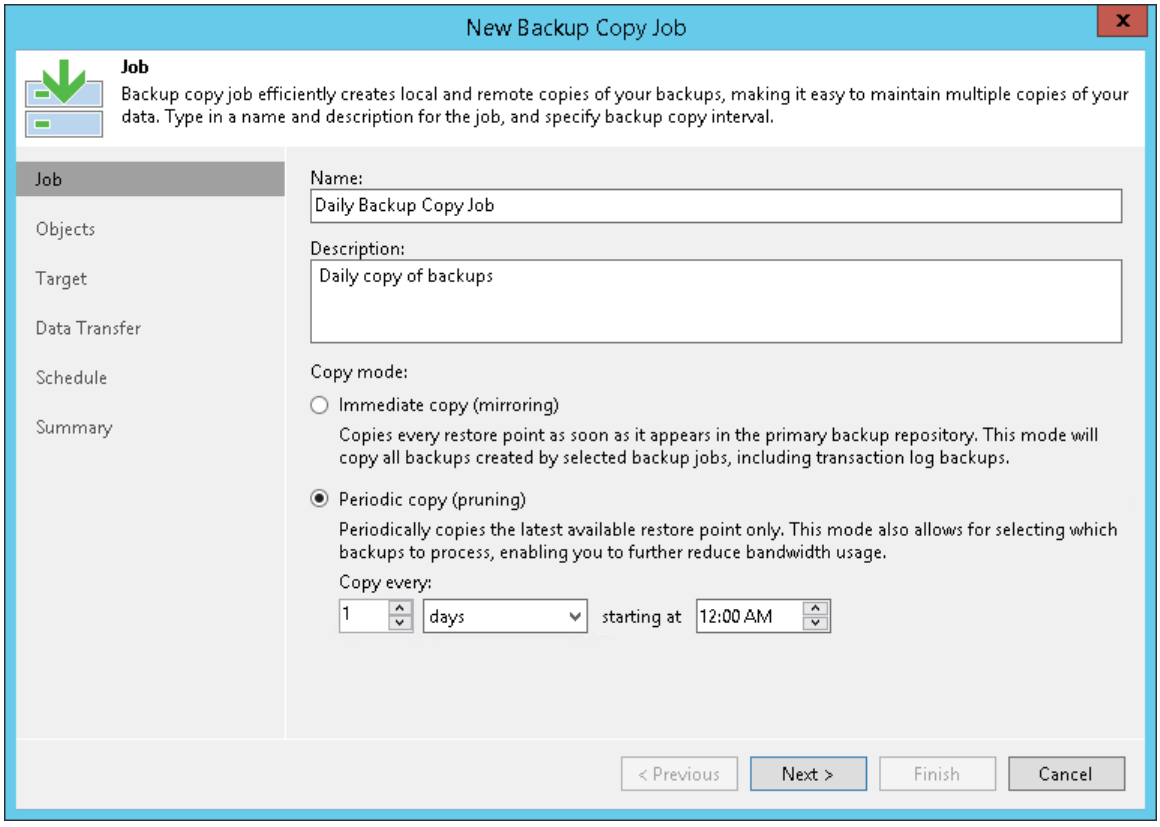


Step 2. Specify Job Name and Copy Mode

At the **Job** step of the wizard, specify basic settings for the backup copy job:

1. In the **Name** field, specify a name for the job.
2. In the **Description** field, provide a description for the job.

   The default description contains information on a user who created the job, date and time when the job was created.
3. Select a backup copy mode. Note that after you finish configuring the backup copy job, you will not be able to change the selected mode. For more information on copy modes and backup types supported in each mode, see Backup Copy Modes.
   - Select **Immediate copy** to copy new restore points and, if required, log backups as soon as they appear.
   - Select **Periodic copy** to copy the most recent restore points once within a backup copy interval.
     - In the **Copy every** field, specify the backup copy interval. For more information on how to choose backup copy interval duration, see Backup Copy Intervals.
     - If you have selected a daily backup copy interval, specify the start time for the interval. By default, the daily backup copy interval starts at 12:00 AM.



Step 3. Select Workloads to Process

At the **Objects** step of the wizard, select workloads whose restore points you want to copy to the target backup repository:

1. Click **Add**.
2. [For the immediate copy mode] Select a type of a source from which you want to copy restore points:
   - **From jobs**. You will see existing backup jobs. Veeam Backup & Replication will copy restore points created by the selected backup jobs.
   - **From repositories**. You will see all backup repositories in the backup infrastructure. Veeam Backup & Replication will copy restore points stored on the selected backup repositories. Veeam Backup & Replication will copy only restore points of workloads of the same platform as you have selected when launching the wizard.

     If you select repositories as sources, and target new jobs to the repositories in future, Veeam Backup & Replication will update backup copy job settings automatically to include these jobs to be copied.
3. [For the periodic copy mode] Select where to browse for workloads whose restore points you want to copy:
   - **From infrastructure**. You will browse for VMs and VM containers in the entire virtual infrastructure. Veeam Backup & Replication will search for restore points of the selected VMs in all backup repositories in the backup infrastructure and will copy the most recent restore points. You can limit the search scope by selecting only specific backup repositories for the backup copy job. For more information, see Select Source Backup Repositories.

     This source is available only for backup copy jobs that process VMware or Hyper-V VMs.
   - **From backups**. You will browse for workloads in existing backups. Veeam Backup & Replication will search for restore points of the selected workloads in all backups of backup jobs created on the backup server and will copy the most recent restore points. You can limit the search scope by selecting only specific backup repositories for the backup copy job. For more information, see Select Source Backup Repositories.

     This source is the only available one for backup copy jobs that process backups of Amazon EC2 instances and Microsoft Azure VMs. You can select instances/VMs only from backups that are stored in external repositories.
   - **From jobs**. You will browse for workloads in existing backup jobs. Veeam Backup & Replication will search for restore points of the selected workloads in backups created by the selected jobs. Note that if multiple jobs process one workload, Veeam Backup & Replication copies only restore points created by the first job in the **Objects to process** list.
4. In the **Add Objects** window, select the necessary sources or workloads.
5. Click **Add**.
6. [For the immediate copy mode] If you have configured processing of transaction log backups in the source backup jobs, and want to copy these log backups to the target repository, select the **Include database transaction log backups** check box.

> **NOTE**
>
> When you copy Veeam Agent backup jobs that process clusters with shared disks, the network traffic will be higher compared to the traffic sent when Veeam Agent backup jobs run. This is because Veeam Agent backup jobs send data of shared disks only with the owner node and then, within the target storage, clone this data to other nodes; whereas backup copy jobs send data as it is stored on the storage — each node with the cloned data.

As an alternative, you can create a backup copy job with an empty source — that is, do not add any workloads at this step of the wizard. In this case, you need to configure a secondary destination for the source backup job and link it to the created backup copy job. For more information, see Linking Backup Jobs to Backup Copy Jobs.

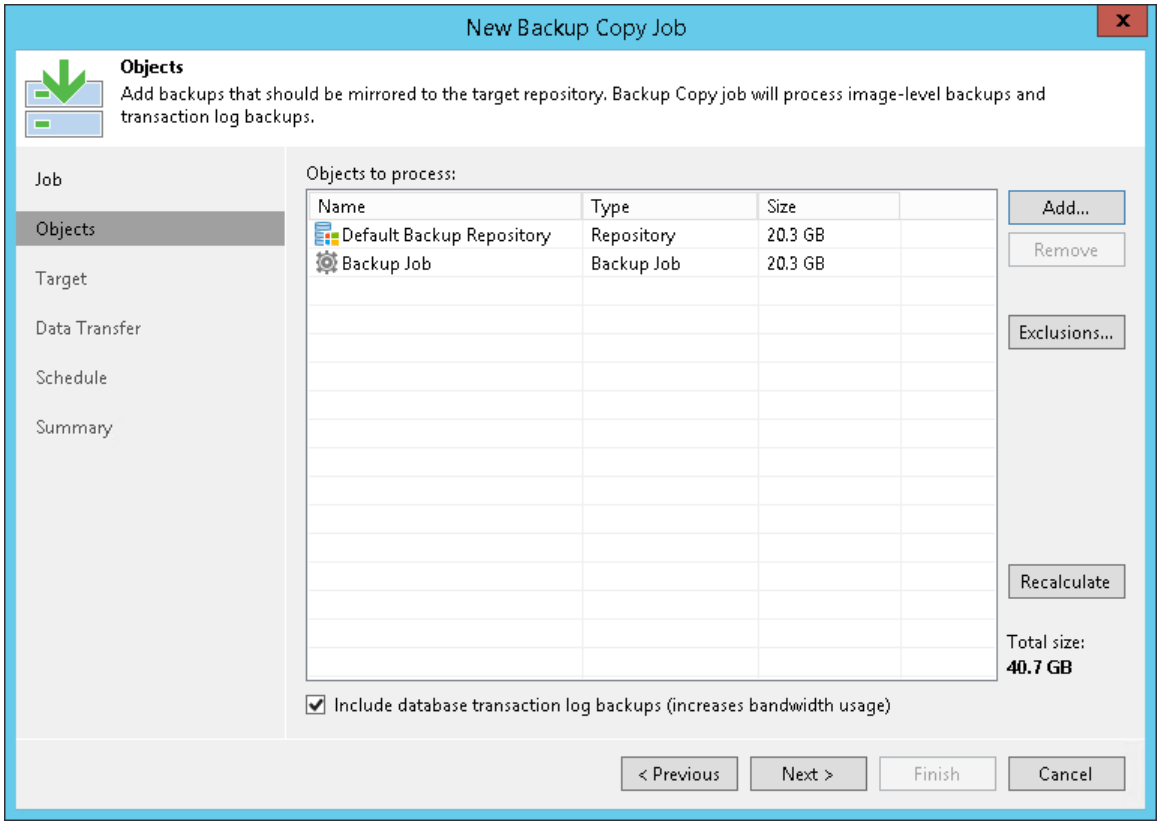
Limitations for Workload Selection

When selecting workloads, consider the following limitations:

- Within one backup copy job, Veeam Backup & Replication processes workloads of one platform only. For example, within one backup copy job, you cannot process backups of VMware vSphere and Microsoft Hyper-V VMs or backups of VMs and physical machines.
- If you use tags to categorize virtual infrastructure objects, check limitations for VM tags. For more information, see VM Tags.
- The following applies to the periodic copy mode:
  - If a workload that you want to add to a backup copy job is processed by multiple source backup jobs that use different block sizes, add this workload **From jobs**.

    When you add a workload **From infrastructure** or **From backups**, Veeam Backup & Replication picks the most recent backup as a data source. If Veeam Backup & Replication picks source backups with different block sizes during different backup copy intervals, the backup copy job will fail.

- If **Use per-machine backup files** option is disabled in a repository that you are planning to use as the target, you can select only workloads that have the same data block size — that is, restore points of all workloads that you select to process must have the same data block size. If the **Use per-machine backup files** option is enabled, you can select workloads that have different block sizes.



Step 4. Exclude Objects from Backup Copy Job

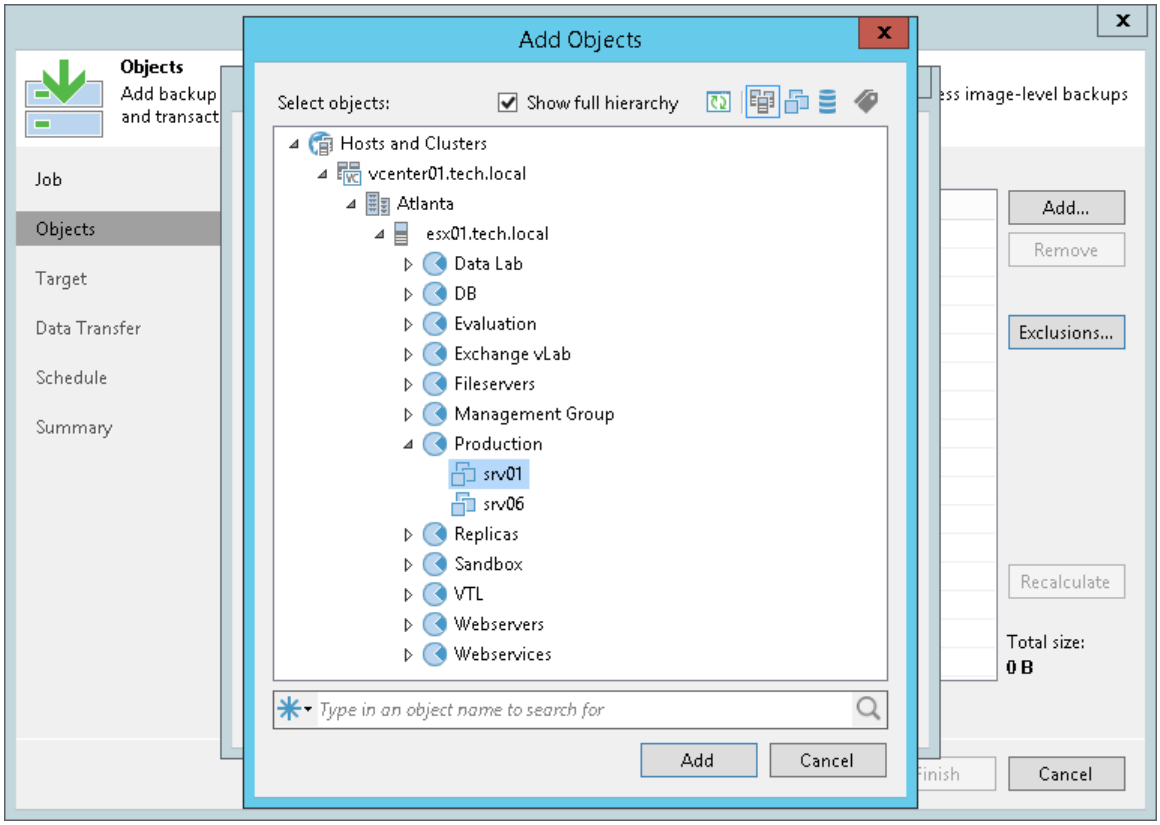
This option is available only for virtual machines.

To specify which objects you want to exclude from the backup copy job:

1. At the **Objects** step of the wizard, click **Exclusions**.
2. In the **Exclusions** window, click **Add**.

   [For the immediate copy mode] If you have added repositories as sources of restore points, you can exclude individual VMs or backup jobs from processing. To exclude individual VMs, click **Add** > **VMs**. To exclude individual jobs, click **Add** > **Jobs**.
3. In the **Add Objects** window, select objects that you want to exclude.

   When you exclude VMs, you can use the **Show full hierarchy** check box to display the hierarchy of all hosts added to Veeam Backup & Replication.
4. Click **Add**.
5. Click **OK**.



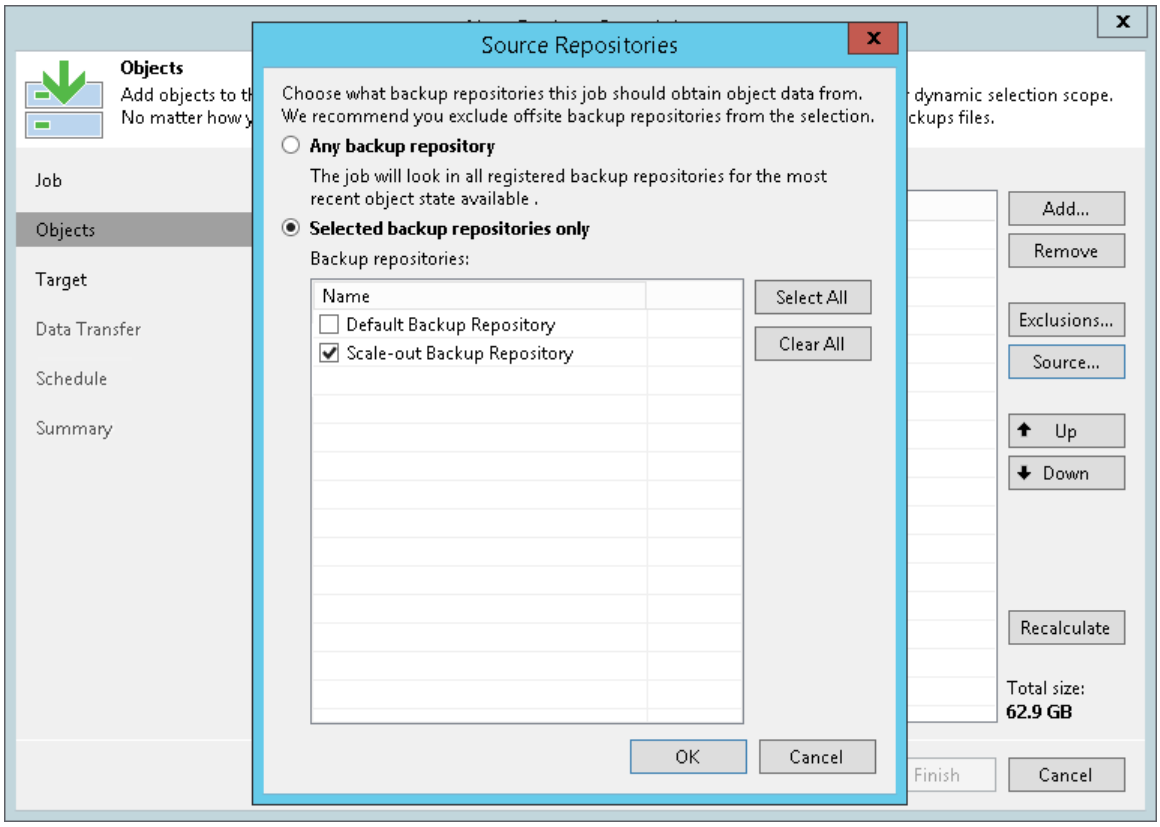
Step 5. Select Source Backup Repositories

This functionality is available if you have selected **Periodic copy** at the Specify Job Name and Description step.

You can select specific backup repositories on which Veeam Backup & Replication must search for restore points:

1. At the **Objects** step of the wizard, click **Source**.
2. Click the **Selected backup repositories only** option and select backup repositories.

> **IMPORTANT**
>
> The **Selected backup repositories only** option does not apply if you have added workloads **From jobs**. In this case, Veeam Backup & Replication copies data from the backup repositories where backups created by source backup jobs reside.



Step 6. Define Processing Order

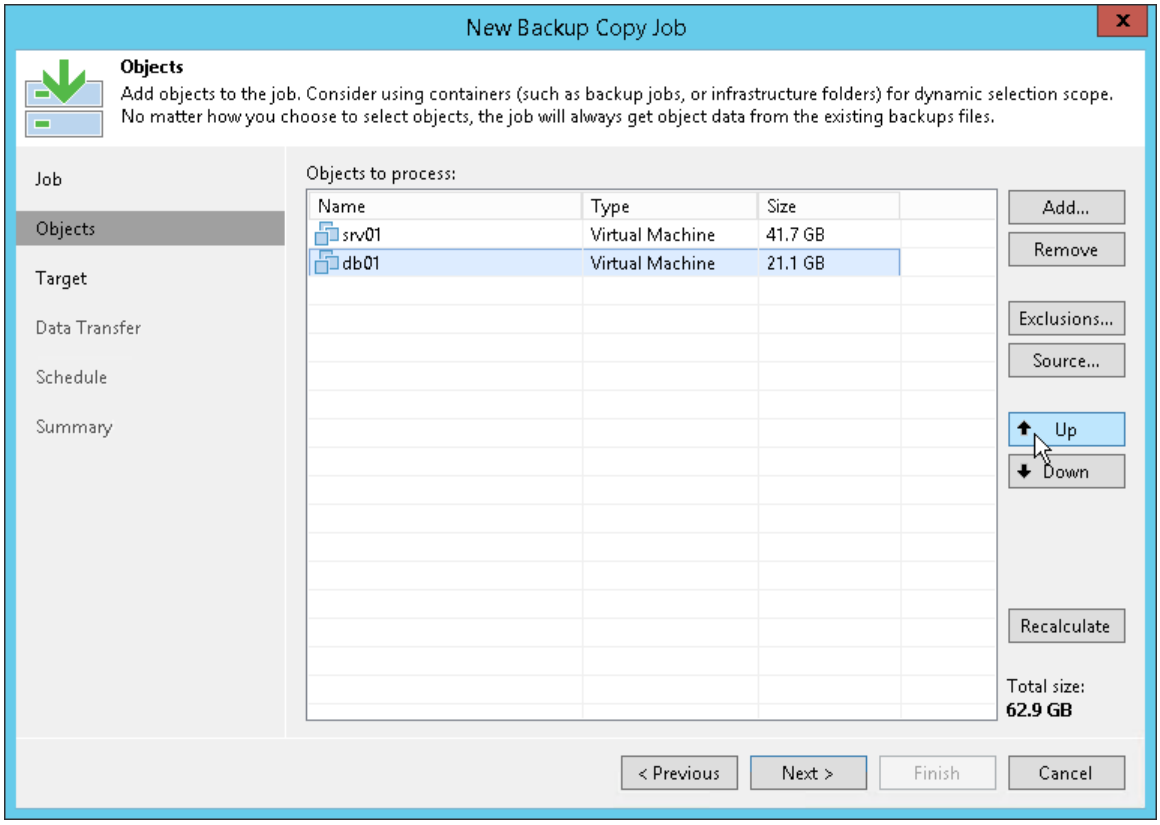
This functionality is available if you have selected **Periodic copy** at the Specify Job Name and Description step.

You can define the order in which the backup copy job must process workloads. Configuring workload order can be helpful, if you want the backup copy job to process mission-critical workloads first. For this, put these workloads higher in the list to ensure that their processing fits the backup window.

Workloads inside a container are processed at random. To ensure that workloads are processed in the defined order, you must add them as standalone workloads, not as part of the container.

To define workload processing order:

1. At the **Objects** step of the wizard, select a workload whose order you want to change.
2. Use the **Up** and **Down** buttons on the right to move the workload up or down in the list.



Step 7. Specify Target Repository and Retention Settings

At the **Target** step of the wizard, define a target backup repository and configure retention policy:

1. From the **Backup repository** list, select a backup repository where copied backups must be stored.
2. In the **Retention Policy** field, configure the short-term retention policy for restore points:
   - If you want to keep the last <N> restore points, select *restore points* from the drop-down list and specify the number of restore points.
   - If you want to keep all restore points created during the last <N> days, select *days* from the drop-down list and specify the number of days.

   When the specified number is exceeded, the earliest restore point will be removed from the backup chain or will be merged with the next closest restore point. For more information on how Veeam Backup & Replication retains the desired number of restore points, see Short-Term Retention Policy.

   > **NOTE**
   >
   > If you enable the GFS retention, the short-term retention policy will not be able to delete and merge the GFS backup files. Thus, the backup copy chain will have more restore points than specified in the short-term retention policy.

3. If you want to create weekly, monthly and yearly full backups, you can configure long-term retention policy (GFS retention policy). GFS full backups will not be deleted or modified until the specified retention period expires. For more information on the GFS retention policy and its limitations, see Long-Term Retention Policy (GFS).

   To configure GFS retention policy, do the following:

   a. Select the **Keep certain full backups longer for archival purposes** check box.

   b. Click **Configure**.

   c. In the **Configure GFS** window, select the necessary GFS backup options. You can configure Veeam Backup & Replication to create weekly, monthly and yearly restore points. For details on settings of the GFS retention, see Backup Copy GFS Cycles.

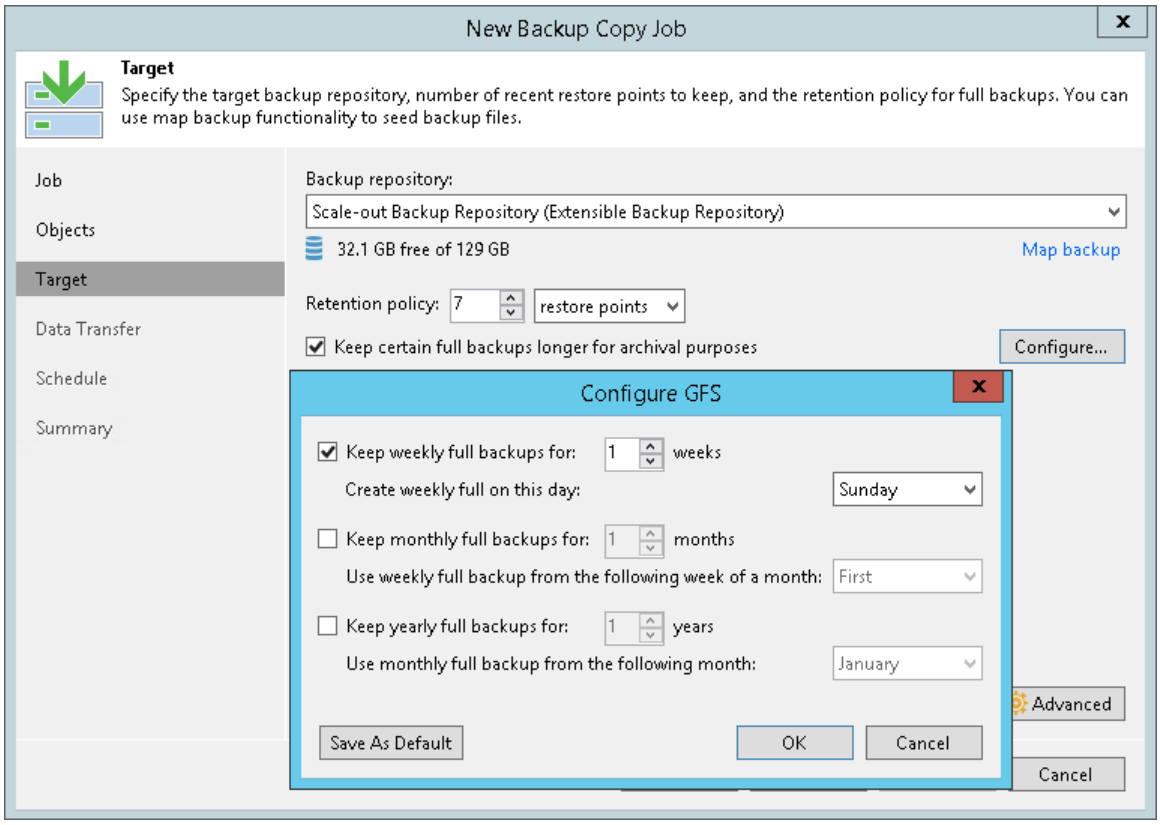
   > **NOTE**
   >
   > Before you implement the GFS retention policy, see Limitations and Considerations.

4. [For the periodic copy mode] You can define a way to create weekly, monthly and yearly full backups:
   - Synthetic Full Method: With this method, during the GFS backup copy creation, Veeam Backup & Replication does not copy data from the source backup repository but synthesizes full backups from backup files that are already stored in the target backup repository. This approach helps to reduce load on the network and production environment.

     The synthetic full method is used by default. To use this method, leave the **Read the entire restore point from source instead of synthesizing it from increments** option unselected.

- Active Full Method: With this method, Veeam Backup & Replication copies data for archive full backups from the source backup repository. This method decreases load on the target repository but increases load on the network and production environment.

  To use this method, select the **Read the entire restore point from source instead of synthesizing it from increments** option.



Step 8. Map Backup File

If the target backup repository already stores a backup of workloads that you want to copy, you can map the backup copy job to this backup. The backup can be created by a backup job or another backup copy job.

The backup copy job will use the backup as a seed. As a result, Veeam Backup & Replication will transfer less data over network. For more information, see Mapping Backup Copy Jobs.

To map the backup copy job to a backup:

1. At the **Target** step of the wizard, click the **Map backup** link.
2. In the **Select Backup** window, select a backup that contains restore points of workloads that you want to copy.

Limitations for Mapping

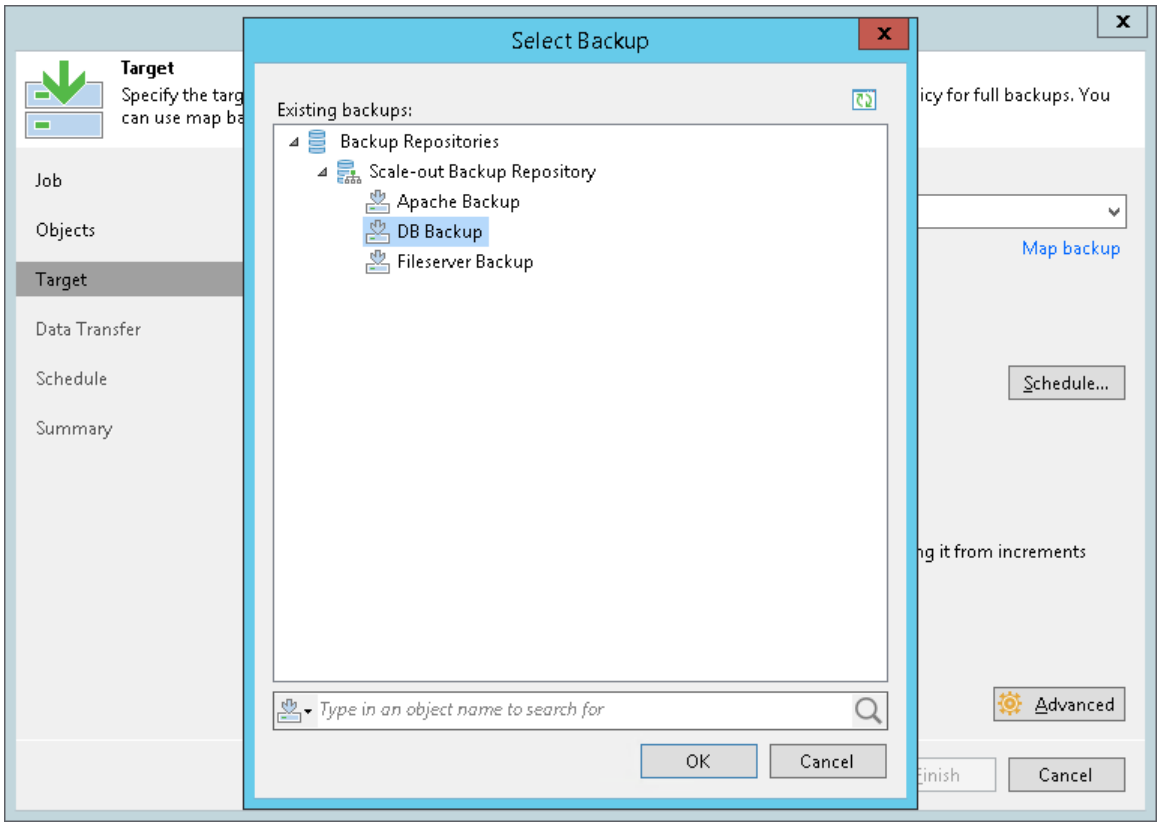
When mapping a backup copy job, mind the following limitations:

- If a backup that you plan to use as a seed is encrypted, you must enable encryption for the backup copy job. The password that you use for the backup copy job can differ from the password used for the initial job.
- The following limitations apply to a backup that can be used as a seed:
  - If you map the backup copy job to a backup created by another backup copy job, the backup must be created with the same copy mode.
  - [For the immediate copy mode] You can map the backup copy job only to a backup created by another backup copy job.
  - [For the periodic copy mode] If you map the backup copy job to a backup created by a backup job, this backup must be created with the incremental backup method only, that is, forever forward or forward incremental.

    > **NOTE**
    >
    > To overcome the listed limitations, you can use a workaround scenario. For more information, see Creating Seed for Backup Copy Job.

  - [For the periodic copy mode] You can map a Veeam Agent backup copy job only to a backup created by the following types of jobs:
    - Backup copy job that processes backups created by Veeam Agent operating in the standalone mode.

- Backup job created in Veeam Agent operating in the standalone mode.



Step 9. Specify Advanced Settings

At the **Target** step of the wizard, you can specify the following settings for the backup copy job:

- Maintenance settings
- Storage settings
- RPO Monitor settings
- Notification settings
- Script settings

> **TIP**
>
> After you specify advanced settings for the backup copy job, you can save them as default settings. For this, click **Save as Default** at the bottom left corner of the **Advanced Settings** window. When you create a new backup copy job, Veeam Backup & Replication will automatically apply the default settings to the new job.

Maintenance Settings

To specify settings for backup files stored in the target backup repository:

1. At the **Target** step of the wizard, click **Advanced**.
2. If you want to periodically perform a health check of the most recent restore point in the backup chain, select the **Perform backup files health check** check box and specify the schedule for the health check.

   By default, the health check is performed on the last Saturday of every month. For more information on the health check, see Health Check.
3. Select the **Remove deleted items data after** check box and specify the retention policy settings for deleted workloads.

   By default, the deleted item retention period is 30 days. It is recommended that you set the retention period to 3 days or more to prevent unwanted data loss. For more information on the retention policy and its limitations, see Deleted Items Retention.
4. To periodically compact a full backup, select the **Defragment and compact full backup file** check box and specify the schedule for the compacting operation.

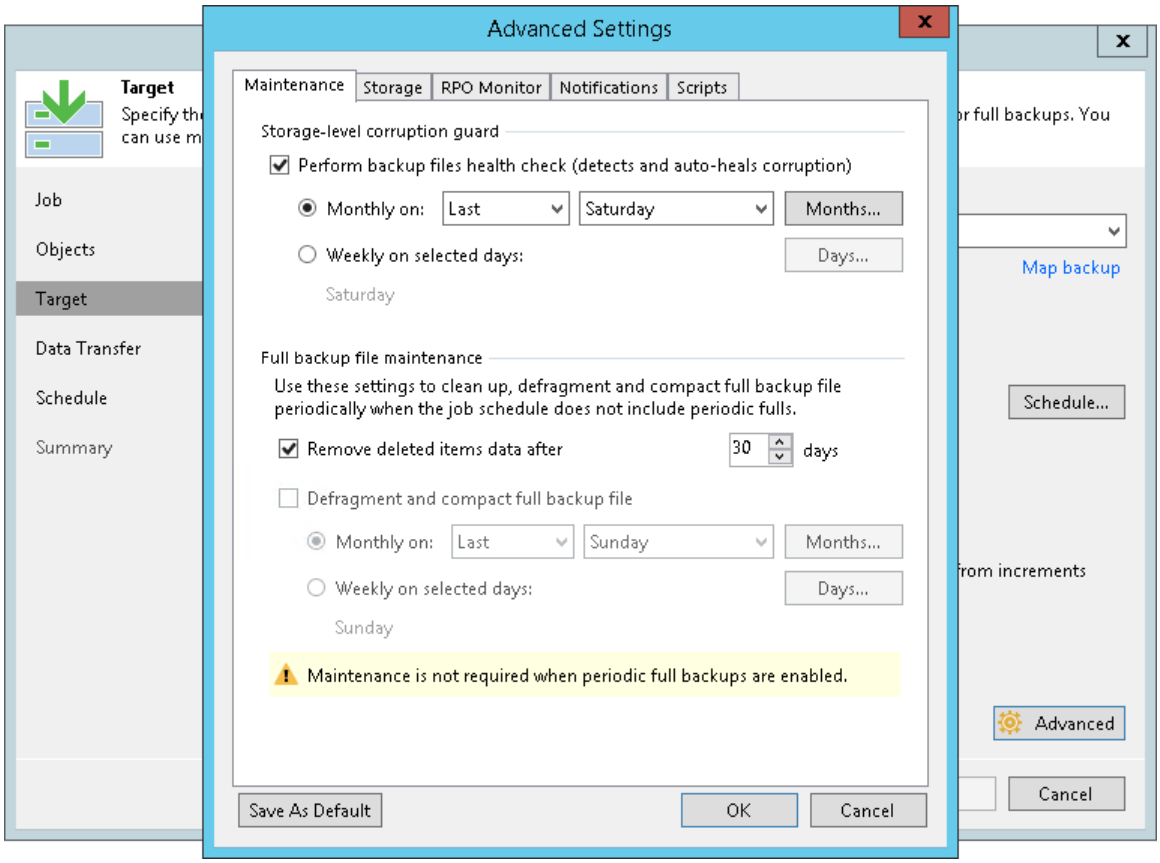
   By default, the compact operation is disabled. For more information on compact of full backup files, see Compact of Full Backup File.

IMPORTANT

When enabling the **Defragment and compact full backup file** option, mind the following:

- The **Defragment and compact full backup file** option can be enabled only if GFS retention policy is disabled.
- The target backup repository must have enough space to store a file of the full backup size. During the compact process, Veeam Backup & Replication creates an auxiliary VBK file that exists in the backup repository until the end of the compact operation.
- If you do not want to copy data for workloads that have only one restore point in the full backup file and this restore point is older than 7 days, check that the following conditions are met: **Remove deleted items data** is disabled; **Use per-machine backup files** is disabled in the settings of the target backup repository.

  Veeam Backup & Replication will extract data for such workloads from the full backup file and write this data to a separate backup file. The file will be displayed under the **Backups > Disk (Imported)** node in the **Home** view.



Storage Settings

To specify compression, deduplication and encryption settings for backup files stored in the target backup repository, do the following:

1. At the **Target** step of the wizard, click **Advanced**.
2. Click the **Storage** tab.

3. In the **Data reduction** section, specify data compression and deduplication settings:
   - By default, Veeam Backup & Replication performs deduplication before storing copied data in the target backup repository. To disable data deduplication, clear the **Enable inline data deduplication** check box.

     For more information on deduplication, see Deduplication.
   - From the **Compression level** list, select a compression level.

     For more information on data compression levels, see Data Compression.
4. In the **Encryption** section, specify encryption settings:
   - To encrypt the backup file created with the backup copy job, select the **Enable backup file encryption** check box.
   - From the **Password** field, select a password that you want to use to encrypt the backup file. If you have not created a password beforehand, click **Add** or use the **Manage passwords** link to specify a new password. For more information, see Managing Passwords for Data Encryption.

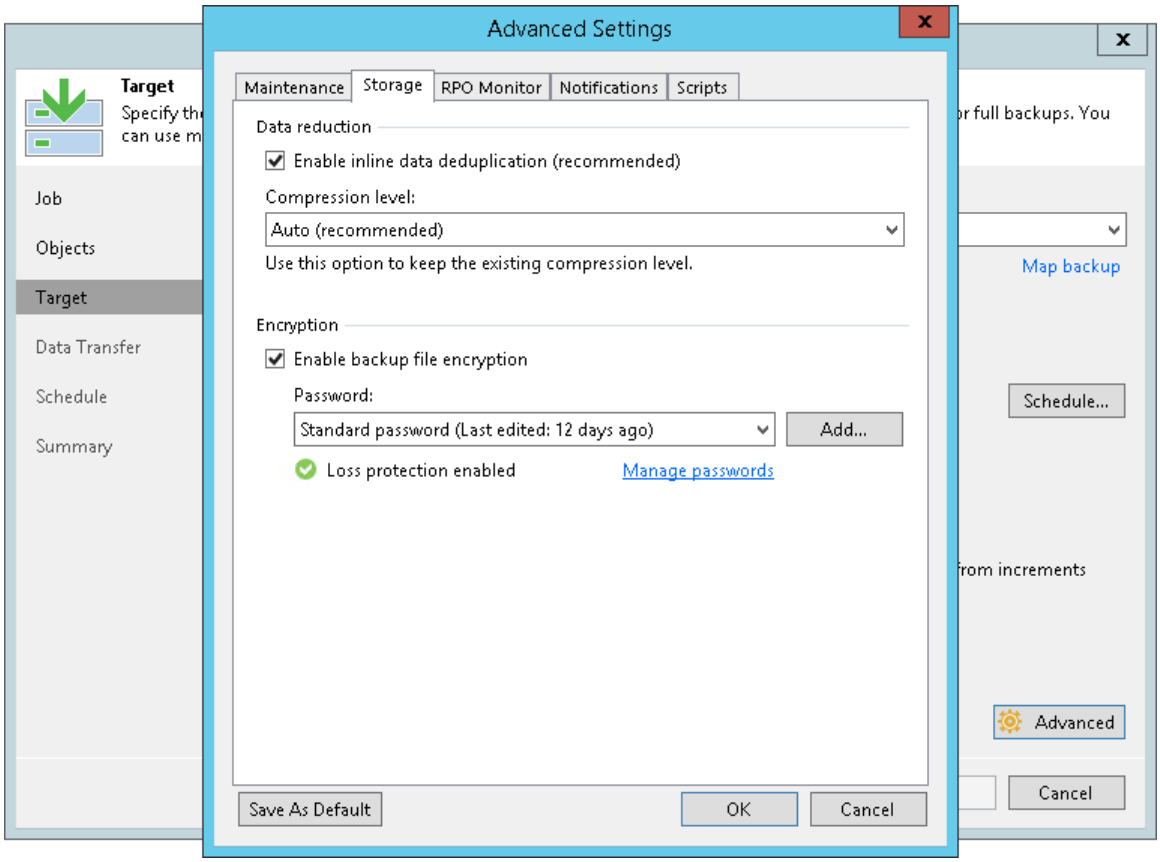
NOTE

When specifying encryption settings, consider the following:

- If you enable encryption for an existing backup copy job, Veeam Backup & Replication applies new settings only starting from the next active full backup (created manually or by the GFS schedule). The active full backup file and subsequent incremental backup files in the backup chain will be encrypted with the specified password.

  Note that if you disable the **Read the entire restore point from source backup instead of synthesizing it from increments** option in the backup copy job, you will have synthetic full backups, not active full backups. For details, see Defining Backup Copy Target.

- Encryption is not retroactive. If you enable encryption for an existing job, Veeam Backup & Replication does not encrypt the previous backup chain created with this job. If you want to start a new chain so that the unencrypted previous chain can be separated from the encrypted new chain, follow the scenario described in this Veeam KB article.



RPO Monitor Settings

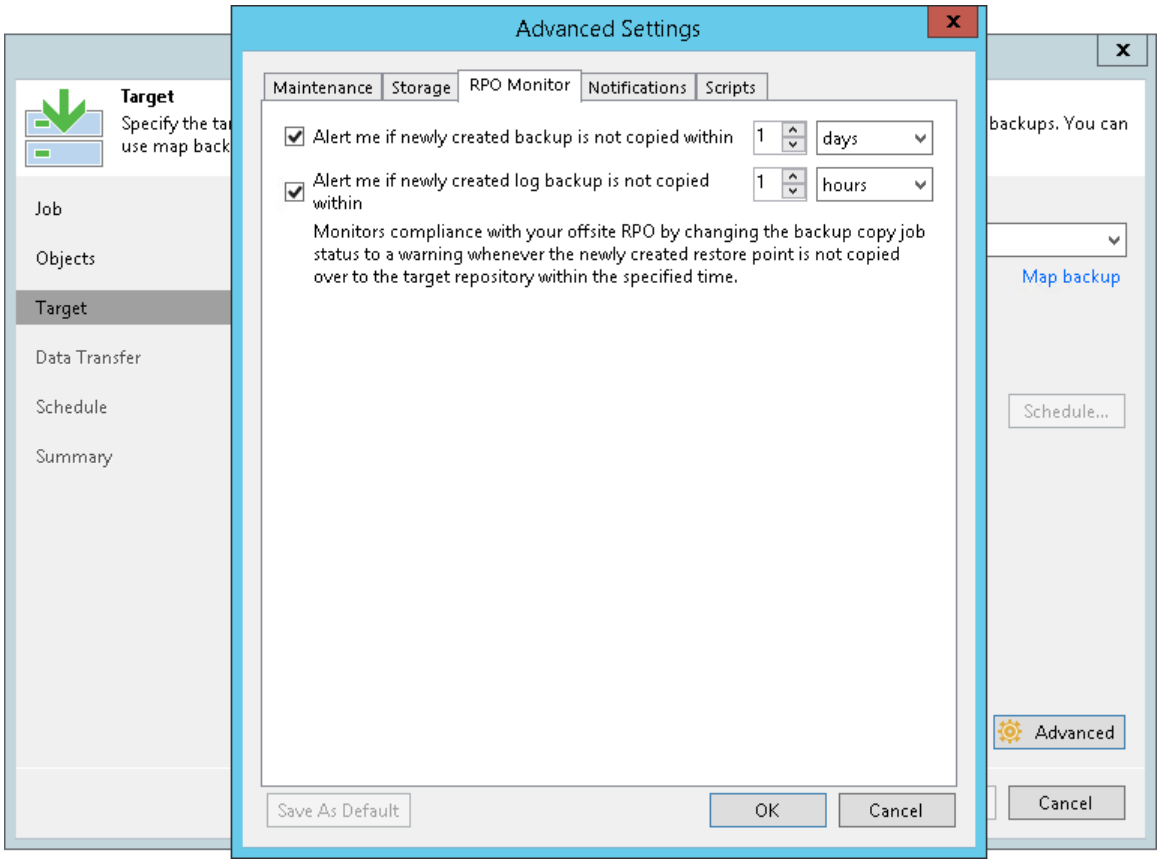
This capability is available if you have selected **Immediate copy** at the Specify Job Name and Description step.

You can instruct a backup copy job to display a warning if a newly created restore point or transaction log is not copied within the desired recovery point objective (RPO). The RPO is counted down from the moment when the source backup job finishes and is ready to be copied.

To mark a job with the *Warning* status when the RPO is exceeded, do the following:

1. At the **Target** step of the wizard, click **Advanced**.
2. Click the **RPO Monitor** tab.

3. Select the **Alert me if newly created backup is not copied within** check box.
4. In the fields on the right, specify the desired RPO in minutes, hours or days.
5. If you have enabled copying of log backups, select the **Alert me if newly created log backup is not copied within** check box.
6. In the fields on the right, specify the desired RPO in minutes, hours or days.



Notification Settings

To specify notification settings for the backup copy job:

1. At the **Target** step of the wizard, click **Advanced**.
2. Click the **Notifications** tab.
3. Select the **Send SNMP notifications for this job** check box if you want to receive SNMP traps when the job completes successfully. In the periodic copy mode, you will receive notifications when the entire backup copy job finishes. In the immediate copy mode — when copying of each source backup job finishes.

   SNMP traps will be sent if you specify global SNMP settings in Veeam Backup & Replication and configure software on the recipient workload to receive SNMP traps. For more information, see Specifying SNMP Settings.

4. Select the **Send email notifications to the following recipients** check box if you want to receive notifications by email in case of job failure or success. In the field below, specify the recipient email address. You can enter several addresses separated by a semicolon.

   Email notifications will be sent if you configure global email notification settings in Veeam Backup & Replication.

   - [For the immediate copy mode] Veeam Backup & Replication sends notifications when copying of each source backup job finishes. For example, if your backup copy job contains two source backup jobs, you will receive two emails.
   - [For the periodic copy mode] Veeam Backup & Replication sends a consolidated email notification once for the specified backup copy interval. Even if the synchronization process is started several times within the interval, for example, due to job retries, only one email notification will be sent.

   For more information on how to configure global notification settings, see Configuring Global Email Notification Settings.

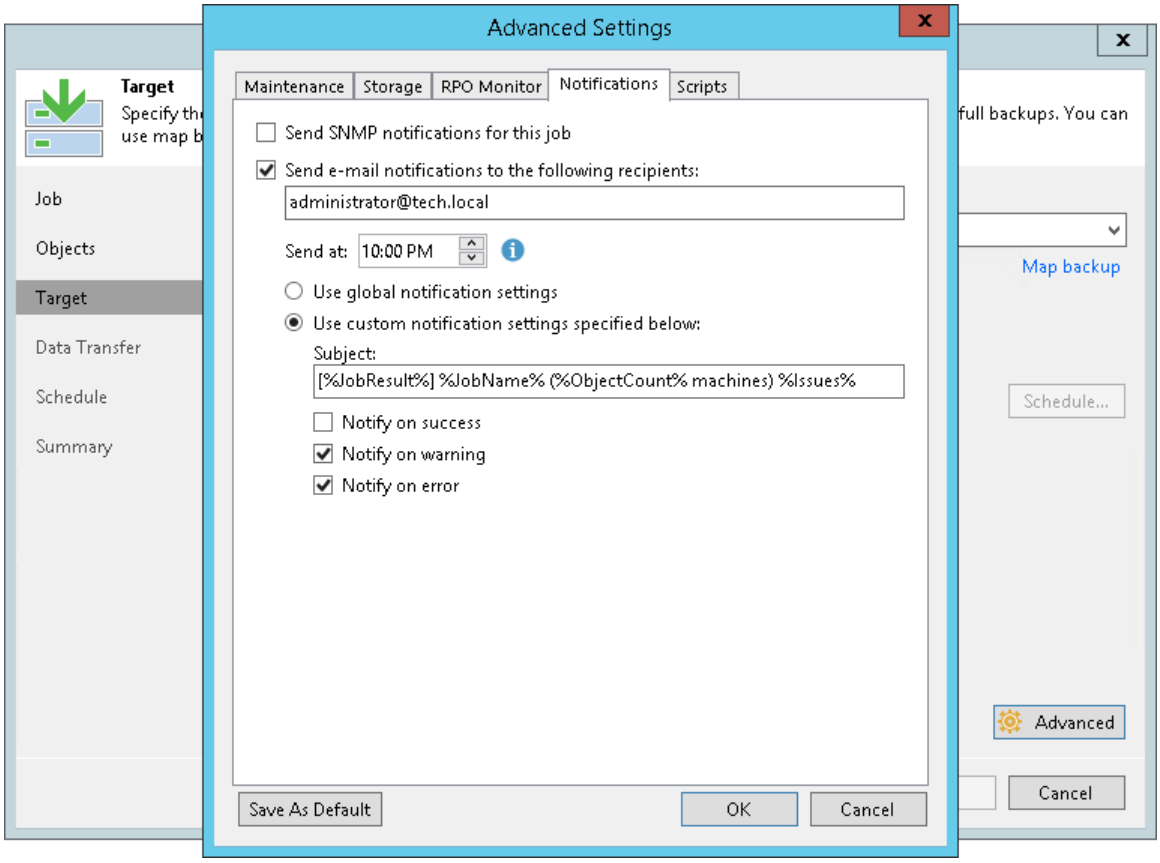
5. [For the immediate copy mode] In the **Send at** field, specify when you want to send notifications about backup copy jobs that process log backups. Veeam Backup & Replication sends a consolidated report once a day at the specified time.

6. You can choose to use global notification settings or specify custom notification settings.

   - To receive a typical notification for the job, select **Use global notification settings**. In this case, Veeam Backup & Replication will apply to the job global email notification settings specified for the backup server. For more information, see Configuring Global Email Notification Settings.
   - To configure a custom notification for a job, select **Use custom notification settings specified below**. You can specify the following notification settings:

     i. In the **Subject** field, specify a notification subject. You can use the following variables in the subject: *%Time%* (completion time), *%JobName%* (the backup copy job name and source backup job name in the *CopyJobName|SourceJobName* format), *%JobResult%*, *%ObjectCount%* (number of workloads in the job) and *%Issues%* (number of workloads in the job that have been processed with the *Warning* or *Failed* status).

ii. Select the **Notify on success**, **Notify on warning** and/or **Notify on error** check boxes to receive email notification if data processing completes successfully, completes with a warning or fails.



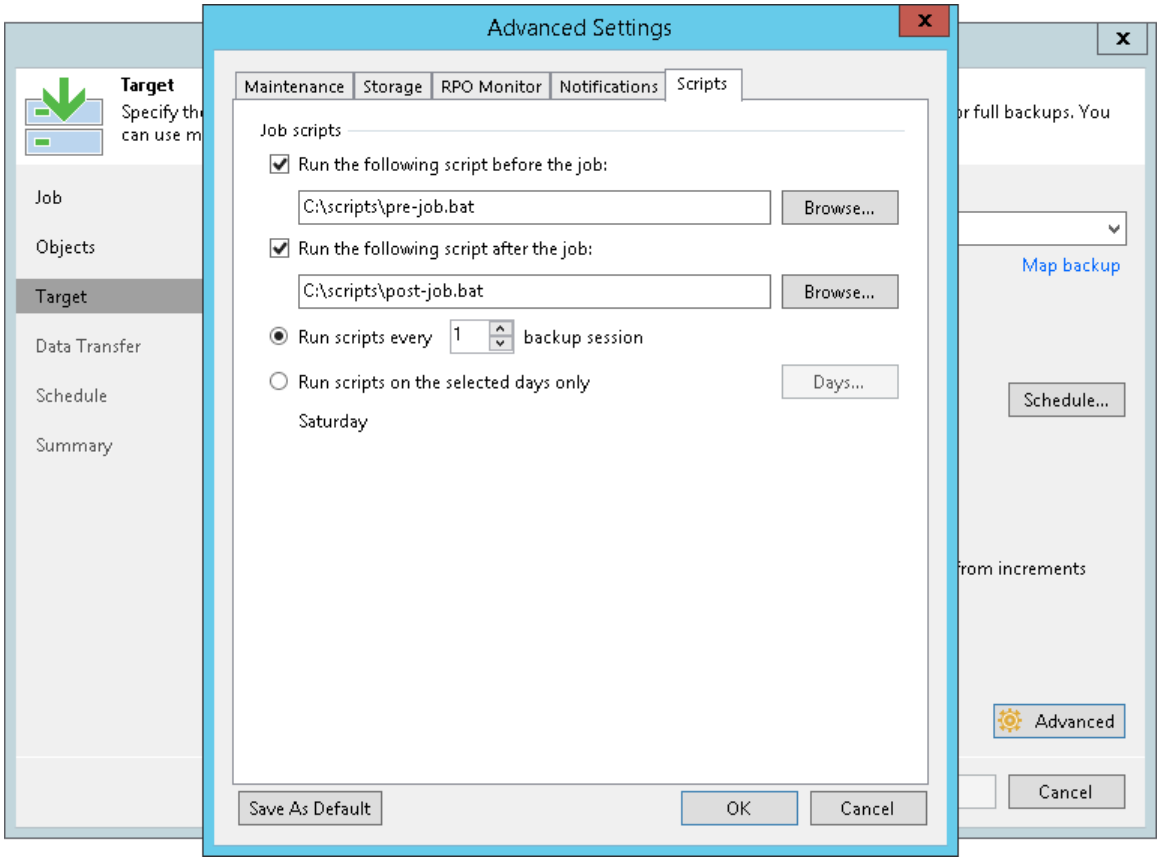
Scripts Settings

To specify script settings for the backup copy job:

1. At the **Target** step of the wizard, click **Advanced**.
2. Click the **Scripts** tab.
3. Select the **Run the following script before the job** and **Run the following script after the job** check boxes to execute custom scripts before and/or after the backup copy job. Note that in the immediate copy mode, scripts are executed for every source backup job.

   Then click **Browse** and select executable files from a local folder on the backup server. The scripts will be executed on the backup server after the transformation processes are completed on the target repository
4. You can change how often the scripts must be executed:
   - To run the scripts after a specific number of backup copy sessions, select **Run scripts every... backup session** option and specify the number of sessions.
   - To run the scripts on specific days, select the **Run scripts on selected days only** option and click the **Days** button to specify week days.

> **NOTE**
>
> If you select the **Run scripts on the selected days only** option, Veeam Backup & Replication executes scripts only once on each selected day — when the job runs for the first time. During subsequent job runs, scripts are not executed.



Step 10. Specify Data Path Settings

The **Data Transfer** step of the wizard is available only if you copy backups of virtual or physical machines created with Veeam Backup & Replication or Veeam Agents.

At this step of the wizard, you can select how Veeam Backup & Replication will transport backed up data — directly or through WAN accelerators. By default, during the backup copy job Veeam Backup & Replication transports data directly from the source backup repository to target backup repository. This type of transport is recommended if you plan to copy backup files over high-speed connections.

If you plan to copy backup files over WAN or slow connections, it is recommended that you configure source and target WAN accelerators in the backup infrastructure and copy backups through these WAN accelerators. For more information, see WAN Acceleration.

To use WAN acceleration for the backup copy job:

1. At the **Data Transfer** step of the wizard, select the **Through built-in WAN accelerators** option.
2. From the **Source WAN accelerator** list, select a WAN accelerator configured in the source site.
3. From the **Target WAN accelerator** list, select a WAN accelerator configured in the target site.

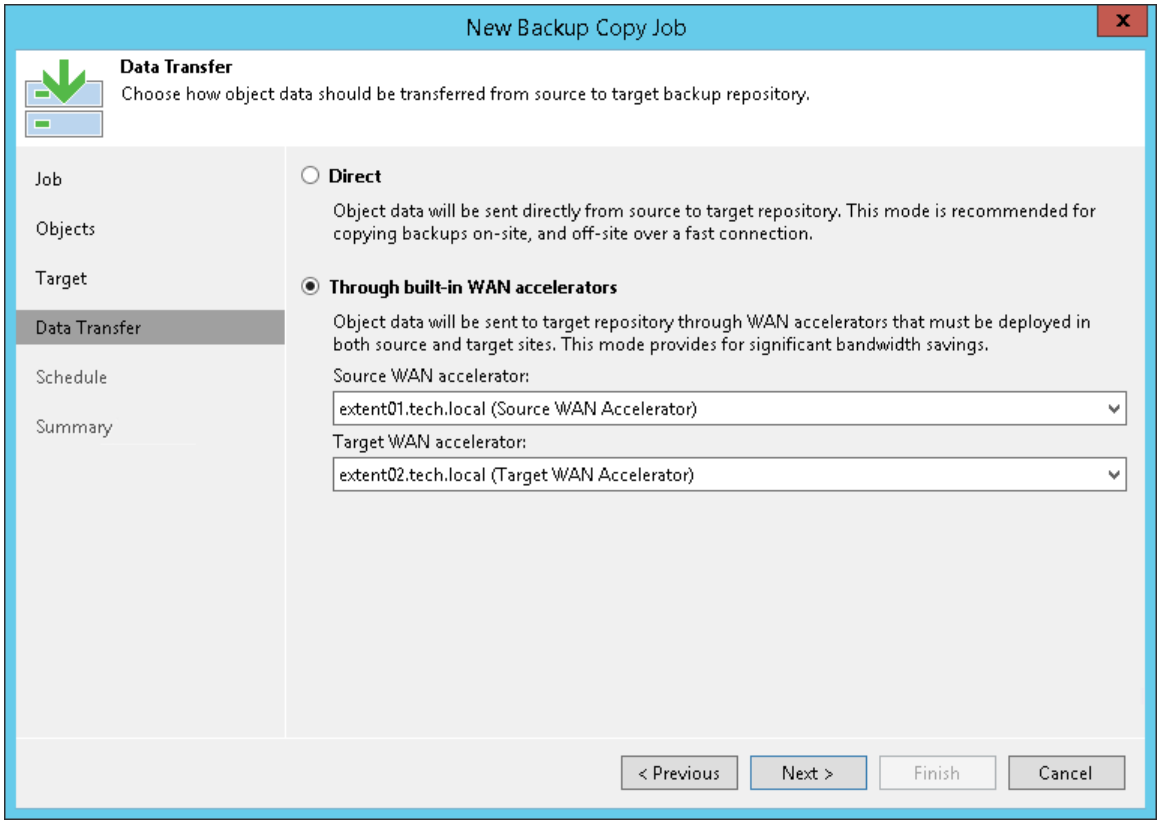
Requirements and Limitations for WAN Accelerators

- You must not assign one source WAN accelerator to several backup copy jobs that you plan to run simultaneously.

  The source WAN accelerator requires a lot of CPU and RAM resources and does not process multiple backup copy tasks in parallel. As an alternative, you can create one backup copy job for all workloads you plan to process over one source WAN accelerator. The target WAN accelerator, however, can be assigned to several backup copy jobs.

- [For WAN accelerators with the high bandwidth mode disabled] It is recommended that you pre-populate the global cache on the target WAN accelerator before you start the backup copy job. Global cache population helps reduce the amount of traffic transferred over WAN. For more information, see Populating Global Cache.

- You cannot use WAN accelerators for backup copy jobs that copy backups of Amazon EC2 instances.

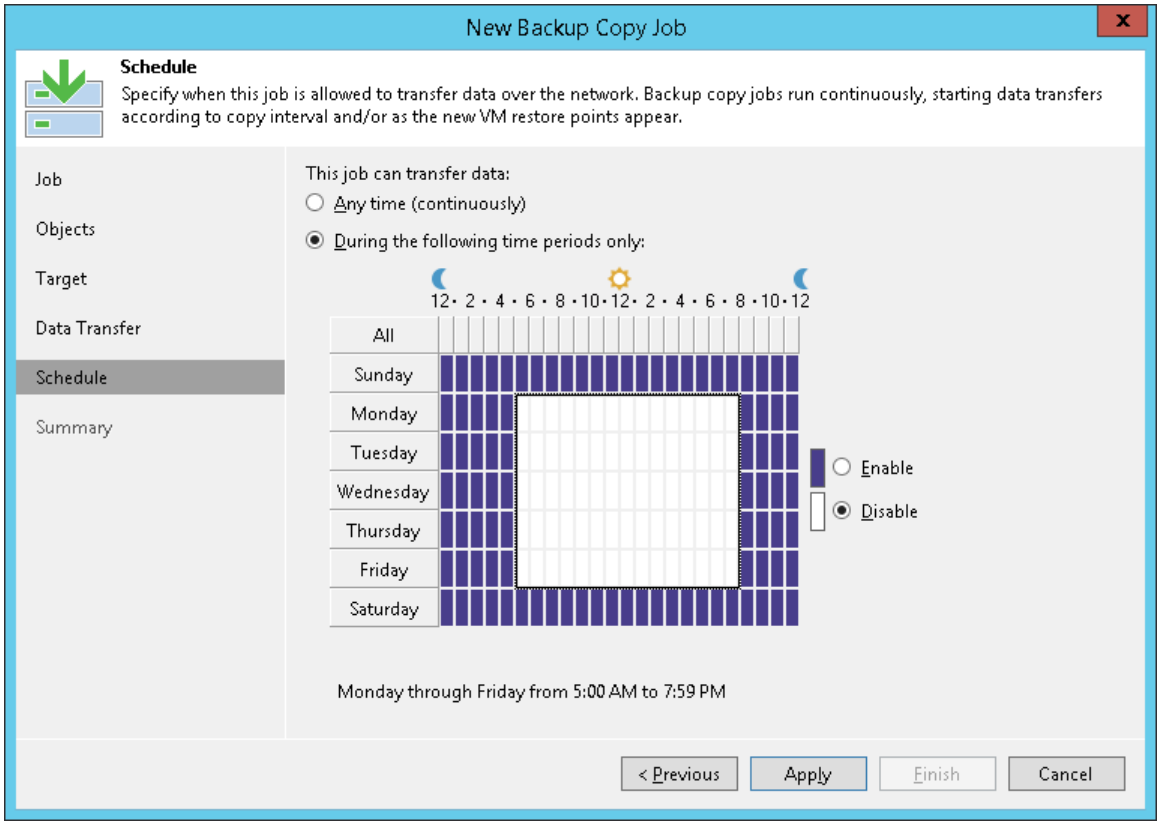


Step 11. Define Backup Copy Window

At the **Schedule** step of the wizard, you can define a time span in which the backup copy job must not transport data between source and target backup repositories. For more information, see Backup Copy Window.

To define a backup window for the backup copy job:

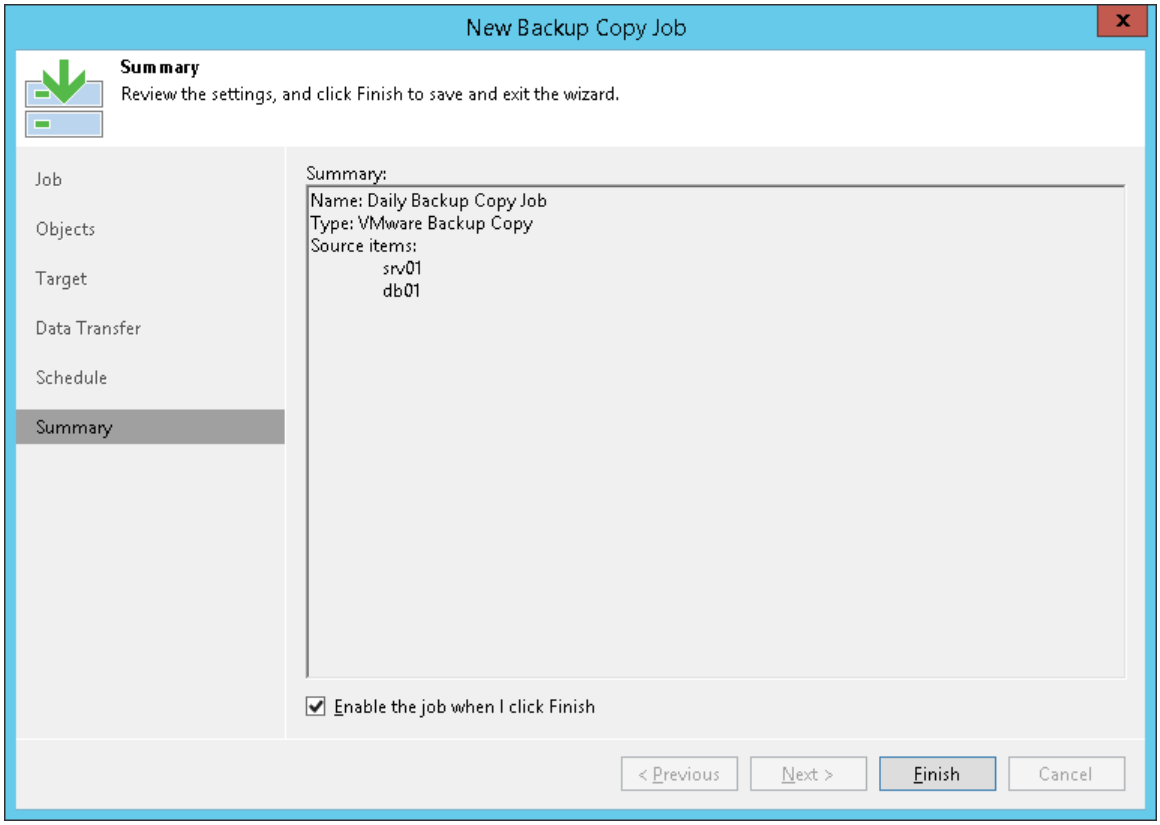
1. Select the **During the following time periods only** option.
2. In the schedule box, select the desired time area.
3. Use the **Enable** and **Disable** options to mark the selected time segments as allowed or prohibited for the backup copy job.



Step 12. Finish Working with Wizard

At the **Summary** step of the wizard, complete the procedure of backup copy job configuration:

1. Review details of the backup copy job.
2. Select the **Enable the job when I click Finish** check box if you want to start the job right after you finish working with the wizard.
3. Click **Finish** to close the wizard.



Creating Backup Copy Jobs for HPE StoreOnce Repositories

To copy backup files between HPE StoreOnce backup repositories, you must configure a backup copy job.

Unlike other backup copy jobs, the backup copy job for HPE StoreOnce mirrors data from the source repository. This backup copy job copies backup files as they are stored in the source repository, without any transformation. To copy the backup files, Veeam Backup & Replication uses the HPE StoreOnce Catalyst Copy technology.

The backup copy job copies only backup files created by backup jobs and other backup copy jobs. The backup files must be of the following types:

- Backup files of VMware vSphere and Microsoft Hyper-V VMs created by Veeam Backup & Replication. Log backup files are not copied.
- Physical machine backup files created by Veeam Agent backup jobs managed by the backup server.
- Backups copied to an HPE StoreOnce repository by other backup copy jobs (with HPE Catalyst Copy disabled).

IMPORTANT

Mind the following limitations:

- The backup jobs must be configured on the same backup server where you configure the backup copy job with HPE Catalyst Copy enabled. Backups created by jobs configured on other backup servers are not copied.
- The backup copy jobs with HPE Catalyst Copy enabled does not copy imported backups.
- To copy backups created by other backup copy jobs (without HPE Catalyst Copy), you must enable the GFS retention for these backup copy jobs. For more information on how to enable the GFS retention, see Specify Target Repository and Retention Settings.

When the backup copy job for HPE StoreOnce runs for the first time, it copies all existing backup files. Then the backup copy job starts each time a new backup file appears in the source repository. In case of a removed backup file, the backup copy job waits 21 days since the backup file creation and after removes the backup file from the target repository. If 21 days have already passed at the moment of removal, the backup copy job removes the backup file immediately. You can change this day limit in the backup copy job settings. For more information, see Maintenance Settings.

Before creating a job, check prerequisites and limitations. Then use the **New Backup Copy Job** wizard to configure the backup copy job.

Before You Begin

Before you create a backup copy job for HPE StoreOnce backup repository, check the following requirements:

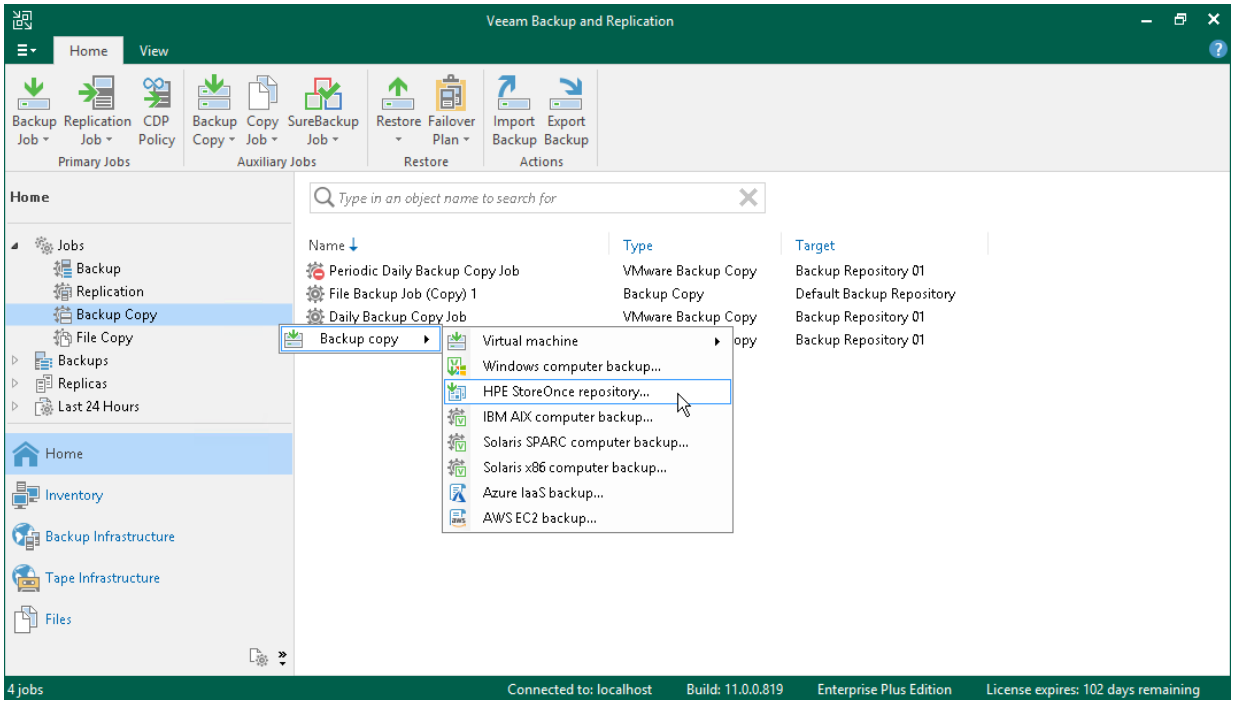
- The minimum supported software versions of HPE StoreOnce are the following:
  - For the third generation, the minimum version is 3.18.7.
  - For the fourth generation, the minimum version is 4.2.0.
- Make sure that all backup infrastructure components that take part in the backup copy process are added to the backup infrastructure. These components include the source and target repositories between which data is copied. For more information on how to add a backup repository, see Adding Backup Repositories.
- Check that repositories between which you plan to copy data have a direct connection to each other. This is required because Veeam Backup & Replication uses the HPE StoreOnce Catalyst Copy technology to copy backup files.

  This direct connection must be of the same type as the connection that you select when adding the target HPE StoreOnce. For example, if you connected the target HPE StoreOnce repository over Fibre Channel, you must connect the source HPE StoreOnce to the target HPE StoreOnce over Fibre Channel.
- HPE StoreOnce repositories connected over Fibre Channel require the two-way connection. Zone the source initiator World Wide Names (WWNs) with the destination target WWNs, and zone the destination initiator WWNs with the source target WWNs.
- If you plan to use pre-job and/or post-job scripts, you must create scripts before you configure the backup copy job.

Step 1. Launch New Backup Copy Job Wizard

To run the **New Backup Copy Job** wizard, do one of the following:

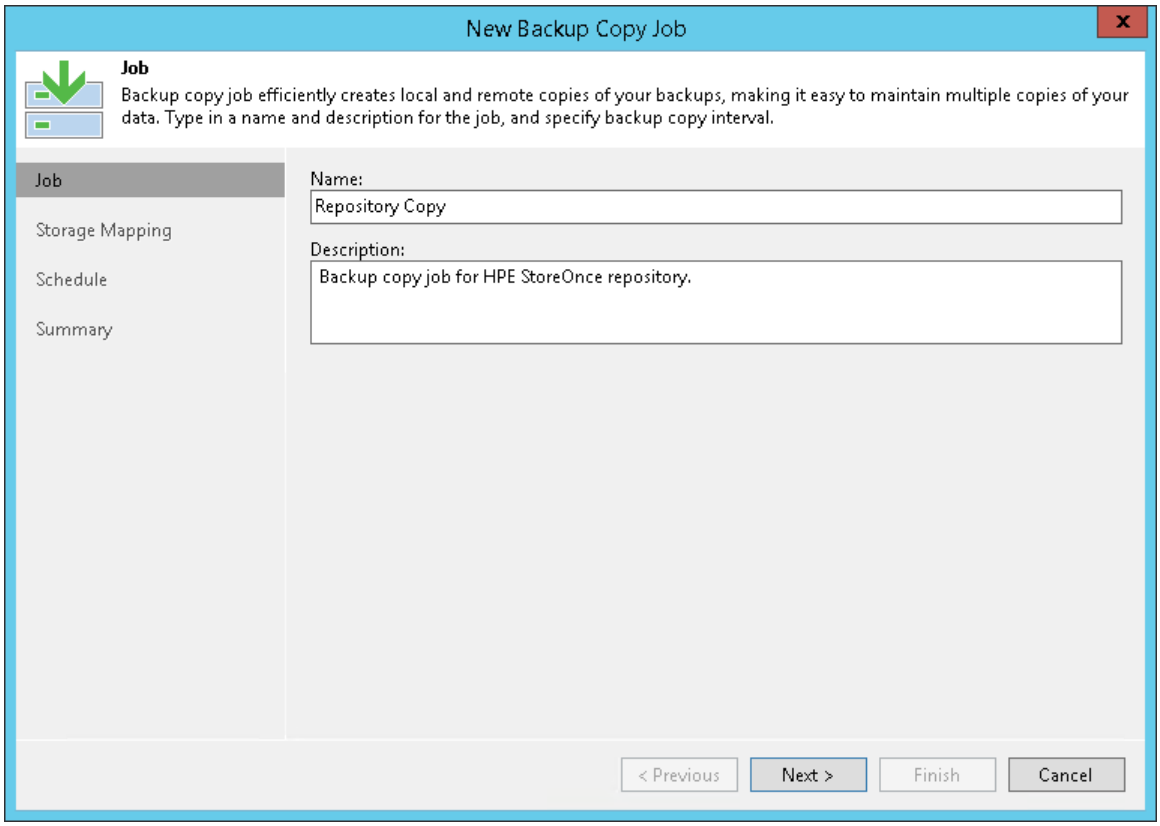
- On the **Home** tab, click **Backup Copy** and select **HPE StoreOnce repository**.
- Open the **Home** view. In the inventory pane, right-click **Jobs** and select **Backup copy** > **HPE StoreOnce repository**.



Step 2. Specify Job Name and Description

At the **Job** step of the wizard, specify name and description for the backup copy job:

1. In the **Name** field, enter a name for the job.
2. In the **Description** field, enter a description for the job. The default description contains information about the user who created the job, date and time when the job was created.



Step 3. Select Source and Target Repositories

At the **Storage Mapping** step of the wizard, select a source repository from which you want to copy backups and a target repository where you want to store the copies.

1. Click **Add** to open the **Add Repository** window.
2. From the **Source backup repository** list, select a backup repository from which you want to copy backup files. The unsupported backup repositories are not shown in the list.
3. From the **Target backup repository** list, select a backup repository where you want to store the copies. The unsupported backup repositories are not shown in the list.

> **IMPORTANT**
>
> Veeam Backup & Replication does not copy all types of backups. For the list of supported backup types, see Creating Backup Copy Jobs for HPE StoreOnce Repositories.

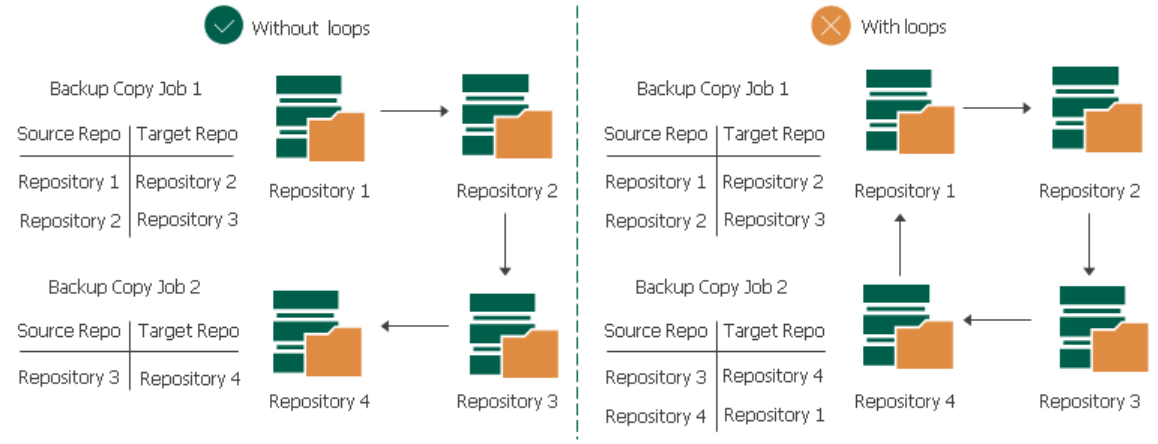
Requirements and Limitations for Source and Target Repositories

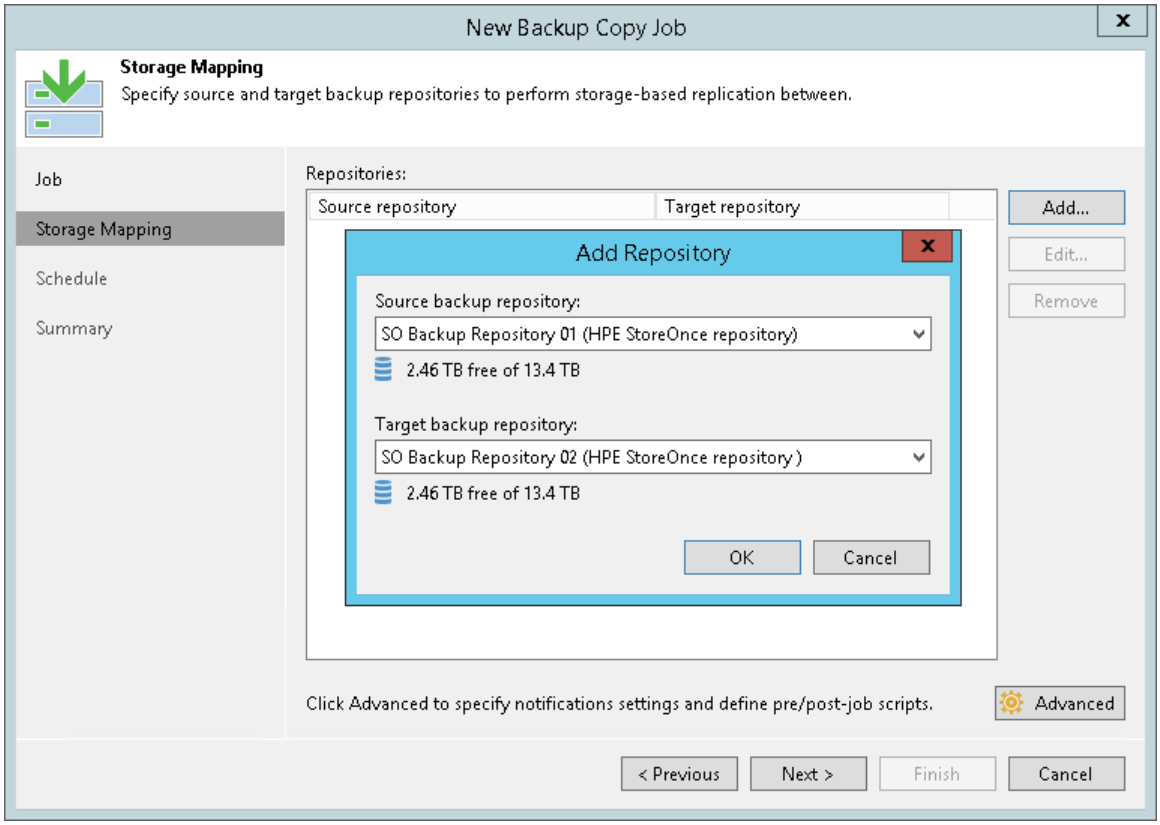
When you create backup copy jobs, check the following prerequisites and limitations for the source and target repositories:

- The source and target backup repositories must be HPE StoreOnce backup repositories or scale-out backup repositories that consist of HPE StoreOnce repositories only.
- [For scale-out backup repositories] The file placement policy must be Data locality. For more information, see Data locality.
- Within one backup copy job, you can use each repository as a source only once.
- You must not create the same pairs of source and target repositories, even in different backup copy jobs.

Requirements for Data Flow

When you create backup copy jobs, check that you do not create loops in data flow across all backup copy jobs. This means that data copied from one repository must not be copied to it again. The following image shows backup copy jobs configured correctly (without loops) and incorrectly (with loops).





Step 4. Specify Advanced Settings

At the **Storage Mapping** step of the wizard, you can specify the following settings for the backup copy job:

- Maintenance Settings
- Notification Settings
- Script Settings

> **TIP**
>
> After you specify necessary settings for the backup copy job, you can save them as default settings. To do this, click **Save as Default** at the bottom left corner of the **Advanced Settings** window. When you create a new backup copy job, Veeam Backup & Replication will automatically apply the default settings to the new job.

Maintenance Settings

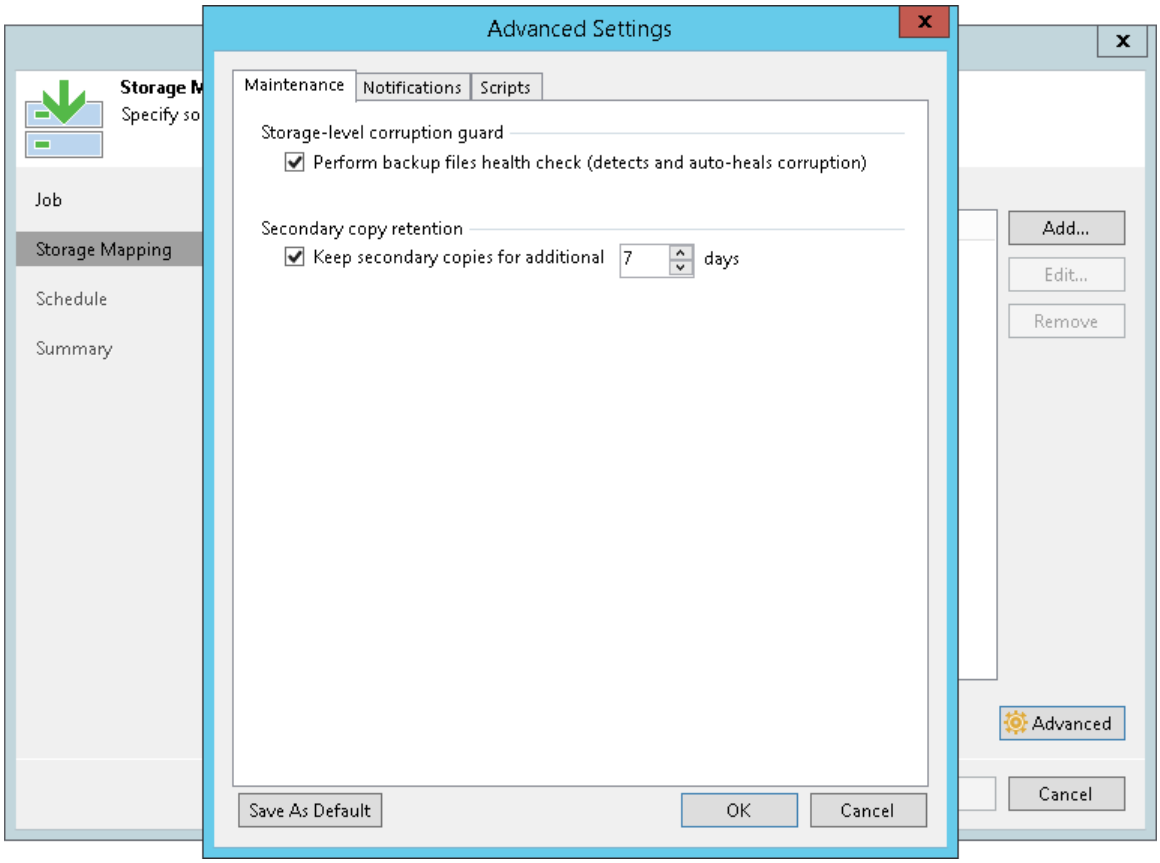
In the maintenance settings, you can configure whether to perform a health check and after which period delete from the target HPE StoreOnce repository files deleted from the source repository. Note that the health check may lower the performance of the target repository. For details, see Health Check for Backup Files.

1. At the **Storage Mapping** step of the wizard, click **Advanced**.
2. To disable the health check, clear the **Perform backup files health check** check box. By default the health check is enabled.
3. In the **Keep secondary copies for additional** field, specify after which period delete files from the target repository after they were deleted from the source repository.

   The backup copy job waits the specified number of days since the backup file creation and after deletes the backup file from the target repository. If the specified number of days has already passed at the moment of deletion, the backup copy job deletes the backup file immediately.

IMPORTANT

Veeam Backup & Replication does not perform the health check for encrypted and compressed backup files.



Notification Settings

To specify notification settings for the backup copy job:

1. At the **Storage Mapping** step of the wizard, click **Advanced**.
2. Click the **Notifications** tab.
3. Select the **Send SNMP notifications for this job** check box if you want to receive SNMP traps when data from each source repository is copied.

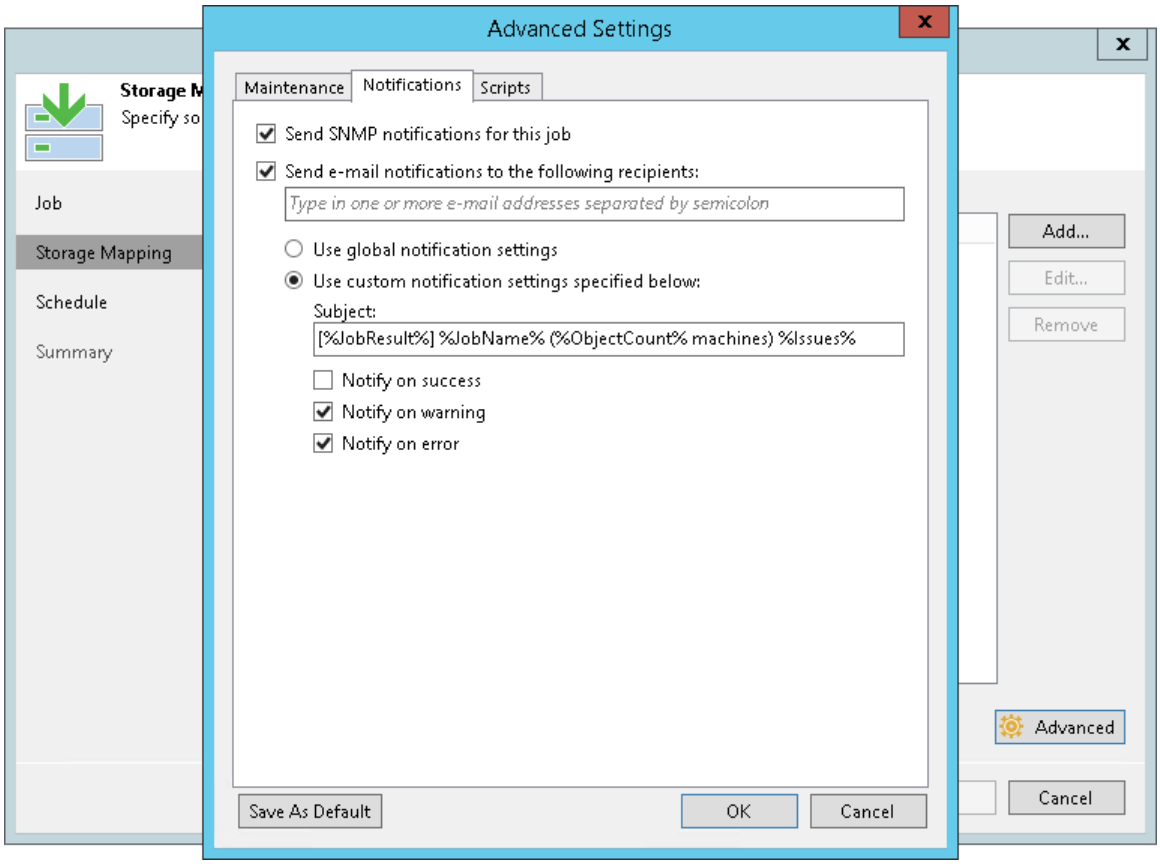
   SNMP traps will be sent if you specify global SNMP settings in Veeam Backup & Replication and configure software on the recipient machine to receive SNMP traps. For more information, see Specifying SNMP Settings.
4. Select the **Send email notifications to the following recipients** check box if you want to receive email notifications when copying data of each source repository finishes with *Success*, *Warning* or *Failed* status.

   Email notifications will be sent if you configure global email notification settings in Veeam Backup & Replication. For more information, see Configuring Global Email Notification Settings.
5. You can choose whether to use global notification settings or specify custom notification settings.
   - To receive typical notifications for the job, select **Use global notification settings**. In this case, Veeam Backup & Replication will apply to the job global email notification settings specified for the backup server.

- To configure custom notifications, select **Use custom notification settings specified below**. You can specify the following notification settings:

  i. In the **Subject** field, specify a notification subject. You can use the following variables in the subject: *%Time%* (completion time), *%JobName%*, *%JobResult%*, *%VmCount%* (number of machines in the job) and *%Issues%* (number of machines in the job that have been processed with the *Warning* or *Failed* status).

  ii. Select the **Notify on success**, **Notify on warning** and/or **Notify on error** check boxes to receive email notifications if data processing completes successfully, fails or completes with a warning.



Script Settings

To specify script settings for the backup copy job:

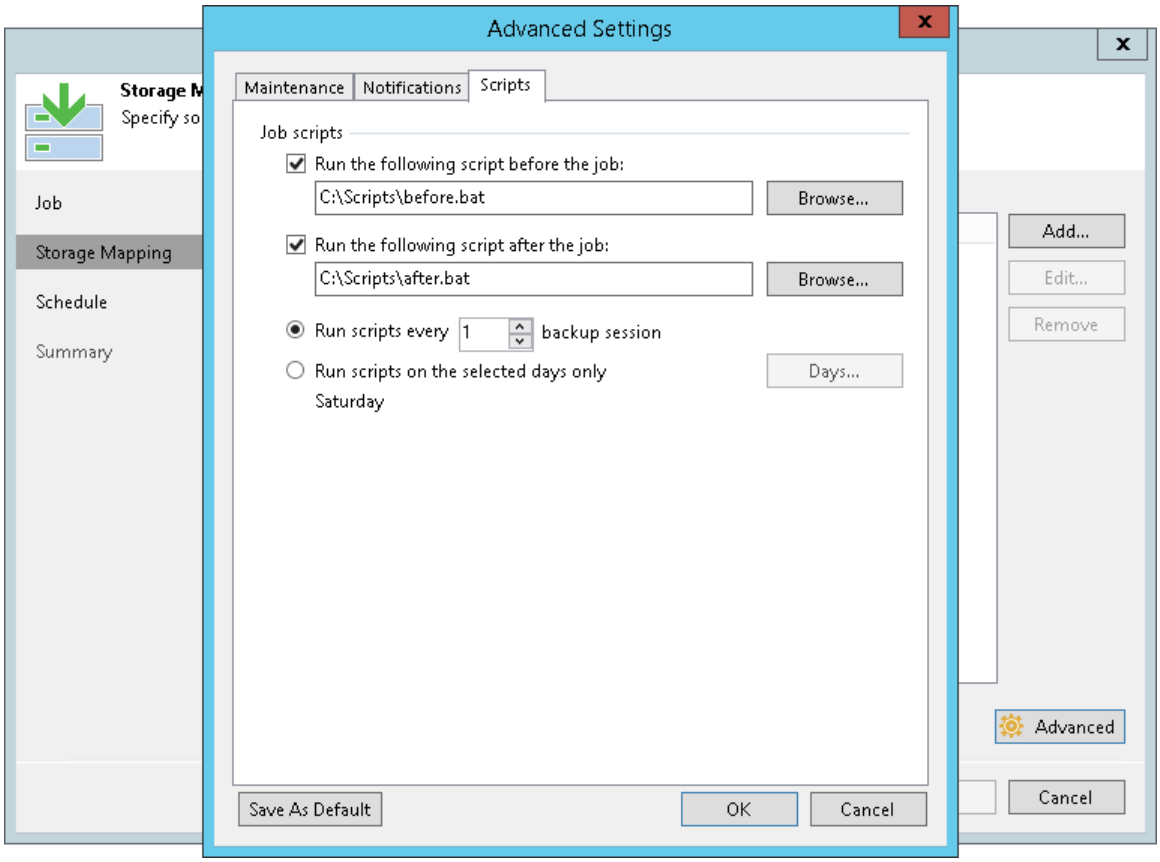
1. At the **Storage Mapping** step of the wizard, click **Advanced**.
2. Click the **Scripts** tab.
3. Select the **Run the following script before the job** and **Run the following script after the job** check boxes to execute custom scripts before and/or after copying data of each source repository finishes.

   Then click **Browse** and select executable files from a local folder on the backup server. The scripts are executed on the backup server after the transformation processes are completed on the target repository.
4. You can change how often the scripts must be executed:

   - To run scripts after a specific number of backup copy sessions, select **Run scripts every... backup session** option and specify the number of sessions.

- To run scripts on specific days, select the **Run scripts on selected days only** option and click the **Days** button to specify week days.

> **NOTE**
>
> If you select the **Run scripts on the selected days only** option, Veeam Backup & Replication executes scripts only once on each selected day — when the job runs for the first time. During subsequent job runs, scripts are not executed.

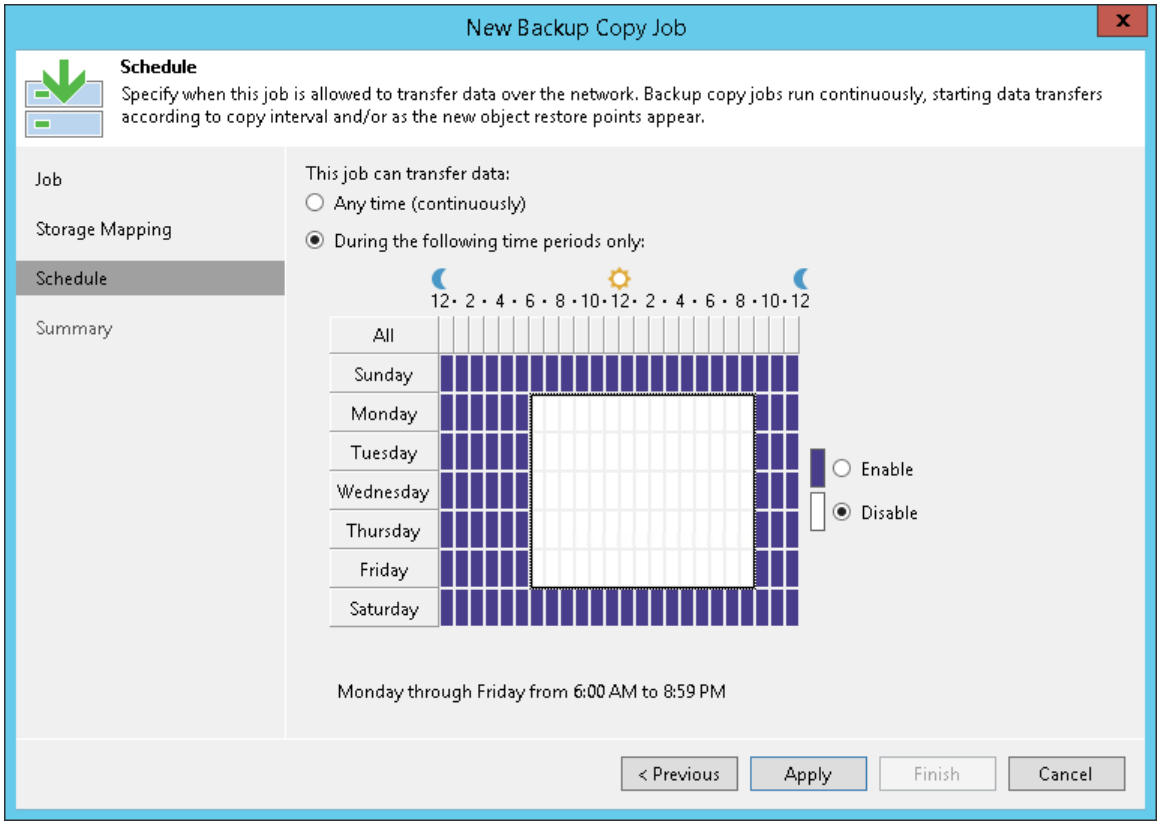


Step 5. Define Backup Copy Window

At the **Schedule** step of the wizard, define the time span in which the backup copy job must not transport data between source and target backup repositories. For more information, see Backup Copy Window.

To define a 'prohibited' period for the backup copy job:

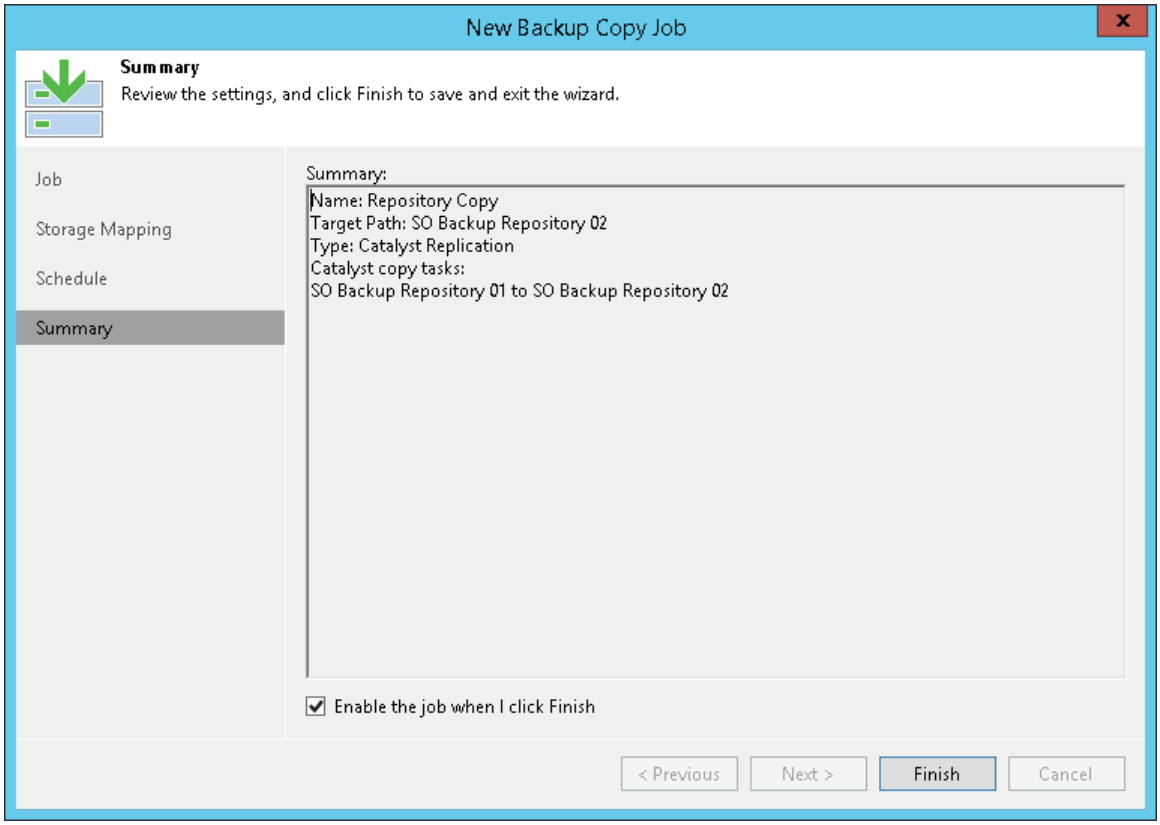
1. Select the **During the following time periods only** option.
2. In the schedule box, select the desired time area.
3. Use the **Enable** and **Disable** options to mark the selected time segments as allowed or prohibited for the backup copy job.



Step 6. Finish Working with Wizard

At the **Summary** step of the wizard, complete the procedure of backup copy job configuration:

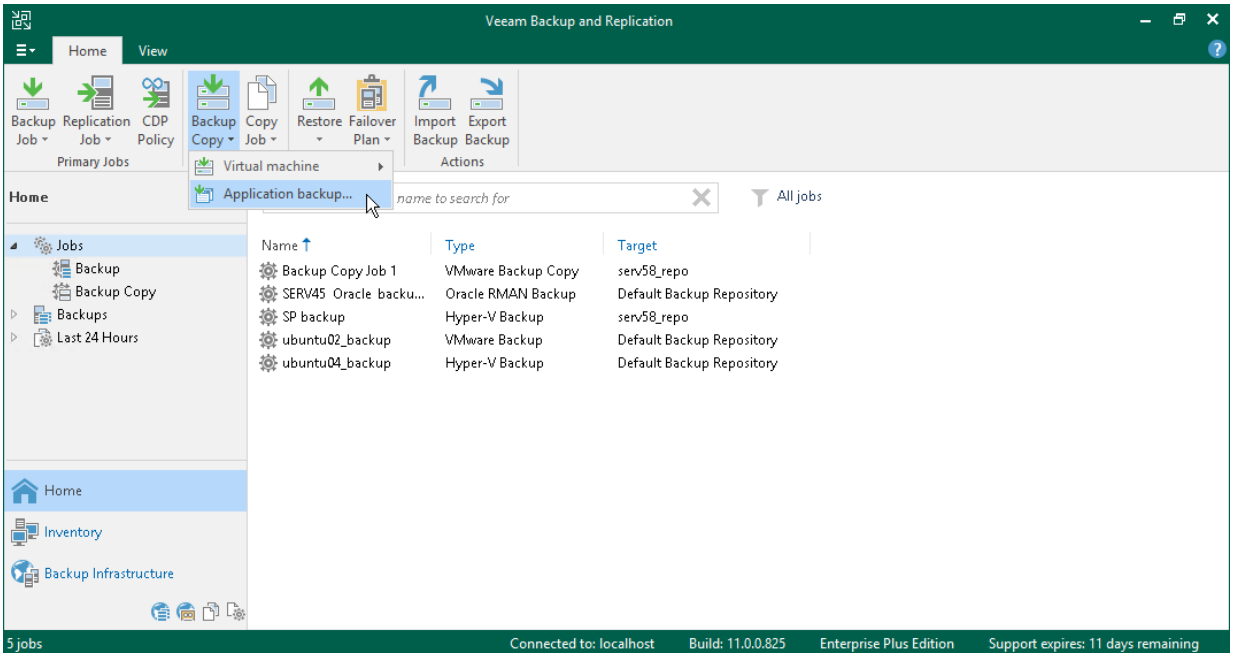
1. Review details of the backup copy job.
2. Select the **Enable the job when I click Finish** check box if you want to start the job right after you finish working with the wizard.
3. Click **Finish** to close the wizard.



Creating Backup Copy Jobs for Oracle and SAP HANA Databases

To create copies of Veeam Plug-in backups of Oracle and SAP HANA databases, you must configure a backup copy job. For more details, see the following sections of the Veeam Plug-ins for Enterprise Applications Guide:

- Backup Copy for Oracle RMAN Backups
- Backup Copy for SAP on Oracle Backups
- Backup Copy for SAP HANA Backups



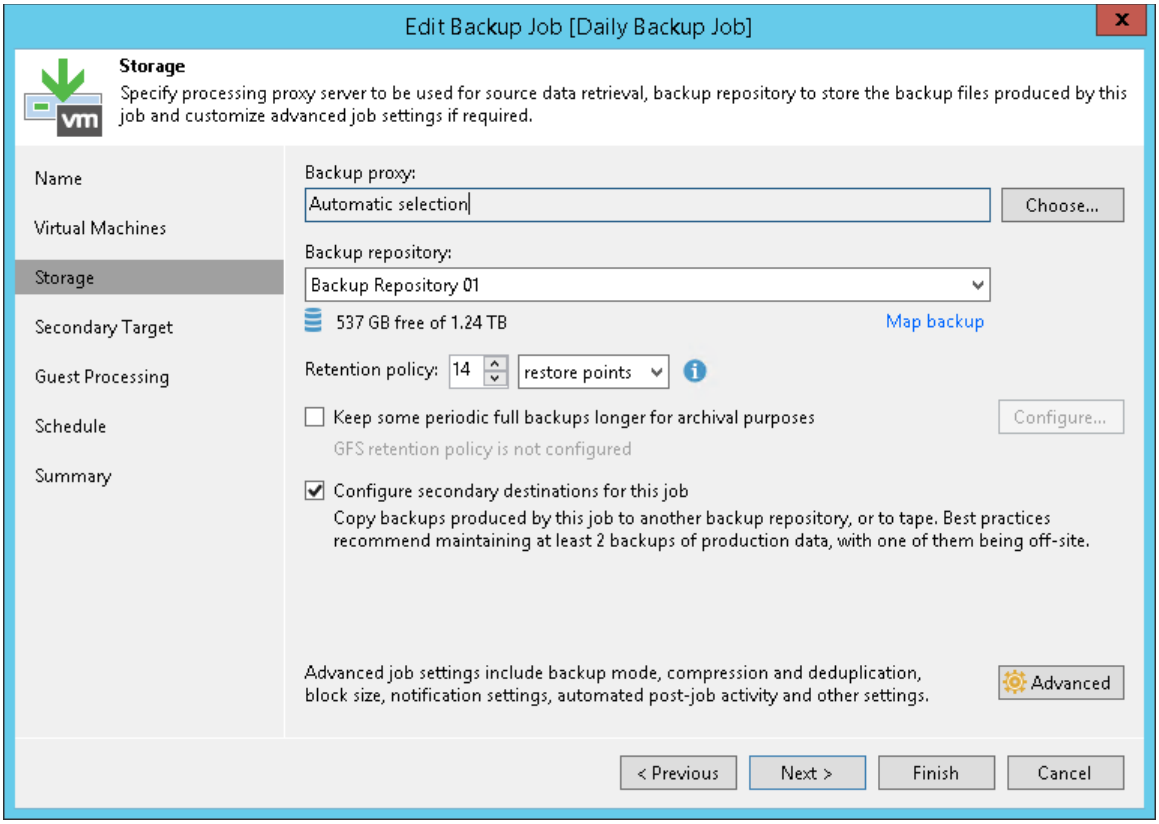
Linking Backup Jobs to Backup Copy Jobs

You can link backup jobs to backup copy jobs. This option lets you create a secondary target for the backup job and store backups created with the backup job in the secondary backup repository.

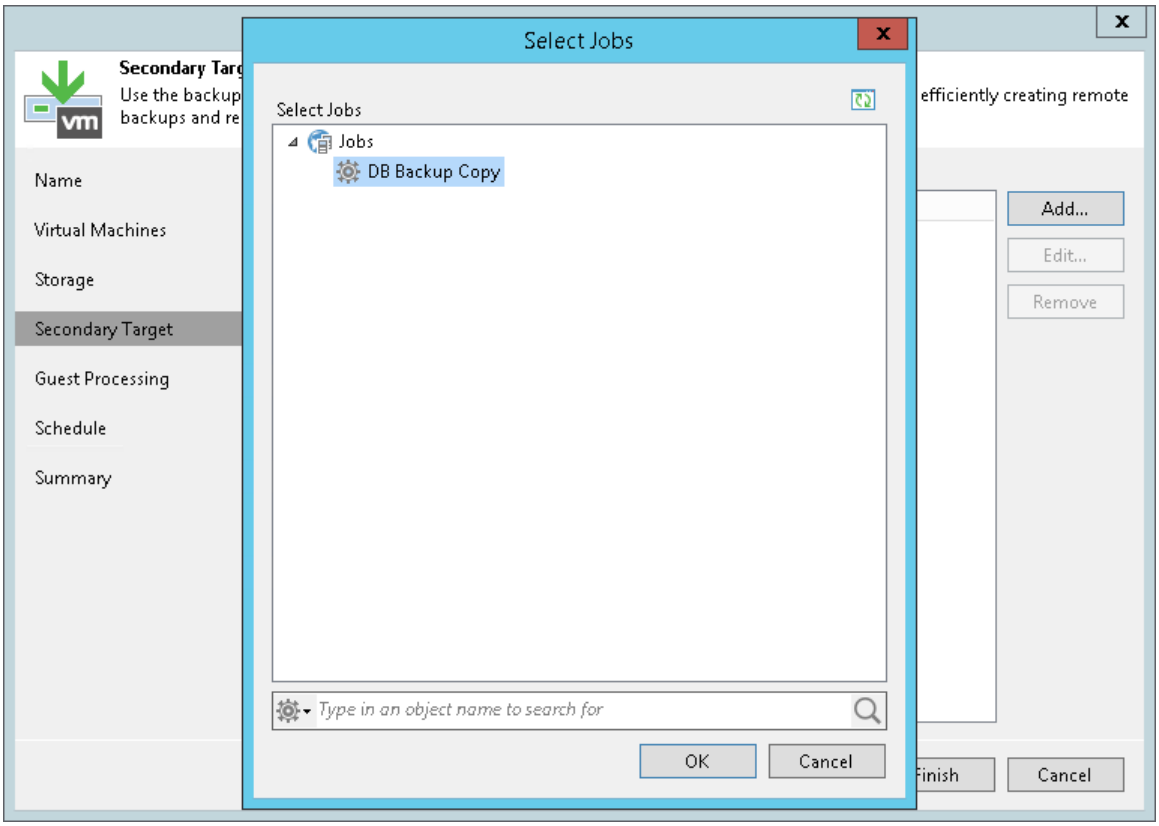
When you link a backup job to the backup copy job, Veeam Backup & Replication automatically updates properties of the backup copy job and adds to it the backup job as a source of data. During every backup copy session, the backup copy job checks the source backup repository for new restore points. As soon as a backup job session is finished and a new restore point appears in the source backup repository, the backup copy job automatically copies this restore point to the target backup repository.

You can link a backup job to an existing backup copy job using the **Backup Job** wizard. To link jobs:

1. Open the backup job settings for editing. For more information, see Editing Job Settings.
2. Navigate to the **Storage** step.
3. Select the **Configure secondary destination for this job** check box.



4. At the **Secondary Target** step of the wizard, click **Add** and choose a backup copy job to which the backup job must be linked. The backup copy job must be already configured on the backup server.



Managing Backups

To view all backups created by backup copy jobs, open the **Home** view and select the **Backups** > **Disk (Copy)** node in the inventory pane. The list of available backups is displayed in the working area. You can view backup properties, remove unnecessary backups and remove missing restore points.

Viewing Backup Properties

You can view summary information about backups created by backup copy jobs. The summary information provides the following data: available restore points, date of restore points creation, compression and deduplication ratios, data size and backup size.

In the summary information, Veeam Backup & Replication displays data about restore points created by the short-term retention scheme and archive restore points created by the GFS retention scheme (if GFS retention is enabled). Archive restore points are marked with the following letters:

- R — full backups created with the short-term retention scheme or active full backups
- W — weekly backups
- M — monthly backups
- Y — yearly backups

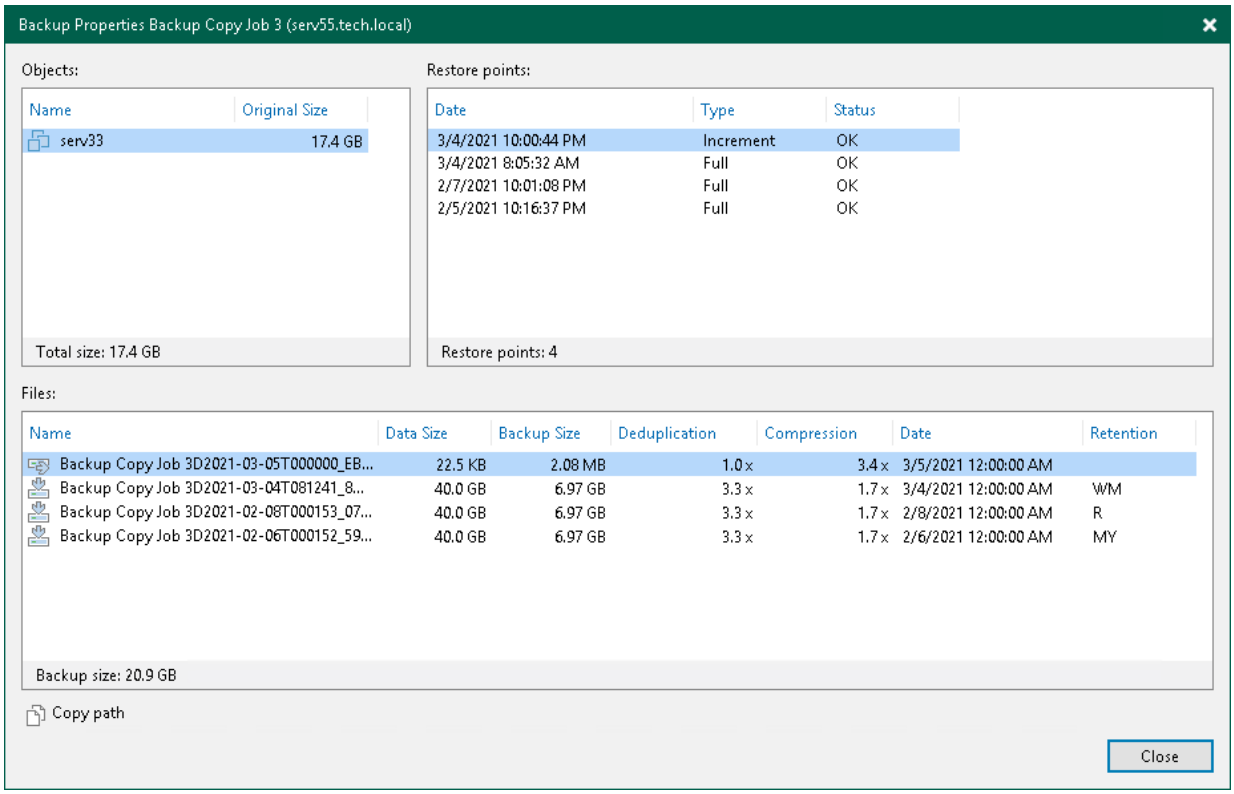
In the summary information, you can also see the following icons:

Icon	State
	Full restore point
	Incremental restore point
	Reverse incremental restore point
	Missing full restore point
	Missing incremental restore point
	Missing reverse incremental restore point

To view summary information for a backup copy:

1. Open the **Home** view.
2. In the inventory pane, select **Backups** > **Disk (copy)**.

3. In the working area, right-click the backup copy and select **Properties**.



Removing Backups from Target Repositories

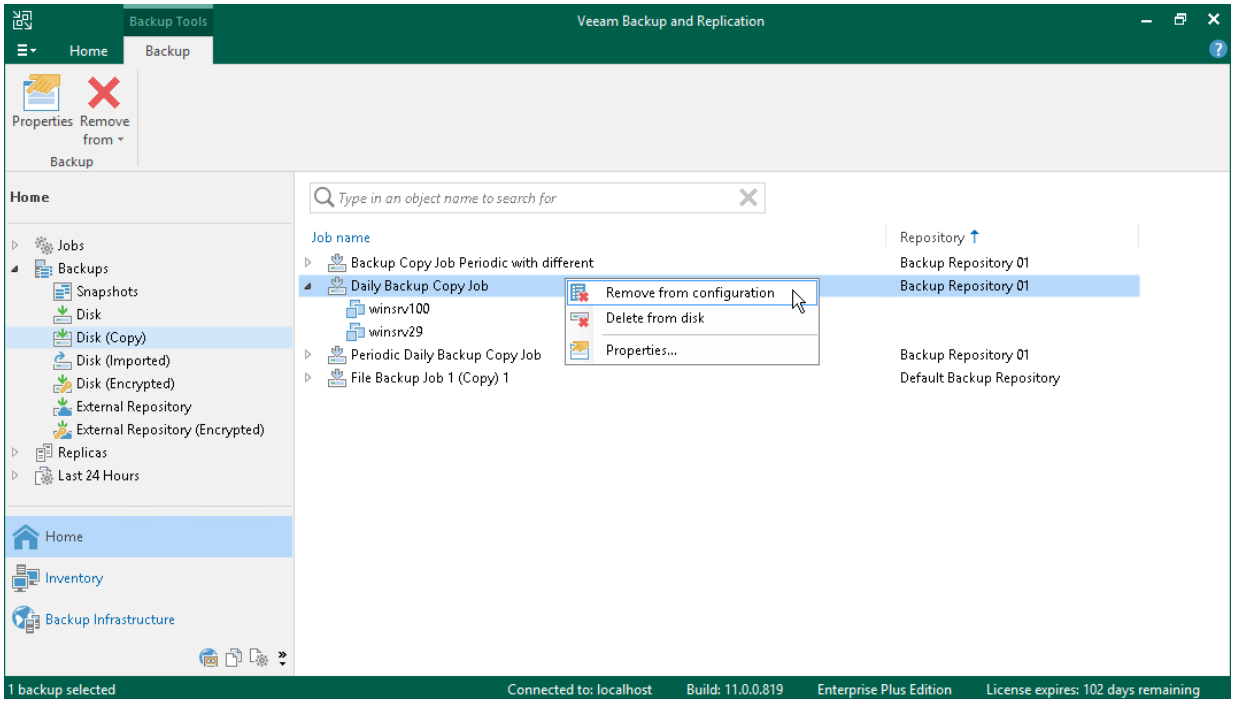
You can remove backups created by backup copy jobs from Veeam Backup & Replication configuration or permanently delete backup chains from the target backup repository.

Removing from Configuration

When you use the **Remove from configuration** option, you delete records about backup copies from the Veeam Backup & Replication console and configuration database. The backup copy job remains in the list of jobs and all backup files remain in the target backup repository. You can easily import backups to the Veeam Backup & Replication console for restore operations at any moment.

To remove from backups:

1. Open the **Home** view.
2. In the inventory pane, select **Backups**.
3. In the working area, right-click the backup copy and select **Remove from configuration**.
4. To remove all weekly, monthly, quarterly and yearly backups created by the job, select the **Include archived full backups** check box and click **Yes**.



Deleting from Disk

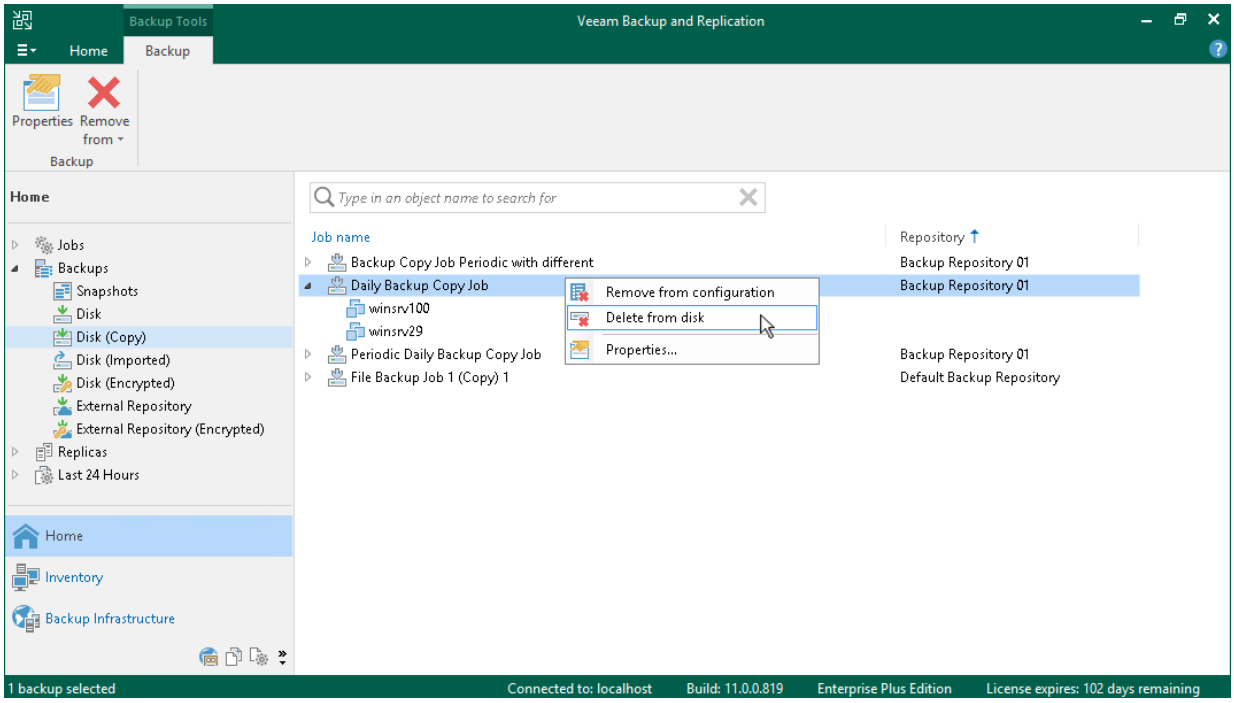
When you use the **Delete from disk** option, you delete records about backup copies from the Veeam Backup & Replication console and configuration database, and delete backup files from the target backup repository. This option can be used for the whole backup copy or for some workloads in the backup copy.

To permanently remove backup copies from the target backup repository:

1. Open the **Home** view.
2. In the inventory pane, select **Backups**.

3. In the working area, right-click the backup copy or a workload in the backup copy and select **Delete from disk**.

4. To remove all weekly, monthly, quarterly and yearly backups from disk, select the **Include archived full backups** check box and click **Yes**.



Removing Missing Restore Points

In some cases, one or more restore points in the backup chain may be inaccessible. This can happen, for example, if the backup repository is put to the maintenance mode (for scale-out backup repositories), the backup repository is not available or some backup file is missing in the backup chain. Backup chains that contain missing restore points get corrupted — you cannot perform backup copy or restore data from the missing restore point, and restore points that depend on the missing restore point.

You can perform the following with missing restore points:

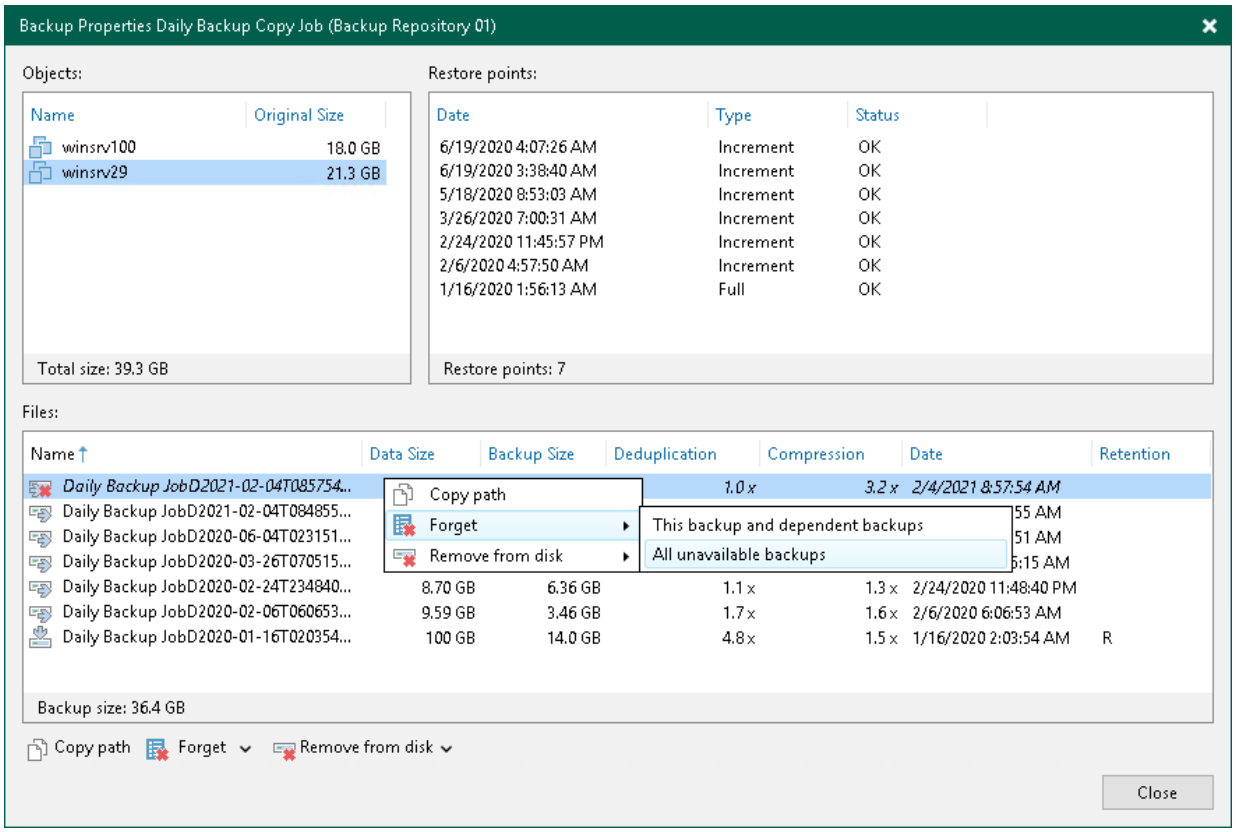
- Forget — you can remove records about missing restore points from the configuration database. Veeam Backup & Replication will ignore the missing restore points and will not display them in the console. The backup files will remain on disk (if backup files are still available).
- Remove — you can remove records about missing restore points from the Veeam Backup & Replication console and configuration database, and delete backup files from disk (if backup files are still available).

> **NOTE**
>
> Consider the following:
>
> - The **Forget** and **Remove from disk** options are available only for restore points that are missing from the backup chain or that depend on missing ones.
> - You can manually update information about missing restore points. For this, disable a backup copy job and rescan the backup repository that is the target for the backup copy job. For more information, see Disabling and Removing Jobs and Rescanning Backup Repositories.
>
>   Manual update can be required because Veeam Backup & Replication requires some time to update information in the configuration database for restore points that were removed from a backup chain or became inaccessible. That is why such restore points may not be displayed in the console as missing restore points.

To remove records about missing restore points from the configuration database:

1. Open the **Home** view.
2. In the inventory pane, select **Disk (copy)** under **Backups**.
3. In the working area, select the backup and click **Properties** on the ribbon or right-click the backup and select **Properties**.
4. In the **Backup Properties** window, right-click the missing restore point and select **Forget**.
   - To remove only the selected restore point and restore points that depend on it (that is, a part of the backup chain starting from this restore point), select **This and dependent backups**.

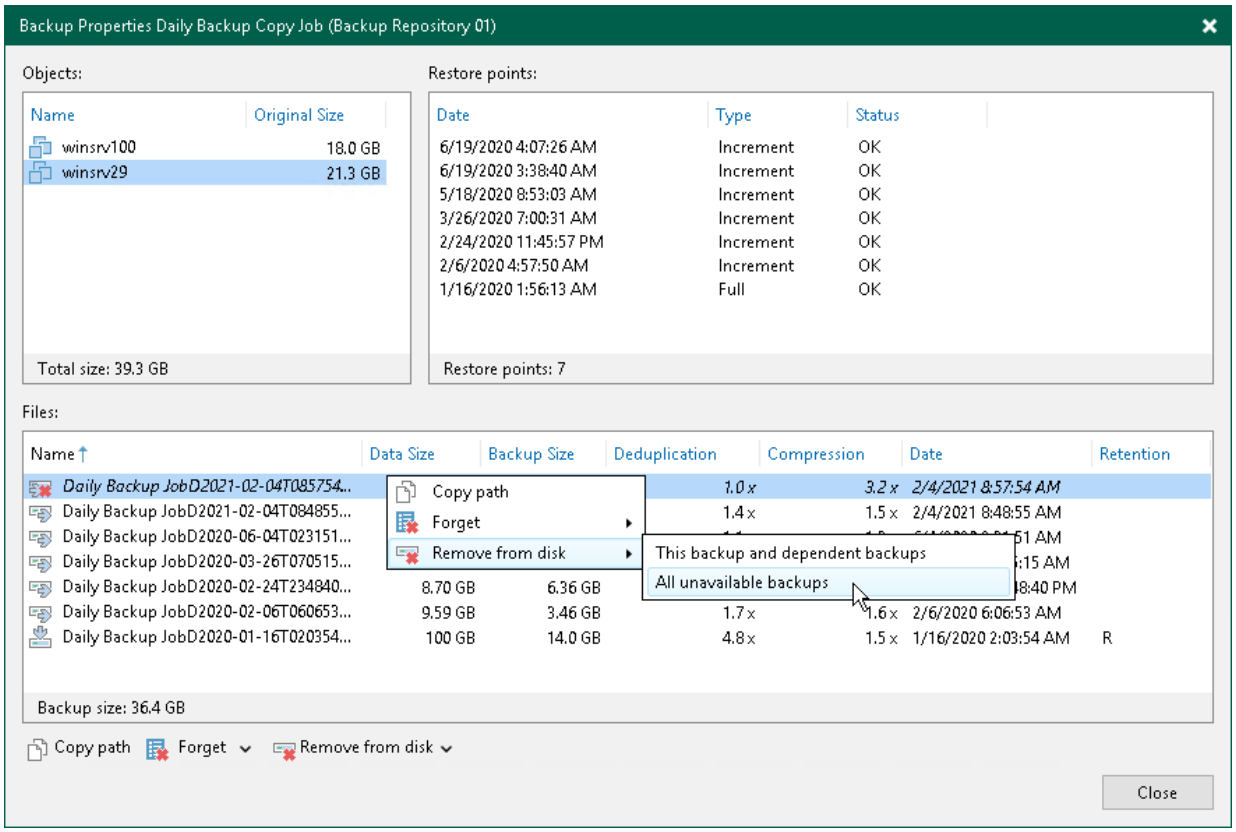
- To remove all missing restore points, select **All unavailable backups**.



To remove missing restore points from the configuration database and disk:

1. Open the **Home** view.
2. In the inventory pane, click **Disk (copy)** under **Backups**.
3. In the working area, select the backup and click **Properties** on the ribbon or right-click the backup and select **Properties**.
4. In the **Backup Properties** window, right-click the missing restore point and select **Remove from disk**.
   - To remove only the selected restore point and restore points that depend on it (that is, a part of the backup chain starting from this restore point), select **This and dependent backups**.

- To remove all missing restore points, select **All unavailable backups**.



Reporting

The way how you perform reporting operations for backup copy jobs and for backup jobs differs but slightly:

- For backup copy jobs in the periodic mode, the process is the same as described in the Backup: Reporting section.
- For backup jobs in the immediate mode, the process differs as described further.

The following list shows the reporting operations and highlights the differences for backup jobs in the immediate mode:

- Viewing real-time statistics

  You can view real-time statistics for the whole job if you select the job in the working area of the **Jobs** node. The whole job report shows general information about the job itself and child jobs — tasks that copy backup jobs added as sources to the backup copy job. You can also view real-time statistics for an individual child job if you select the child job in the working area of the **Last 24 Hours** or **Running** node. This statistics shows detailed information about the selected child job including the processed VMs.

- Viewing job session results

  You select and view session results for a child job.

- Viewing job reports

  The job report shows results for the last job run and does not provide details on child jobs. If you want to get reports once a child job finishes, configure notifications. For more information, see Notification Settings.

- Viewing session reports

  You can view session reports only if you configured notifications for a job. In this case, you get reports once a child job finishes. For more information, see Notification Settings.

Managing Jobs

To view all configured jobs, open the **Home** view and select the **Jobs** > **Backup Copy** node in the inventory pane. The list of available jobs is displayed in the working area. You can edit job properties, start and stop jobs, and delete unnecessary jobs.

Editing Backup Copy Jobs

You can edit backup copy job settings after you create it. For example, you may want to change scheduling settings for the job or add some machines to the job.

To edit job settings:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs** > **Backup Copy**.
3. In the working area, select the job and click **Edit** on the ribbon or right-click the job and select **Edit**.

You will follow the same steps as you have followed when creating the job and can change job settings as required.

Related Topics

- Creating Backup Copy Jobs for VMs and Physical Machines
- Creating Backup Copy Jobs for HPE StoreOnce Repositories
- Creating Backup Copy Jobs for Oracle and SAP HANA Databases

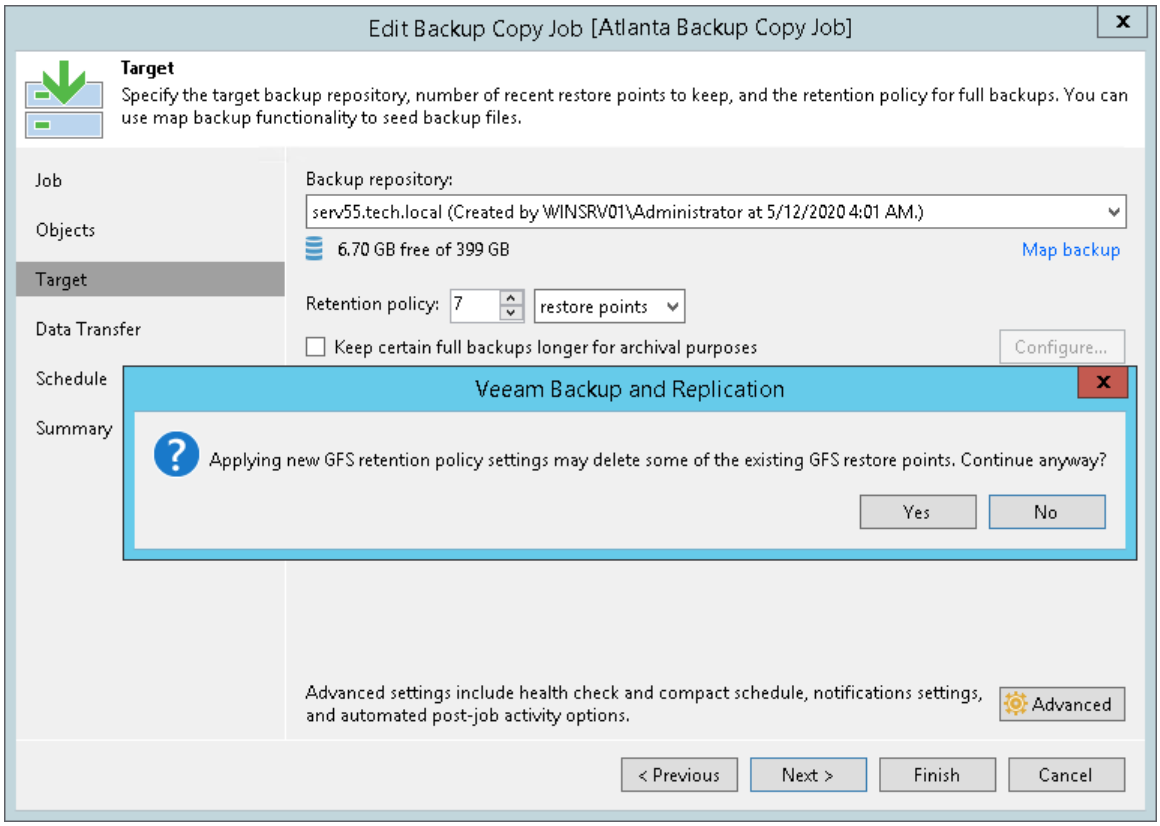
Disabling GFS Scheme

If you disable the **Keep the following restore points as full backups for archival purposes** option, and you already have archive full backups in the target backup repository, Veeam Backup & Replication will offer you to remove existing archive full backups.

- Click **Yes** to remove archive full backups from the target backup repository. Archive full backups will be removed during the next retention cycle (next backup copy session). The backup copy job will not create archive full backups.
- Click **No** to keep archive full backups in the target backup repository. Archive full backups will be displayed under the **Backups** > **Disk (Imported)** node in the Veeam Backup & Replication console. The backup copy job will not create archive full backups.

NOTE

If you disable the **Keep the following restore points as full backups for archival purposes** option and enable it again later, archive full backups that remained on disk will not be linked to the backup copy job. They will still be displayed under the **Backups** > **Disk (Imported)** node in the Veeam Backup & Replication console.



Starting Backup Copy Jobs Manually

You can manually start a backup copy job. This procedure differs for the periodic and immediate copy modes. You can find details in the following sections:

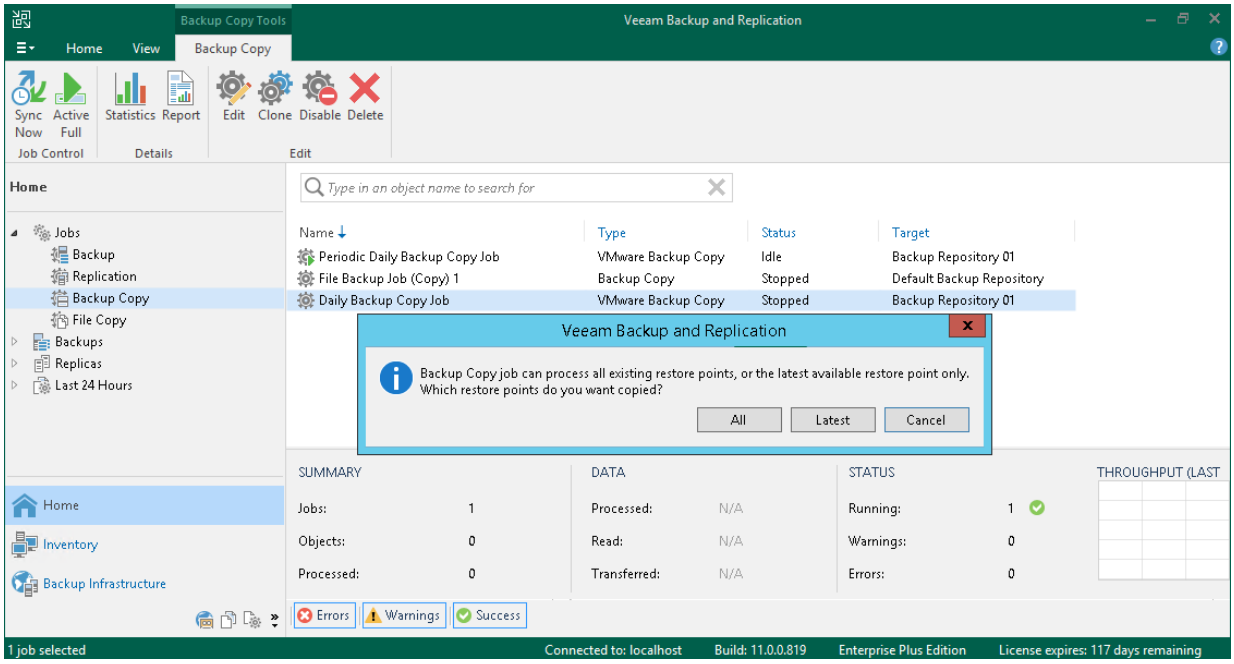
- Starting Jobs in Immediate Copy Mode
- Starting Jobs in Periodic Copy Mode

Starting Jobs in Immediate Copy Mode

Manual start can be helpful if the backup copy job was disabled for some time and restore points were not copied to the target repository.

To start the backup copy job manually:

1. Open the **Home** view.
2. In the inventory pane, select the **Backup Copy** node.
3. In the working area, select the backup copy job and click **Sync now** on the ribbon or right-click the backup copy job and select **Sync now**.
4. In the opened window, do the following:
   - If you want to copy all restore points created by the source jobs but that were not copied since the last backup copy job session, click **All**.
   - If you want to copy only the latest restore point for each source job, click **Latest**.



Starting Jobs in Periodic Copy Mode

Manual start of the backup copy job can be helpful if a new restore point has already appeared in the source backup repository but the previous backup copy interval has not finished yet.

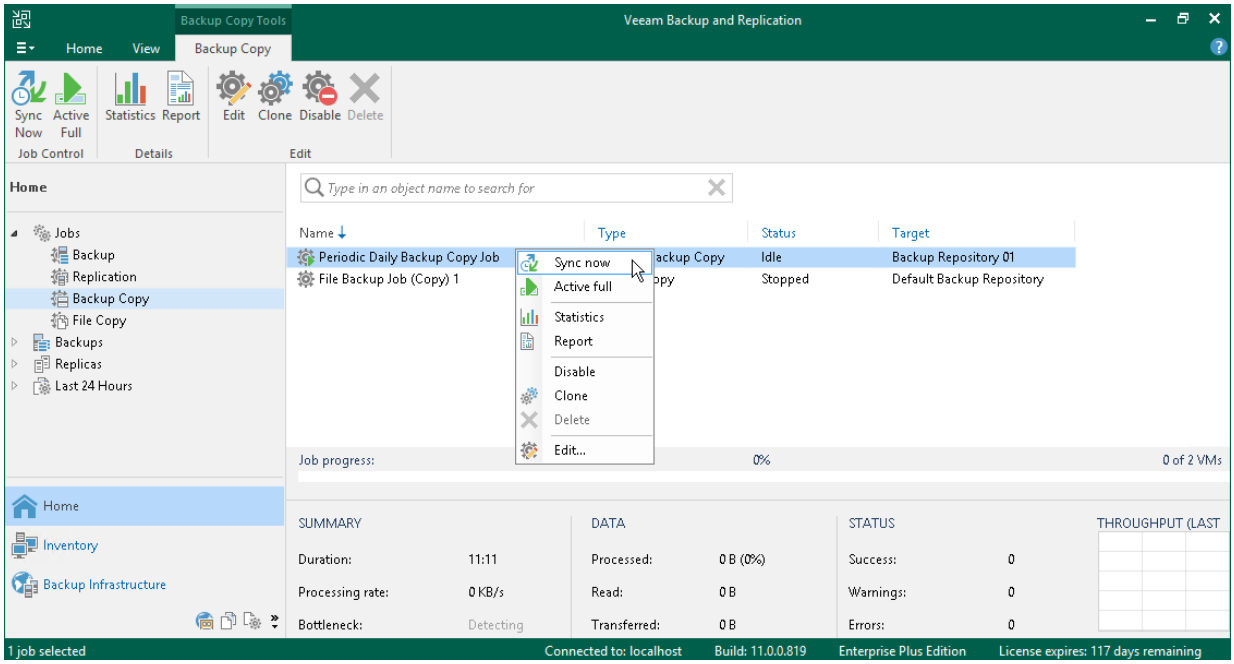
When you manually start the synchronization process, Veeam Backup & Replication creates a new backup copy interval.

- In case of backup copy jobs with minutely and hourly intervals, the backup copy interval is equal to those that are created automatically by the schedule. As a result, the start time of backup copy processing shifts forward.
- In case of backup copy jobs with intervals equal to one or several days, the day of the next backup copy interval shifts forward for the number of days equal to the interval. The start time of the backup copy interval, however, remains the same.

  For example, you configure a backup copy job to copy data every 30 days. The backup copy interval starts at 2:00 AM. The manual backup copy interval starts on May 1 at 1:00 PM. The manual backup copy interval will work from 1:00 PM on May 1 till 2:00 AM on May 31. On May 31 at 2:00 AM Veeam Backup & Replication will automatically start a new 30-day backup copy interval.

To start the backup copy job manually:

1. Open the **Home** view.
2. In the inventory pane, select the **Backup Copy** node under **Jobs**.
3. In the working area, select the backup copy job and click **Sync now** on the ribbon or right-click the backup copy job and select **Sync now**.

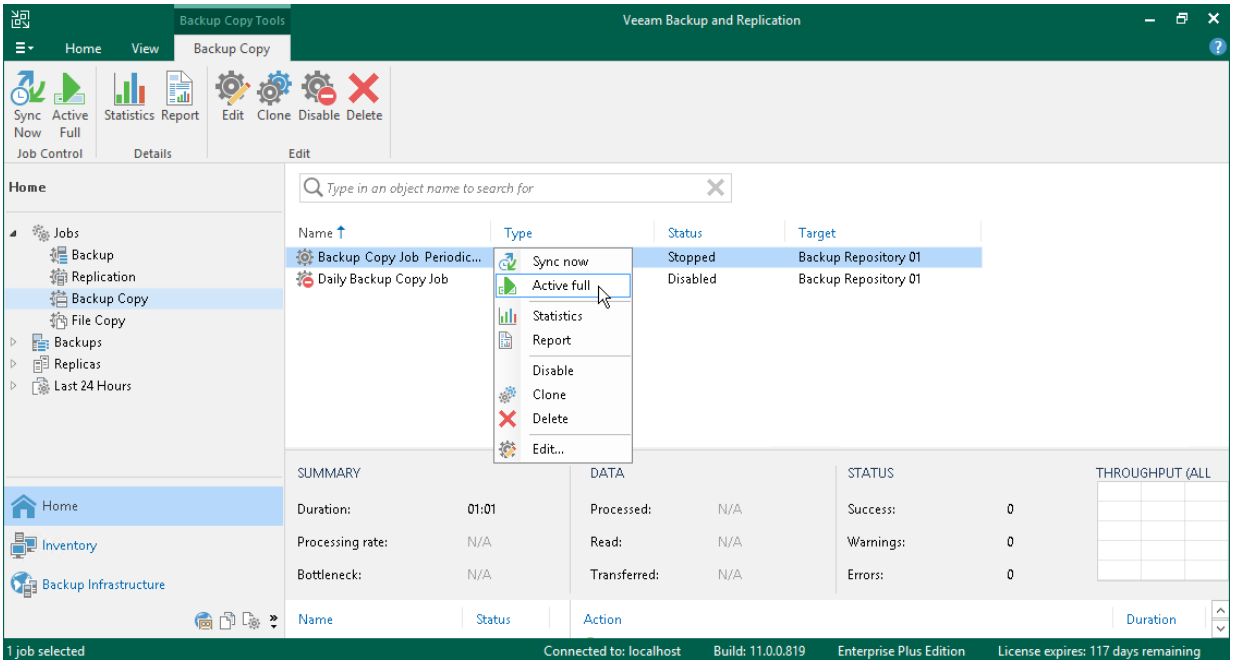


Creating Active Full Backups

You can manually create an ad-hoc full backup — active full backup, and add it to the backup chain in the target backup repository. Active full backup can be helpful if you want to change backup copy job settings, for example, enable or disable encryption. Veeam Backup & Replication will apply new settings starting from this full backup.

To create an active full backup manually:

1. Open the **Home** view.
2. In the inventory pane, select the **Backup Copy** node under **Jobs**.
3. In the working area, select the backup copy job and click **Active full** on the ribbon or right-click the backup copy job and select **Active full**. Veeam Backup & Replication will start a new backup copy session, copy data from the source backup repository and save it in a full backup file in the target backup repository.



Disabling and Deleting Jobs

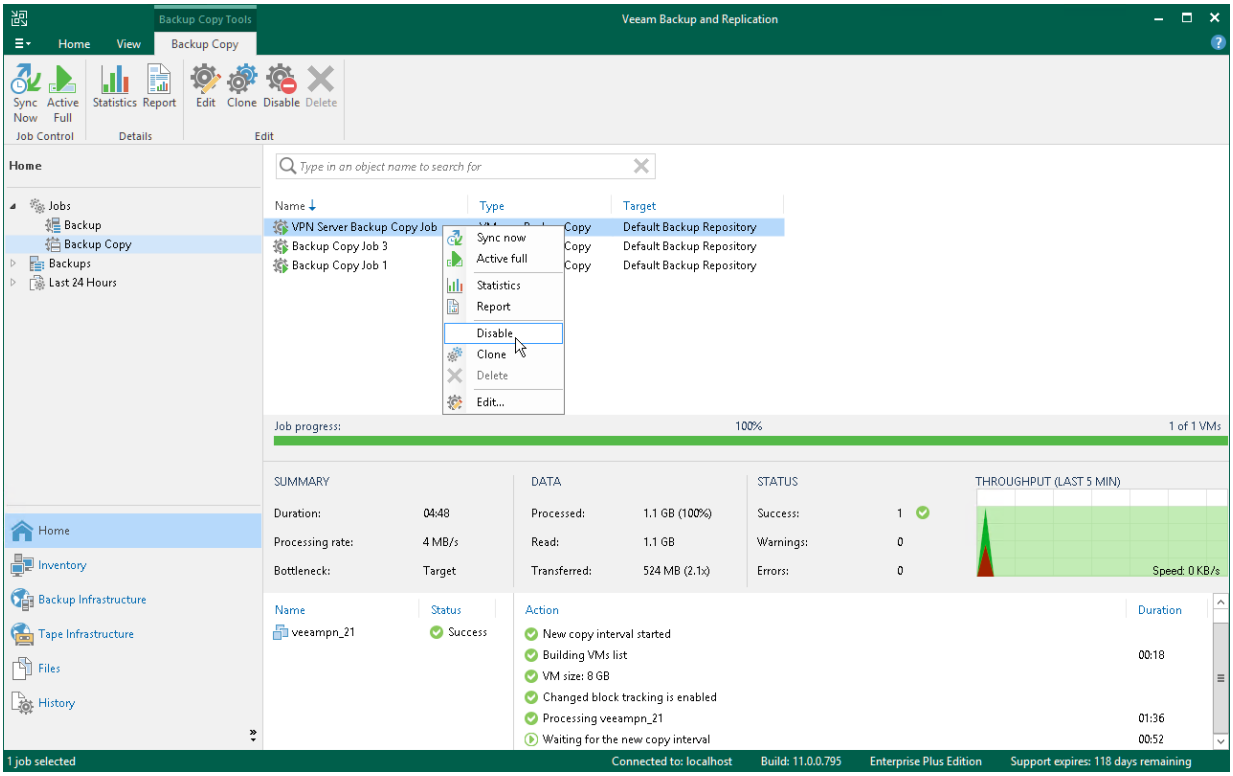
You can temporary disable backup copy jobs. The disabled job is paused for some period of time and is not run by the specified schedule. You can enable a disabled job at any time. You can also permanently delete a job from Veeam Backup & Replication and from the configuration database.

Disabling Job

To disable a job:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs** > **Backup Copy**.
3. In the working area, select the job and click **Disable** on the ribbon or right-click the job and select **Disable**.

To enable a disabled job, select the job in the list and click **Disable** once again.



Deleting Job

To delete a job:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs** > **Backup Copy**.
3. In the working area, select the job, click **Delete** on the ribbon or right-click the job and select **Delete**.

> **NOTE**
>
> If you want to delete a periodic backup copy job, you must first stop the synchronization process. To do this, disable the backup job. After the job is disabled, you can delete it.

After the job is deleted, the backups created by this job are displayed under the **Backups** > **Disk (Orphaned)** node. If the backup files created by this job were also stored in an object storage repository, they will also be displayed under the **Backups** > **Object Storage (Orphaned)** node.

VM Copy

With Veeam Backup & Replication, you can run a VM copy job to create an independent fully-functioning copy of a VM or VM container (host, cluster, folder, resource pool, VirtualApp, datastore or tag) on the selected storage. VM copying can be helpful if you want to move your datacenter, create a test lab and so on.

The produced copy of a VM is stored decompressed, in a native VMware vSphere format, so it can be started right away. Although VM copy is similar to replication in many respects, there are several important differences.

- VM copy is a single-use process (that is, every run of a VM copy job mirrors a VM in its latest state). Due to their nature, VM copy jobs do not support incremental runs.
- Veeam Backup & Replication does not create and maintain restore points for VM copies. If you schedule to run a VM copy job periodically, every new run will overwrite the existing copy.
- With the VM copy job, all VM disks are copied as thick, while replication allows you to preserve the format of disks or convert the disk format on the fly.
- There are no failover or failback possibilities for a VM copy.

VM copy jobs use the same infrastructure components as backup jobs (for details, see Backup Architecture). In addition to available scenarios, you can also copy VMs to a target folder on any server or host connected to the backup server.

Copying VMs

With VM copy jobs you can create a fully-functional copy of a VM and store this copy in the backup repository or storage device. VM copying may be helpful if you want to move your datacenter to another location, archive a VM before decommissioning and so on.

To create a VM copy, you must configure a VM copy job. One job can be used to process one VM or more VMs.

You can configure a job and start it immediately or save the job to start it later. Jobs can be started manually or scheduled to run automatically at specific time.

Before you create a VM copy job, check prerequisites. Then use the **New VM Copy Job** wizard to configure a VM copy job.

Before You Begin

Before you create a VM copy job, check the following prerequisites:

- Backup infrastructure components that will take part in the VM copying process must be added to the backup infrastructure and properly configured. These include the source ESXi host and server or backup repository on which you plan to store the VM copy.
- The target storage device must have enough free space to store created VM copies. To receive alerts about low space on the storage device, configure global notification settings. For more information, see Specifying Other Notification Settings.
- If you plan to use pre-freeze and/or post-thaw scripts, you must create scripts before you configure the VM copy job.

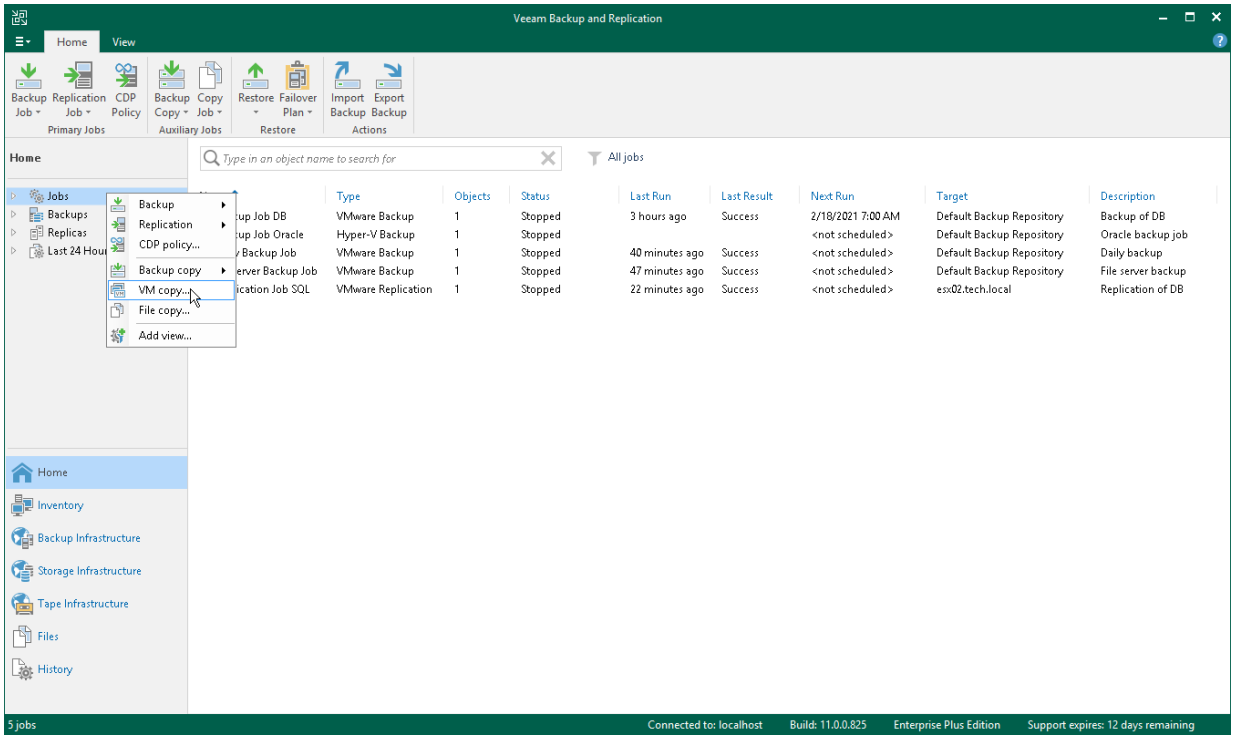
Mind the following limitations:

- Due to Microsoft limitations, you cannot use Microsoft Azure Active Directory credentials to perform application-aware processing on VMs running Microsoft Windows 10 (or later).
- If you use tags to categorize virtual infrastructure objects, check limitations for VM tags. For more information, see VM Tags.

Step 1. Launch VM Copy Job Wizard

To run the **VM Copy Job** wizard, do either of the following:

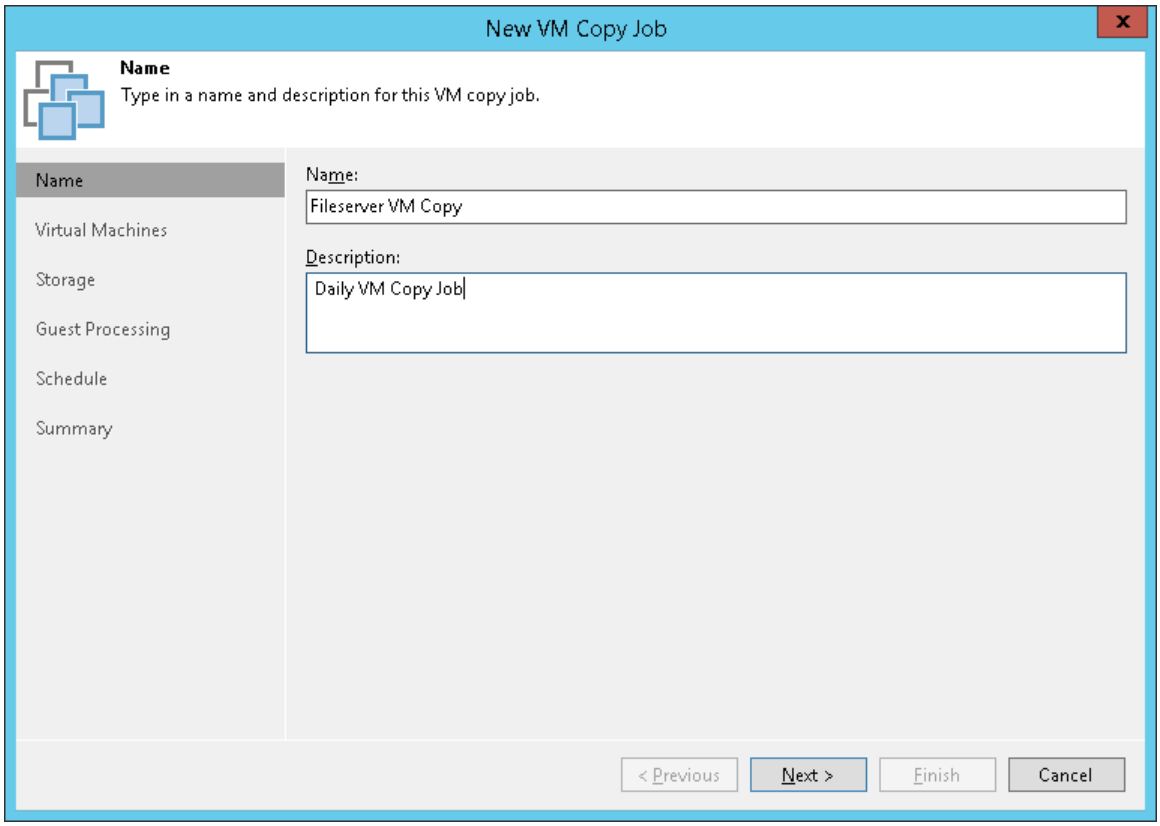
- On the **Home** tab, click **Copy Job** > **Virtual machine**.
- Open the **Home** view. In the inventory pane, right-click **Jobs** and select **VM Copy**.
- Open the **Inventory** view, in the working area select the VMs, click **Add to VM Copy** on the ribbon and select **New job** or right-click the VMs area and select **Add to VM copy job** > **New job**. In this case, the selected VMs will be automatically added to the VM copy job. You can add other VMs to the job when passing through the wizard steps.
- You can quickly add the VMs to an already existing job. To do this, open the **Inventory** view, in the working area select the VMs and click **Add to VM Copy** > *name of the job* on the ribbon or right-click the VMs and select **Add to VM copy job** > *name of the job*.



Step 2. Specify Job Name and Description

At the **Name** step of the wizard, specify a name and description for the VM copy job.

1. In the **Name** field, enter a name for the VM copy job.
2. In the **Description** field, provide a description for future reference. The default description contains information about the user who created a job, date and time when the job was created.



Step 3. Select VMs to Copy

At the **Virtual Machines** step of the wizard, select VMs and VM containers (hosts, clusters, folders, resource pools, VirtualApps, datastores or tags) that you want to copy.

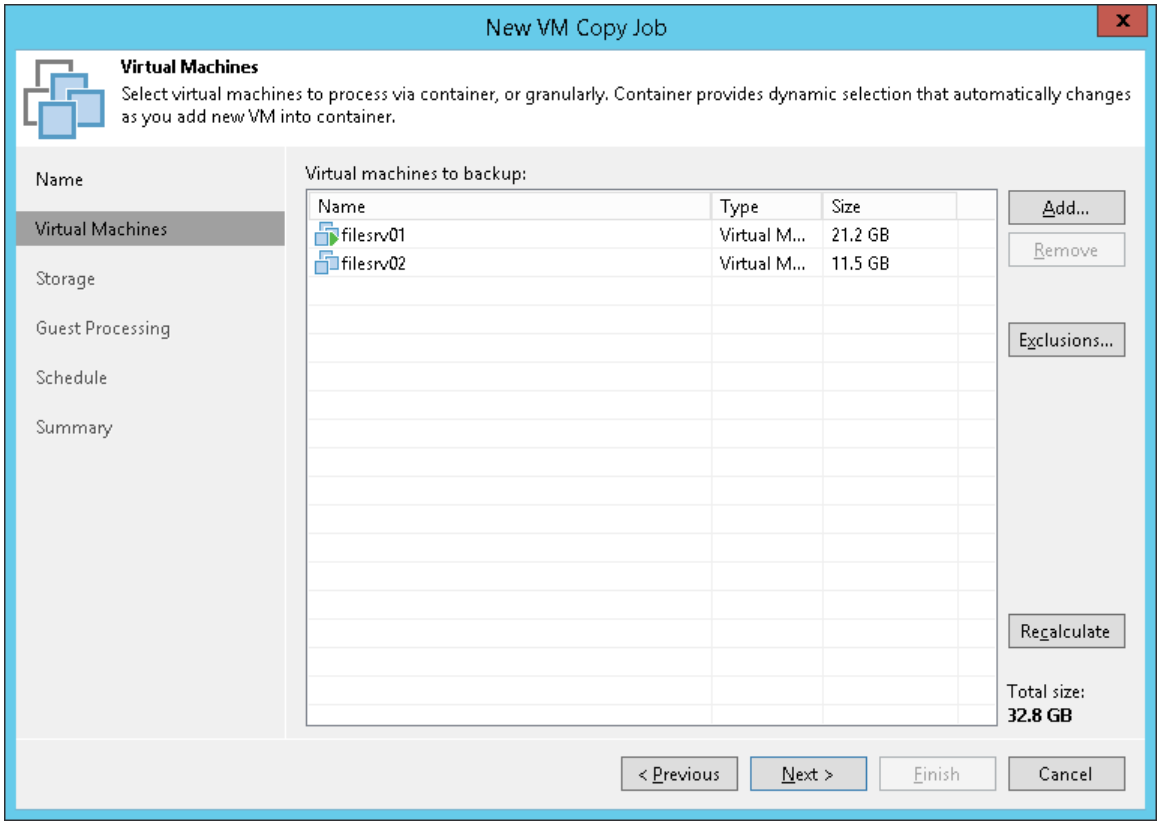
Jobs with VM containers are dynamic in their nature. If a new VM is added to the container in the virtual infrastructure after the VM copy job is created, Veeam Backup & Replication will automatically update the job settings to include the added VM.

1. Click **Add**.
2. Use the toolbar at the top right corner of the window to switch between views: **Hosts and Clusters**, **VMs and Templates**, **Datastores and VMs** and **Tags**. Depending on the view you select, some objects may not be available. For example, if you select the **VMs and Templates** view, no resource pools, hosts or clusters will be displayed in the tree.
3. Select the object and click **Add**.

To quickly find the necessary object, you can use the search field at the bottom of the **Add Objects** window.

1. Click the button to the left of the search field and select the necessary type of object to search for: *Everything*, *Folder*, *Cluster*, *Host*, *Resource pool*, *VirtualApp* or *Virtual machine*.
2. Enter the object name or a part of it in the search field.
3. Click the **Start search** button on the right or press [Enter] on the keyboard.

The initial size of VMs and VM containers added to the VM copy job is displayed in the **Size** column in the list. The total size of objects is displayed in the **Total size** field. Use the **Recalculate** button to refresh the total size value after you add a new object to the job.



Step 4. Exclude Objects from VM Copy Job

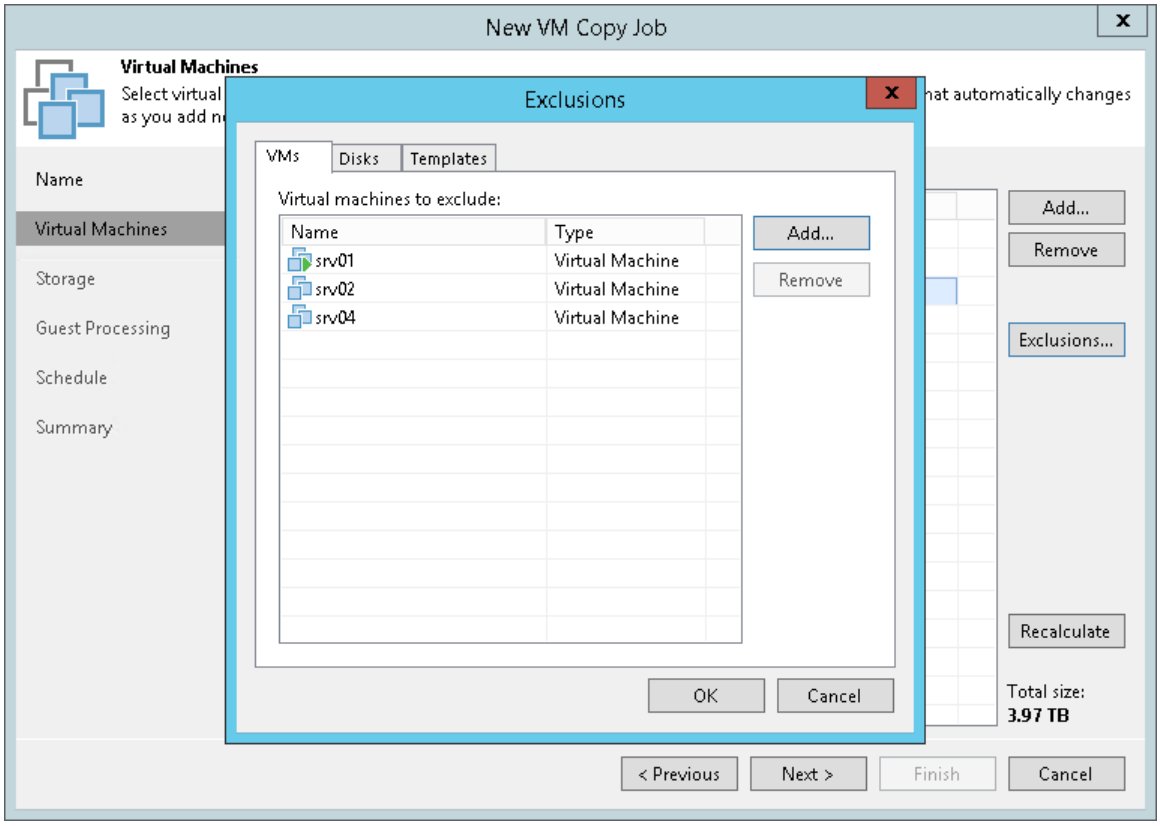
After you have added VMs and VM containers to the job, you can specify which objects you want to exclude from the VM copy. You can exclude the following types of objects:

- VMs from VM containers
- Specific VM disks
- VM templates

> **NOTE**
>
> Veeam Backup & Replication automatically excludes VM log files from VM copies to make copying process faster and reduce the size of the resulting file.

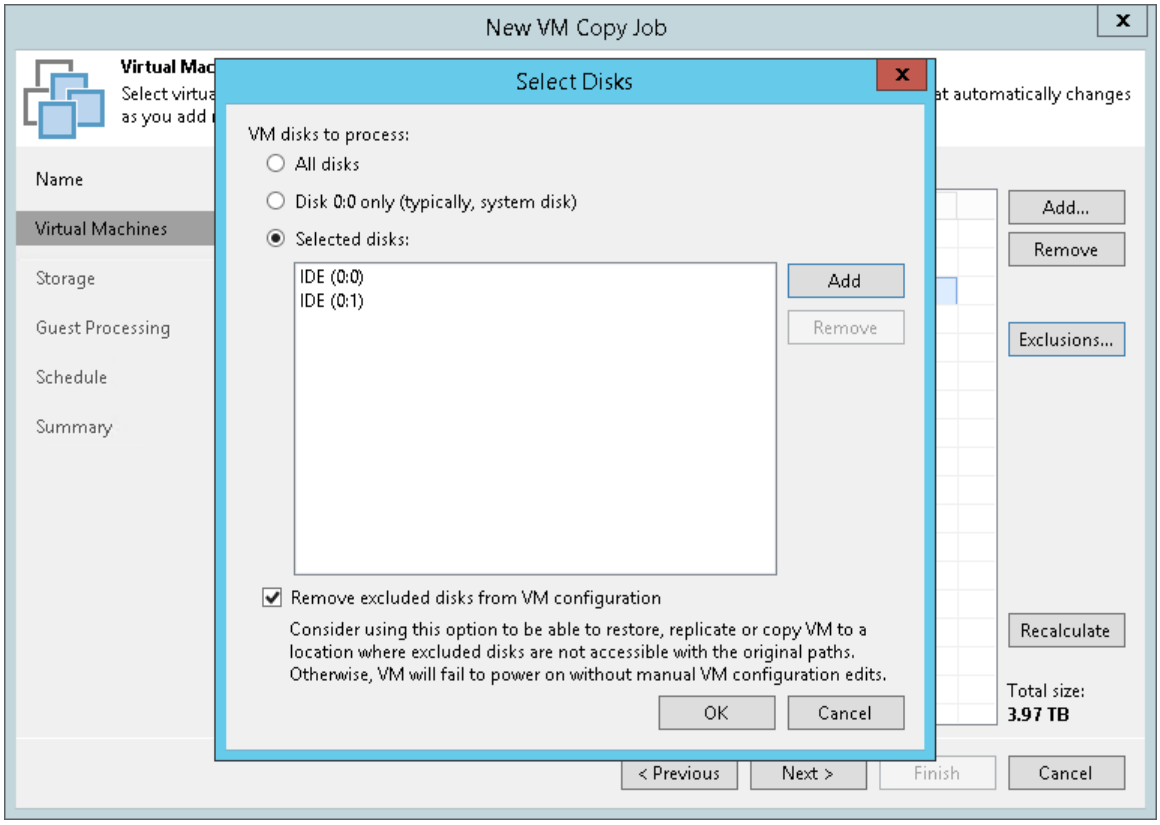
To exclude VMs from a VM container:

1. At the **Virtual Machines** step of the wizard, select a VM container added to the job and click **Exclusions**.
2. Click the **VMs** tab.
3. Click **Add**.
4. Use the toolbar at the top right corner of the window to switch between views: **Hosts and Clusters**, **VMs and Templates**, **Datastores and VMs** and **Tags**. Depending on the view you select, some objects may not be available. For example, if you select the **VMs and Templates** view, no resource pools, hosts or clusters will be displayed in the tree.
5. In the displayed tree, select the necessary object and click **Add**. Use the **Show full hierarchy** check box to display the hierarchy of all VMware Servers added to Veeam Backup & Replication.
6. Click **OK**.



To exclude VM disks:

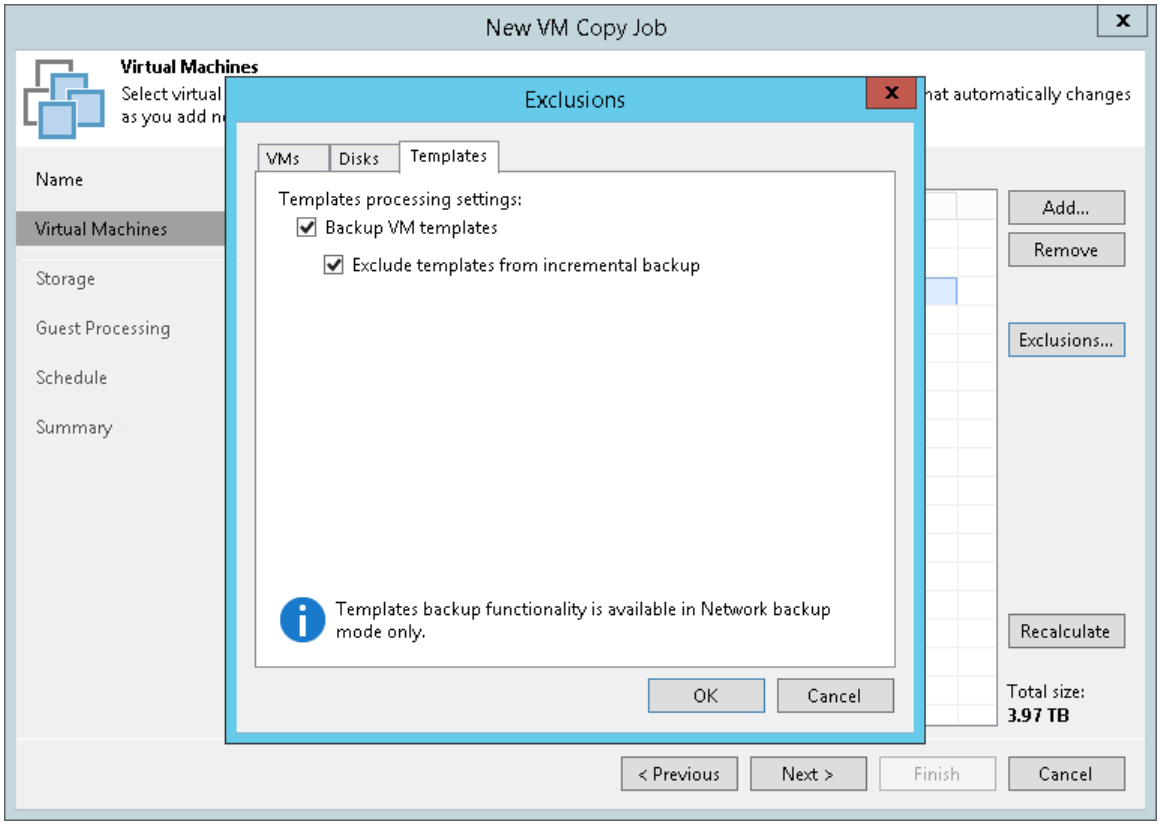
1. At the **Virtual Machines** step of the wizard, select a VM or VM container added to the job and click **Exclusions**.
2. Click the **Disks** tab.
3. Select the VM in the list and click **Edit**. If you want to exclude disks of a VM added as a part of the container, click **Add** to include the VM in the list as a standalone object.
4. Choose disks that you want to copy. You can choose to process all disks, 0:0 disks (typically, the system disks) or add to the list custom IDE, SCSI or SATA disks.
5. Select the **Remove excluded disks from VM configuration** check box. Veeam Backup & Replication will modify the VMX file of a copied VM to remove excluded disks from the VM configuration. If you use the VM copy to register the VM in a location where excluded disks are not accessible with the original paths, you will not have to manually edit the VM configuration file to be able to power on the VM.



To exclude VM templates:

1. At the **Virtual Machines** step of the wizard, select a VM or VM container added to the job and click **Exclusions**.
2. Click the **Templates** tab.
3. Clear the **Backup VM templates** check box.

4. If you want to include VM templates into the full VM copy only, leave the **Backup VM templates** check box selected and select the **Exclude templates from incremental backup** check box.



Step 5. Specify Copy Destination

At the **Storage** step of the wizard, select which backup proxy must be used for VM data transporting and specify the destination for the VM copy.

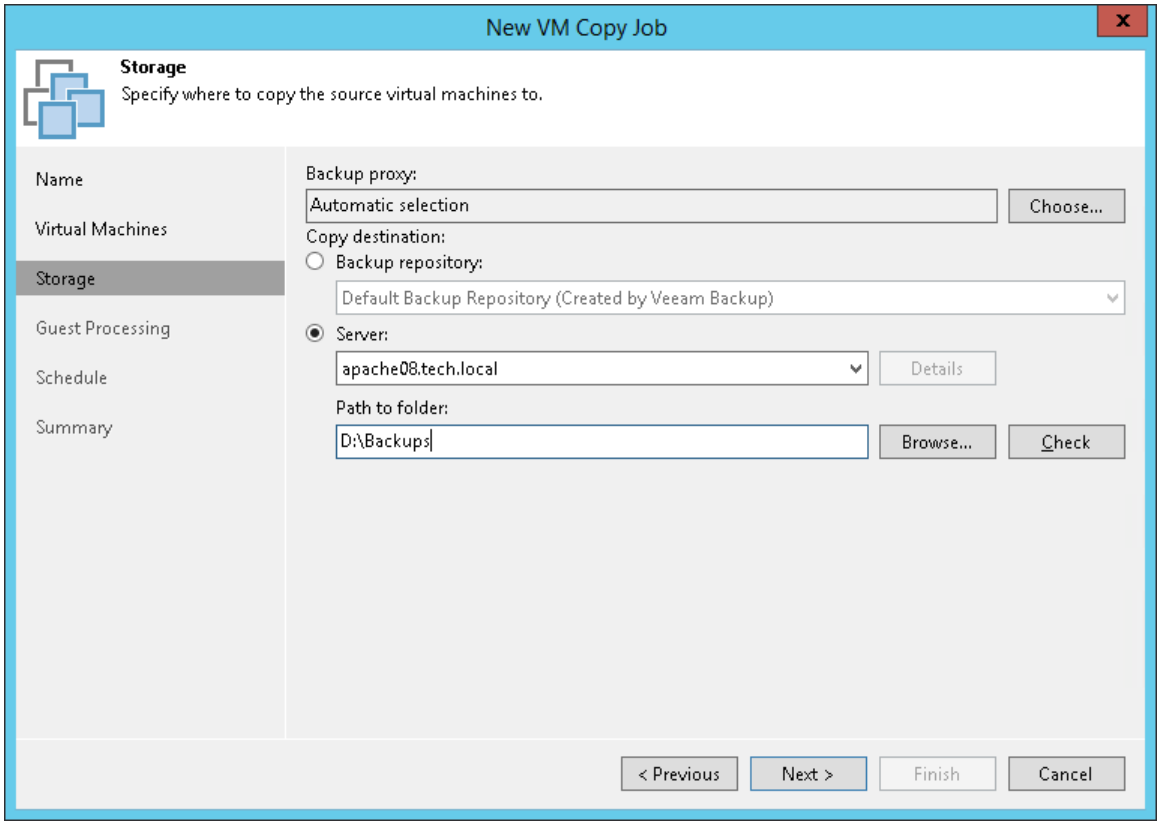
1. Click **Choose** next to the **Backup proxy** field to select a backup proxy.
   - If you choose **Automatic selection**, Veeam Backup & Replication will detect backup proxies that have access to the source datastore and automatically assign an optimal backup proxy for processing VM data.

     Veeam Backup & Replication assigns backup proxies to VMs included in the VM copy job one by one. Before processing a new VM in the VM list, Veeam Backup & Replication checks available backup proxies. If more than one backup proxy is available, Veeam Backup & Replication analyzes transport modes that the backup proxies can use for data retrieval and the current workload on the backup proxies to select the most appropriate one for VM processing.
   - If you choose **Use the selected backup proxy servers only**, you can explicitly select backup proxies that the job must use. It is recommended that you select at least two backup proxies to ensure that the VM copy job starts if one of the proxies fails or loses its connectivity to the source datastore.
2. In the **Copy destination** section, select a location where the created VM copy must be stored.
   - Select a backup repository from the list if you want to create a VM copy in the backup repository configured in the backup infrastructure. When you select a backup repository, Veeam Backup & Replication automatically checks how much free space is available on it.

     > **IMPORTANT**
     >
     > You cannot use a scale-out backup repository as the copy destination.

- Select **Server** if you want to store the VM copy on a disk or storage device attached to the server. From the **Server** list, select a server added to the backup infrastructure. In the **Path to folder** field, specify a folder on the server where the created VM copy must be stored.

  Use the **Check** button to see how much free space is available in the target location.



Step 6. Specify Guest Processing Settings

At the **Guest Processing** step of the wizard, you can enable the following settings for VM guest OS processing:

- Application-aware processing
- Microsoft SQL Server transaction log settings
- Oracle archived log settings
- Use of pre-freeze and post-thaw scripts

To coordinate guest processing activities, Veeam Backup & Replication deploys non-persistent runtime components or uses (if necessary, deploys) persistent agent components on the VM guest OS.

The non-persistent runtime components run only during guest processing and are stopped immediately after the processing is finished (depending on the selected option, during the VM copy job session or after the replication job completes).

You must specify a user account that will be used to connect to the VM guest OS and deploy the non-persistent runtime components or connect to (if necessary, deploy) persistent agent components:

1. From the **Guest OS credentials** list, select a user account that has enough permissions. For more information on the permissions and requirements for the user account, see Required Permissions for Guest Processing.

   If you have not set up credentials beforehand, click the **Manage accounts** link or click the **Add** button to add credentials. For more information on adding credentials, see the Credentials Manager section.

2. By default, Veeam Backup & Replication uses the same credentials for all VMs in the job. If some VM requires a different user account, click **Credentials** and enter custom credentials for the VM.

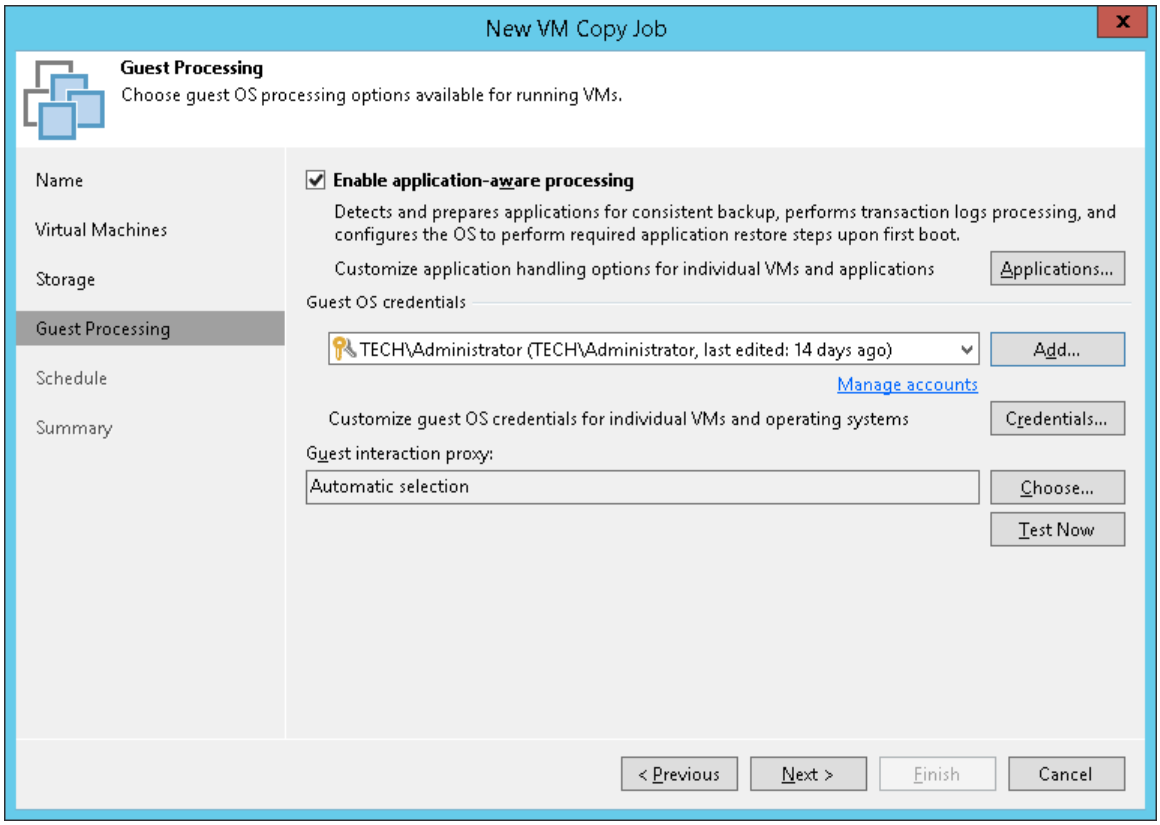
> **IMPORTANT**
>
> Credentials for application-aware processing and guest OS file indexing for Microsoft Windows VMs must be specified in the following format:
>
> - For Active Directory accounts — *DOMAIN|Username*
> - For local accounts — *Username* or *HOST|Username*

3. If you have added Microsoft Windows VMs to the job, specify which guest interaction proxy Veeam Backup & Replication can use to deploy the non-persistent runtime components or connect to (if necessary, deploy) persistent agent components on the VM guest OS. On the right of the **Guest interaction proxy** field, click **Choose**.
   - Leave **Automatic selection** to let Veeam Backup & Replication automatically select the guest interaction proxy.
   - Select **Use the selected guest interaction proxy servers only** to explicitly define which servers will perform the guest interaction proxy role. The list of servers contains Microsoft Windows servers added to the backup infrastructure.

To check if Veeam Backup & Replication can communicate with VMs added to the job and deploy the non-persistent runtime components or connect to (if necessary, deploy) persistent agent components on their guest OSes, click **Test Now**. Veeam Backup & Replication will use the specified credentials to connect to all VMs in the list.

NOTE

The guest interaction proxy functionality is included in the Veeam Universal License. When using a legacy socket-based license, Enterprise or higher edition is required.



Application-Aware Processing

If you add to the VM copy job VMs running VSS-aware applications, you can enable application-aware processing to create a transactionally consistent VM copy. The transactionally consistent VM copy guarantees proper recovery of applications on VMs without data loss.

To enable application-aware processing:

1. Select the **Enable application-aware processing** check box.
2. Click **Applications**.
3. In the displayed list, select the VM and click **Edit**.

   To define custom settings for a VM added as a part of the VM container, you must include the VM in the list as a standalone object. To do this, click **Add** and choose a VM whose settings you want to customize. Then select the VM in the list and define the necessary settings.
4. On the **General** tab, in the **Applications** section specify the VSS behavior scenario:
   - Select **Require successful processing** if you want Veeam Backup & Replication to stop the VM copy process if any VSS errors occur.
   - Select **Try application processing, but ignore failures** if you want to continue the VM copy process even if VSS errors occur. This option is recommended to guarantee completion of the job. The created VM image will not be transactionally consistent but crash consistent.

- Select **Disable application processing** if you do not want to enable quiescence for the VM.

5. [For Microsoft Exchange, Microsoft SQL and Oracle VMs] In the **Transaction logs** section, specify if Veeam Backup & Replication must process transaction logs or copy-only VM copies must be created.

   a. Select **Process transaction logs with this job** if you want Veeam Backup & Replication to process transaction logs.

      [For Microsoft Exchange VMs] With this option selected, the non-persistent runtime components or persistent components running on the VM guest OS will wait for the VM copy job to complete successfully and then trigger truncation of transaction logs. If the VM copy job fails, the logs will remain untouched on the VM guest OS until the next start of the non-persistent runtime components or persistent components.

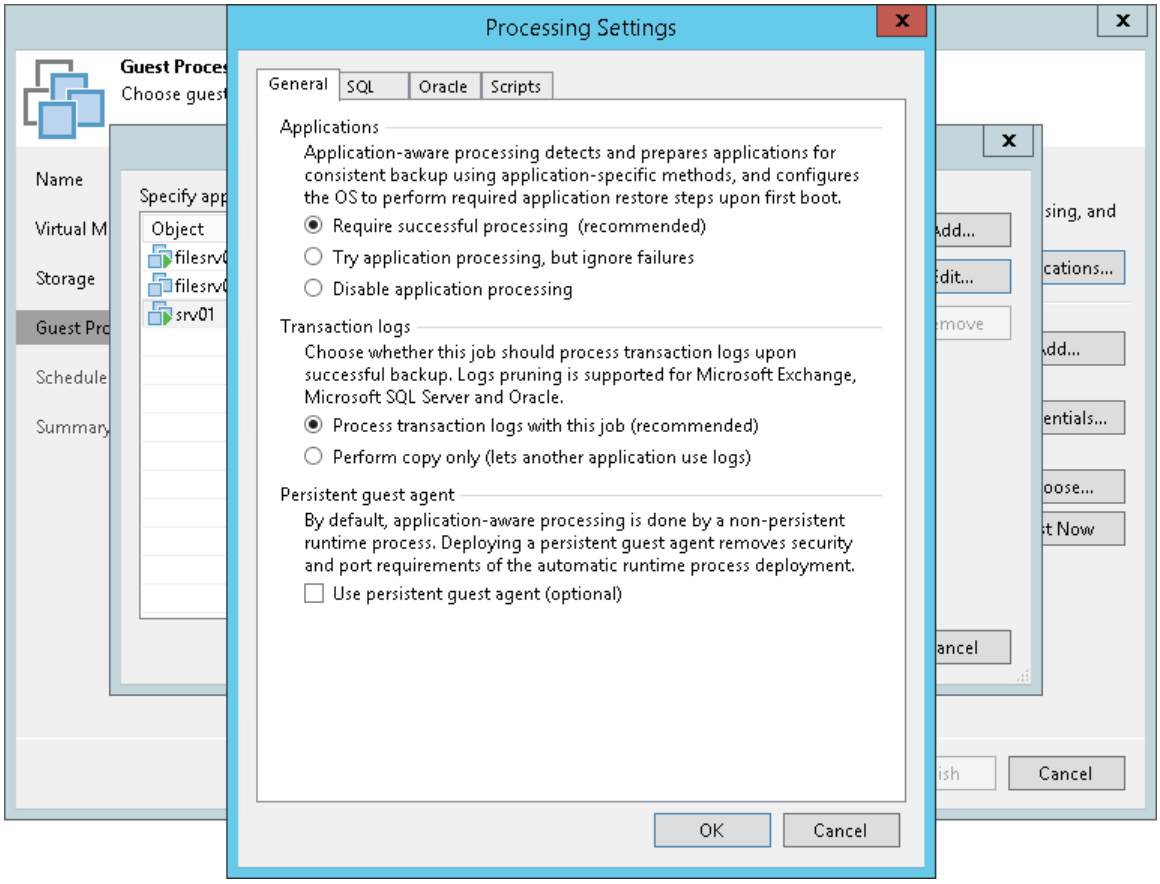
      [For Microsoft SQL Server VMs and Oracle VMs] You will have to specify settings for transaction log handling on the **SQL** and **Oracle** tabs of the **VM Processing Settings** window. For more information, see Microsoft SQL Server Transaction Log Settings and Oracle Archived Log Settings.

   b. Select **Perform copy only** if you use another backup tool to perform VM guest level backup or replication, and this tool maintains consistency of the database state. Veeam Backup & Replication will create a copy-only VM image for the selected VMs. The copy-only VM image preserves the chain of full/differential backup files and transaction logs on the VM. For more information, see Microsoft Docs.

6. In the **Persistent guest agent** section, specify if Veeam Backup & Replication must use persistent guest agents on each protected VM for application-aware processing.

   By default, Veeam Backup & Replication uses non-persistent runtime components.
   Veeam Backup & Replication deploys runtime components on each protected VM when the backup job starts, and removes the runtime components as soon as the backup job finishes.

Select the **Use persistent guest agent check** box to enable persistent agent components for guest processing. For more information, see Non-Persistent Runtime Components and Persistent Agent Components.

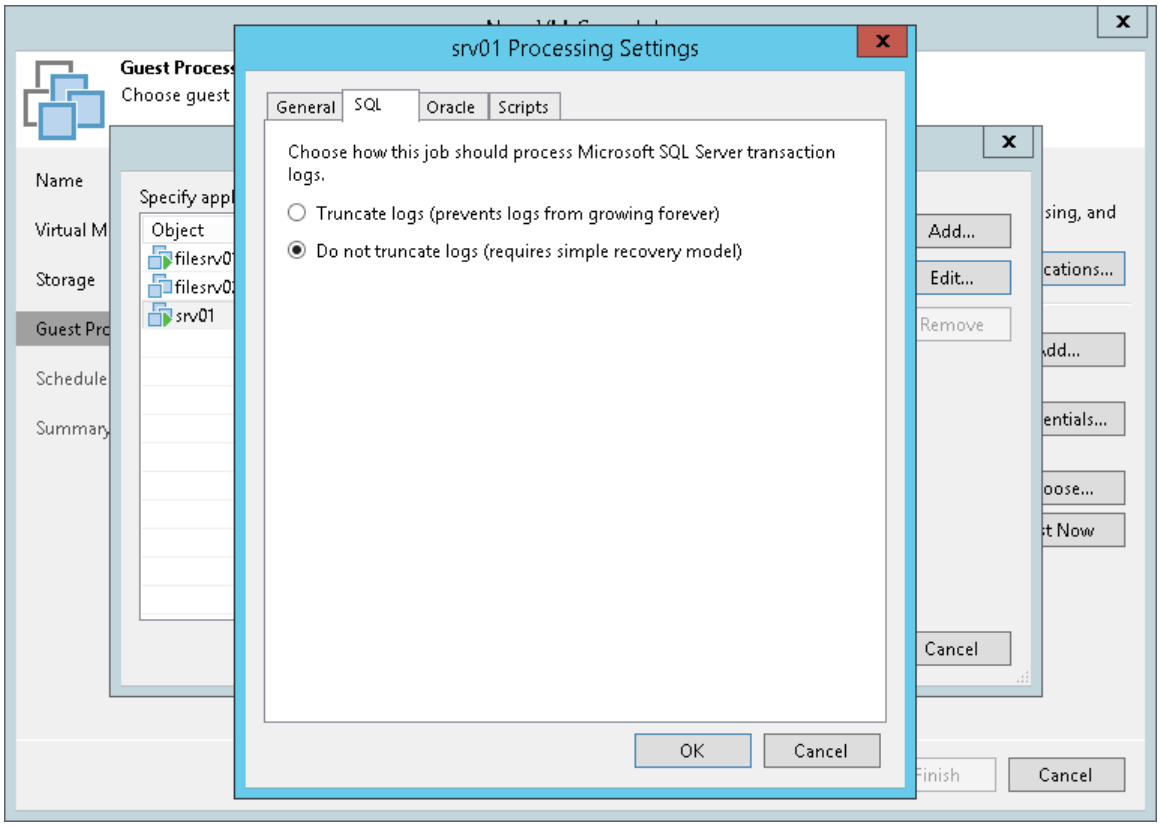


Microsoft SQL Server Transaction Log Settings

If you copy a Microsoft SQL VM, you can specify how Veeam Backup & Replication must process transaction logs:

1. At the **Guest Processing** step of the wizard, select the **Enable application-aware processing** check box.
2. Click **Applications**.
3. In the displayed list, select the Microsoft SQL Server VM and click **Edit**.
4. In the **Transaction logs** section, select **Process transaction logs with this job**.
5. In the **VM Processing Settings** window, click the **SQL** tab.
6. Specify how transaction logs must be processed:
   - Select **Truncate logs** if you want Veeam Backup & Replication to trigger truncation of transaction logs only after the job completes successfully. In this case, the non-persistent runtime components or persistent components will wait for the job to complete and then trigger truncation of transaction logs. If the VM copy job fails, the logs will remain untouched on the VM guest OS until the next start of the non-persistent runtime components or persistent components.

- Select **Do not truncate logs** if you do not want Veeam Backup & Replication to truncate logs at all. This option is recommended if you are using another backup tool to perform VM guest-level backup or replication, and this tool maintains consistency of the database state. In such scenario, Veeam Backup & Replication will not trigger transaction log truncation. After you fail over to the necessary restore point of the VM copy, you will be able to apply transaction logs to get the database system to the necessary point in time between VM copy job sessions.



Oracle Archived Log Settings

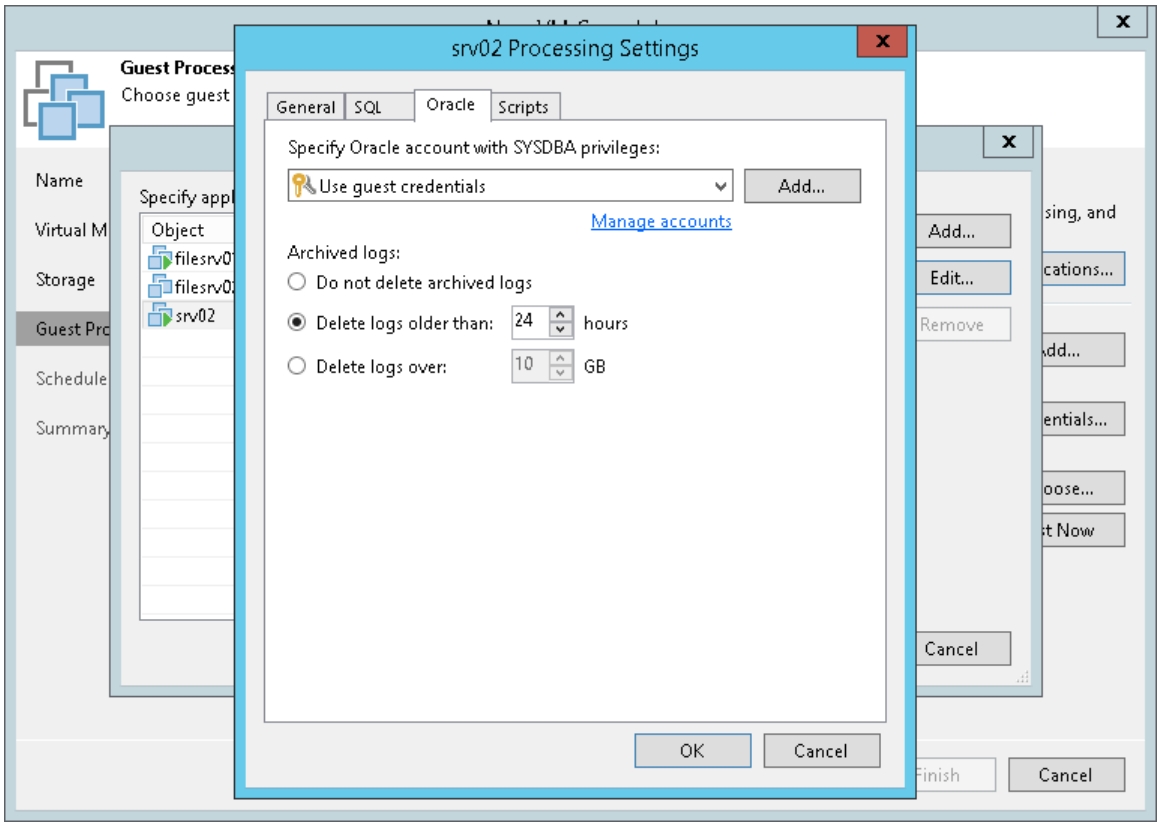
If you copy an Oracle VM, you can specify how Veeam Backup & Replication must process transaction logs:

1. At the **Guest Processing** step of the wizard, select the **Enable application-aware processing** check box.
2. Click **Applications**.
3. In the displayed list, select the Oracle VM and click **Edit**.
4. In the **Transaction logs** section, select **Process transaction logs with this job**.
5. In the **VM Processing Settings** window, click the **Oracle** tab.
6. In the **Specify Oracle account with SYSDBA privileges** section, specify a user account that Veeam Backup & Replication will use to connect to the Oracle database. The account must have SYSDBA rights on the Oracle database.

   You can select **Use guest credentials** in the list of user accounts. In this case, Veeam Backup & Replication will use the account specified at the **Guest Processing** step of the wizard to access the VM guest OS and connect to the Oracle database.

7. In the **Archived logs** section, specify if Veeam Backup & Replication must truncate transaction logs on the Oracle VM:
   - Select **Do not truncate archived logs** if you want Veeam Backup & Replication to preserve archived logs on the VM guest OS. When the VM copy job completes, the non-persistent runtime components or persistent components will not truncate transaction logs.

     It is recommended that you select this option for databases for which the ARCHIVELOG mode is turned off. If the ARCHIVELOG mode is turned on, transaction logs on the VM guest OS may grow large and consume all disk space. In this case, the database administrators must take care of transaction logs themselves.
   - Select **Truncate logs older than <N> hours** or **Truncate logs over <N> GB** if you want Veeam Backup & Replication to truncate archived logs that are older than <N> hours or larger than <N> GB. The non-persistent runtime components or persistent components running on the VM guest OS will wait for the VM copy job to complete successfully and then trigger transaction logs truncation using Oracle Call Interface (OCI). If the job does not manage to copy the Oracle VM, the logs will remain untouched on the VM guest OS until the next start of the non-persistent runtime components or persistent components.



Pre-Freeze and Post-Thaw Scripts

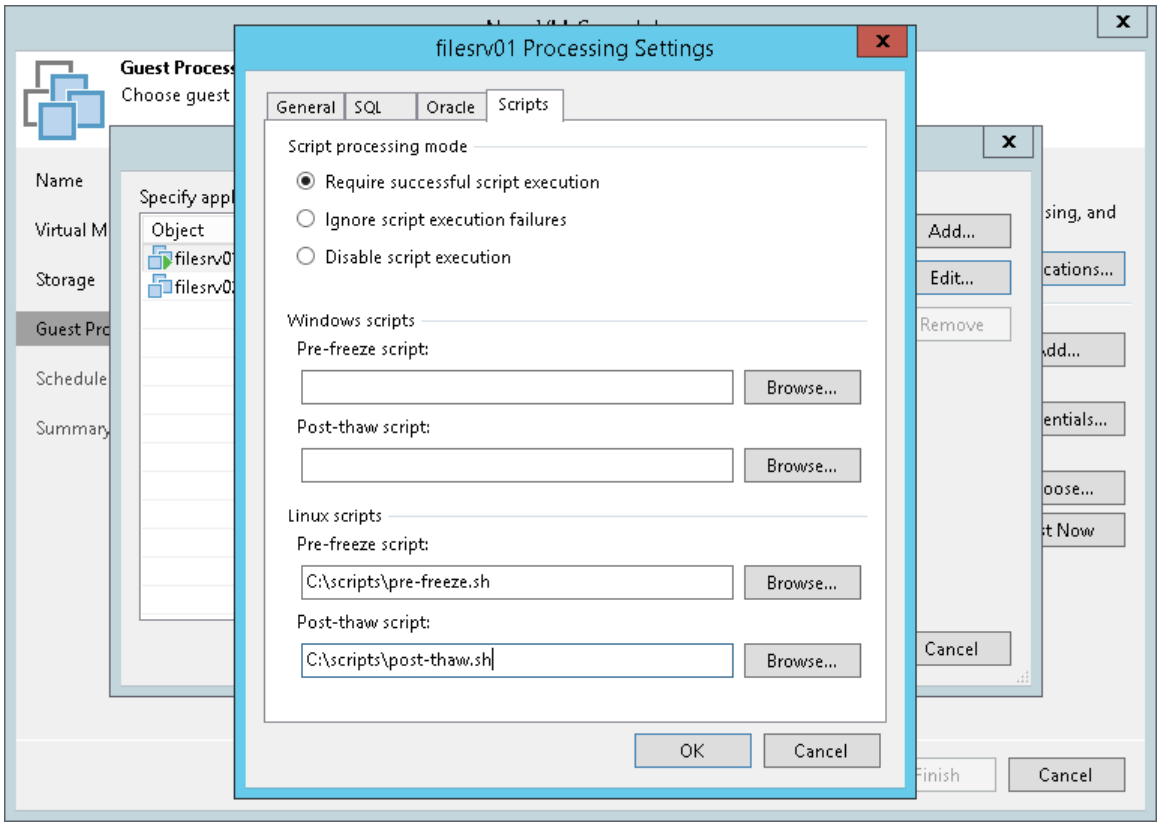
If you plan to copy VMs running applications that do not support VSS, you can instruct Veeam Backup & Replication to run custom pre-freeze and post-thaw scripts for these VMs. The pre-freeze script quiesces the VM file system and application data to bring the VM to a consistent state before Veeam Backup & Replication triggers a VM snapshot. After the VM snapshot is created, the post-thaw script brings the VM and applications to their initial state.

To specify pre-freeze and post-thaw scripts for the job:

1. At the **Guest Processing** step, click **Applications**.

2. In the displayed list, select the VM and click **Edit**.
3. Click the **Scripts** tab.
4. In the **Script processing mode** section, specify the scenario for scripts execution:
   - Select **Require successful script execution** if you want Veeam Backup & Replication to stop the VM copy process if the script fails.
   - Select **Ignore script execution failures** if you want to continue the VM copy process even if script errors occur.
   - Select **Disable script execution** if you do not want to run scripts for the VM.
5. In the **Windows scripts** section, specify paths to pre-freeze and post-thaw scripts for Microsoft Windows VMs. For the list of supported script formats, see Pre-Freeze and Post-Thaw Scripts.
6. In the **Linux scripts** section, specify paths to pre-freeze and/or post-thaw scripts for Linux VMs. For the list of supported script formats, see Pre-Freeze and Post-Thaw Scripts.

   If you have added to the job a VM container with Microsoft Windows and Linux VMs, you can select to execute both Microsoft Windows and Linux scripts for the VM container. When the job starts, Veeam Backup & Replication will automatically determine what OS type is installed on the VM and apply corresponding scripts to quiesce this VM.



Step 7. Define Job Schedule

At the **Schedule** step of the wizard, select to run the VM copy job manually or schedule the job to run on a regular basis.

To specify the job schedule:

1. Select the **Run the job automatically** check box. If this check box is not selected, you will have to start the job manually to perform VM replication.
2. Define scheduling settings for the job:
   - To run the job at specific time daily, on defined week days or with specific periodicity, select **Daily at this time**. Use the fields on the right to configure the necessary schedule.
   - To run the job once a month on specific days, select **Monthly at this time**. Use the fields on the right to configure the necessary schedule.

   > **NOTE**
   >
   > When you configure the job schedule, keep in mind possible date and time changes (for example, related to daylight saving time transition).

   - To run the job repeatedly throughout a day with a set time interval, select **Periodically every**. In the field on the right, select the necessary time unit: *Hours* or *Minutes*. Click **Schedule** and use the time table to define the permitted time window for the job. In the **Start time within an hour** field, specify the exact time when the job must start.

     A repeatedly run job is started by the following rules:

     - Veeam Backup & Replication always starts counting defined intervals from 12:00 AM. For example, if you configure to run a job with a 4-hour interval, the job will start at 12:00 AM, 4:00 AM, 8:00 AM, 12:00 PM, 4:00 PM and so on.
     - If you define permitted hours for the job, after the denied interval is over, Veeam Backup & Replication will immediately start the job and then run the job by the defined schedule.

     For example, you have configured a job to run with a 2-hour interval and defined permitted hours from 9:00 AM to 5:00 PM. According to the rules above, the job will first run at 9:00 AM, when the denied period is over. After that, the job will run at 10:00 AM, 12:00 PM, 2:00 PM and 4:00 PM.

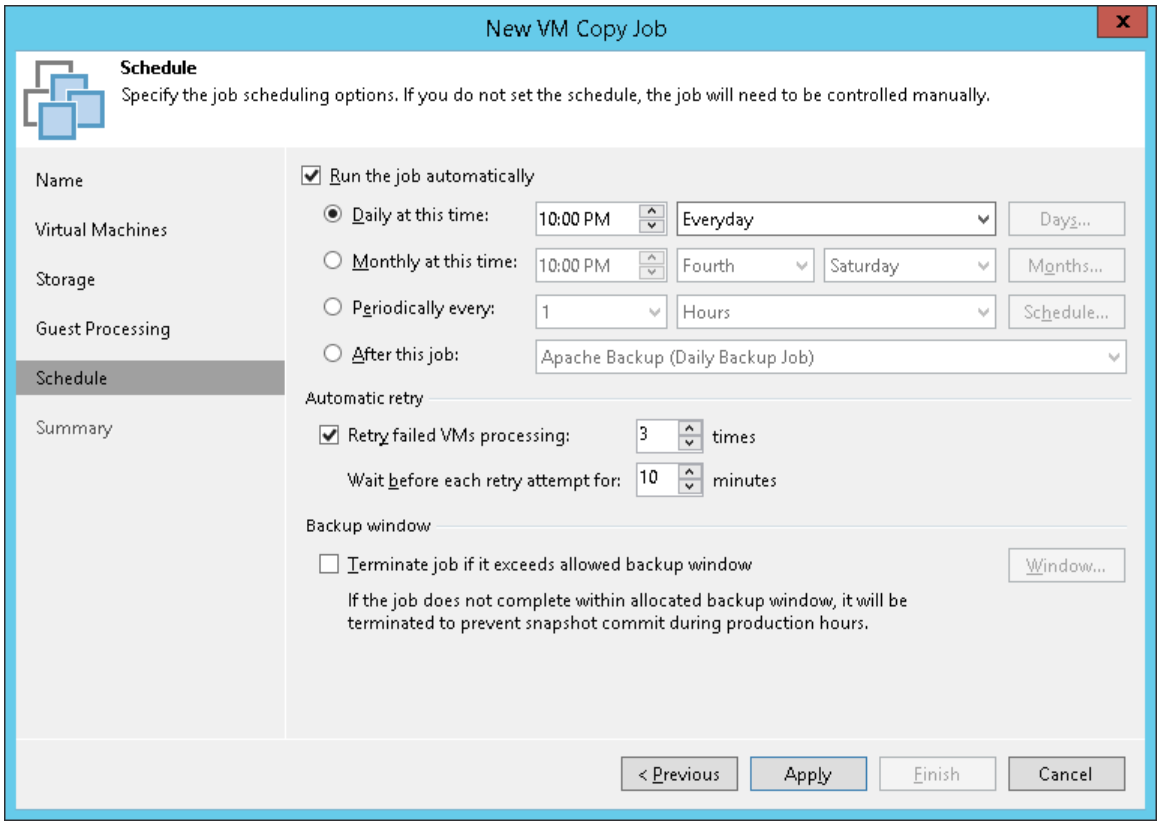
   - To run the job continuously, select the **Periodically every** option and choose **Continuously** from the list on the right.
   - To chain jobs, use the **After this job** field. In the common practice, jobs start one after another: when job *A* finishes, job *B* starts and so on. If you want to create a chain of jobs, you should define the time schedule for the first job in the chain. For the rest of the jobs in the chain, select the **After this job option** and choose the preceding job from the list.
3. In the **Automatic retry** section, define whether Veeam Backup & Replication must attempt to run the job again if the job fails for some reason. During a job retry, Veeam Backup & Replication processes failed VMs only. Enter the number of attempts to run the job and define time spans between them. If you select continuous schedule for the job, Veeam Backup & Replication will retry the job for the defined number of times without any time intervals between the job sessions.

4. In the **Backup window** section, determine a time interval within which the job must be completed. The backup window prevents the job from overlapping with production hours and ensures it does not provide unwanted overhead on your production environment. To set up a backup window for the job:

   a. Select the **Terminate job if it exceeds allowed backup window** check box and click **Window**.

   b. In the **Time Periods** window, define the allowed hours and prohibited hours for VM copying. If the job exceeds the allowed window, it will be automatically terminated.

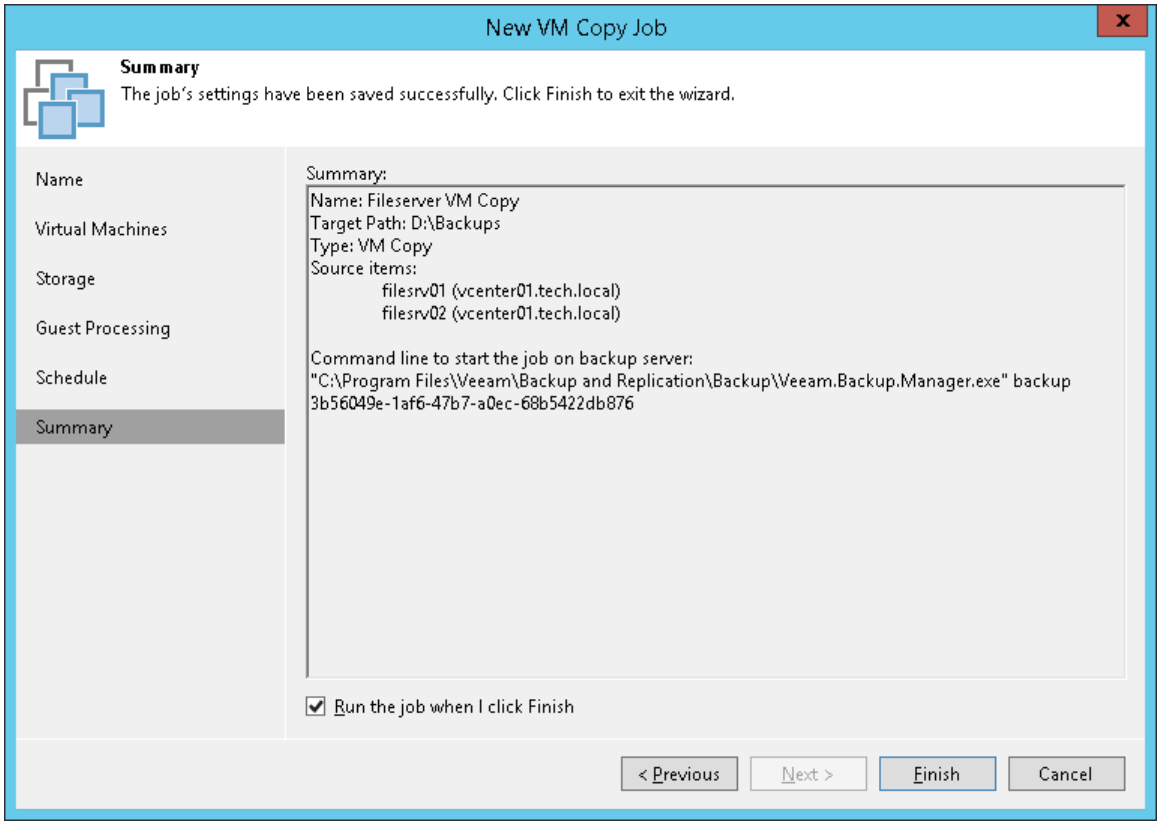
> **NOTE**
>
> The **After this job** function will only start a job if the first job in the chain is started automatically by schedule. If the first job is started manually, jobs chained to it will not be started.



Step 8. Finish Working with Wizard

At the **Summary** step of the wizard, complete the procedure of VM copy job configuration.

1. Review details of the VM copy job.
2. Select the **Run the job when I click Finish** check box if you want to start the job right after you finish working with the wizard.
3. Click **Finish** to close the wizard.



File Copy

You can copy and move files and folders between servers and hosts added to the backup infrastructure. For file copying operations, Veeam Backup & Replication offers a Windows Explorer-like user interface familiar to a Microsoft Windows user. You can copy files manually or schedule file copy jobs to run automatically by the defined schedule.

The file copy functionality is not intended for creating backups of VM guest OS files. Use backup jobs to create VM image-level backups instead.

Creating File Copy Jobs

To schedule a copying process for files and folders, you must configure a file copy job. You can run the file copy job immediately after its creation, schedule or save the job.

File copy jobs let you copy files between the following backup infrastructure objects:

- Virtualization hosts
- Microsoft Windows servers
- Linux servers
- ExaGrid storage appliances used as backup repositories

Before you configure a file copy job, check prerequisites. Then use the **New File Copy Job** wizard to create a job.

Before You Begin

Before you configure a file copy job, check the following prerequisites:

Backup infrastructure components that will take part in the file copying process must be added to the backup infrastructure and properly configured. These include a source and target host or server between which files and folders will be copied.

Mind the following limitations:

- File copy is not supported for Unix systems, for example, Solaris, FreeBSD and AIX.
- Veeam Backup & Replication does not preserve the Access Control List (ACL) settings for copied guest OS folders. The ACL settings are preserved for files only.

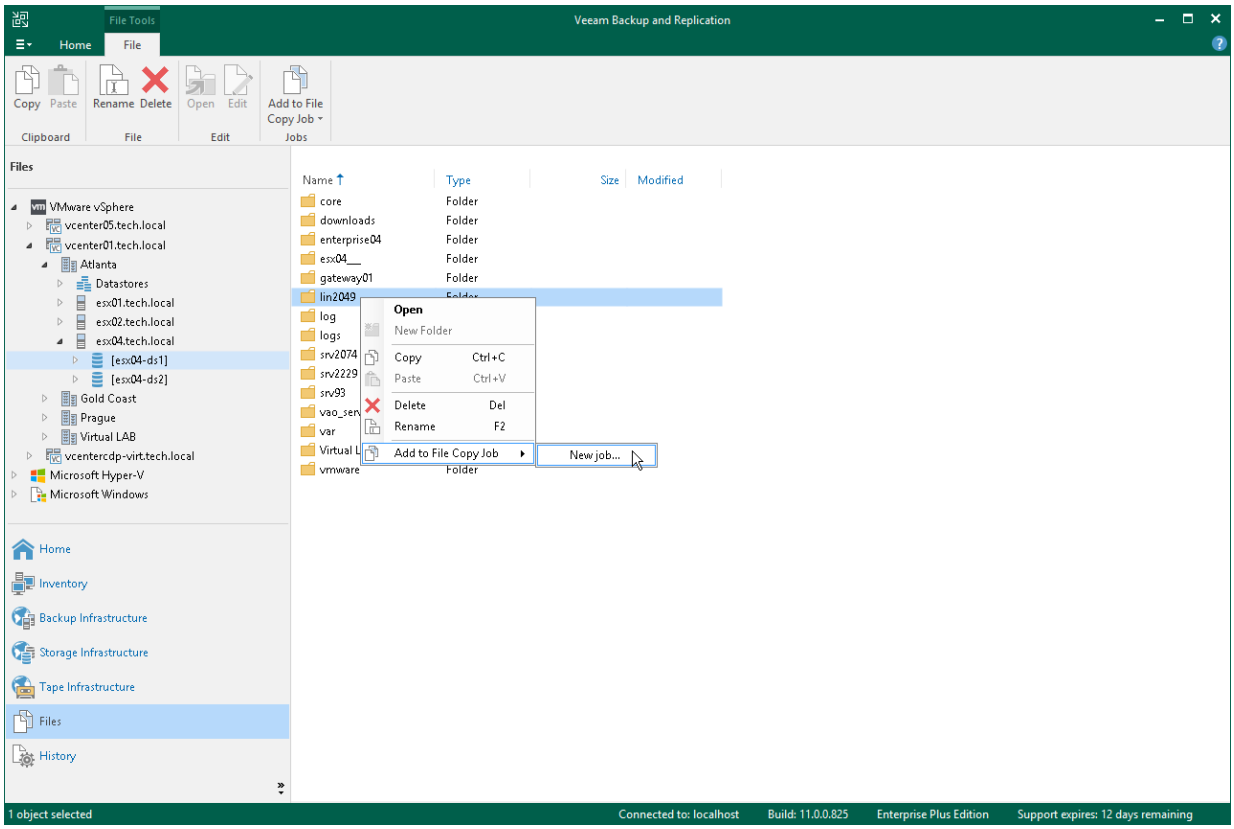
> **TIP**
>
> You can restore the ACL settings for recovered guest OS files and folders using Instant File-Level Restore.

Step 1. Launch New File Copy Job Wizard

To launch the **New File Copy Job** wizard, do either of the following:

- On the **Home** tab, click **Copy Job > File**.
- Open the **Files** view, in the working area right-click the necessary files and folders and select **Add to File Copy Job > New job**. Veeam Backup & Replication will start the **New File Copy Job** wizard and add selected files and folders to this job. You can add other files and folders to the job later on, when you pass through the wizard steps.

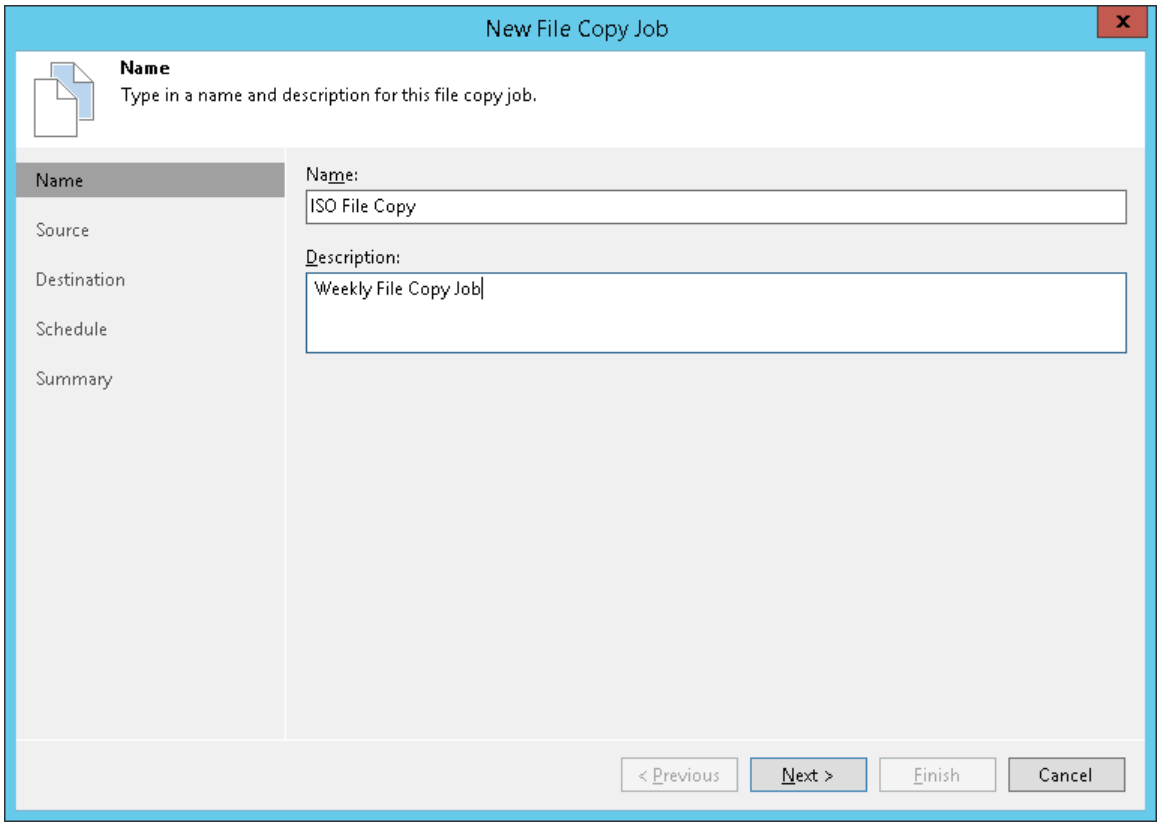
You can add files and folders to already existing jobs. To do this, open the **Files** view, in the working area right-click necessary objects and select **Add to file copy job > name of the job**.



Step 2. Specify Job Name and Description

At the **Name** step of the wizard, enter a name and description of the created job.

1. In the **Name** field, enter a name for the file copy job.
2. In the **Description** field, provide a description for future reference. The default description contains information about the user who created a job, date and time when the job was created.



Step 3. Select Files and Folders to Be Copied

At the **Source** step of the wizard, select files and folders that you want to copy.

You can use the following sources for the file copy job:

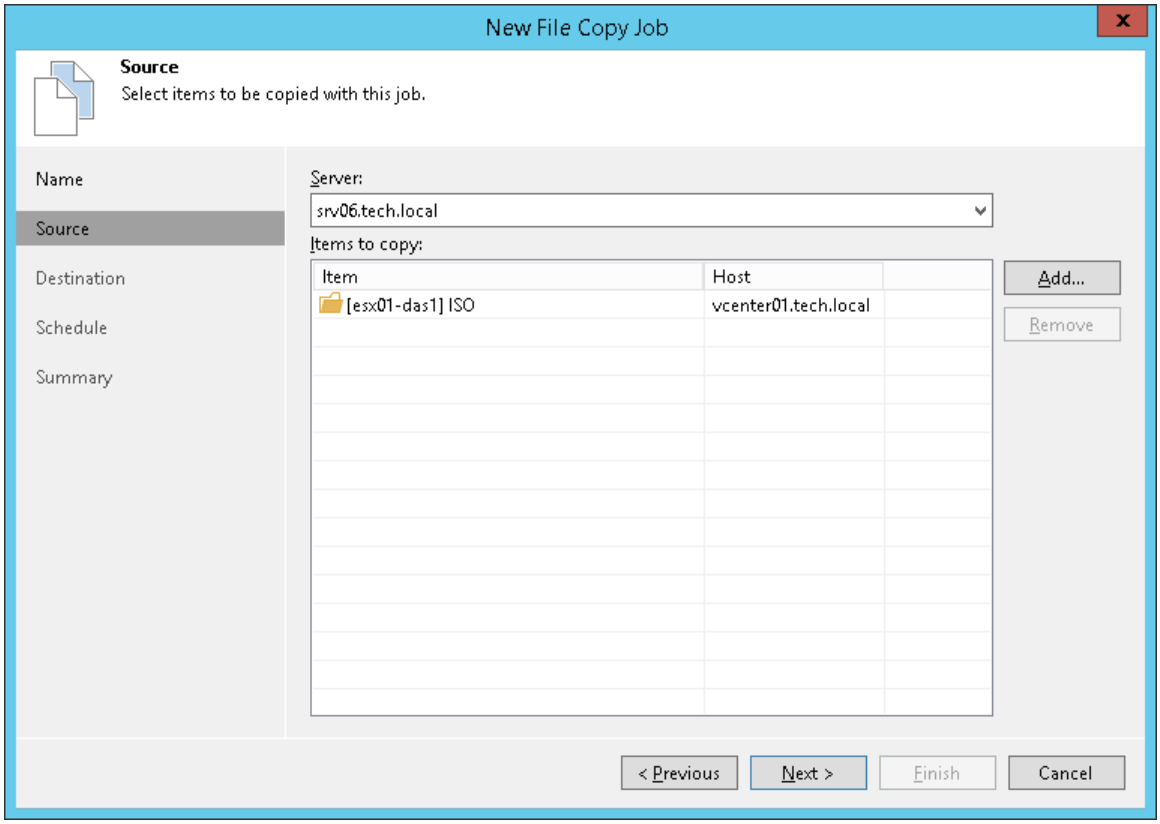
- Virtualization hosts
- Microsoft Windows servers
- Linux servers
- ExaGrid storage appliances used as backup repositories

To select files and folders that you want to copy:

1. From the **Host** list, choose a host or server on which files or folders that you want to copy reside.
2. Click **Add** and select files or folders that must be copied. The selected items will be added to the list.

> **IMPORTANT**
>
> If the list contains files/folders with the same names and extensions, Veeam Backup & Replication copies only one instance of a file/folder. This limitation applies even if you add files/folders from different hosts or servers. To avoid this limitation, you can rename files/folders on the source or add parent folders to the list.

To remove a file or folder from the list, select it and click **Remove**.



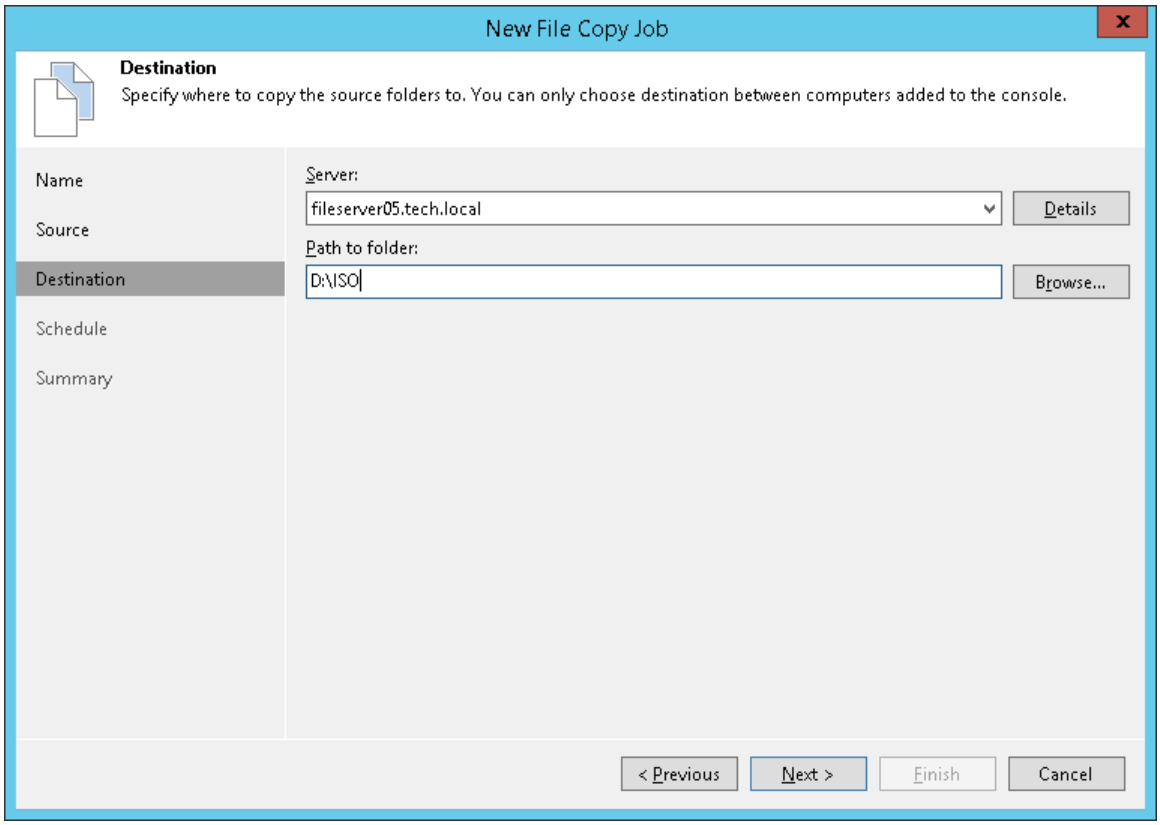
Step 4. Select Destination for Copying

At the **Destination** step of the wizard, select a destination to which files or folders must be copied.

1. From the **Server** list, select the target host or server.
2. Click **Details** on the right of the **Server** field to view or edit server properties.
3. Click **Browse** next to the **Path to folder** field and select a folder where copied items must be stored. To create a dedicated folder for copied files or folders, use the **New Folder** button at the bottom of the **Select Folder** window.

IMPORTANT

If the target folder already contains files/folders with the same names and extensions as the files/folder that must be copied, Veeam Backup & Replication will replace files/folders in the target folder with new files/folders.



Step 5. Define Job Schedule

At the **Schedule** step of the wizard, you can select to run the file copy job manually or schedule the job to run on a regular basis.

To specify the job schedule:

1. Select the **Run the job automatically** check box. If this check box is not selected, you will have to start the job manually to copy files or folders.
2. Define scheduling settings for the job:
   - To run the job at specific time daily, on defined week days or with specific periodicity, select **Daily at this time**. Use the fields on the right to configure the necessary schedule.
   - To run the job once a month on specific days, select **Monthly at this time**. Use the fields on the right to configure the necessary schedule.

     NOTE

     When you configure the job schedule, keep in mind possible date and time changes (for example, related to daylight saving time transition).

   - To run the job repeatedly throughout a day with a set time interval, select **Periodically every**. In the field on the right, select the necessary time unit: *Hours* or *Minutes*. Click **Schedule** and use the time table to define the permitted time window for the job. In the **Start time within an hour** field, specify the exact time when the job must start.

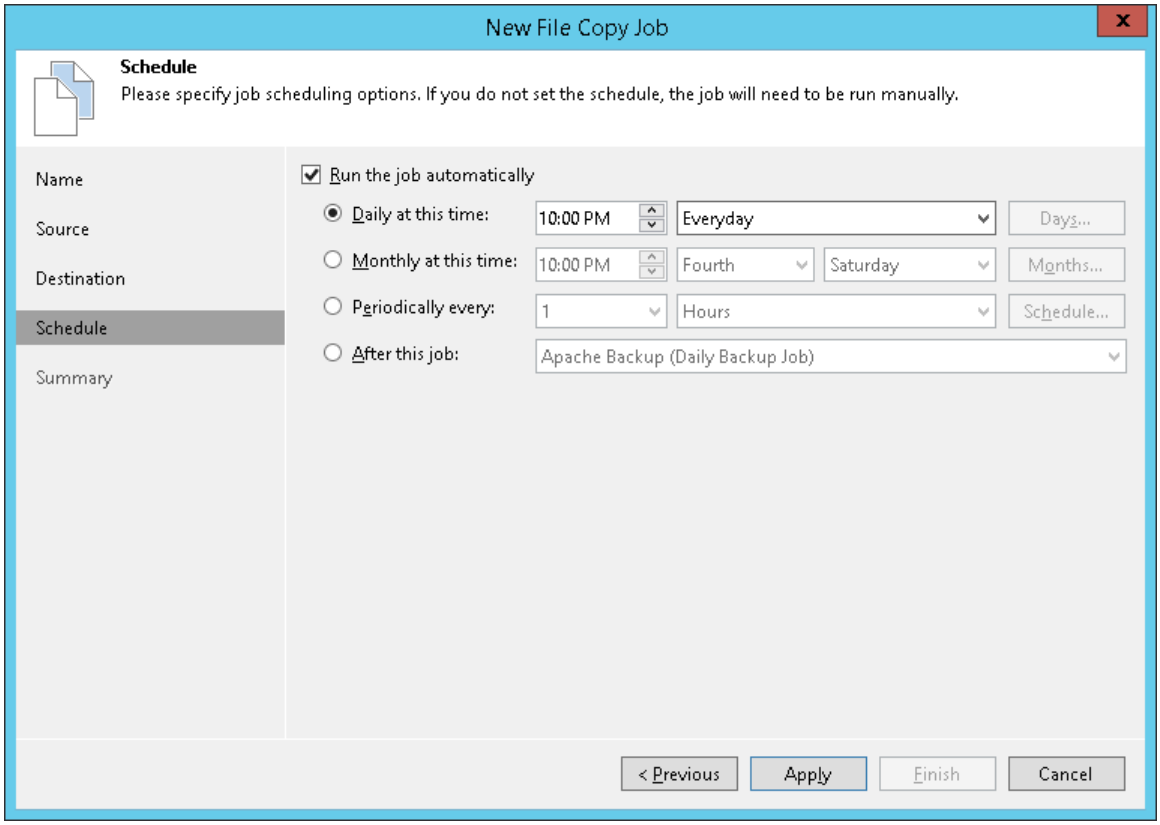
     A repeatedly run job is started by the following rules:

     - Veeam Backup & Replication always starts counting defined intervals from 12:00 AM. For example, if you configure to run a job with a 4-hour interval, the job will start at 12:00 AM, 4:00 AM, 8:00 AM, 12:00 PM, 4:00 PM and so on.
     - If you define permitted hours for the job, after the denied interval is over, Veeam Backup & Replication will immediately start the job and then run the job by the defined schedule.

     For example, you have configured a job to run with a 2-hour interval and defined permitted hours from 9:00 AM to 5:00 PM. According to the rules above, the job will first run at 9:00 AM, when the denied period is over. After that, the job will run at 10:00 AM, 12:00 PM, 2:00 PM and 4:00 PM.

   - To run the job continuously, select the **Periodically every** option and choose **Continuously** from the list on the right.

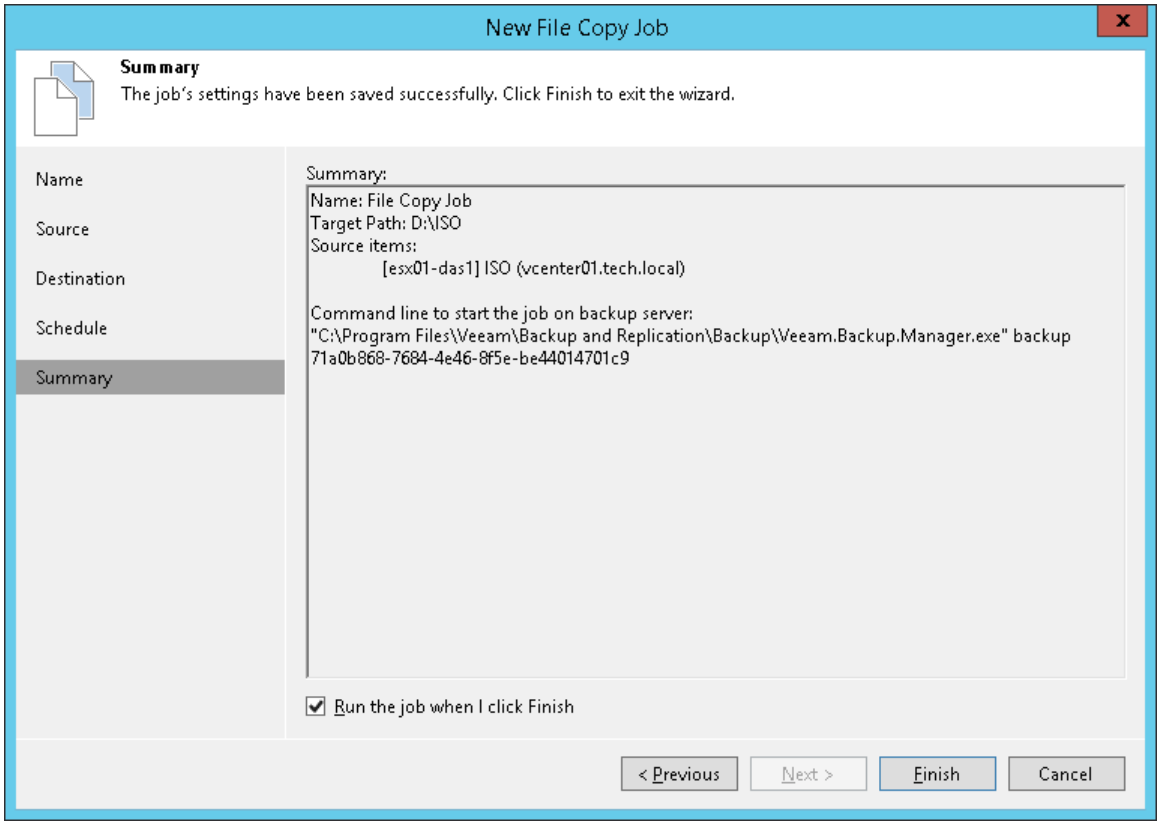
- To chain jobs, use the **After this job** field. In the common practice, jobs start one after another: when job *A* finishes, job *B* starts and so on. If you want to create a chain of jobs, you should define the time schedule for the first job in the chain. For the rest of the jobs in the chain, at the **Schedule** step of the wizard, select the **After this job option** and choose the preceding job from the list.



Step 6. Finish Working with Wizard

At the **Summary** step of the wizard, complete the procedure of file copy job configuration.

1. Review details for the created file copy job.
2. Select the **Run the job when I click Finish** check box if you want to start the job right after you finish working with the wizard.
3. Click **Finish** to close the wizard.



Copying Files and Folders Manually

You can manually copy and move files and folders between servers and hosts added to the backup infrastructure.

Veeam Backup & Replication lets you copy files manually between the following backup infrastructure objects:

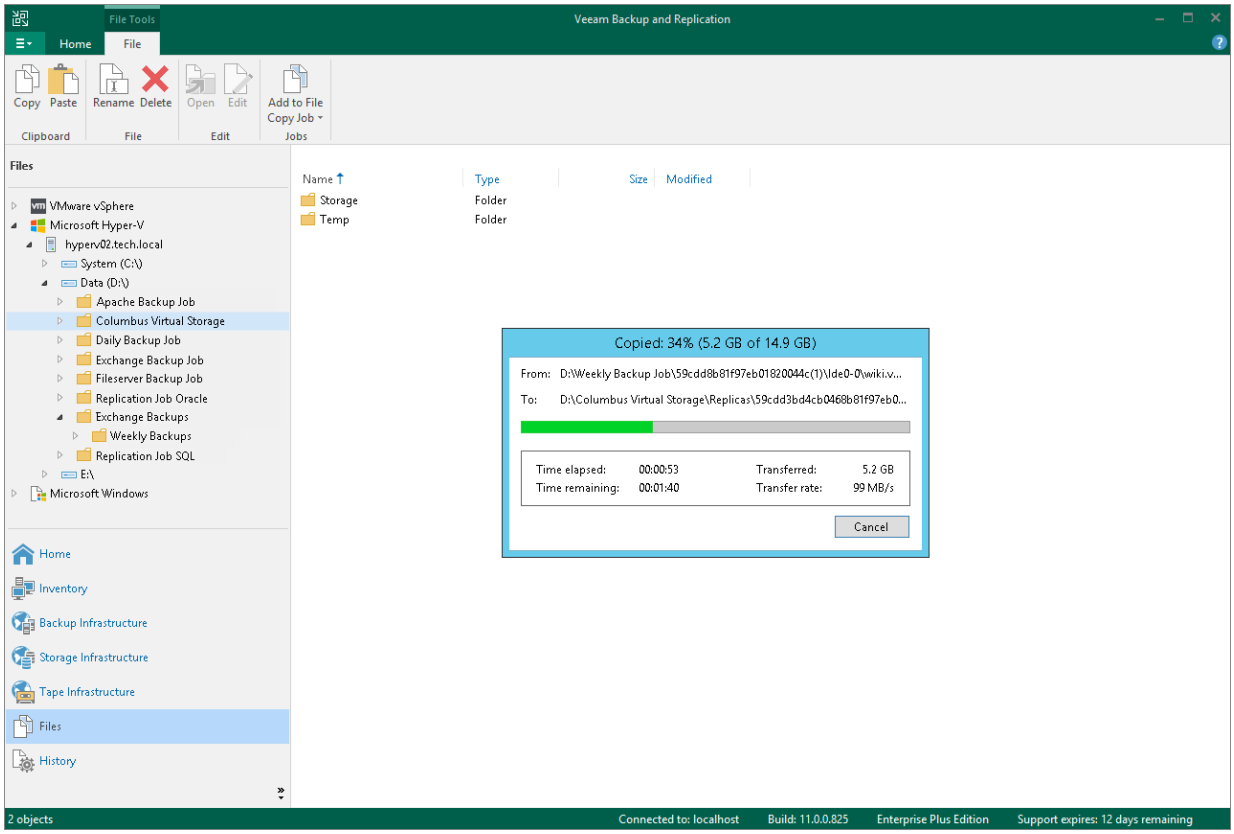
- Virtualization hosts
- Microsoft Windows servers
- Linux servers
- Deduplicating storage appliances used as backup repositories

> **IMPORTANT**
>
> You cannot copy backup files (VBK, VIB and VRB) to HPE StoreOnce storage appliances used as backup repositories. To copy such files, use backup copy jobs.

To copy files and folders:

1. Open the **Files** view.
2. In the inventory pane, expand the file tree of the source server or host.
3. Right-click files and folders that you want to copy and select **Copy**.
4. In the inventory pane, expand the file tree of the target server or host.
5. Right-click a destination folder and select **Paste**.

You can also use a drag-n-drop operation to copy files and folders between the source and target hosts or servers.



Managing Folders

You can create, rename and delete folders in the **Files** view of Veeam Backup & Replication.

To create a folder:

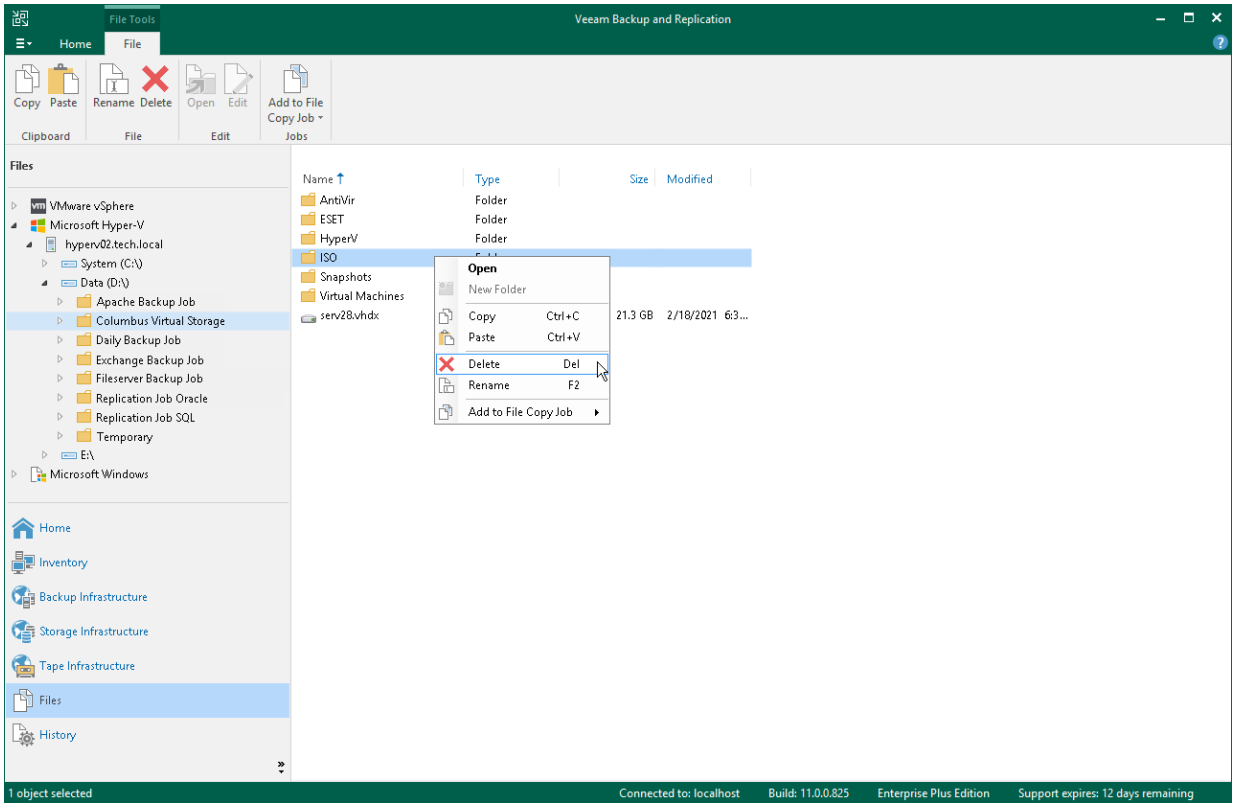
1. Open the **Files** view.
2. In the inventory pane, expand the file tree of the necessary server or host.
3. In the working area, right-click anywhere on the blank area and select **New Folder**.

To rename a folder:

1. Open the **Files** view.
2. In the inventory pane, expand the file tree of the necessary server or host.
3. In the working area, select the folder and click **Rename** on the ribbon or right-click the folder and select **Rename**.
4. Enter a new name for the folder and press [Enter] on the keyboard.

To remove a folder:

1. Open the **Files** view.
2. In the inventory pane, expand the file tree of the necessary server or host.
3. In the working area, select the folder and click **Delete** on the ribbon or right-click the folder and select **Delete**.



Editing and Deleting Files

You can edit files and delete them in the **Files** view of Veeam Backup & Replication. For example, you may want to edit a configuration file of the VM (VMX) or need to delete from the storage files of unused VMs.

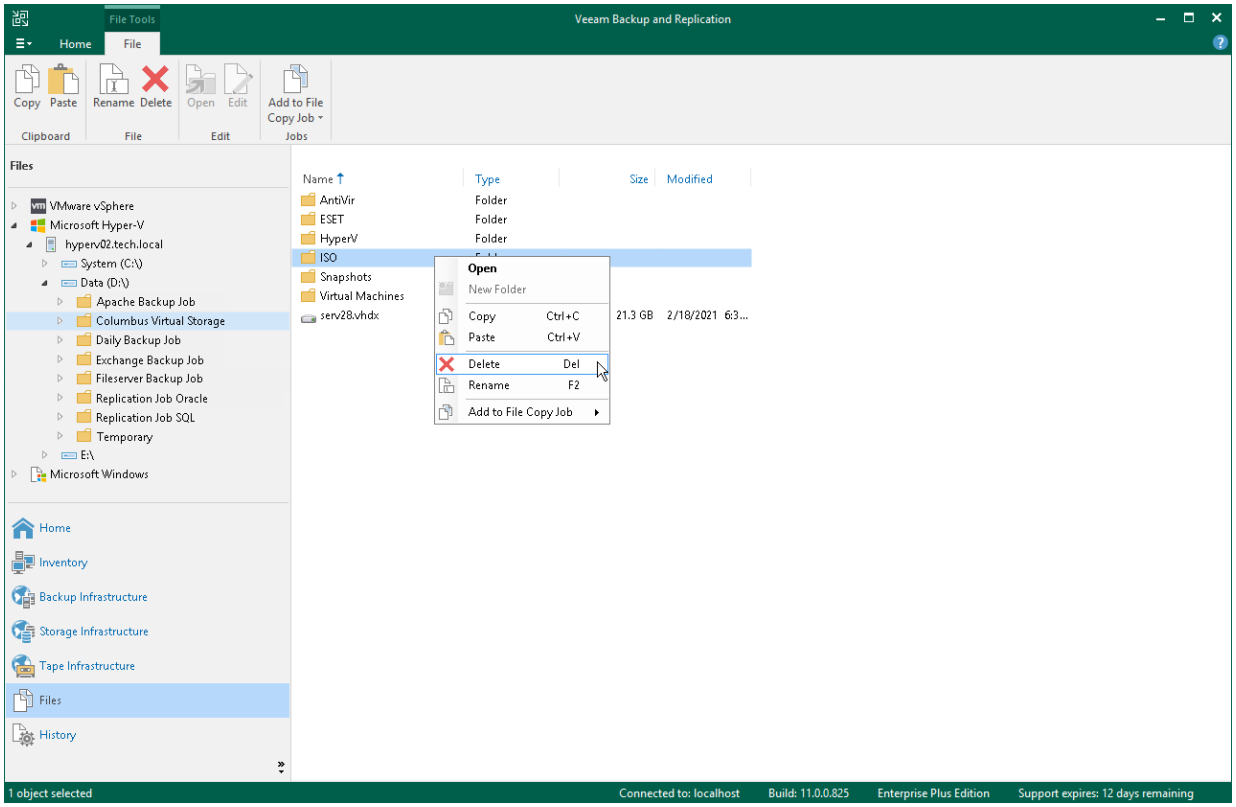
To edit a file:

1. Open the **Files** view.
2. In the inventory pane, expand the file tree of the necessary server or host.
3. In the working area, select the file and click **Edit** on the ribbon or right-click the folder and select **Edit**.
4. Veeam Backup & Replication will open the selected file in the editor. Edit the file as required and click **Save** on the file editor toolbar or press [Ctrl+S] on the keyboard.

To delete a file:

1. Open the **Files** view.
2. In the inventory pane, expand the file tree of the necessary server or host.
3. In the working area, select the file and click **Delete** on the ribbon or right-click the folder and select **Delete**.

> **NOTE**
>
> To delete a folder on VSAN, you must remove a real folder, not a symbolic link to this folder. The real folder is named with GUID, for example, *c07a2953-8096-5b20-a11a-002590c5857c*, while the symbolic link contains the folder name, for example, *srv02_vm*. If you delete the folder symbolic link, the delete operation will fail, and the folder will not be removed.



Quick Migration

Quick Migration allows you to migrate VMs or virtual disks between ESXi hosts and datastores.
Veeam Backup & Replication supports migration of VMs or their disks in any state with minimum disruption to business operations. You can use Quick Migration as a self-contained capability or as a way to finalize the Instant Recovery and Instant Disk Recovery processes.

When you perform Quick Migration, Veeam Backup & Replication analyzes your virtual environment, its configuration, the state of VMs and selects the most appropriate VM relocation method:

- **vMotion and Storage vMotion**

  vMotion and Storage vMotion are native migration mechanisms of VMware vCenter.
  Veeam Backup & Replication uses these methods whenever it is possible.

- **Veeam Quick Migration**

  Veeam Quick Migration is the Veeam Backup & Replication proprietary technology.
  Veeam Backup & Replication uses this method when VMware vCenter methods cannot be used. For example, if your VMware vSphere license does not provide support for vMotion and Storage vMotion, or you need to migrate VMs from one standalone ESXi host to another.

  Veeam Quick Migration supports two modes of VM migration:

  - **SmartSwitch**

    With SmartSwitch, Veeam Backup & Replication suspends a VM, then moves the VM configuration file and copies changes made to the VM disk after snapshot creation to the target host. After the migration is completed, the VM is resumed on the target host.

  - **ColdMigration**

    With ColdMigration, Veeam Backup & Replication stops the VM, then copies changes made to the VM disk after snapshot creation to the new host. After, the VM is started on the target host.

Veeam Quick Migration of VMs

Migration of a VM using the Veeam Quick Migration method includes the following steps:

1. Veeam Backup & Replication copies VM configuration file (.VMX) to the target host and registers the VM.
2. Veeam Backup & Replication triggers a VM snapshot creation and copies VM disk content to the new destination.
3. Veeam Backup & Replication uses different modes when moving the VM between hosts with compatible and non-compatible CPUs.
   - If you move a VM between two hosts with compatible CPUs, Veeam Backup & Replication uses the SmartSwitch mode.
   - If you move a VM between two hosts with non-compatible CPUs or VM RAM is more than 8 GB, Veeam Backup & Replication uses the ColdMigration mode.

Veeam Quick Migration of Virtual Disks

Migration of a VM disk using the Veeam Quick Migration method includes the following steps:

1. A temporary VM is created on the target datastore.

2. Veeam Backup & Replication copies the disk from the original datastore to the temporary VM on the target datastore.
3. If the original VM is powered on, Veeam Backup & Replication suspends it.
4. If changes were made to the original disk during the copy process performed at step 3, Veeam Backup & Replication copies these changes to the disk of the temporary VM.
5. On the original VM, Veeam Backup & Replication replaces old path to the disk with the path to the disk of the temporary VM.
6. If the original VM was suspended during the migration, Veeam Backup & Replication powers the VM on.

IMPORTANT

Before the migration, VM disks must be recovered using Instant Disk Recovery.

Veeam Quick Migration of First Class Disks (FCDs)

Migration of First Class Disks (FCDs) supports only the vMotion method and includes the following steps:

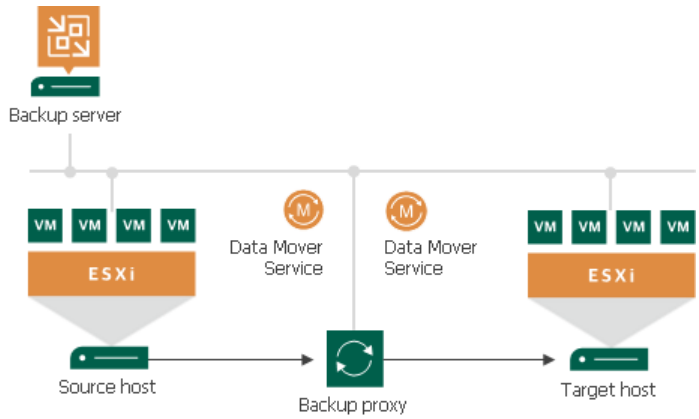
1. Veeam Backup & Replication gets information about the target datastore.
2. Veeam Backup & Replication connects to one of the ESXi hosts of a cluster to which the target datastore is mounted.
3. Veeam Backup & Replication migrates FCDs from the source datastore to the target datastore.
4. If redo logs are stored on a custom datastore, Veeam Backup & Replication will delete FCDs snapshots after migration.
5. Veeam Backup & Replication assigns storage policies to FCDs.

Quick Migration Architecture

Quick Migration architecture in a VMware vSphere environment comprises the following components:

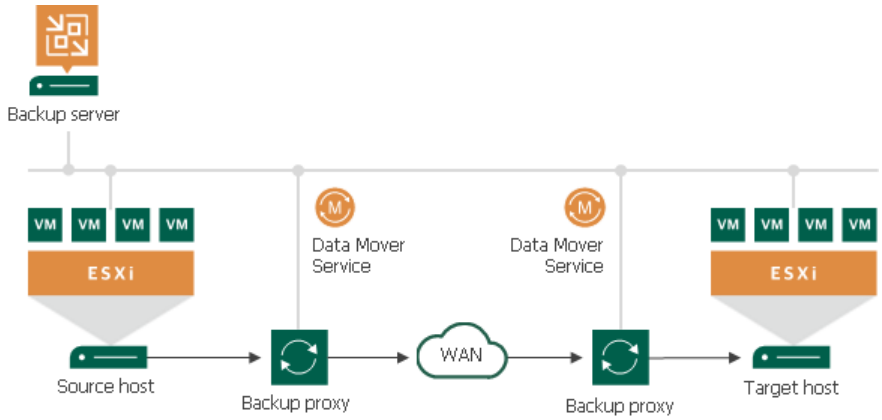
- Source host and target host with associated datastores
- One or two backup proxy servers

Similar to the backup architecture, Quick Migration uses two-service architecture: the source-side Veeam Data Mover interacts with the source host, and the target-side Veeam Data Mover interacts with the target host. To perform on-site migration, you can deploy one backup proxy for data processing and transfer. This backup proxy must have access to the source host and to the target host at the same time. In this scenario, the source-side Veeam Data Mover and the target-side Veeam Data Mover are started on the same backup proxy.



The common requirement for off-site migration is that one Veeam Data Mover runs in the production site (closer to the source host and datastore), and the other Veeam Data Mover runs in the remote target site (closer to the target host and datastore). During backup, Veeam Data Movers maintain a stable connection, which allows for uninterrupted operation over WAN or slow links.

For off-site migration, you need to deploy at least one local backup proxy in each site: a source backup proxy in the production site, and a target backup proxy in the remote target site.



Migrating VMs

You can relocate one or more VMs with quick migration. Quick migration can be used to move VMs from one ESXi host to another one. You can perform "hot" quick migration for running VMs or "cold" quick migration for VMs that are powered off.

Quick migration is not job-driven: it cannot be saved as a job or scheduled to run later.
Veeam Backup & Replication will start relocating VMs immediately after you finish working with the **Quick Migration** wizard.

Before you start quick migration, check prerequisites. Then use the **Quick Migration** wizard to migrate VMs.

Before You Begin

Before you perform quick migration, check the following prerequisites and limitations:

- Backup infrastructure components that will take part in quick migration must be added to the backup infrastructure and properly configured. These include the source and target ESXi hosts.
- The target datastore must have enough free space to store disks of the migrated VMs. To receive alerts about low space on the target datastore, configure global notification settings. For more information, see Specifying Other Notification Settings.
- If you want to use VMware vSphere vMotion to relocate VMs between hosts and/or VMware vSphere Storage vMotion to relocate VM disks between datastores, make sure that you have a VMware vSphere license covering these features.
- If you use tags to categorize virtual infrastructure objects, check limitations for VM tags. For more information, see VM Tags.

Encryption

Veeam Backup & Replication does not keep encryption settings if a VM is migrated with VMware vMotion. After the migration process is finished, you will need to enable encryption for the migrated VM manually.

Integration with Instant Recovery

When you restore a VM using Instant Recovery, Veeam Backup & Replication starts the VM directly from a compressed and deduplicated backup file. To finalize recovery of a VM, you still need to move it to a new location. Moving the VM with VMware Storage vMotion or hot replication may require a lot of time and resources, or it may cause loss of valuable data.

Veeam Quick Migration was designed to complement Instant Recovery. Instead of pulling data from vPower NFS datastore, Quick Migration registers the VM on the target host, restores the VM contents from the backup file located in the backup repository and synchronizes the VM restored from backup with the running VM.

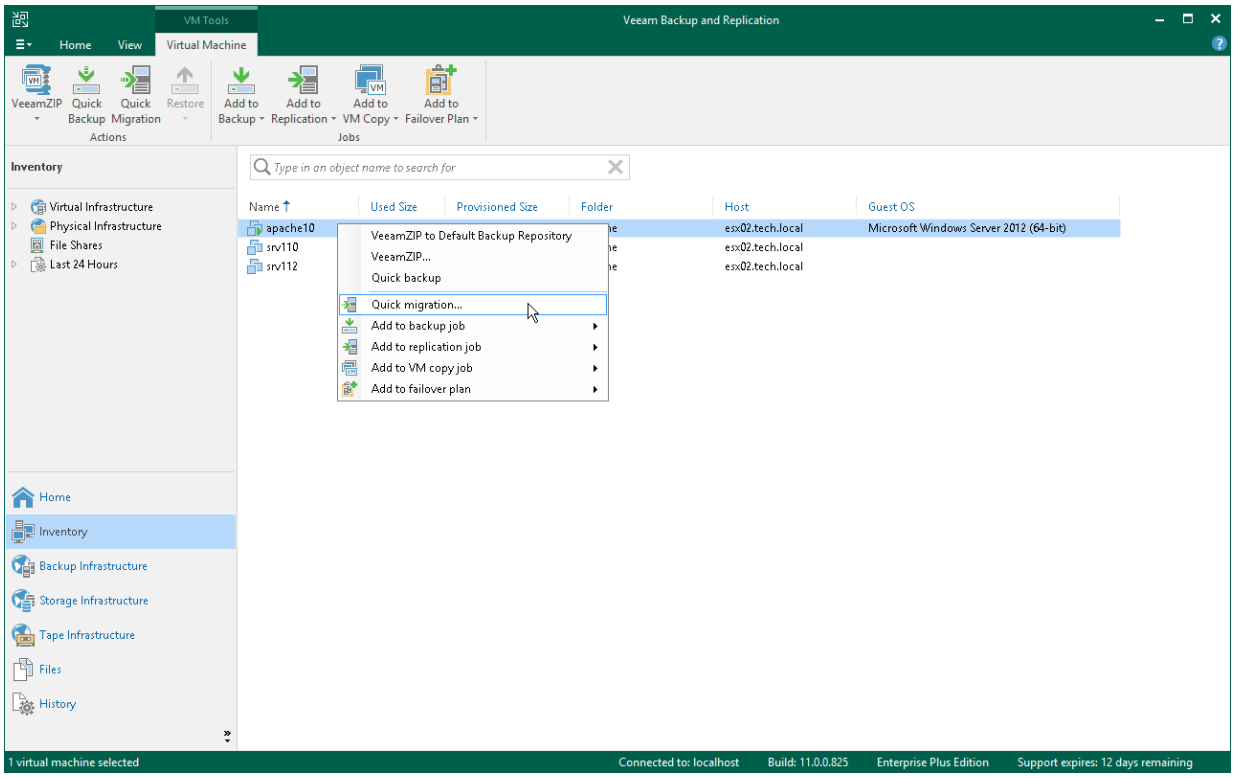
For more information, see Instant Recovery to VMware vSphere.

Step 1. Launch Quick Migration Wizard

If you use Quick Migration as a way to finalize the Instant Recovery or Instant Disk Recovery processes, you must launch Quick Migration as described in Finalizing Instant Recovery to VMware vSphere and Finalizing Instant Disk Recovery.

To launch Quick Migration as a self-contained capability:

1. Open the **Inventory** view.
2. In the infrastructure tree, select a host or VM container (host, cluster, folder, resource pool, VirtualApp, datastore or tag) in which the VMs that you want to relocate reside.
3. In the working area, select the VM and click **Quick Migration** on the ribbon or right-click the VMs and select **Quick Migration**.



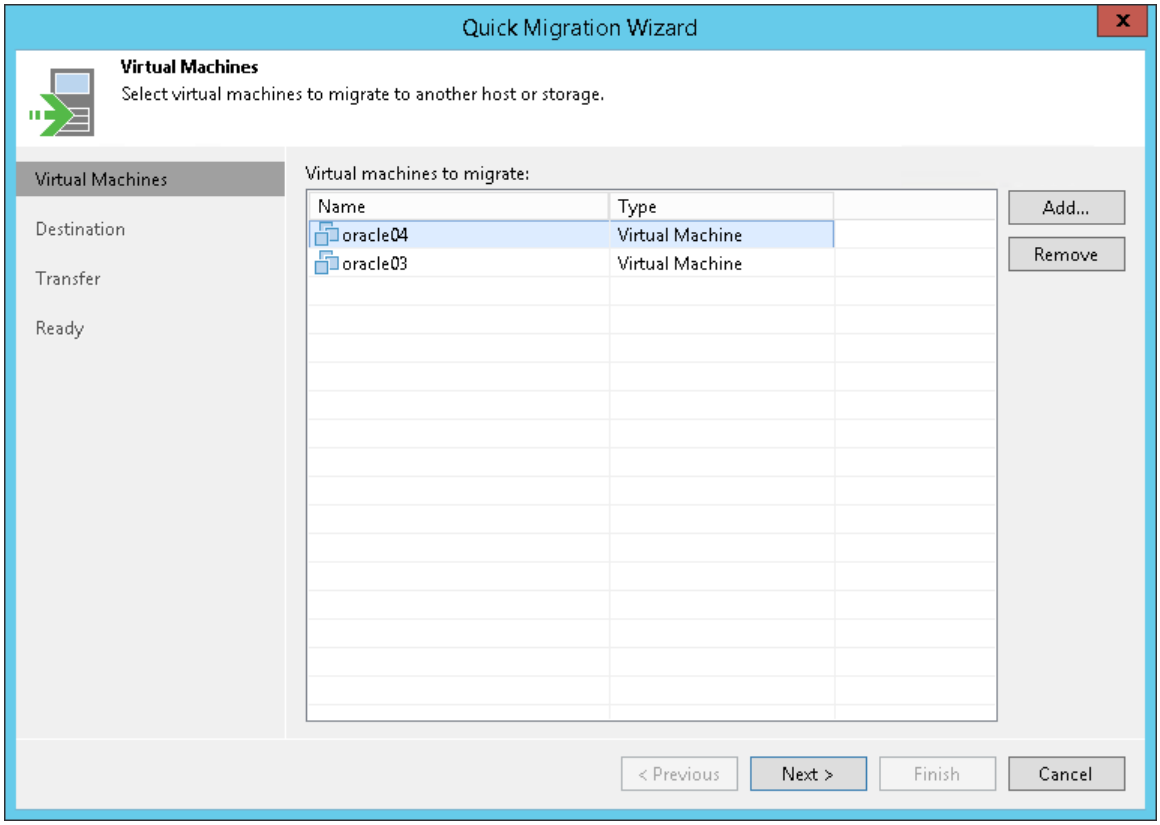
Step 2. Select VMs to Relocate

At the **Virtual Machines** step of the wizard, select the VMs and VM containers that you want to relocate:

1. Click **Add**.
2. Use the toolbar at the top right corner of the window to switch between views: **Hosts and Clusters**, **VMs and Templates**, **Datastores and VMs** and **Tags**. Depending on the view you select, some objects may not be available. For example, if you select the **VMs and Templates** view, no resource pools, hosts or clusters will be displayed in the tree.
3. Select the necessary object and click **Add**.

To quickly find the necessary object, you can use the search field at the bottom of the **Add Objects** window.

1. Click the button to the left of the search field and select the necessary type of object to search for: *Everything*, *Folder*, *Cluster*, *Host*, *Resource pool*, *VirtualApp* or *Virtual machine*.
2. Enter the object name or a part of it in the search field.
3. Click the **Start search** button on the right or press [Enter] on the keyboard.



Step 3. Specify VM Destination

At the **Destination** step of the wizard, select the destination to which the selected VMs must be relocated:

1. Click **Choose** next to the **Host or cluster** field and select an ESXi host or cluster where the relocated VM must be registered.
2. If all or majority of relocated VMs must belong to the same resource pool, click **Choose** next to the **Resource pool** field and select the target resource pool.

   If you want to place relocated VMs to different resource pools:

   a. Click the **Pick resource pool for selected VMs** link.

   b. In the **Choose Resource Pool** window, click **Add VM** on the right and select the VMs.

   c. Select the added VM in the **VM resource pool** list and click **Resource Pool** at the bottom of the window.

   d. From the list of available resource pools, select the target resource pool.
3. If all or majority of relocated VMs must be placed to the same folder, click **Choose** and select the folder.

   If you want to place relocated VMs to different folders:

   a. Click the **Pick VM folder for selected VMs** link.

   b. In the **Choose Folder** window, click **Add VM** on the right and select the VMs.

   c. Select the added VM in the **VM folder** list and click **VM Folder** at the bottom of the window.

   d. From the list of available folders, select the target folder.

   The **VM folder** section is disabled if you selected a standalone ESXi host as a target for VM relocation.
4. If all or majority of relocated VMs must be stored on the same datastore, click **Choose** and select the datastore. Veeam Backup & Replication displays only those datastores that are accessible by the selected ESXi host. If you have chosen relocate VMs to a cluster, Veeam Backup & Replication will display only shared datastores.

   > **IMPORTANT**
   >
   > For Instant Recovery finalization, mind the following:
   >
   > - If you migrate VMs to the same datastore cluster that is used as the destination for redirecting virtual disk updates (the Datastore step of the **Instant Recovery to VMware** wizard), you must enable the **Force Veeam transport usage** check box at the Transfer step of the **Quick Migration** wizard. Otherwise, your migrated VM may be deleted.
   > - If you migrate workloads to the same datastore that is used as the destination for redirecting virtual disk updates, Veeam Backup & Replication uses Veeam Quick Migration instead of Storage vMotion. Such behavior prevents data loss due to a bug in VMware Storage vMotion. For more information on migration, see the Quick Migration section.

   If you want to place relocated VMs to different datastores:

   a. Click the **Pick datastore for selected virtual disks** link.

   b. In the **Choose VM Files Location** window, click **Add VM** on the right and select the VM that must be placed on datastores.

   c. Select the added VM in the **Files location** list and click **Datastore** at the bottom of the window.

d. From the list of available datastores, select the target datastore.

You can choose to store VM configuration files and disk files in different locations.

a. Add the VM to the **Files location** list, expand the VM and select the required type of files.

b. Click **Datastore** at the bottom of the window and choose the destination for the selected type of files.

5. By default, Veeam Backup & Replication saves disks of relocated VMs in the thin format. If necessary, you can change the disk format. For example, if the original VM uses thick disks, you can change the disk format of the relocated VM to thin provisioned and save on disk space required to store VM data.

   Disk format change is available only for VMs using virtual hardware version 7 or later.

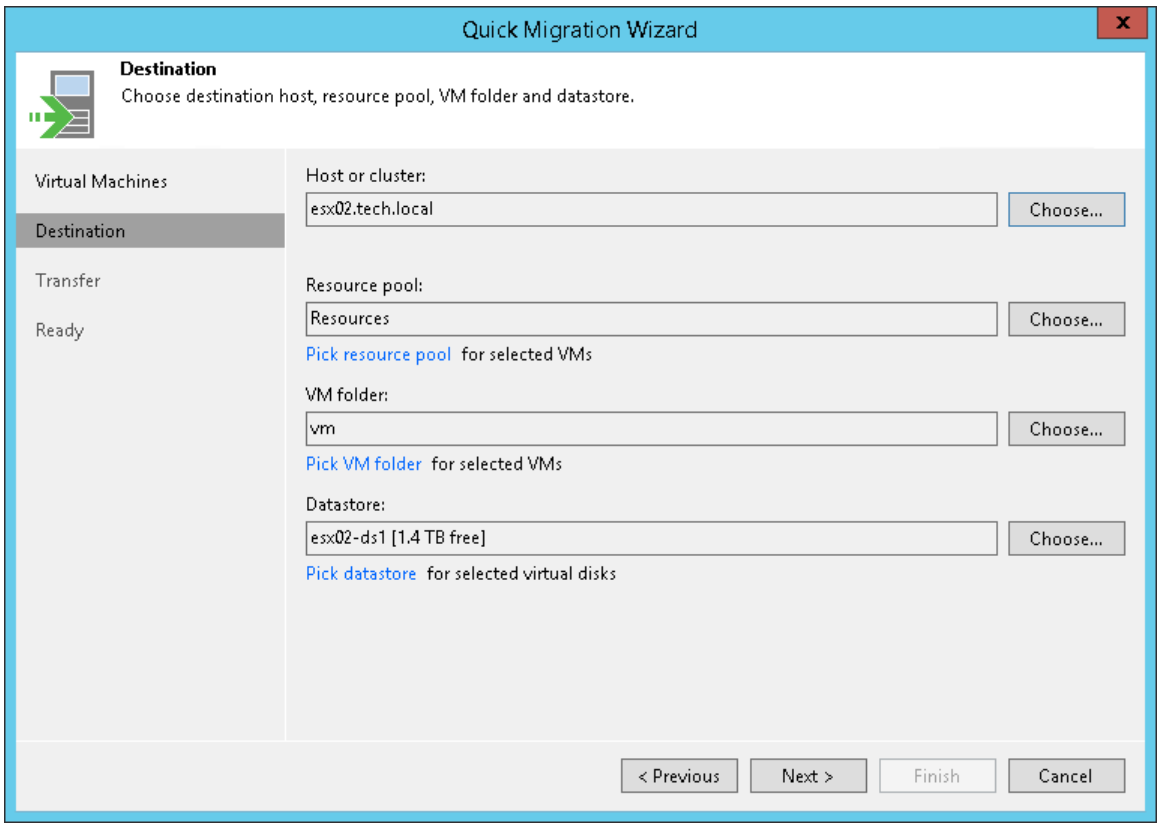
   To change VM disk format:

   a. Click the **Pick datastore for selected virtual disks** link.

   b. In the **Choose VM Files Location** window, click **Add VM** on the right and select the VM whose disk format you want to change.

   c. Select the added VM in the list and click **Disk type** at the bottom of the window.

   d. In the **Disk Type Settings** section, choose the format that will be used to restore VM disk files: same as the source disk, thin or thick.



Step 4. Select Infrastructure Components for Data Transfer

At the **Transfer** step of the wizard, assign infrastructure components to relocate the VMs:

1. In the **Data transfer** section, select backup proxies that must be used to transfer VM data from source to target.

   If you plan to migrate VMs within one site, the same backup proxy can act as the source backup proxy and target backup proxy. For off-site migration, you must deploy at least one backup proxy in each site to establish a stable connection across the sites for data transfer.

   Click **Choose** next to the **Source proxy** and **Target proxy** fields to select backup proxies for migration. In the **Backup Proxy** window, you can choose automatic proxy selection or assign proxies explicitly.

   - If you choose **Automatic selection**, Veeam Backup & Replication will detect backup proxies that have access to the source datastore and will automatically assign optimal proxy resources for processing VM data.

     Migrated VMs are processed one by one. Before processing a new VM in the VM list, Veeam Backup & Replication checks available backup proxies. If more than one backup proxy is available, Veeam Backup & Replication analyzes transport modes that the backup proxies can use for data retrieval and the current workload on the backup proxies to select the most appropriate resource for VM processing.

   - If you choose **Use the selected backup proxy servers only**, you can explicitly select backup proxies that must be used to perform migration.

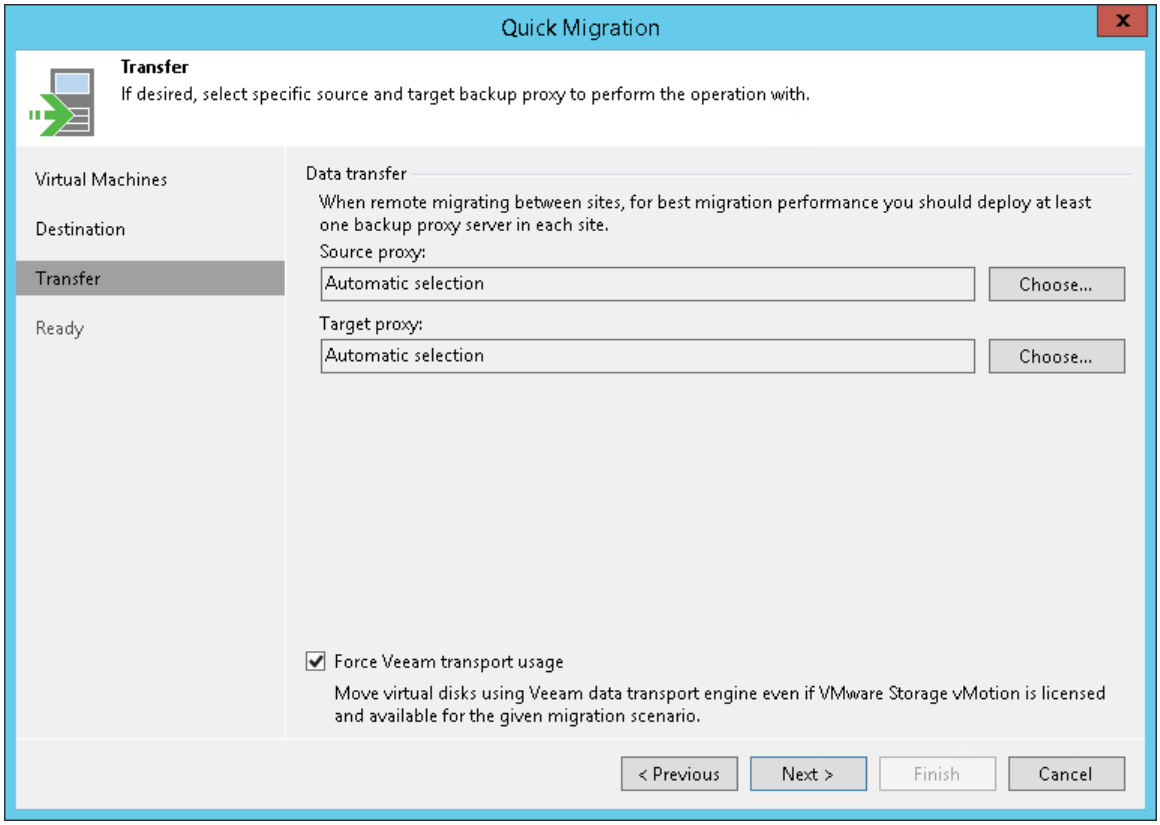
2. Select which migration mechanism to use: VMware vMotion or Veeam Quick Migration. Veeam Backup & Replication can use VMware vMotion only if your VMware license covers this functionality.

   - If you want to use VMware vMotion to relocate the VMs, leave the **Force Veeam transport usage** check box not selected. Veeam Backup & Replication will attempt to use the VMware vMotion mechanism to migrate the selected VMs. If VMware vMotion cannot be used for some reason (for example, if using it can cause data loss or if you do not have a VMware vSphere license for this functionality), Veeam Backup & Replication will fail over to its native migration mechanism.

   - If you do not want to use VMware vMotion, select the **Force Veeam transport usage** check box. Veeam Backup & Replication will use its native migration mechanism.

IMPORTANT

If you use a native Veeam mechanism to relocate a VM, Veeam Backup & Replication suspends the initial VM on the source ESXi host (SmartSwitch) or powers off the initial VM (cold switch) for a short period of time during quick migration. For more information, see Quick Migration.



Step 5. Finish Working with Wizard

At the **Ready** step of the wizard, you actions differ depending on the method you have chosen to use for migration.

If VMware vMotion is used for migration, review details and click **Finish**. In this case, all existing jobs to which the source VM is added will switch to the migrated VM. The backup chain will be continued, thus, the next job session for the VM will be incremental.

If Veeam Quick Migration is used, you must also choose whether you want to delete files of the VM for which the migration was launched after Veeam Backup & Replication receives a heartbeat signal from the VM on the target host (migrated VM):

- If you want to delete the VM for which the migration was launched, leave the **Delete source VM files upon successful migration** check box selected. All jobs to which the VM is added will switch to the migrated VM. The backup chain will be continued, thus, the next job session for the VM will be incremental.

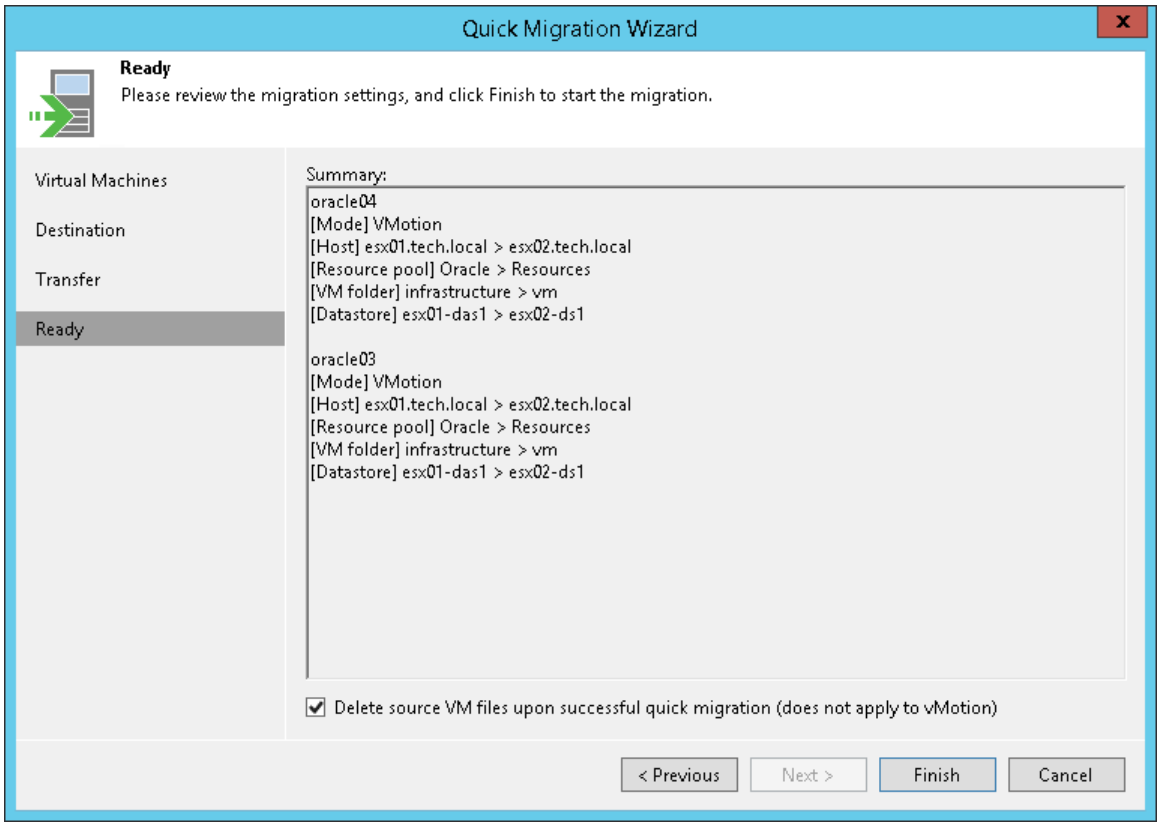
  If the heartbeat signal is not received from the migrated VM, the source VM will not be deleted and the migrated VM will not be added to any jobs. To protect the VM, you must add it to a backup job manually.

- If you want to leave the VM for which the migration was launched, clear the selection of the **Delete source VM files upon successful migration** check box. In this case, the VM will not be deleted. All jobs to which the VM is added will continue to process this VM. To protect the migrated VM, you must add it to a backup job manually.

NOTE

The **Delete source VM files upon successful migration** option applies if the following conditions are met:

- VMware vMotion is not used.
- VMs are powered on before migration.
- VMware Tools are installed on VMs.



Migrating First Class Disks (FCDs)

You can relocate First Class Disks (FCDs) between datastores using Quick Migration. Migration of FCDs supports only the vMotion method. Depending on whether FCDs that you want to migrate are attached to a VM or not, Veeam Backup & Replication defines the migration scenario:

- If FCDs are not attached, Veeam Backup & Replication will migrate only the FCDs that you have selected.
- If FCDs are attached to a VM, Veeam Backup & Replication will also look for other FCDs attached to this VM and will perform migration of these FCDs as well.

Quick migration is not a job-driven process: it cannot be saved as a job or scheduled to run later. To initiate a quick migration of FCDs, you must complete FCD Quick Migration wizard and launch the **FCD Quick Migration** wizard.

Before You Begin

Before you perform FCD quick migration, check the following prerequisites and limitations.

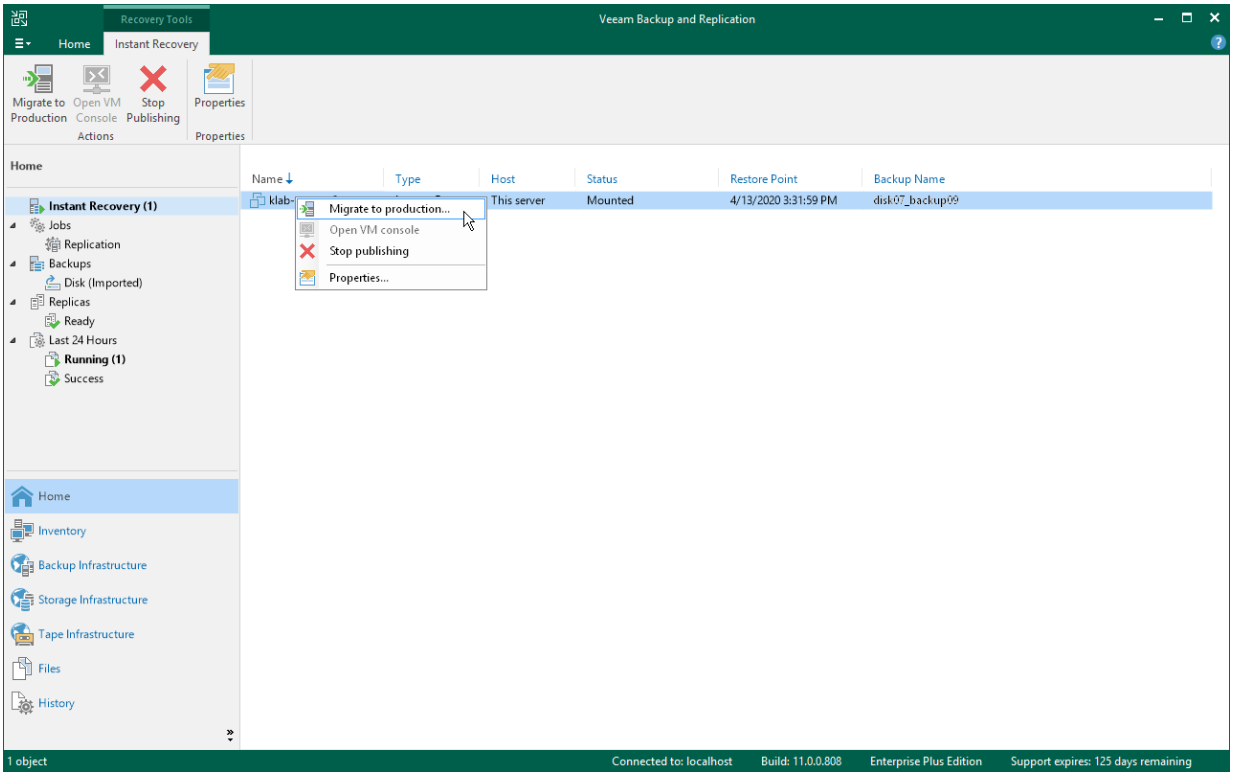
- All FCDs that you plan to migrate, must be located on a vPower NFS datastore that is mounted to the target cluster.
- If redo logs are stored on a custom datastore, Veeam Backup & Replication checks that these redo logs are available on this datastore.
- The target datastore to which you want to migrate FCDs must meet the following requirements:
  - The datastore must be added to the infrastructure of the vCenter which you specified to perform FCD quick migration.
  - The datastore must be mounted to all ESXi hosts of a cluster which you have selected to perform FCD quick migration.
  - The datastore must have enough free space for all migrated FDCs.
  - You can not migrate FCDs to the vPower NFS datastore.
- You can not migrate FCDs to the same datastore where they have already been registered.

Step 1. Launch FCD Quick Migration Wizard

To launch the **FCD Quick Migration** wizard

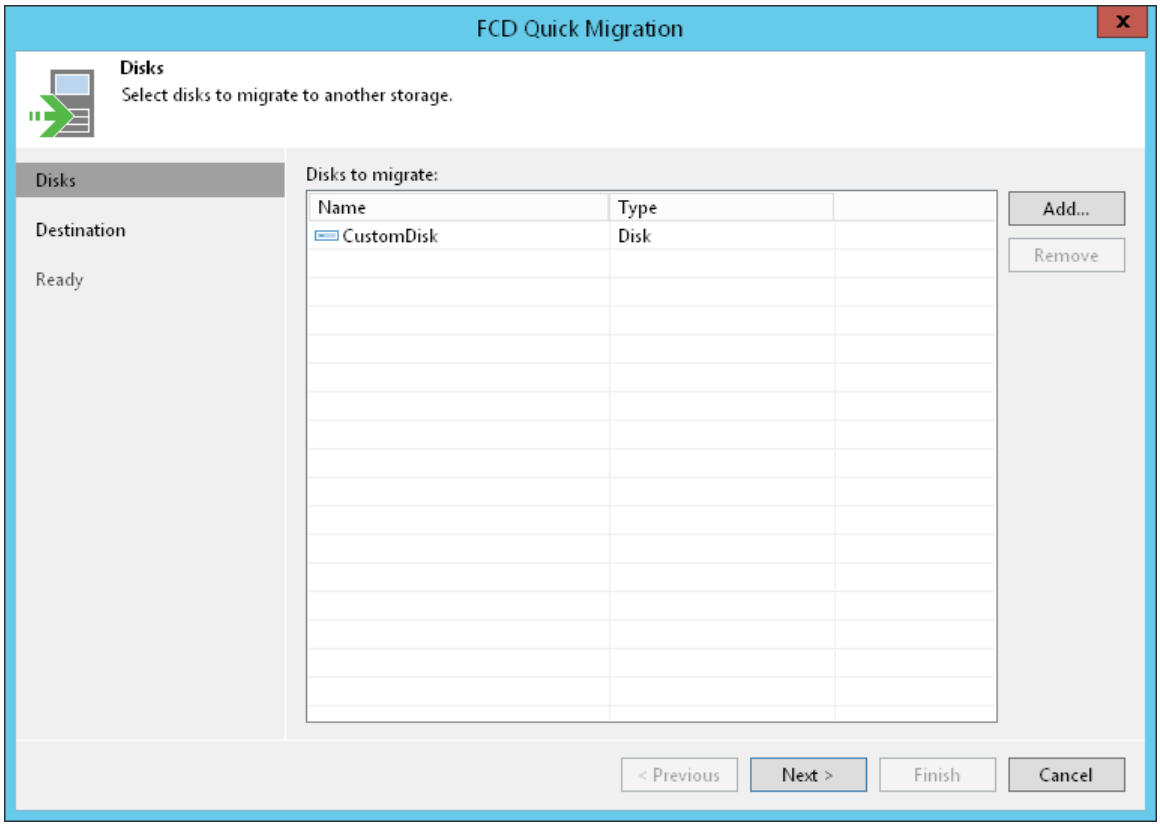
1. Open the **Home** view.
2. In the inventory pane, select the **Instant Recovery** node.
3. In the working area, right-click the VM to which disks were recovered and select **Migrate to production**. Veeam Backup & Replication will launch the FCD Quick Migration wizard.

After you finish working with the wizard, Veeam Backup & Replication migrates the disks with all changes made after the disk recovery and before its migration.



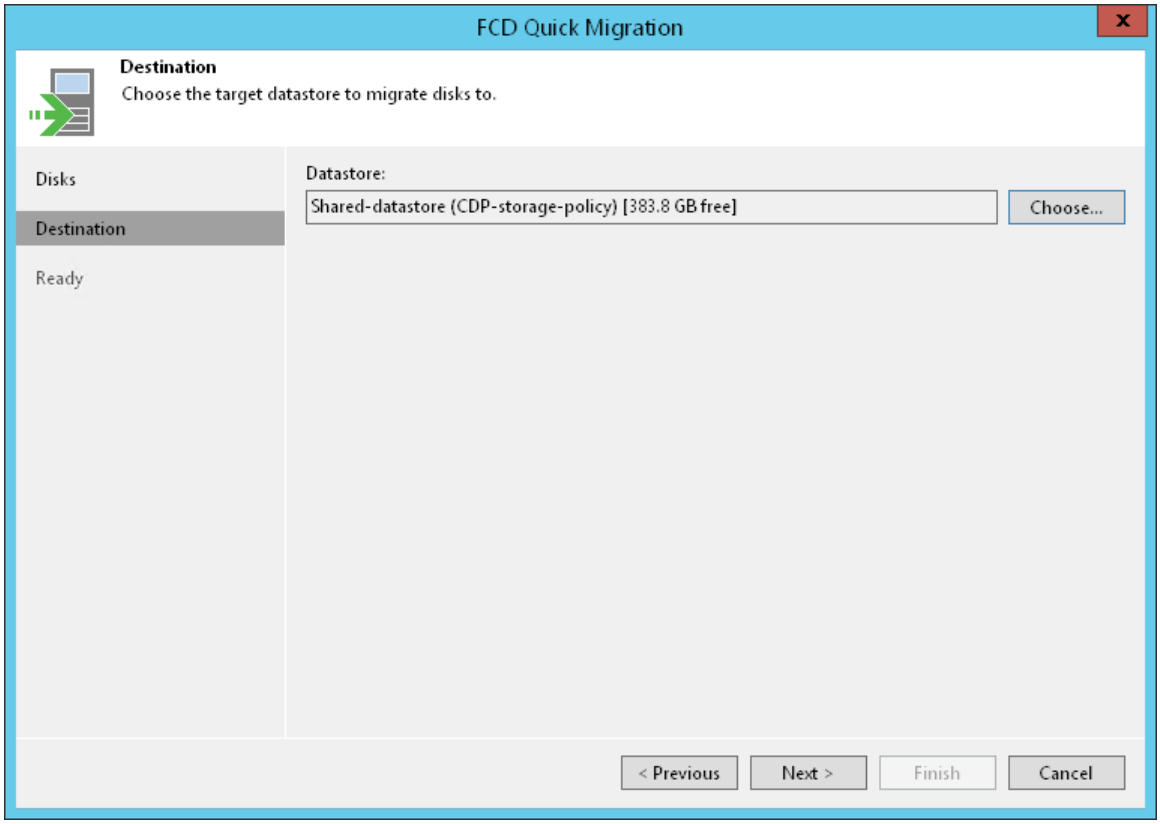
Step 2. Select FCDs to Migrate

At the **Disks** step of the wizard, specify the FCDs that you want to migrate. To add FCDs, click **Add** and select necessary FCDs from the virtual environment. If you want to exclude specific FDCs from migration, select the necessary FCD and click **Remove**.



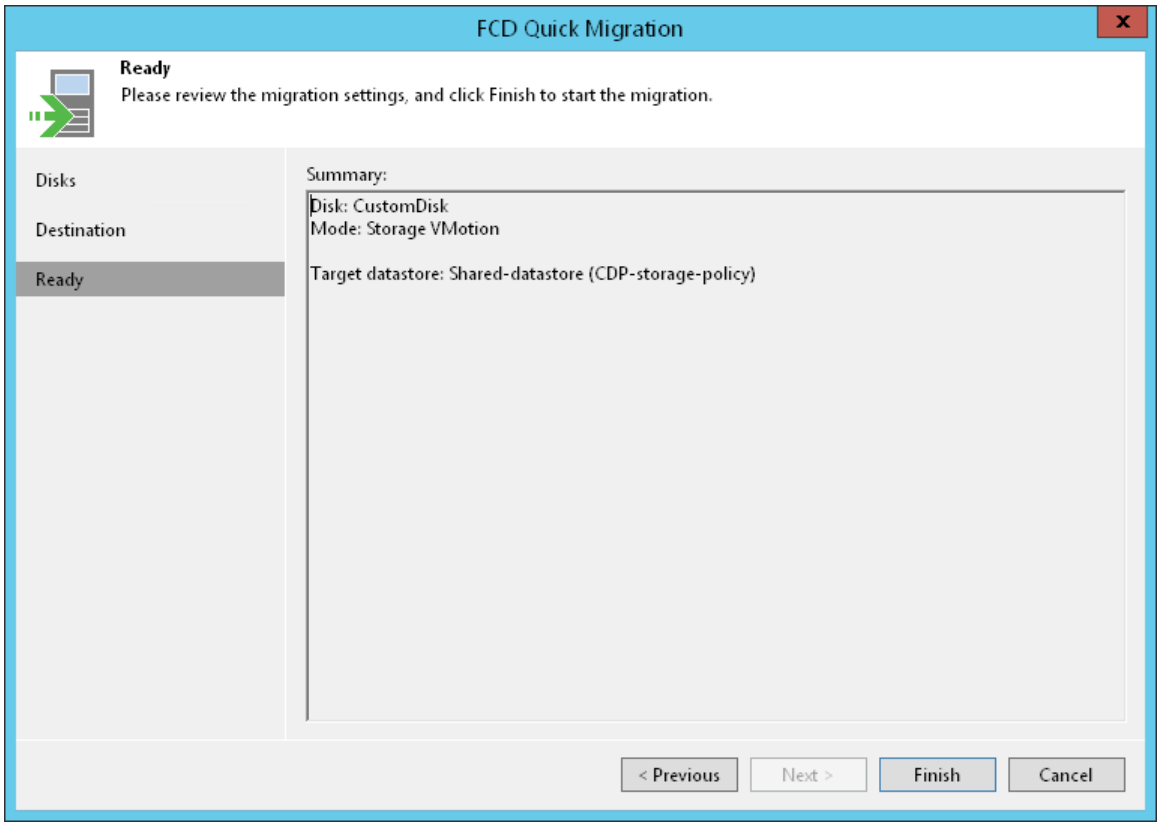
Step 3. Specify FCD Destination

At the **Destination** step of the wizard, specify the datastore to which you want to migrate FCDs. You can migrate different FCDs to various datastores.



Step 4. Finish Working with Wizard

At the **Ready** step of the wizard, review details on FCDs quick migration and click **Finish** to exit the wizard.



Recovery Verification

Veeam Backup & Replication offers two technologies to verify recoverability of VM backups and replicas:

- SureBackup
- SureReplica

IMPORTANT

The recovery verification functionality is included in the Veeam Universal License.

When using a legacy socket-based license, Enterprise or higher edition is required. If you use the Standard edition, you can manually verify VM backups with Instant Recovery. For more information, see Manual Recovery Verification.

SureBackup

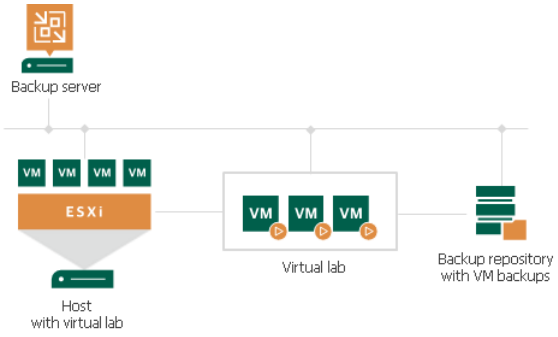
SureBackup is the Veeam technology that allows you to test VM backups and check if you can recover data from them. You can verify any restore point of a backed-up VM.

During a SureBackup job, Veeam Backup & Replication performs "live" verification: scans the backed-up data for malware, boots the VM from the backup in the isolated environment, runs tests for the VM, powers the VM off and creates a report on recovery verification results.

How SureBackup Works

For SureBackup, Veeam Backup & Replication uses a regular image-based backup. During recovery verification, Veeam Backup & Replication performs the following actions:

1. If the SureBackup job is configured to perform malware scan, Veeam Backup & Replication scans data of the VMs from the application group with antivirus software.
2. Veeam Backup & Replication publishes the VMs from the application group and the VM under verification in the isolated environment — virtual lab. VMs are started directly from compressed and deduplicated backup files that reside in the backup repository. To achieve this, Veeam Backup & Replication utilizes the Veeam vPower NFS Service.
3. If the SureBackup job is configured to perform malware scan, Veeam Backup & Replication scans data of the VM under verification with antivirus software.
4. Veeam Backup & Replication performs a number of tests against the VMs in the application group and VM under verification: heartbeat test, ping test and application test.
5. If the SureBackup job is configured to validate backup files, Veeam Backup & Replication performs a a cyclic redundancy check for the backup file from which the VM under verification is started and, optionally, for backup files from which the VMs in the application group are started. The backup file validation is performed after all verification tests are complete.
6. When the recovery verification process is over, Veeam Backup & Replication unpublishes VMs and creates a report on their state. The report is sent to the backup administrator by email.



During verification, a backed-up VM image remains in read-only state. All changes that take place when the VM is running are written to redo log files that are stored on the datastore selected in the virtual lab settings. When the recovery process is complete, the redo logs are removed.

To perform recovery verification, you need to create the following objects:

1. Application group. During recovery verification, the verified VM may need to be started with a group of VMs on which it is dependent. The application group enables full functionality of applications running inside the VM and lets you run these applications just like in the production environment.
2. Virtual lab. The virtual lab is the isolated virtual environment in which the verified VM and VMs from the application group are started and tested.
3. SureBackup job. The SureBackup job is a task to perform recovery verification. You can run the SureBackup job manually or schedule it to run automatically by schedule.

Backup Recovery Verification Tests

To verify VMs with a SureBackup job, you can instruct Veeam Backup & Replication to run the following backup recovery verification tests:

- Predefined tests
- Custom verification scripts

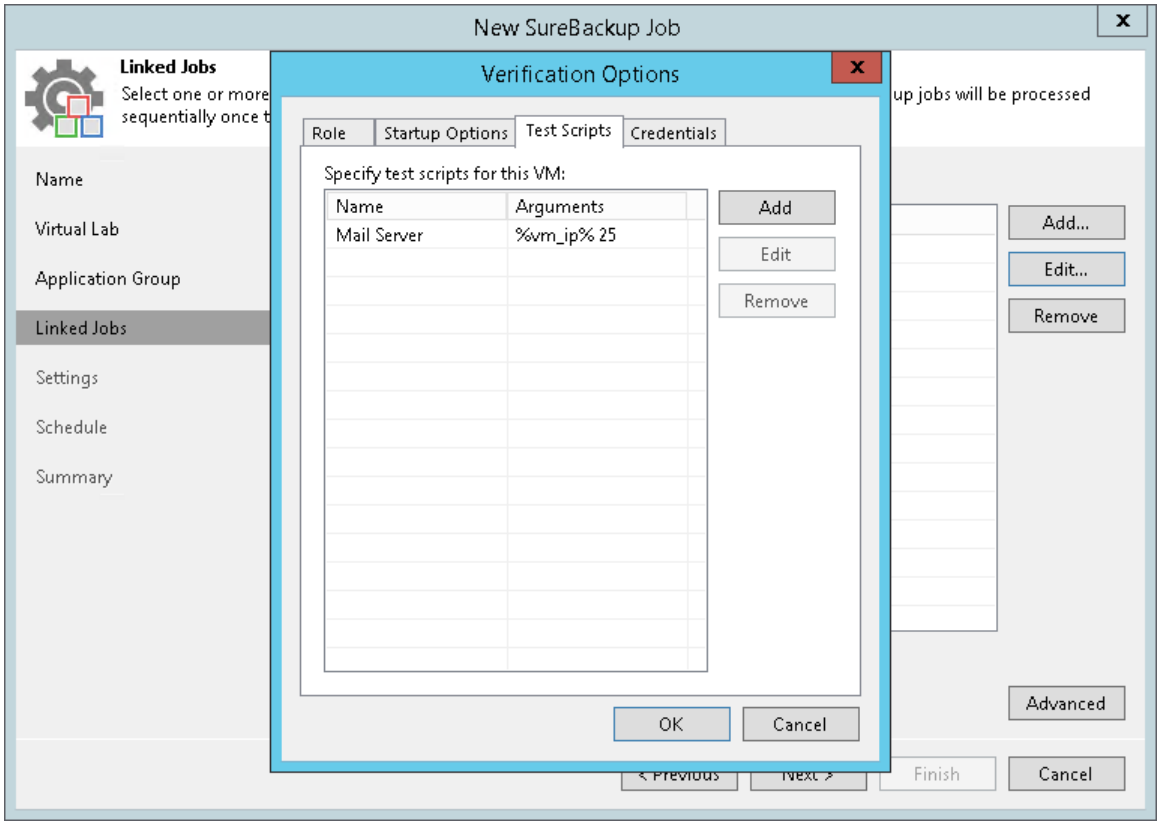
Predefined Tests

Veeam Backup & Replication can verify VMs with the following predefined tests:

- **Heartbeat test**. When the VM starts, Veeam Backup & Replication performs a heartbeat test. It waits for a heartbeat signal from VMware Tools installed inside the VM to determine that the VM guest OS is running. If the signal comes regularly at specific time intervals, the test is passed.
- **Ping test**. Veeam Backup & Replication sends ping requests to the VM from the backup server and checks if the VM can respond to them. If the VM responds to ping requests, the test is passed.
- **Application test**. Veeam Backup & Replication waits for applications inside the VM to start and runs a script against these applications. Veeam Backup & Replication uses two types of predefined scripts:
  - For DNS servers, domain controllers, Global Catalog servers, mail servers and web servers, Veeam Backup & Replication uses a script that probes an application-specific port. For example, to verify a domain controller, Veeam Backup & Replication probes port 389 for a response. If the response is received, the test is passed.
  - For Microsoft SQL Server, Veeam Backup & Replication uses a script that attempts to connect to instances and databases on the Microsoft SQL Server. For more information, see Microsoft SQL Server Checker Script.

> **NOTE**
>
> To run the heartbeat and ping tests, you must have VMware Tools installed on the VM. If VMware Tools are not installed, these tests will be skipped.

You can run verification tests for VMs added to the application group or processed with a linked SureBackup job. You can specify and customize settings for verification tests in the application group or SureBackup job settings.



Custom Verification Scripts

Veeam Backup & Replication can verify VMs with the following custom verification scripts:

- Microsoft SQL Server Checker script
- Backup file validation

Microsoft SQL Server Checker Script

If you need to verify a virtualized Microsoft SQL Server, you can instruct Veeam Backup & Replication to run the Microsoft SQL Server Checker script against it during the SureBackup job. The script sequentially performs the following operations:

1. Connects to Microsoft SQL Server instances.
2. Enumerates databases on these instances.
3. Employs the USE SQL statement to connect to databases and check their availability.

The script is located on the backup server in the Veeam Backup & Replication product folder. The path by default: `C:\Program Files\Veeam\Backup and Replication\Backup\Veeam.Backup.SqlChecker.vbs`.

Credentials for Script Execution

To execute the script, Veeam Backup & Replication connects to Microsoft SQL Server. By default, Veeam Backup & Replication uses the account under which the Veeam Backup Service is running. If you need to run the script under another account, you can specify credentials for this account. The script supports Microsoft Windows and SQL Server authentication methods.

For the Microsoft Windows authentication mode, you can specify credentials for the account on the Credentials tab in the application group or SureBackup job settings.

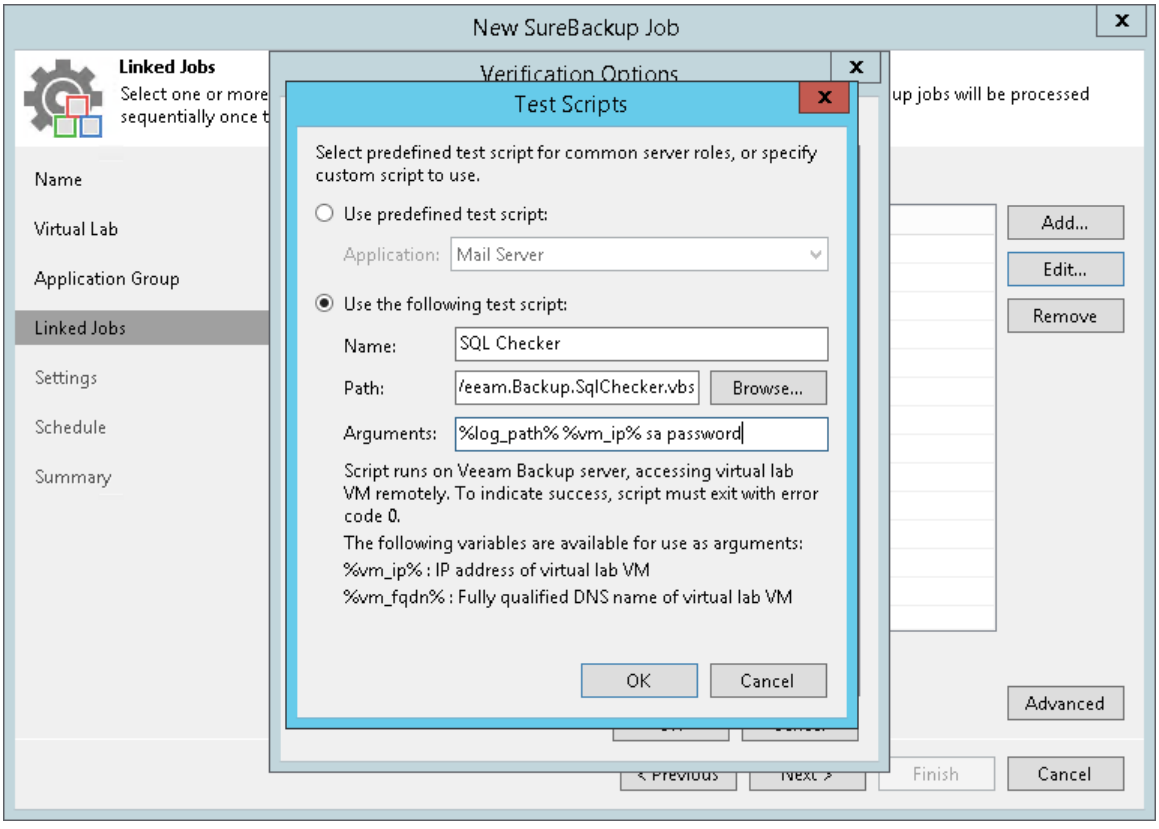
For the SQL Server authentication mode, you must pass credentials of the account as arguments to the script. You can do it using the command line interface or the UI.

To pass credentials using the command line, run the script from the command line with the following parameters:

```
cscript Veeam.Backup.SqlChecker.vbs [logs folder] <sql server[\instance]> <user
name> <password>
```

To pass credentials using the UI:

1. In the application group or SureBackup job settings, select to use a custom script.
2. Specify a path to the Microsoft SQL Server Checker script. Paths by default: `C:\Program Files\Veeam\Backup` and `Replication\Backup\Veeam.Backup.SqlChecker.vbs`.
3. Specify the user name and password in the **Arguments** field.



IMPORTANT

If you use the Microsoft SQL Server authentication mode, you may need to specify credentials of the account to connect to the machine on which Microsoft SQL Server is installed. To do this, use the **Credentials** tab in the application group or SureBackup job settings.

Database Exclusion

By default, Veeam Backup & Replication verifies all databases on all instances of Microsoft SQL Server. However, you can exclude specific databases from verification. For example, vCenter Server database. To exclude an instance or a database, you must open the script in the text editor and edit the **Settings** section in the following way:

- To exclude a specific databases, uncomment the `gDBsToExclude.Push "dbname"` line in the script and specify the names of a database that you want to exclude. To exclude several databases, specify a separate line for each database.

```
gDBsToExclude.Push "dbname1"
gDBsToExclude.Push "dbname2"
```

- To exclude a specific instance, uncomment the `gInstancesToExclude.Push "instancename"` line in the script and specify the name of an instance that you want to exclude. To exclude several instances, specify a separate line for each instance.

```
gInstancesToExclude.Push "instancename1"
gInstancesToExclude.Push "instancename2"
```

- To exclude the default instance, uncomment the `gInstancesToExclude.Push "MSSQLSERVER"` line.

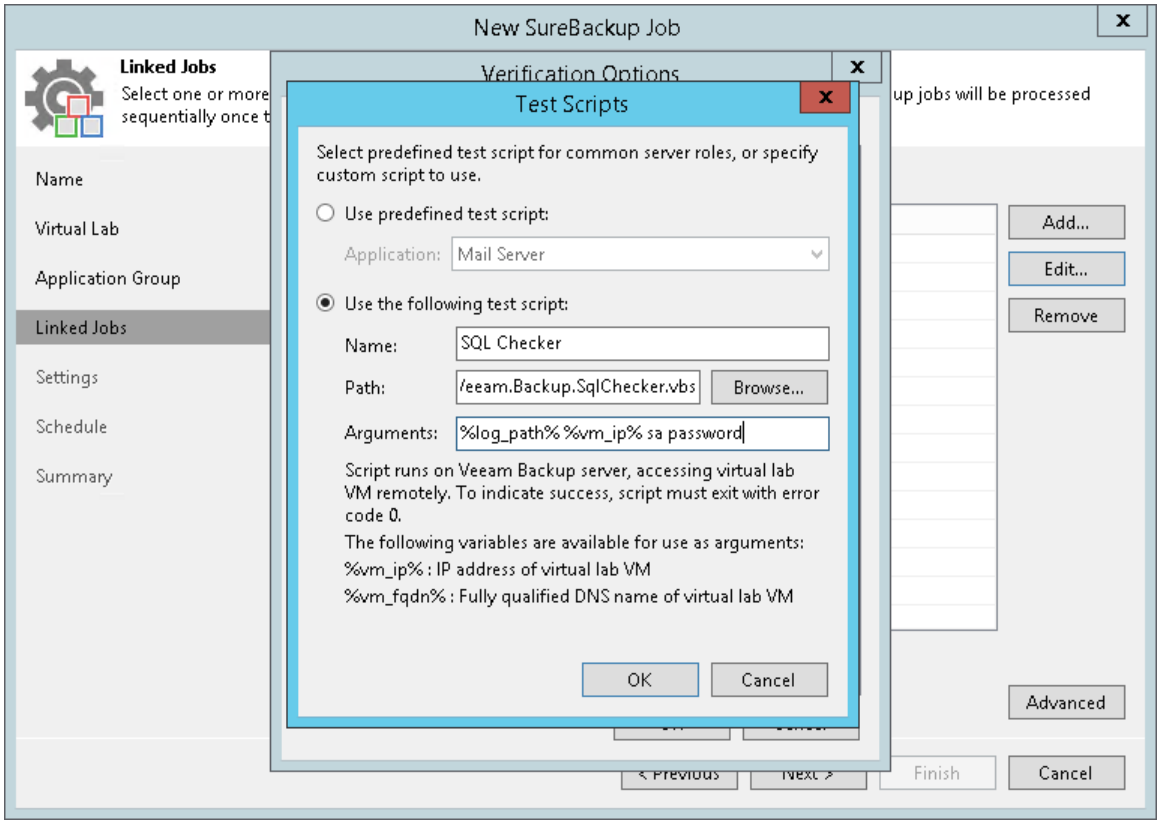
> **IMPORTANT**
>
> Instance and database names are case sensitive.

Logging

To define whether the script has completed successfully or not, Veeam Backup & Replication publishes the following return codes in the SureBackup job session statistics:

- 0 – test is passed successfully.
- 1 – you use a wrong syntax for the script command.
- 2 – Veeam Backup & Replication is unable to connect to Microsoft SQL Server.
- 3 – all instances are excluded from the check.
- 4 – error occurred while Veeam Backup & Replication was getting the list of databases.
- 5 – unknown error
- 6 – one or more databases are not accessible.

Results of script execution are written to the log file located by the following path: `%programdata%\Veeam\Backup\<name of the job>\<VM name>_SQLChecker.log`. If necessary, you can change the log file location. To do this, pass a new path to the log file in the `%log_path%` argument in the application group or SureBackup job settings.



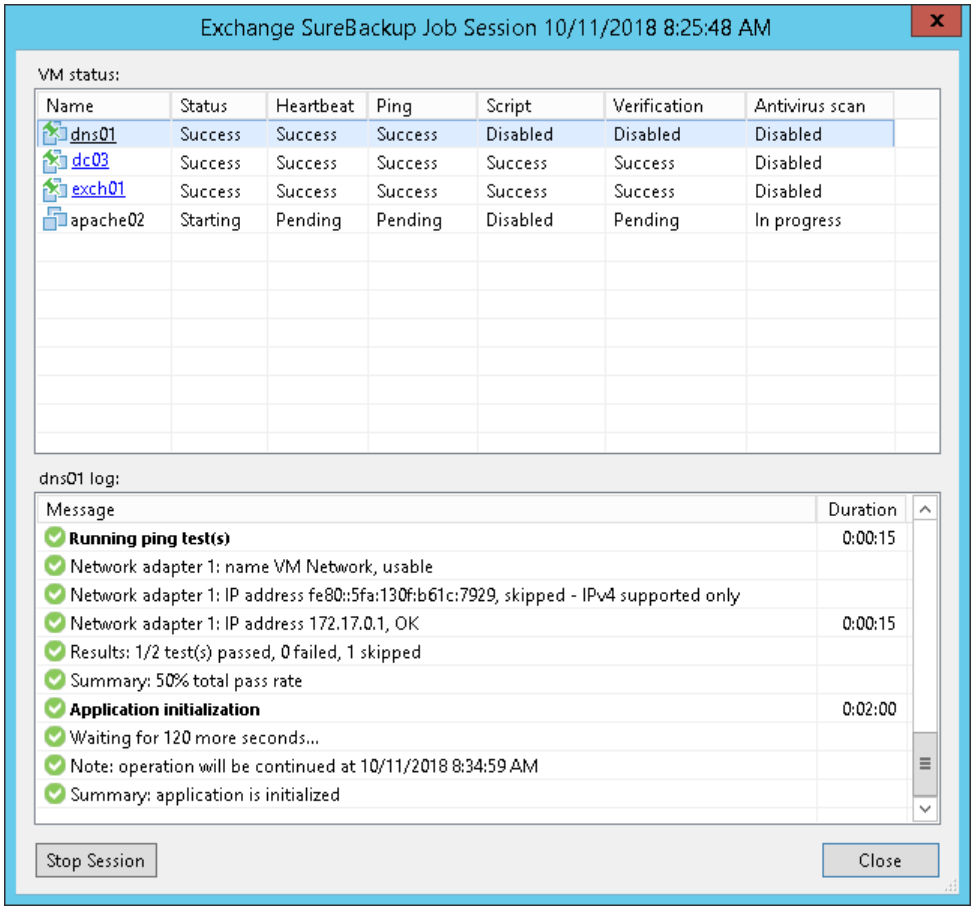
Backup File Validation

In addition to recovery verification tests, Veeam Backup & Replication allows you to perform backup file validation. For backup file validation, Veeam Backup & Replication performs a CRC check for backup files of VMs verified by the SureBackup job. You can also validate backup files for VMs from the application group with this test.

To validate the backup file, Veeam Backup & Replication uses the checksum algorithm. When Veeam Backup & Replication creates a backup file for a VM, it calculates a checksum for every data block in the backup file and stores this data in the backup file, together with VM data. During the backup file validation test, Veeam Backup & Replication decompresses the backup file, recalculates checksums for data blocks in the decompressed backup file and compares them with initial checksum values. If the results match, the test is passed.

The backup file validation test is started after recovery verification tests. As soon as Veeam Backup & Replication completes all "live" verification for all VMs in the SureBackup job, it unpublishes VMs and starts the backup file validation test.

The result of the backup file validation test impacts the state of the SureBackup job session. If the verification tests are completed successfully but the backup validation is not passed, Veeam Backup & Replication marks the SureBackup job session with the *Failed* status.



Application Group

In most cases, a VM works not alone but in cooperation with other services and components. To verify such VM, you first need to start all services and components on which this VM is dependent. To this aim, Veeam Backup & Replication uses the application group.

The application group creates the "surroundings" for the verified VM. The application group contains one or several VMs on which the verified VM is dependent. These VMs run applications and services that must be started to enable fully functional work of the verified VM. Typically, the application group contains at least a domain controller, DNS server and DHCP server.

When you set up an application group, you specify a role of every VM, its boot priority and boot delay. Additionally, you can specify what tests must be performed to verify VMs in the application group.

When a SureBackup job is launched, Veeam Backup & Replication first starts in the virtual lab VMs from the application group in the required order and performs necessary tests against them. This way, Veeam Backup & Replication creates the necessary environment for the verified VM. Only after all VMs from the application group are started and tested, Veeam Backup & Replication starts the verified VM in the virtual lab.

For example, if you want to verify a Microsoft Exchange Server, you need to test its functionality in cooperation with the domain controller and DNS server. Subsequently, you must add to the application group a virtualized domain controller and DNS server. When Veeam Backup & Replication runs a SureBackup job, it will first start and verify the domain controller and DNS server in the virtual lab to make verification of the Microsoft Exchange Server possible.

> **NOTE**
>
> All VMs added to the application group must belong to the same platform — VMware or Hyper-V. Mixed application groups are not supported.

Creating Application Groups

Before you create an application group, check prerequisites. Then use the **New Application Group** wizard to create an application group.

Before You Begin

Before you create an application group, check the following prerequisites:

- A valid license for Enterprise edition of Veeam Backup & Replication must be installed on the backup server.
- All applications and services on which verified VMs are dependent must be virtualized in your environment.
- If you plan to scan VM data for malware, check requirements and limitations.
- If you plan to verify VMs with a ping test, the firewall on tested VMs must allow ping requests.
- If you plan to verify VMs with a heartbeat test, VMware Tools must be installed in tested VMs.
- [For storage snapshots] The storage system must be added to the backup infrastructure.

Mind the following limitations:

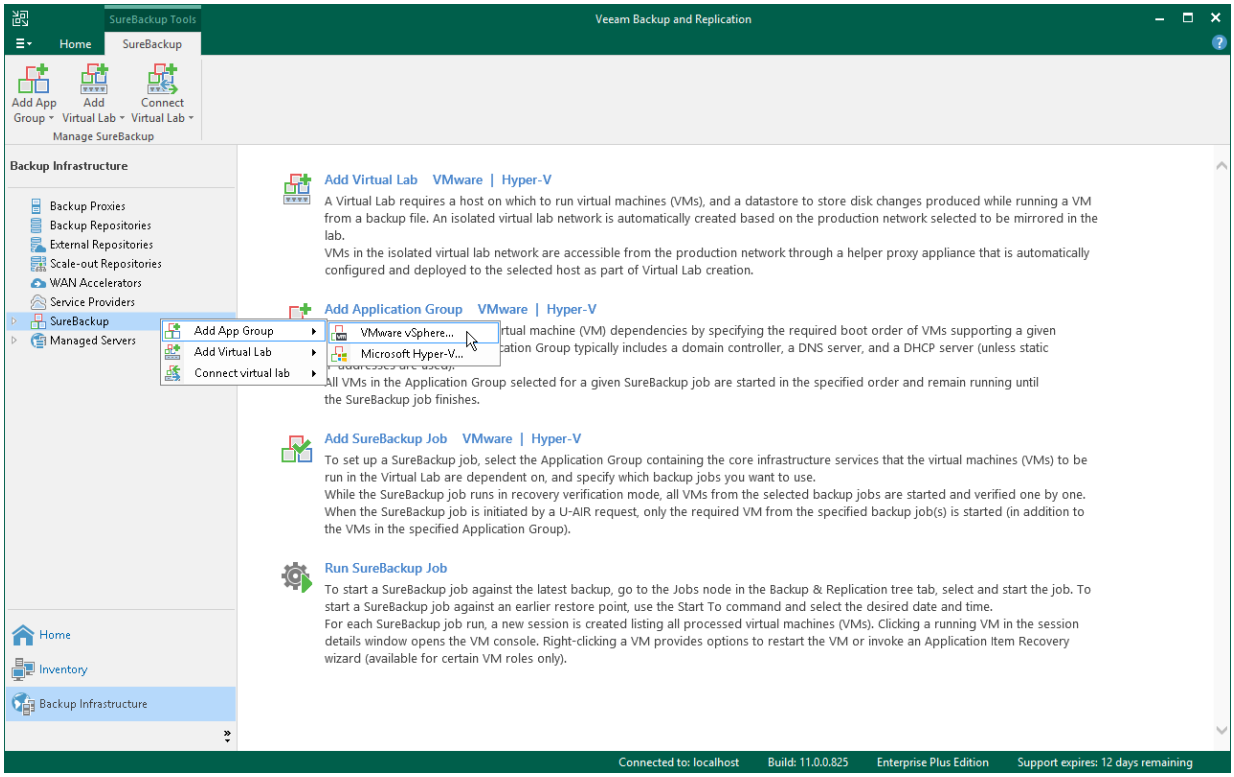
- VM replicas must be in the *Ready* state. If a VM replica is in the *Failover* or *Failback* state, you will not be able to add it to the application group.

- You cannot add to application groups VMs from backups of vCloud Director VMs, backups created with backup copy jobs and backups stored in cloud backup repositories.

Step 1. Launch New Application Group Wizard

To launch the **New Application Group** wizard, do one of the following:

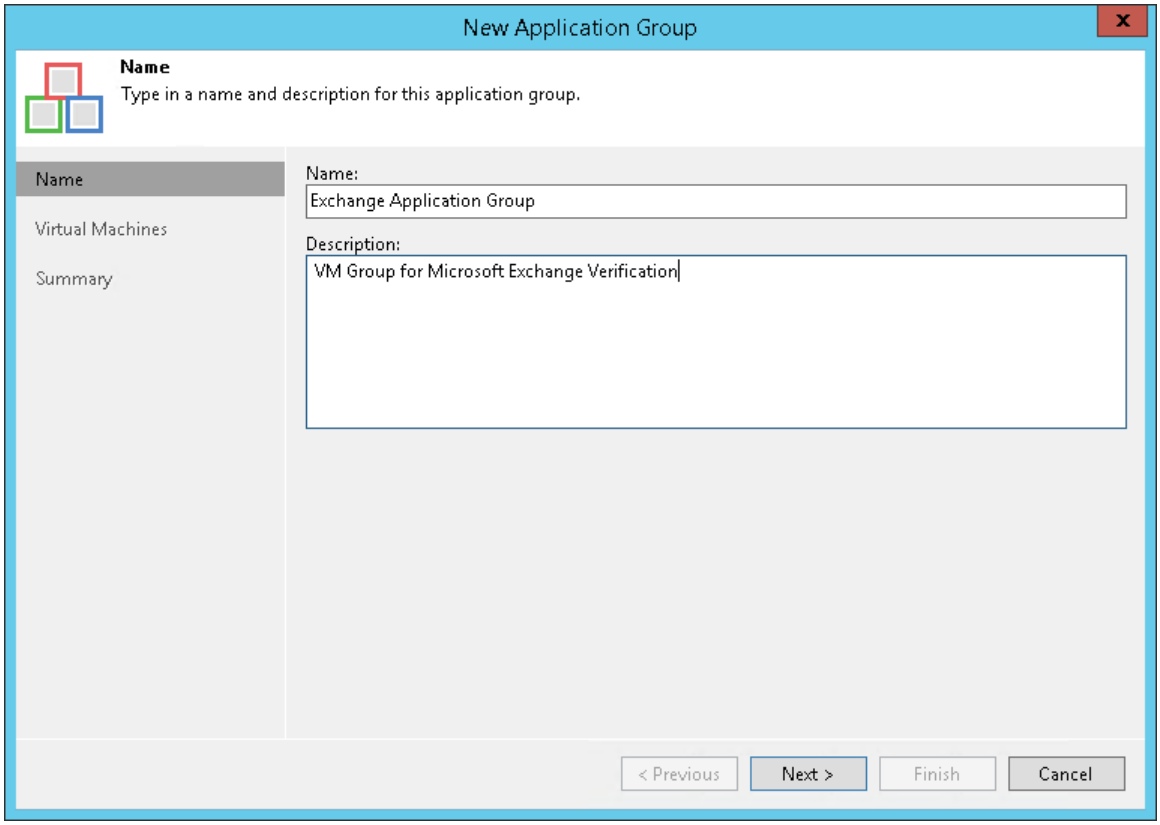
- Open the **Backup Infrastructure** view, in the inventory pane select **SureBackup**. In the working area, click **Add Application Group > VMware**.
- Open the **Backup Infrastructure** view, in the inventory pane select **Application Groups** under **SureBackup** and click **Add Group** > **VMware** on the ribbon.
- Open the **Backup Infrastructure** view, in the inventory pane right-click **Application Groups** under **SureBackup** and select **Add App Group** > **VMware**.



Step 2. Specify Application Group Name and Description

At the **Name** step of the wizard, specify a name and description for the application group.

1. In the **Name** field, enter a name for the application group.
2. In the **Description** field, provide a description for future reference. The default description contains information about the user who created the group, date and time when the group was created.



Step 3. Add VMs to Application Group

At the **Virtual Machines** step of the wizard, add VMs to the created application group. You can add VMs from different sources:

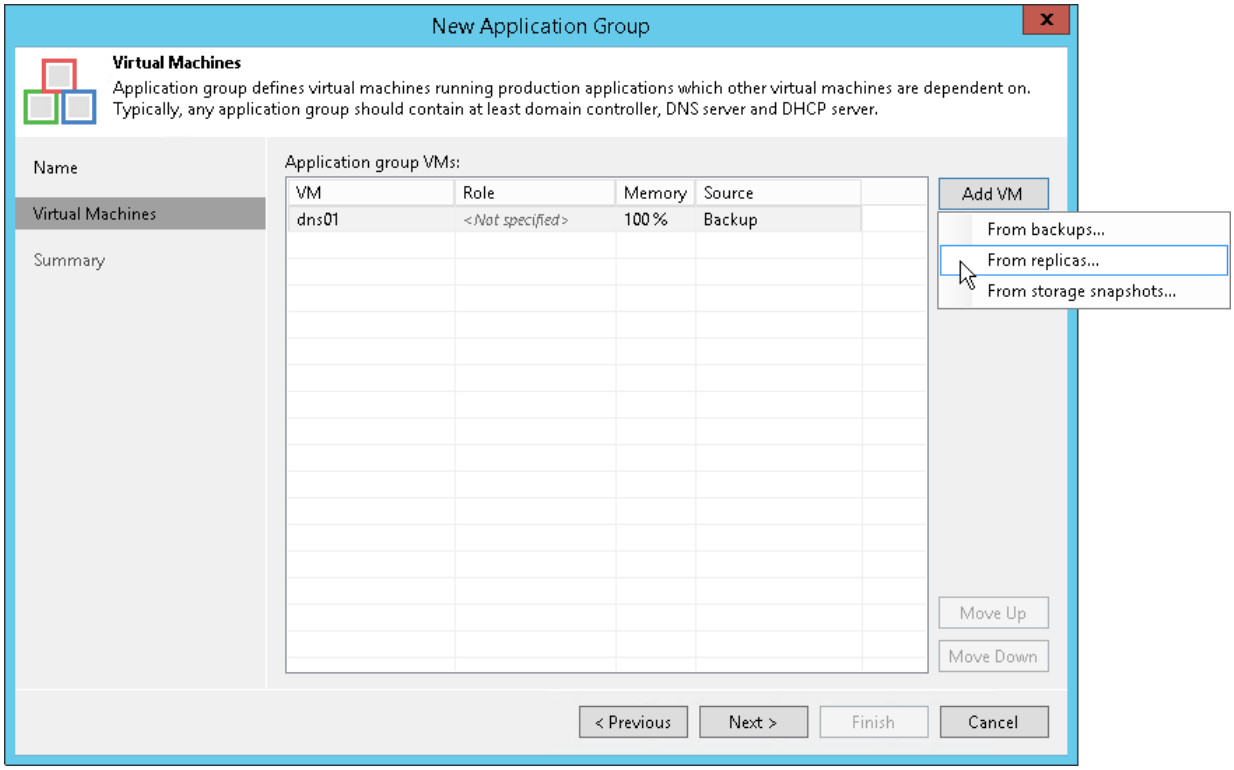
- VM backups
- VM replicas
- Storage snapshots

You can add VMs from backups, storage snapshots and VM replicas to the same application groups. Keep in mind the following limitations:

- VMs must belong to the same platform — VMware vSphere or Microsoft Hyper-V.
- VMs must have at least one valid restore point or must reside on a storage snapshot.
- You cannot add the same VM twice. For example, if you add a VM from the storage snapshot, you will not be able to add the same VM from the backup.

To add VMs to the application group:

1. Click **Add VM** and select **From backups**, **From replicas** or **From storage snapshots**.
2. In the displayed window, expand the job or storage snapshot, select the VM and click **Add**.
3. VMs in the list are specified in the order of their boot priority. To move a VM up and down in the list, select it and click **Move Up** or **Move Down**.



Step 4. Specify Recovery Verification Options and Tests

You must specify verification options for every VM in the application group:

- Select a role that the VM performs
- Configure startup settings
- Select tests that must be performed for the VM
- Specify credentials for running the verification script

To specify recovery verification options:

1. At the **Virtual Machines** step of the wizard, select the VM in the list.
2. Click **Edit** on the right.
3. Use the **Verification Options** window to specify verification options.

Role Settings

On the **Role** tab, select a role that the VM performs. Veeam Backup & Replication offers the following predefined roles for VMs:

- DNS Server
- Domain Controller (Authoritative Restore). In the Authoritative Restore mode, Veeam Backup & Replication starts a domain controller in the virtual lab and marks it as being authoritative to its replication partners. When other domain controllers (replication partners) are started in the virtual lab, they replicate data from the domain controller started in the Authoritative Restore mode.
- Domain Controller (Non-Authoritative Restore). In the Non-Authoritative Restore mode, Veeam Backup & Replication restores a domain controller in the virtual lab and marks it as being non-authoritative to its replication partners. Non-authoritative domain controllers then replicate data from a domain controller started in the Authoritative Restore mode.
- Global Catalog
- Mail Server
- SQL Server
- Veeam Backup for Microsoft Office 365 (machine on which Veeam Backup for Microsoft Office 365 is installed)
- Web Server

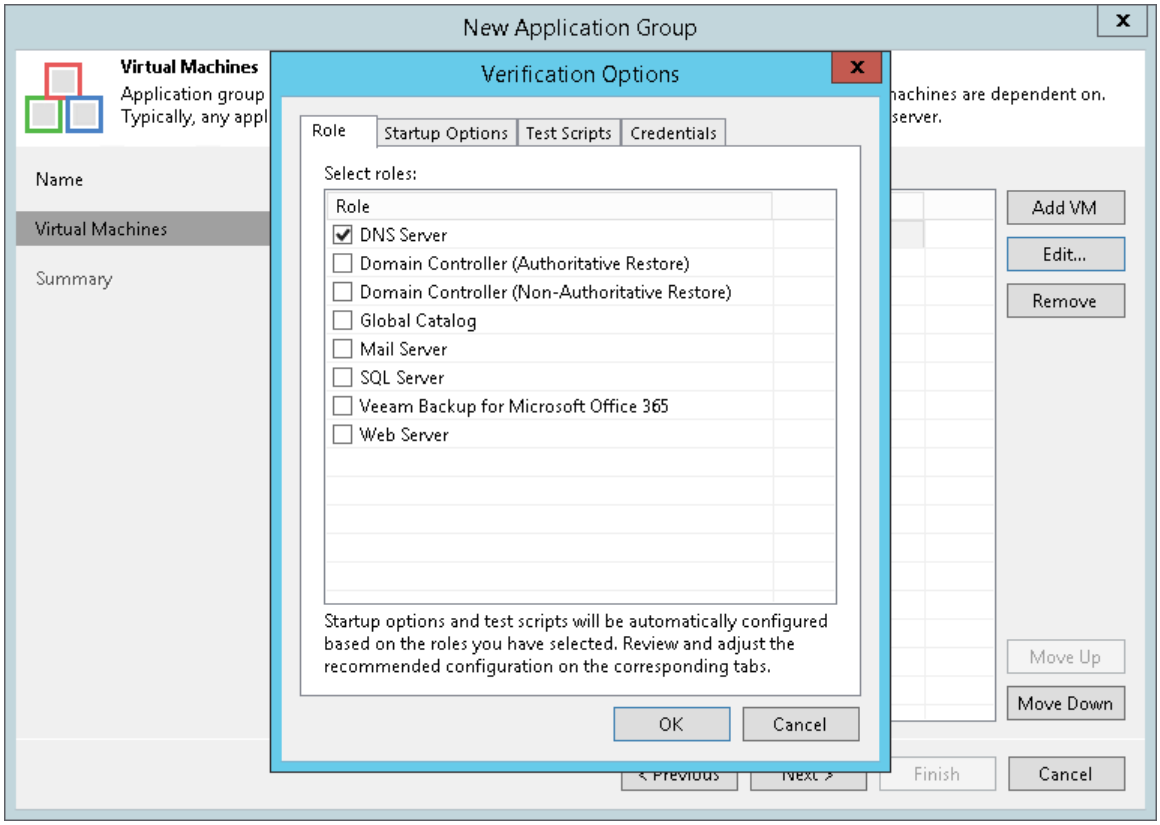
VM roles are described in XML files stored in the *%ProgramFiles%|Veeam|Backup and Replication|Backup|SbRoles* folder on the backup server. You can create your own roles. To do this, you must create a new XML file and specify role and test scripts settings in it. For more information, see Creating XML files with VM Roles Description.

After you select a role for the VM, Veeam Backup & Replication will automatically configure startup options and assign predefined test scripts for the chosen role. You can use these settings or specify custom settings on the **Startup Options** and **Test Scripts** tabs.

To verify VMs that perform roles other than those specified in the list, you will have to manually configure startup options and specify test scripts that must be run for these VMs.

IMPORTANT

If you want to add several domain controllers to the application group, you must assign the **Domain Controller (Authoritative Restore)** role to the first domain controller started in the virtual lab. Other domain controllers must have the **Domain Controller (Non-Authoritative Restore)** role.



Startup Settings

To configure VM startup settings:

1. In the **Verification Options** window, click the **Startup Options** tab.
2. In the **Memory** section, specify the amount of memory that you want to pre-allocate to the VM when this VM starts. The amount of pre-allocated memory is defined in percent. The percentage rate is calculated based on the system memory level available for the production VM. For example, if 1024 MB of RAM is allocated to the VM in the production environment and you specify 80% as a memory rate, 820 MB of RAM will be allocated to the verified VM on startup.

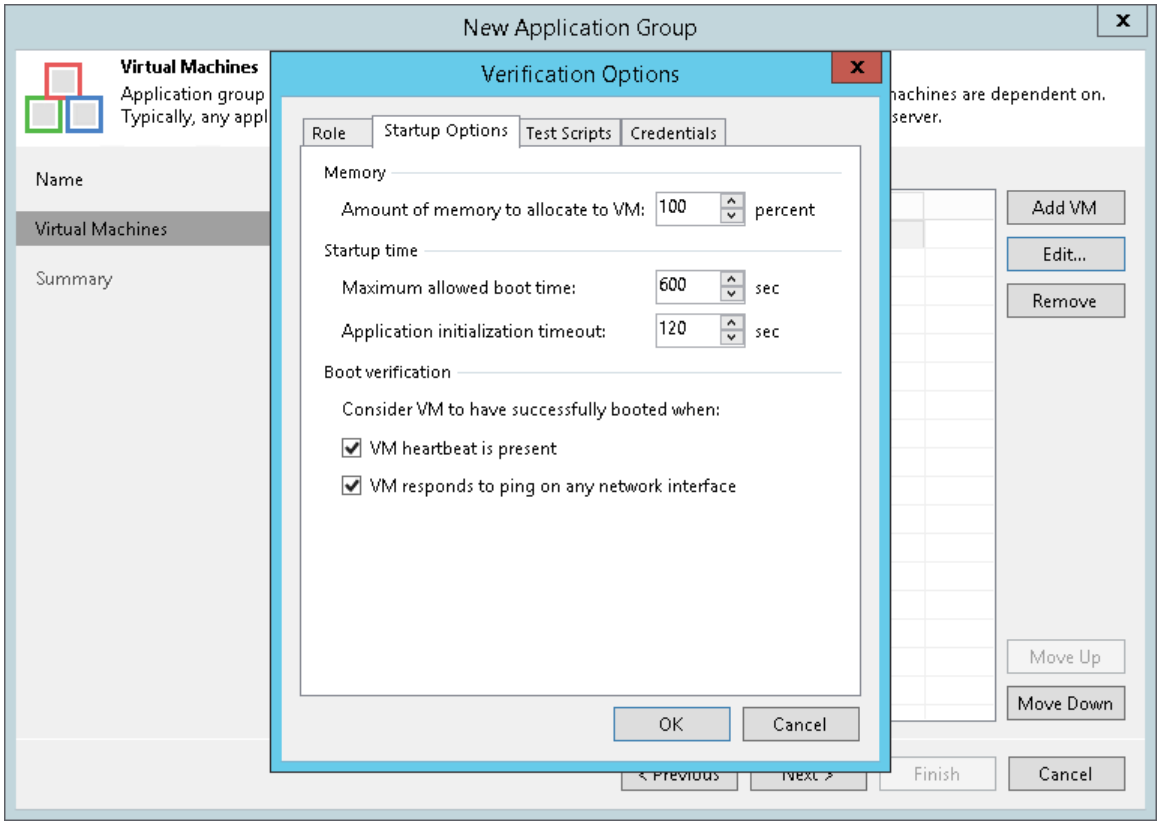
   Veeam Backup & Replication does not allow you to change VM CPU manually, it does this automatically. If the VM has more CPU than the host can provide, Veeam Backup & Replication scales down the CPU of the VM.
3. In the **Startup time** section, specify the allowed boot time for the VM and timeout to initialize applications on the VM.

   Be careful when specifying the **Maximum allowed boot time** value. Typically, a VM started by the SureBackup job requires more time to boot than a VM started in the production environment. If an application is not initialized within the specified interval of time, the recovery verification process fails with the timeout error. If such error occurs, you need to increase the **Maximum allowed boot time** value and run the SureBackup job again.

4. In the **Boot verification** section, specify when the VM must be considered to have been booted successfully:
   - **VM heartbeat is present**. If you enable this option, Veeam Backup & Replication will perform a heartbeat test for the verified VM.
   - **VM responds to ping on any network interface**. If you enable this option, Veeam Backup & Replication will perform a ping test for the verified VM.

   If you enable both options, Veeam Backup & Replication will require that both tests are completed successfully: heartbeat test and ping test.

> **NOTE**
>
> Veeam Backup & Replication performs a heartbeat test only if a VM has VMware Tools are installed. If VMware Tools are not installed, the VM will be started but the test will not be performed. VMs without VMware Tools can still be used as auxiliary VMs: they can be started to enable proper work of other VMs. In this case, you do not need to select any role for such VMs.



Test Script Settings

When you select a VM role, Veeam Backup & Replication automatically assigns a predefined script that must be run to verify applications inside this VM. If you want to verify a VM that has some other role not listed on the **Role** tab:

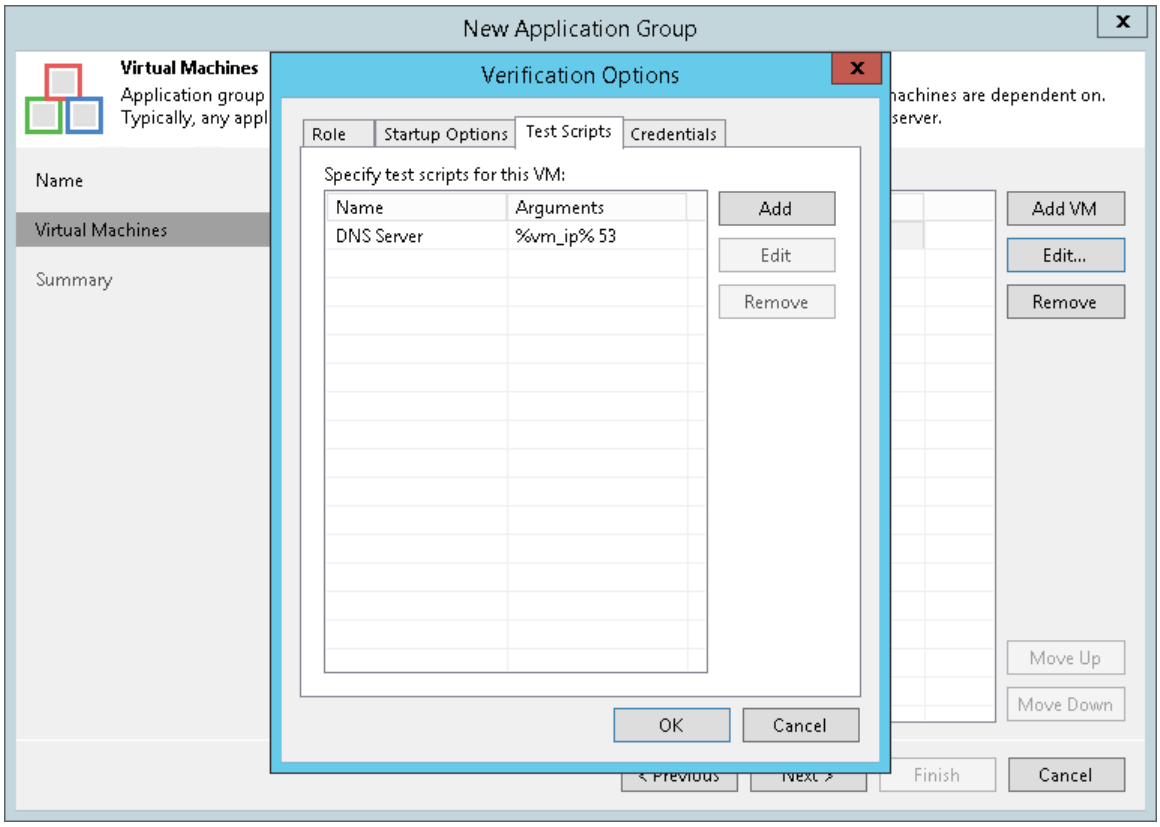
1. In the **Verification Options** window, click the **Test Scripts** tab.
2. Click **Add**.
3. In the **Test Scripts** window, select **Use the following test script**.
4. In the **Name** field, specify a name for the script.

5. In the **Path** field, define a path to an executable script file that must be run to verify the VM. You can do one of the following:
   - If you have your own custom script, define a path to it in the **Path** field.
   - If you do not have a custom script, you can use a standard utility by Veeam, *Veeam.Backup.ConnectionTester.exe*, that probes application communication ports. The utility is located in the installation folder of Veeam Backup & Replication: *%ProgramFiles%|Veeam|Backup and Replication|Backup|Veeam.Backup.ConnectionTester.exe*. Specify this path in the **Path** field.
6. In the **Arguments** field, specify an IP address of the verified VM and the port that you want to probe (if the selected test probes the port). You can use the *%vm_ip%* variable to define the VM IP address or the *%vm_fqdn%* variable to define the VM fully qualified domain name.

   For Microsoft SQL Server, you can also specify a path to the log file in the `%log_path%` argument. For more information, see Backup Recovery Verification Tests.
7. Click **OK** to add the configured test.

To edit test settings, select the test in the list and click **Edit**. To delete a test, select the test in the list and click **Remove**.

> **NOTE**
>
> If a VM performs several roles and runs a number of applications, you can add several verification scripts to verify work of these applications. It is recommended that you specify the maximum startup timeout value and allocate the greatest amount of memory for such VMs.

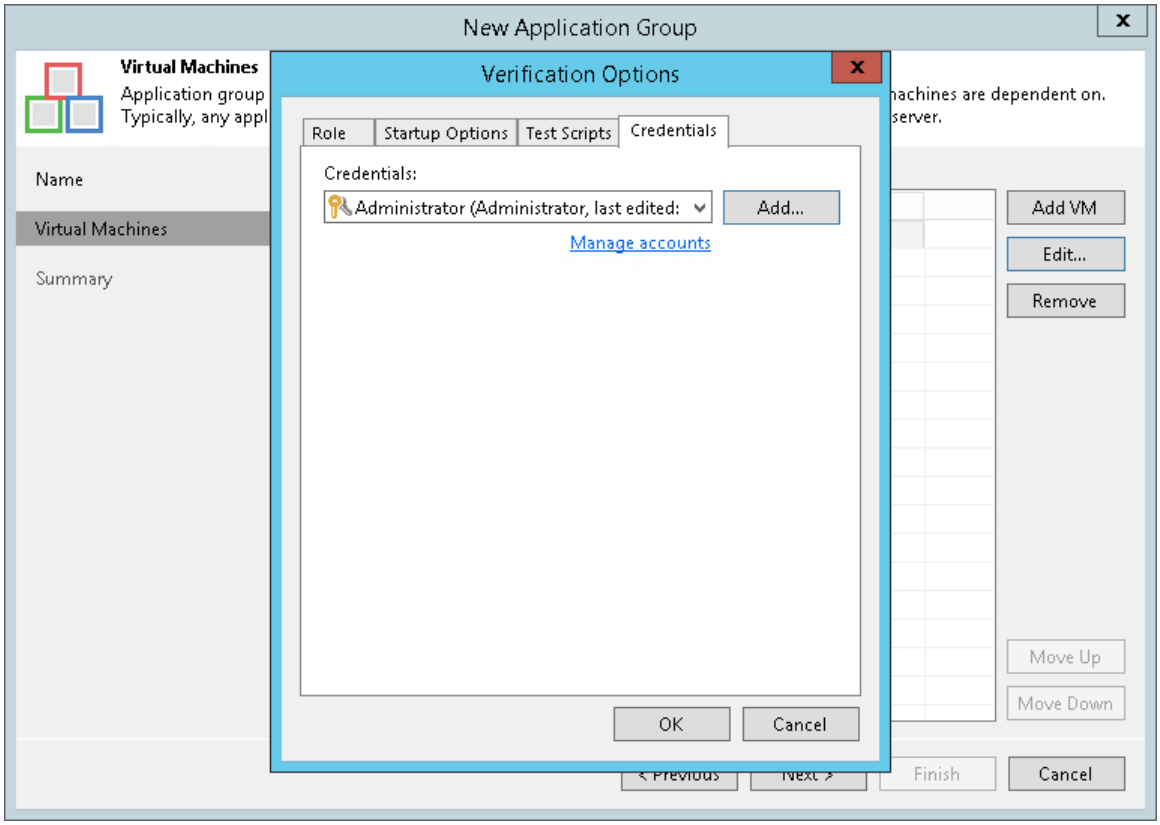


Credentials Settings

By default, to run the verification script Veeam Backup & Replication uses the account under which the Veeam Backup Service is running. If you need to run the script under some other account, you can specify credentials for this account in the application group settings.

1. Click the **Credentials** tab.
2. From the **Credentials** list, select credentials for the account under which you want to run the script.

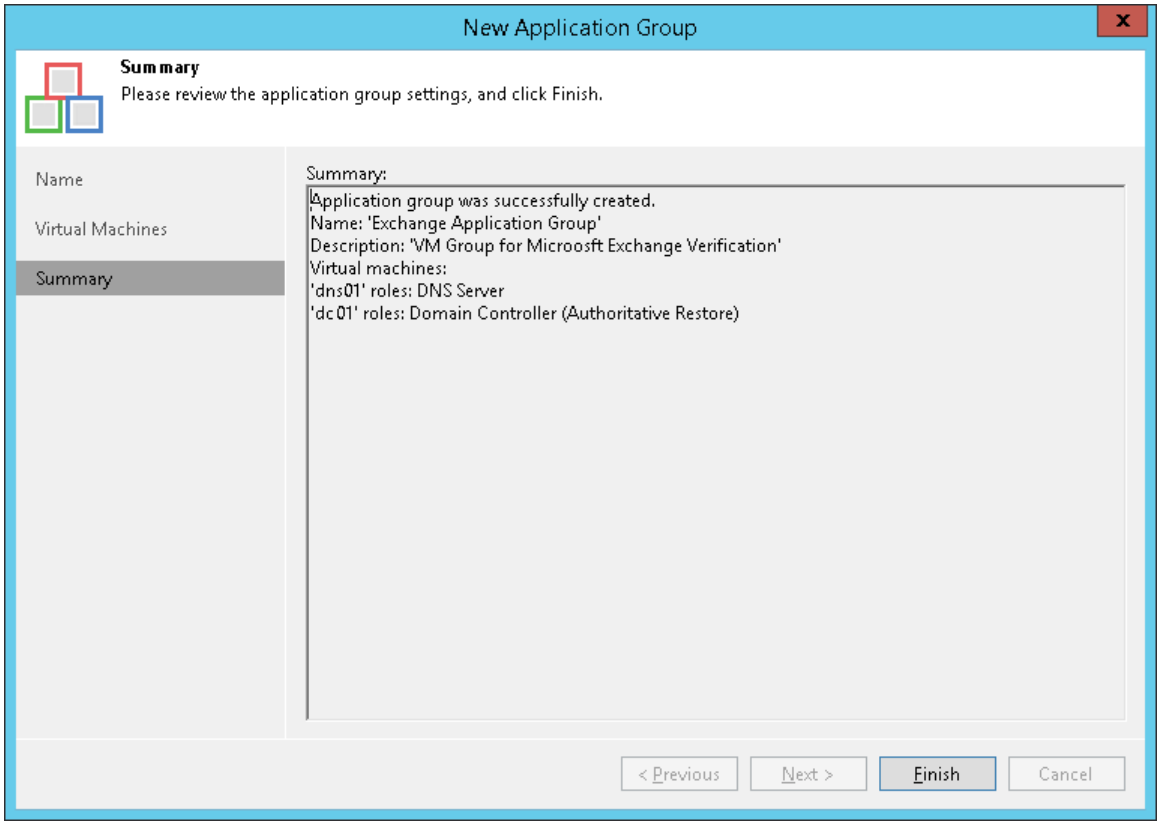
   If you have not set up credentials beforehand, click the **Manage accounts** link or click **Add** on the right to add the credentials. For more information, see Managing Credentials.



Step 5. Review Application Group Settings and Finish Working with Wizard

At the **Summary** step of the wizard, complete the procedure of application group configuration.

1. Review details of the application group.
2. Click **Finish** to save the application group settings and close the wizard.



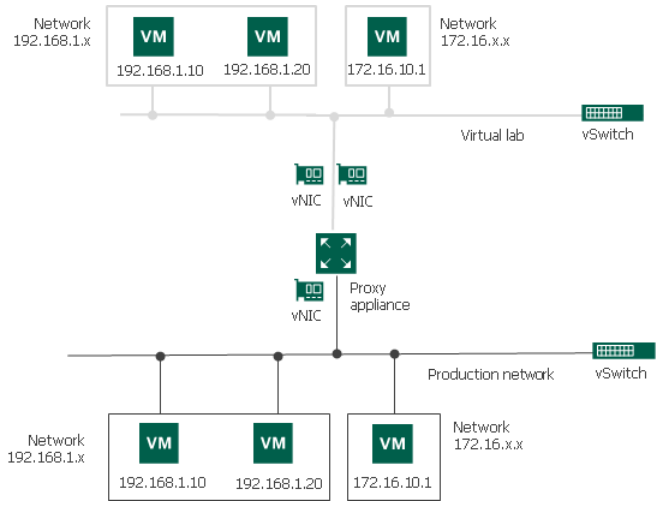
Virtual Lab

The virtual lab is an isolated virtual environment in which Veeam Backup & Replication verifies VMs. In the virtual lab, Veeam Backup & Replication starts VMs from the application group and the verified VM. The virtual lab is used not only for the SureBackup verification procedure, but also for U-AIR, On-Demand Sandbox and staged restore.

The virtual lab itself does not require that you provision extra resources for it. However, VMs running in the virtual lab consume CPU and memory resources of the ESXi host where the virtual lab is deployed. All VM changes that take place during recovery verification are written to redo log files. By default, Veeam Backup & Replication stores redo logs on the datastore selected in the virtual lab settings and removes redo logs after the recovery process is complete.

The virtual lab is fully fenced off from the production environment. The network configuration of the virtual lab mirrors the network configuration of the production environment. For example, if verified VMs and VMs from the application group are located in two logical networks in the production environment, the virtual lab will also have two networks. The networks in the virtual lab will be mapped to corresponding production networks.

VMs in isolated networks have the same IP addresses as in the production network. This lets VMs in the virtual lab function just as if they function in the production environment.



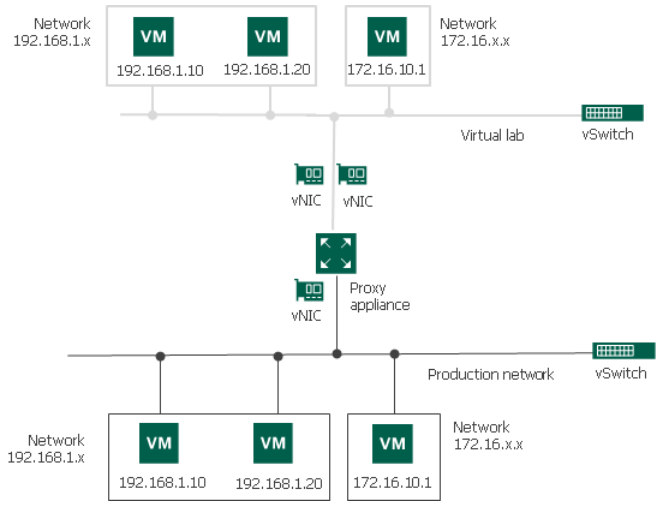
Proxy Appliance

To enable communication between the production environment and isolated networks in the virtual lab, Veeam Backup & Replication uses a proxy appliance. The proxy appliance is an auxiliary Linux-based VM that is deployed on the ESXi host where the virtual lab is created. The proxy appliance VM is assigned an IP address from the production network and placed to the dedicated virtual lab folder and resource pool on the ESXi host.

The proxy appliance is connected to the production network and to the isolated network, so that it has visibility of the production environment and virtual lab. In essence, the proxy appliance acts as a gateway between the two networks — it routes requests from the production environment to VMs in the virtual lab.

To connect to isolated networks, the proxy appliance uses network adapters. Veeam Backup & Replication adds to the proxy appliance one network adapter per every isolated network. For example, if there are two networks in the virtual lab, Veeam Backup & Replication will add two network adapters to the proxy appliance. The network adapter gets an IP address from the isolated network. Typically, this IP address is the same as the IP address of the default gateway in the corresponding production network.

The proxy appliance is an optional component for recovery verification. Technically, you can create a virtual lab without the proxy appliance. However, in this case, you will not be able to perform automatic recovery verification of VMs. VMs will be started from backups in the virtual lab; you will have to access them using the VM console and perform necessary tests manually.

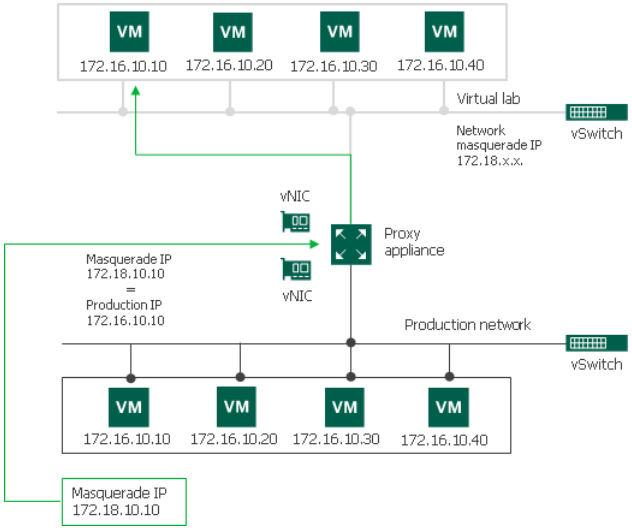


IP Masquerading

To let the traffic into the virtual lab, Veeam Backup & Replication uses masquerade IP addressing.

Every VM in the virtual lab has a masquerade IP address, along with the IP address from the production network. The masquerade IP address resembles the IP address in the production network. For example, if the IP address of a VM is 172.16.1.13, the masquerade IP address may be 172.18.1.13.

The masquerade IP address is an "entry point" to the VM in the virtual lab from the production environment. When you want to access a specific VM in the virtual lab, Veeam Backup & Replication addresses it by its masquerade IP address.



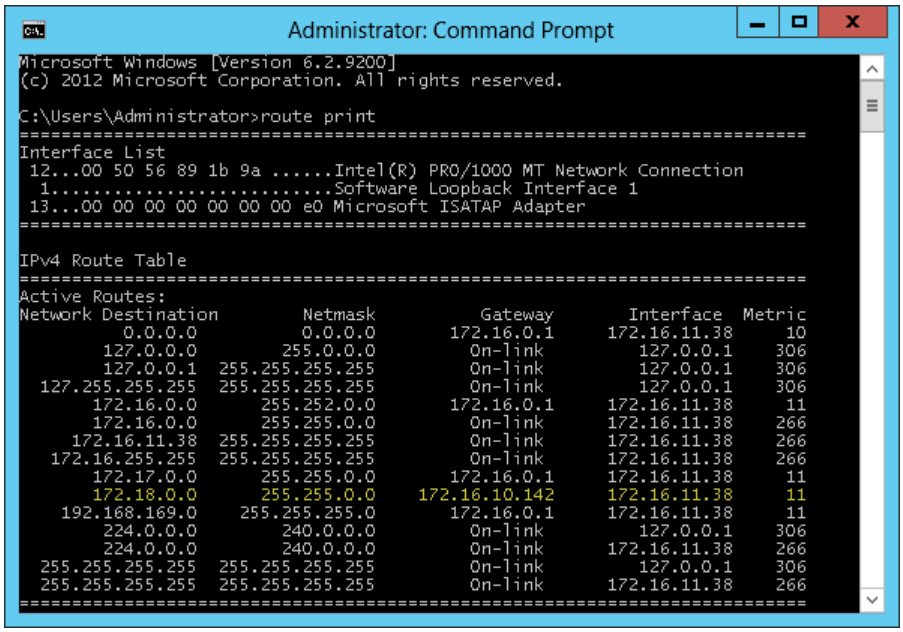
The rules that route requests to VMs in the virtual lab are specified in the routing table on the server from which you want to access VMs in the virtual lab. The routing table can be updated on the following servers:

- **Backup server**. Veeam Backup & Replication automatically creates the necessary static routes in the routing table on the backup server at the moment you launch a SureBackup job and Veeam Backup & Replication starts the virtual lab.

- **Client machine**. If you want to provide your users with access to VMs in the virtual lab, you need to manually update routing tables on their machines and add new static routes. For more information, see Static IP Mapping.

The added static route destines the masquerade network traffic to the proxy appliance. The proxy appliance acts as a NAT device: it resolves the masquerade IP address, replaces it with "real" IP address of the VM from the production network and then directs the request to the necessary VM in the virtual lab. The static route is non-persistent: when you power off the virtual lab, the route is removed from the routing table on the backup server or client machine.

For example, when trying to access a VM with IP address 172.16.10.10 in the isolated network, Veeam Backup & Replication sends a request to the masquerade IP address 172.18.10.10. According to the routing rule added to the IP routing table, all requests are first sent to the next hop — the proxy appliance. The proxy appliance performs address translation, substitutes the masquerade IP address with the IP address in the isolated network, and forwards the request to the necessary VM in the isolated network — in this example, to 172.16.10.10.

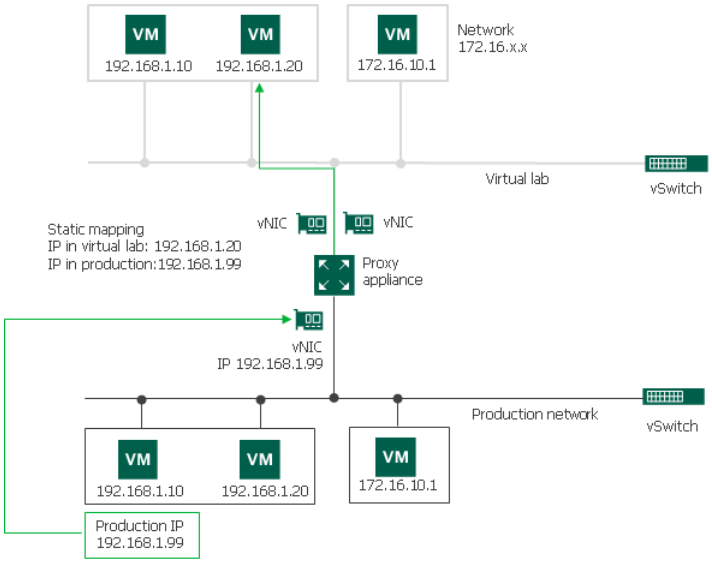


Static IP Mapping

Sometimes it is necessary to provide many clients with access to a restored VM, which is especially the case for user-directed application item-level recovery. For example, you may want to provide users with access to the Microsoft Exchange Server started in the virtual lab using web-based access (like Outlook Web Access). Technically, you may update the routing table on every client machine; however, this will demand a lot of administrative effort.

For such situations, Veeam Backup & Replication lets you get access to a VM in the virtual lab directly from the production environment. To access to a VM in the virtual lab, you must reserve a static IP address in the pool of production IP addresses and map this IP address to the IP address of a VM in the virtual lab.

The static IP address is assigned to the proxy appliance network adapter connected to the production network. IP traffic directed to the specified static IP address is routed by the proxy appliance to the VM in the isolated network.



For example, for a VM with IP address 192.168.1.20 in the isolated network, you can reserve IP address 192.168.1.99 (a free IP address from the production network). As a result, you will be able to use IP address 192.168.1.99 to access the VM in the virtual lab from the production side.

You can also register an alias record in the production DNS server for the reserved IP address. For example, you can register backup.exchange.local as an alias for the IP address 192.168.1.99, and users will be able to access Microsoft Exchange Server by this alias.

Virtual Lab Configuration

For SureBackup recovery verification, Veeam Backup & Replication offers two types of the virtual lab configuration:

- Basic single-host virtual lab
- Advanced single-host virtual lab

> **NOTE**
>
> You can also verify VM backups in Advanced Multi-Host virtual labs with DVS. This scenario can be helpful if you want to test VM backups and VM replicas in the same virtual lab or want to add verified VM backups and replicas to the same SureBackup job.
> For more information, see Advanced Multi-Host Virtual Labs.

Basic Single-Host Virtual Labs

The basic single-host virtual lab can be used if all VMs that you want to verify, VMs from the application group and the backup server are connected to the same network.

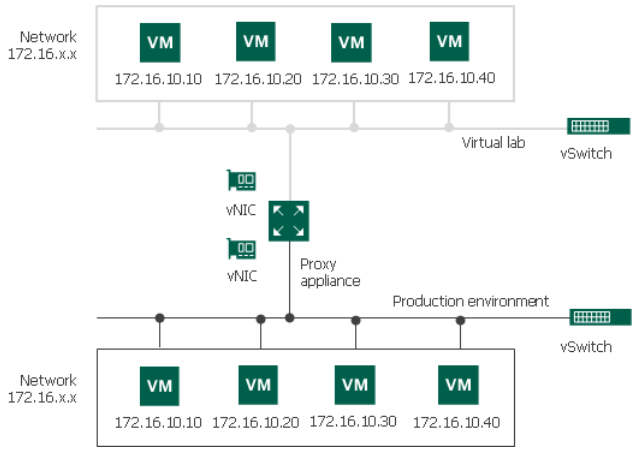
For the basic single-host virtual lab, Veeam Backup & Replication creates one virtual network that is mapped to the corresponding production network. Veeam Backup & Replication automatically adds a number of new objects on the ESXi host where the virtual lab is created:

- A resource pool

- A VM folder
- A standard vSwitch

The vSwitch is only used by the VMs started in the virtual lab. There is no routing outside the virtual lab to other networks.

Veeam Backup & Replication automatically configures all settings for the basic single-host virtual lab. The proxy appliance is also created and configured automatically on the ESXi host where the virtual lab is created.



Advanced Single-Host Virtual Labs

The advanced single-host virtual lab can be used if VMs that you want to verify and/or VMs from the application group are connected to different networks.

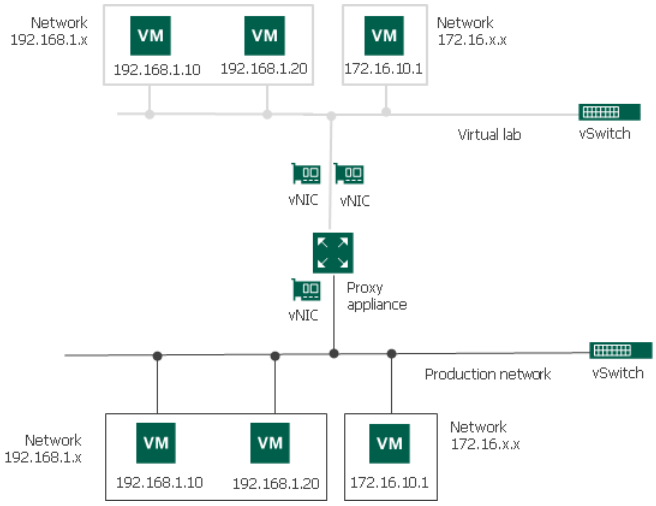
In the advanced single-host virtual lab, Veeam Backup & Replication creates several virtual networks for the virtual lab. The number of virtual networks corresponds to the number of production networks to which verified VMs are connected. Networks in the virtual lab are mapped to production networks.

Veeam Backup & Replication automatically adds a number of new VMware objects on the ESXi host where the virtual lab is created:

- A resource pool
- A VM folder
- A standard vSwitch

The vSwitch is only used by the VMs started in the virtual lab. There is no routing outside the virtual lab to other networks.

When you create an advanced single-host virtual lab, Veeam Backup & Replication configures basic settings for networks that are created in the virtual lab. You need to review these settings and manually adjust them.



Creating Virtual Lab

Before you create a new virtual lab, check prerequisites. Then use the **New Virtual Lab** wizard to create a virtual lab.

Before You Begin

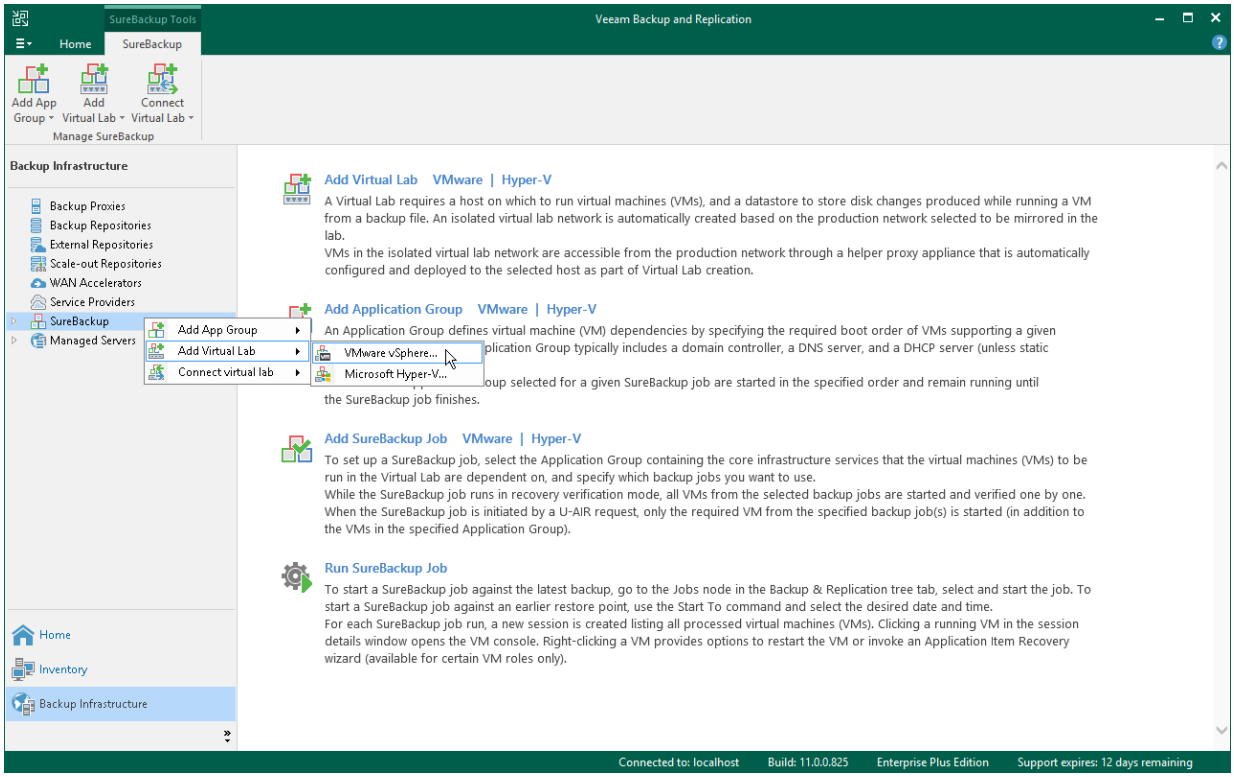
Before you create a virtual lab, check the following prerequisites:

- A valid license for Enterprise edition of Veeam Backup & Replication must be installed on the backup server.
- The ESXi host on which you plan to deploy a virtual lab must have a VMkernel interface. Otherwise the vPower NFS datastore will not be mounted on the ESXi host. For more information, see Veeam vPower NFS Service.
- If you plan to use the advanced multi-host networking mode for VM replicas verification, you must configure a DVS beforehand. For more information, see Advanced Multi-Host Virtual Labs.

Step 1. Launch New Virtual Lab Wizard

To launch the **New Virtual Lab** wizard, do one of the following:

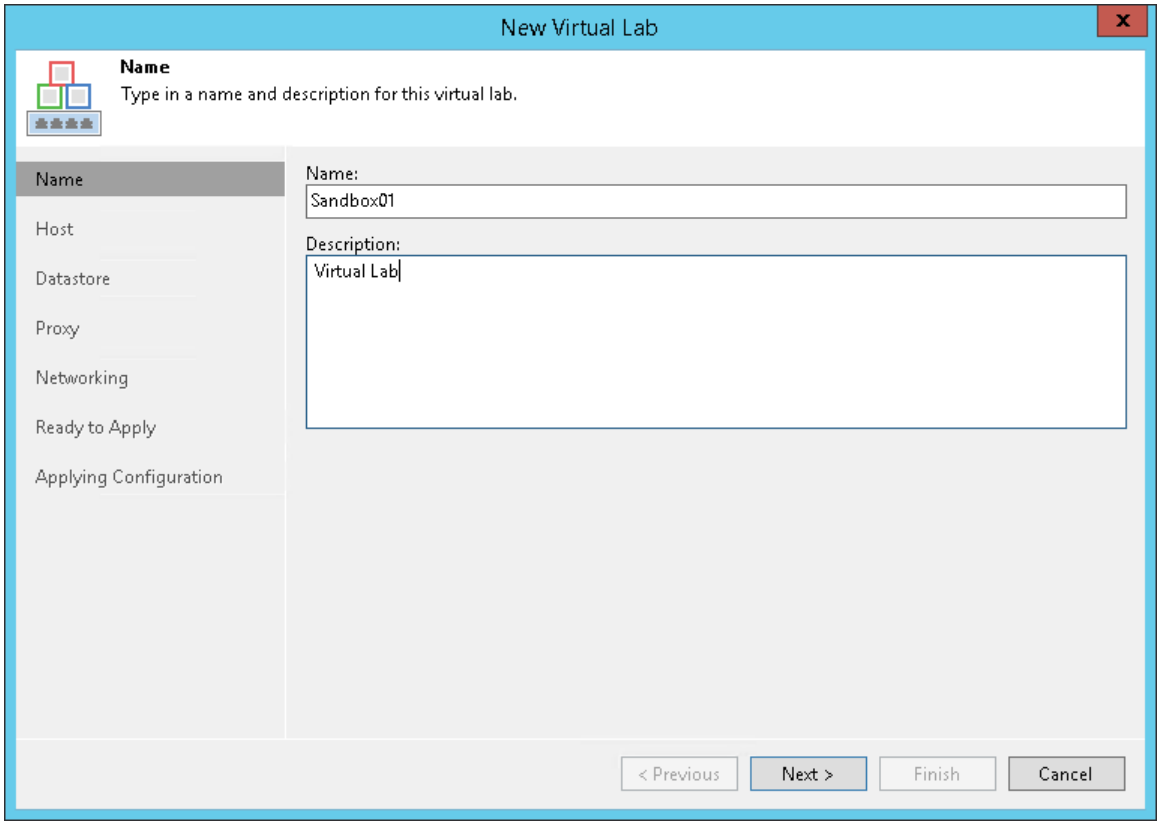
- Open the **Backup Infrastructure** view, in the inventory pane select **SureBackup**. In the working area, click **Add Virtual Lab** > **VMware**.
- Open the **Backup Infrastructure** view, in the inventory pane select **Virtual Labs** node under **SureBackup** and click **Add Virtual Lab** > **VMware** on the ribbon.
- Open the **Backup Infrastructure** view, in the inventory pane right-click **Virtual Labs** under **SureBackup** and select **Add Virtual Lab** > **VMware**.



Step 2. Specify Virtual Lab Name and Description

At the **Name** step of the wizard, specify a name and description for the virtual lab.

1. In the **Name** field, enter a name for the virtual lab.
2. In the **Description** field, provide a description for future reference. The default description contains information about the user who created the virtual lab, date and time when the lab was created.



Step 3. Select Host

At the **Host** step of the wizard, select an ESXi host on which the virtual lab must be created.

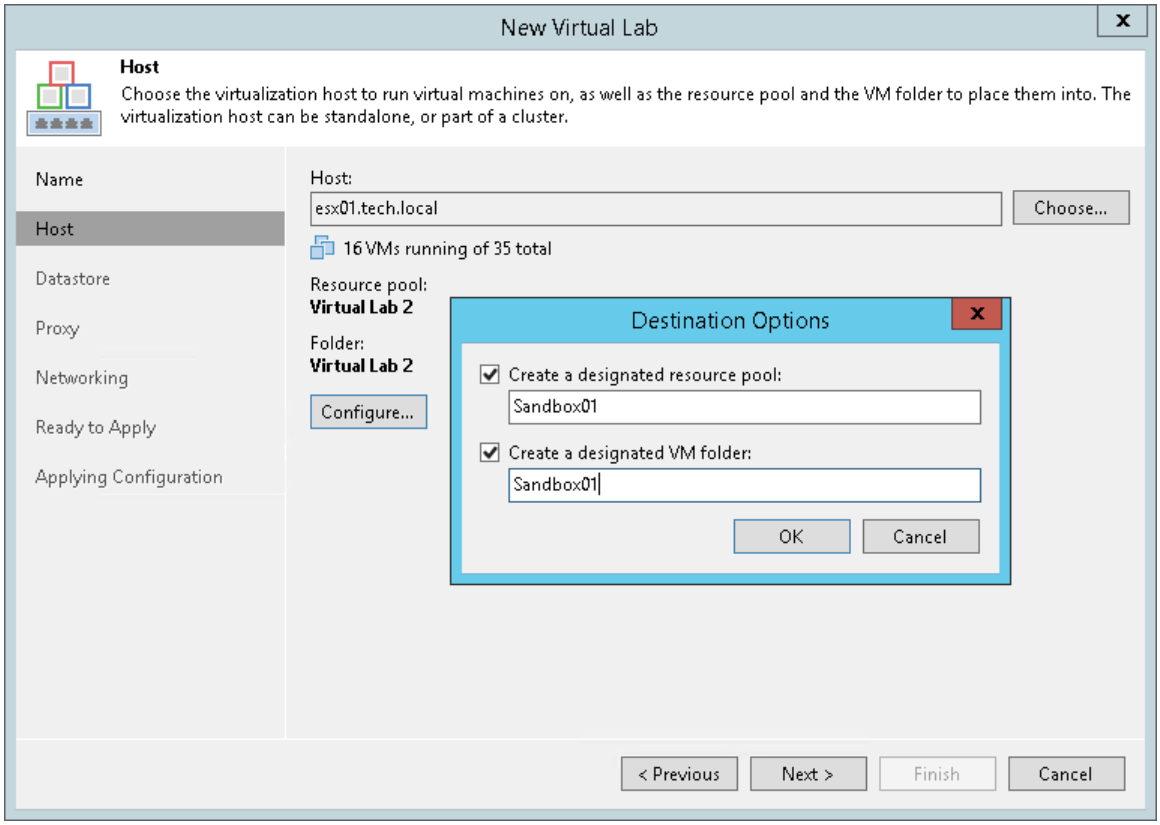
To select an ESXi host:

1. Click **Choose**.
2. Select an ESXi host on which the new virtual lab must be created. You can select a standalone ESXi host or an ESXi host being part of a cluster or vCenter Server hierarchy.
3. For every new virtual lab, Veeam Backup & Replication creates a dedicated folder and resource pool on the ESXi host. By default, the folder and pool have the same name as the virtual lab. To change the name of the folder and/or resource pool, click **Configure**. In the **Destination Options** window, enter the necessary names.

IMPORTANT

You cannot create resource pools in clusters with disabled DRS. If the target host is a part of such a cluster, the **Create a designated resource pool** option will be disabled in the **Destination Options** window. For more information, see this VMware KB article.

You cannot create folders on standalone ESXi hosts or ESXi hosts that are managed by the vCenter Severs but are added to Veeam Backup & Replication as standalone hosts. To overcome this situation, add the corresponding vCenter Server to Veeam Backup & Replication.



Selecting an ESXi Host for VM Replicas Verification

When you select an ESXi host for the virtual lab where VM replicas will be verified, mind the location of verified VM replicas and VM replicas added to the application group:

- If verified VM replicas and VM replicas from the application group are located on the same ESXi host, you must select the ESXi host on which these VM replicas are registered. Verified VM replicas and VMs from the application group will be started on the selected ESXi host. If the application group contains VMs added from VM backups or storage snapshots, these VMs will also be started on the selected ESXi host.

  For this type of virtual lab configuration, you need to choose one of single-host networking modes: Basic single-host or Advanced single-host. For more information, see Selecting a Networking Mode.

- If verified VM replicas and/or VM replicas from the application group are located on different ESXi hosts, you can select any ESXi host in your virtual environment. Veeam Backup & Replication will create the virtual lab on the selected ESXi host. Verified VM replicas and VM replicas from the application group will be started on ESXi hosts where they are registered and connected to the virtual lab with the help of VMware DVS technology.

  The ESXi host on which the virtual lab is created must meet the following requirements:

  - The ESXi host must be located in the same datacenter where VM replicas are registered.
  - The ESXi host must have enough CPU and RAM resources. If the application group contains VMs that are started from backups or storage snapshots, these VMs will be started on the same ESXi host where the virtual lab is located, which will require a lot of resources.
  - For this type of virtual lab configuration, you must use the Advanced multi-host networking mode. For more information, see Selecting a Networking Mode.

Step 4. Select Datastore

At the **Datastore** step of the wizard, you can select where redo logs for verified VMs must be stored. Redo logs are auxiliary files used to keep changes that take place when VMs run in the virtual lab. By default, redo logs are stored on the vPower NFS server. However, you can store redo logs on any datastore in the virtual environment. Redirecting redo logs improves verification performance. As soon as a recovery verification job completes, Veeam Backup & Replication deletes redo logs.

To redirect redo logs:

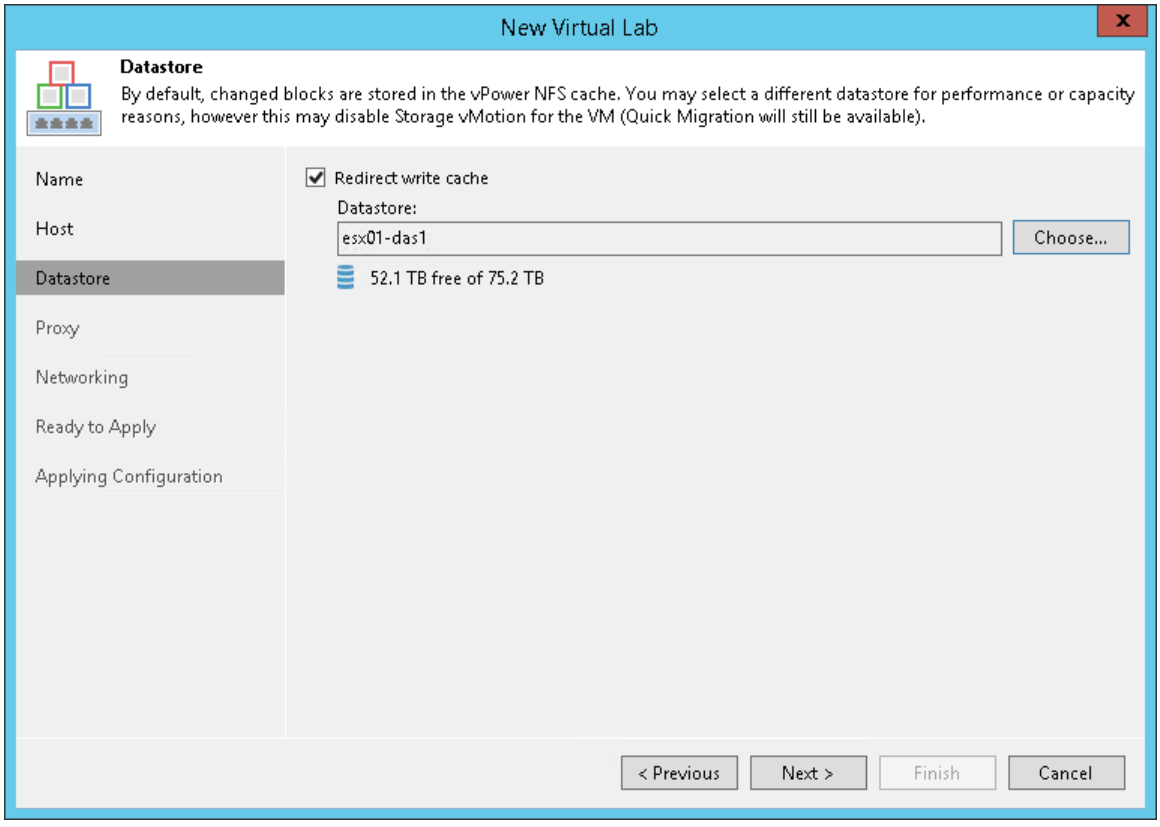
1. Select the **Redirect write cache** check box.

   If you perform staged restore, the **Redirect write cache** option allows you to select a datastore where VM delta files will be stored. Delta files are auxiliary files that keep changes made to a VM during script execution. For more information, see Staged Restore.

2. Click **Choose** and select a datastore from the list.

IMPORTANT

If disks of verified VMs are greater than 2 TB, you must not place redo logs on a VSAN datastore. Otherwise, Veeam Backup & Replication will fail to create snapshots for verified VMs. For more information, see VMware Docs.



Step 5. Set Up Proxy Appliance

At the **Proxy** step of the wizard, configure proxy appliance settings.

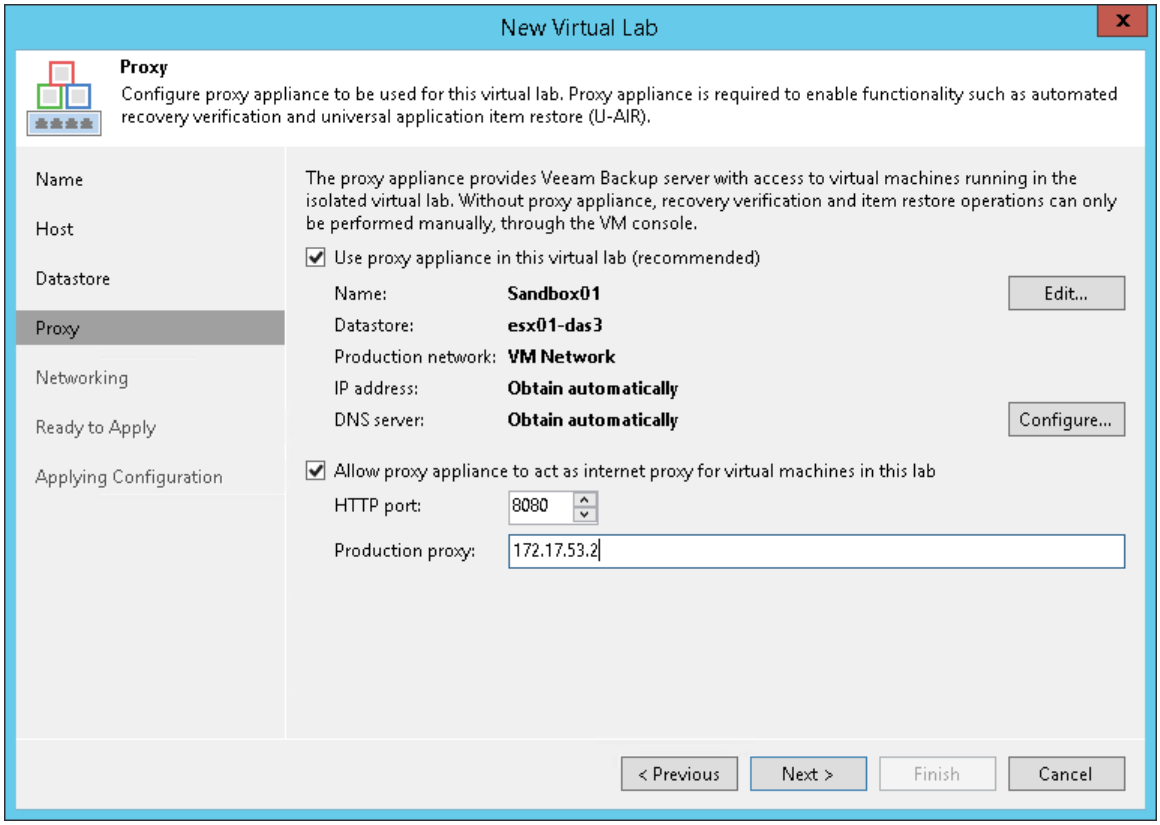
1. Select the **Use proxy appliance in this virtual lab** check box to enable automatic recovery verification of VMs. The proxy appliance acts as a gateway that provides access from the backup server to VMs in the virtual lab. If you do not select this check box, during recovery verification Veeam Backup & Replication will only start VMs in the virtual lab and perform the heartbeat test for VMs. You will have to manually test VMs or perform manual item-level restore over the VM console.
2. By default, the proxy appliance is placed on a datastore with the maximum amount of free space. The default name of the proxy appliance is the virtual lab name that you have specified at the **Name** step of the wizard. To change a name or a datastore for the proxy appliance, click **Edit** and specify a new name or choose a different datastore.
3. Click **Configure** and select a production network in which the proxy appliance will be created. Specify an IP address for the proxy appliance in the production network and settings of the DNS server to be used. You can choose to automatically obtain an IP address for the backup proxy and DNS server settings or set them manually.

> **IMPORTANT**
>
> Mind the following:
>
> - If you assign to the proxy appliance an IP address from the same network where the backup server is located, Veeam Backup & Replication will automatically add a new route to the routing table on the backup server. If you assign to the proxy appliance an IP address from a different network, you will have to manually add a new route to the routing table on the router in the production network. If you do not add a new route, tests and application scripts will fail and you will not be able to access VMs in isolated networks.
> - When Veeam Backup & Replication starts a virtual lab, it verifies if the proxy appliance is available by sending a ping request to it. If the corresponding route is not added to the routing table, the SureBackup job will fail.

4. By default, VMs in the virtual lab work in the isolated environment and do not have access to internet. If you want to let VMs in the virtual lab access the Internet, select the **Allow proxy appliance to act as internet proxy for virtual machines in this lab** check box. In the **Port** field, specify a port for HTTP traffic. By default, port 8080 is used. In the **Production proxy** field, you can optionally specify an IP address or a fully qualified domain name of an Internet-facing proxy server that VMs must use to access internet.
5. On every VM that you plan to start in the virtual lab, adjust proxy settings in the Internet options. To do this, on the VM open **Internet Options** > **Connections** > **LAN Settings** > **Proxy server** and specify an IP address of the proxy appliance on the isolated network and port number.

NOTE

When you allow the proxy appliance to act as an Internet proxy, you enable the HTTP(S) Internet access for VMs in the virtual lab. The proxy appliance does not proxy other protocols (such as ICMP protocol used for ping tests) for VMs in the virtual lab.



Step 6. Select Networking Mode

At the **Networking** step of the wizard, select the type of network settings configuration. The virtual lab configuration depends on objects that you plan to verify in the virtual lab:

- Backups
- Replicas
- VMs from storage snapshots

Selecting Networking Mode for Verifying Backups

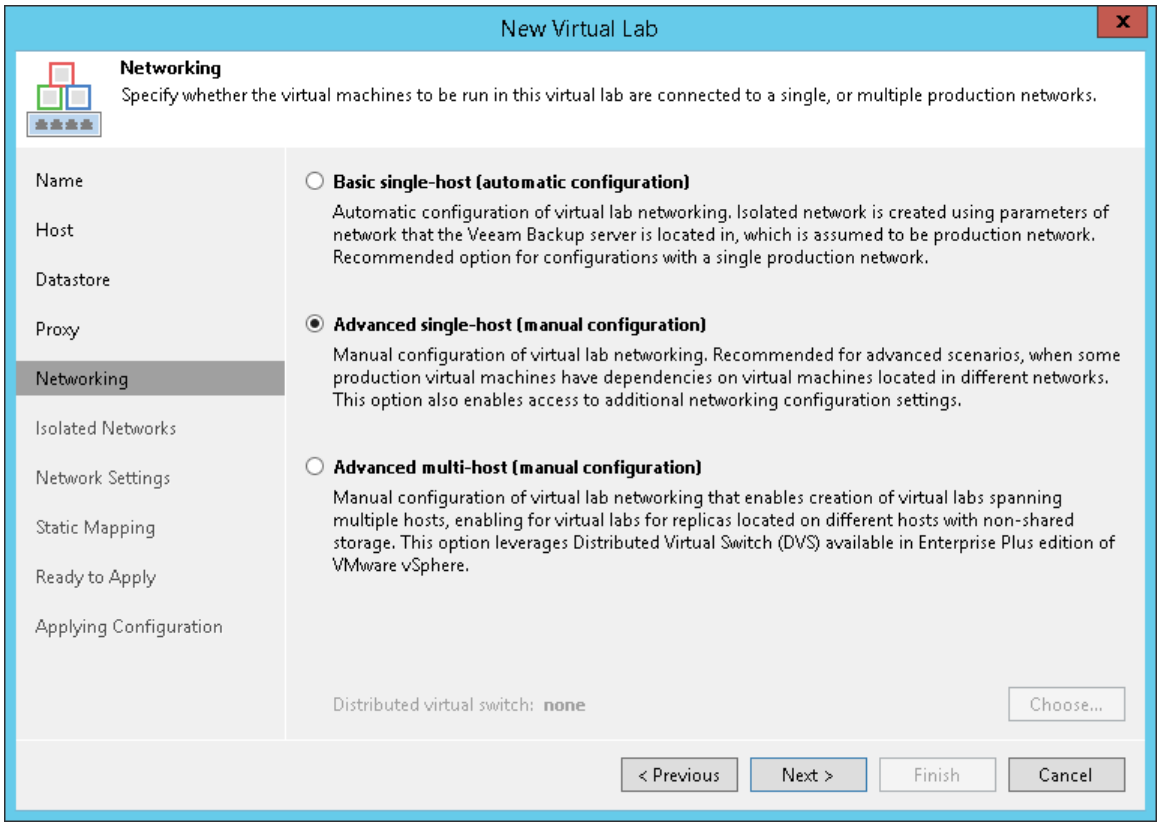
Veeam Backup & Replication offers two networking modes for the virtual lab in which VMs from backups can be verified:

- **Basic single-host**. This networking mode is recommended if all VMs that you plan to verify, VMs from the application group and the backup server are located in the same production network. In this case, Veeam Backup & Replication will automatically define all networking settings for the virtual lab.
- **Advanced single-host**. This networking mode is recommended if VMs that you plan to verify and/or VMs from the application group are located in different networks. In this case, you will have to manually define settings for isolated networks in the virtual lab.

If you select the **Advanced single-host** option, the **New Virtual Lab** wizard will include additional steps for customizing network settings.

NOTE

You can also verify VM backups in Advanced Multi-Host virtual labs with DVS. This scenario can be helpful if you want to test VM backups and replicas in the same virtual lab or want to add verified VM backups and replicas to the same SureBackup job.
For more information, see Advanced Multi-Host Virtual Labs.



Selecting Networking Mode for Verifying Replicas

Veeam Backup & Replication offers three networking modes for the virtual lab in which VM replicas are verified:

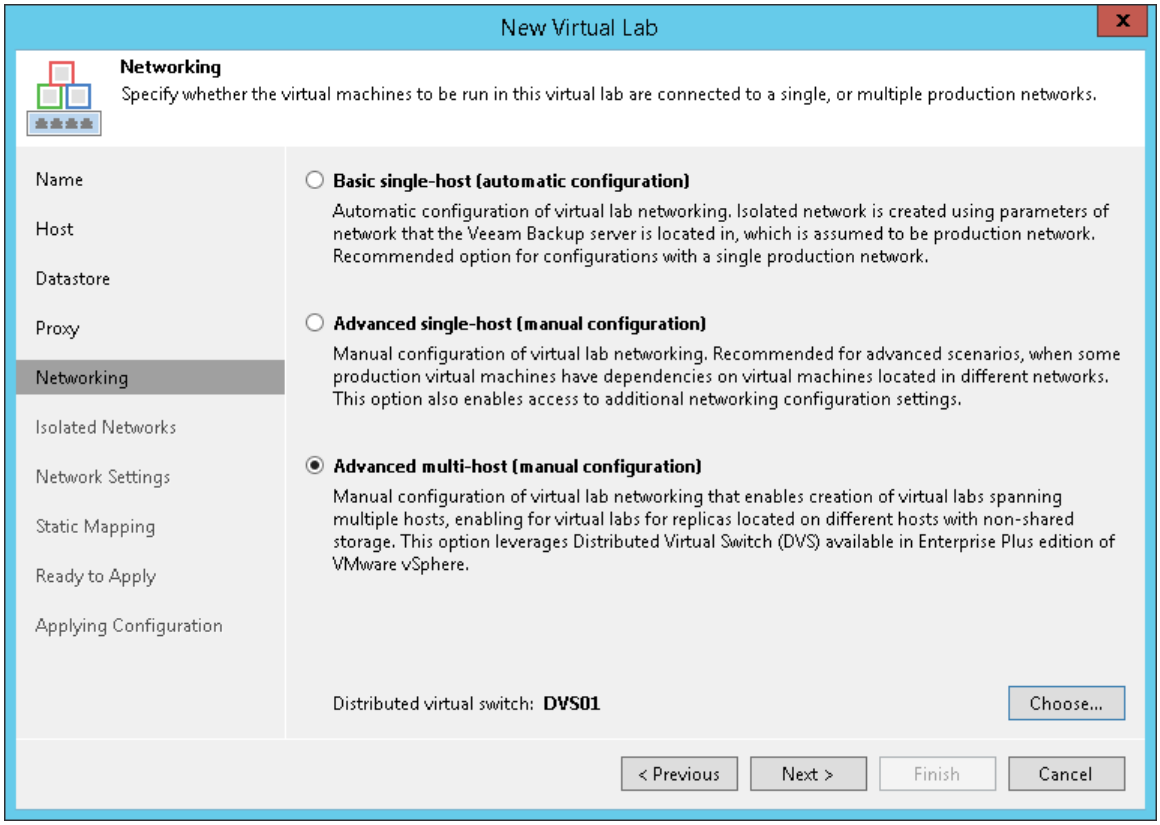
- **Basic single-host**. This type of networking is recommended if VM replicas that you plan to verify are located on the same ESXi host and are connected to the same production network. The backup server must also be located in this network. In this case, Veeam Backup & Replication will automatically define all networking settings for the virtual lab.
- **Advanced single-host**. This type of networking is recommended if VM replicas that you plan to verify are located on the same ESXi host but connected to different networks. In this case, you will have to manually define settings for isolated networks in the virtual lab.
- **Advanced multi-host**. This type of networking is recommended if VM replicas that you plan to verify are located on the different ESXi hosts. For multi-host configuration of the virtual lab, Veeam Backup & Replication uses VMware DVS technology.

  If you select the **Advanced multi-host** option, click **Choose** and select the necessary DVS in your virtual environment. Note that Veeam Backup & Replication does not configure a DVS automatically: you must configure it beforehand.

If the **Advanced single-host** or **Advanced multi-host** option is selected, the **New Virtual Lab** wizard will include additional steps for customizing network settings.

IMPORTANT

For every isolated network in the virtual lab, Veeam Backup & Replication adds a new port group to the DVS. If you use a production DVS, you must isolate port groups created with Veeam Backup & Replication from the production environment. For more information, see Isolated Networks on DVS.



Selecting Network Mode for Verifying VMs on Storage Snapshots

For verifying VMs from storage snapshots, you can select any networking mode.

Step 7. Create Isolated Networks

The **Isolated Networks** step of the wizard is available if you have selected the advanced networking option at the Networking step of the wizard.

At the **Isolated Networks** step of the wizard, you must configure isolated networks to which verified VMs and VMs from the application group will be connected and map these networks to production networks where original VMs are located.

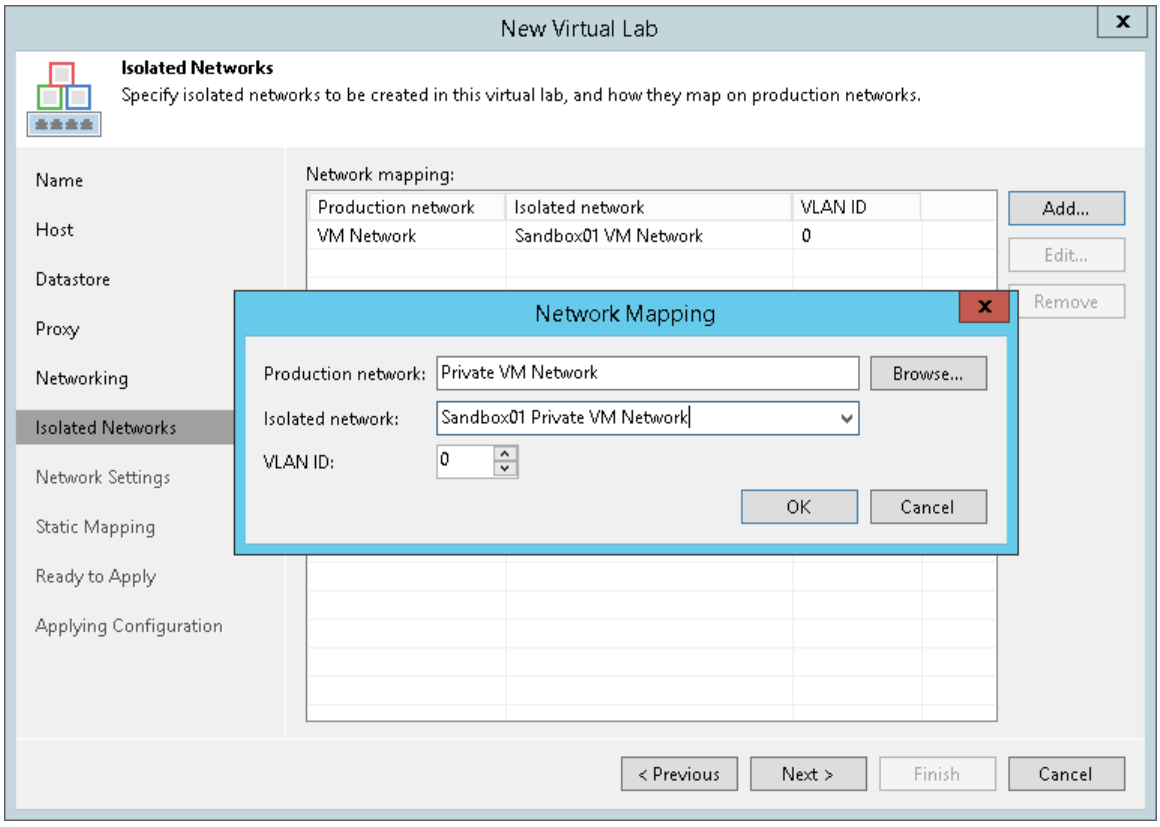
To add a network:

1. Click **Add**.
2. From the **Production network** list, select a production network in which a VM from the application group or verified VM resides.
3. In the **Isolated network** field, specify a name for the isolated network that must be mapped to the selected production network.
4. In the **VLAN ID** field, enter an ID for the created network. In the advanced multi-host virtual lab, VLAN IDs help ensure that the created network is isolated from the production environment. Alternatively, you can manually connect the DVS that you plan to use to the isolated network. For more information, see Advanced Multi-Host Virtual Lab.

IMPORTANT

Mind the following:

- You can map several production networks to the same isolated network. The production networks that you plan to map must have the same network masks and pools of IP addresses.
- You cannot map one production network to several isolated networks.



Step 8. Specify Network Settings

The **Network Settings** step of the wizard is available if you have selected the advanced networking option at the Networking step of the wizard.

At the **Network Settings** step of the wizard, you must specify settings for every created isolated network and define how production networks map to isolated networks in the virtual lab.

Communication between production networks and isolated networks is carried out through vNIC adapters on the proxy appliance. A new vNIC adapter must be added for every isolated network.

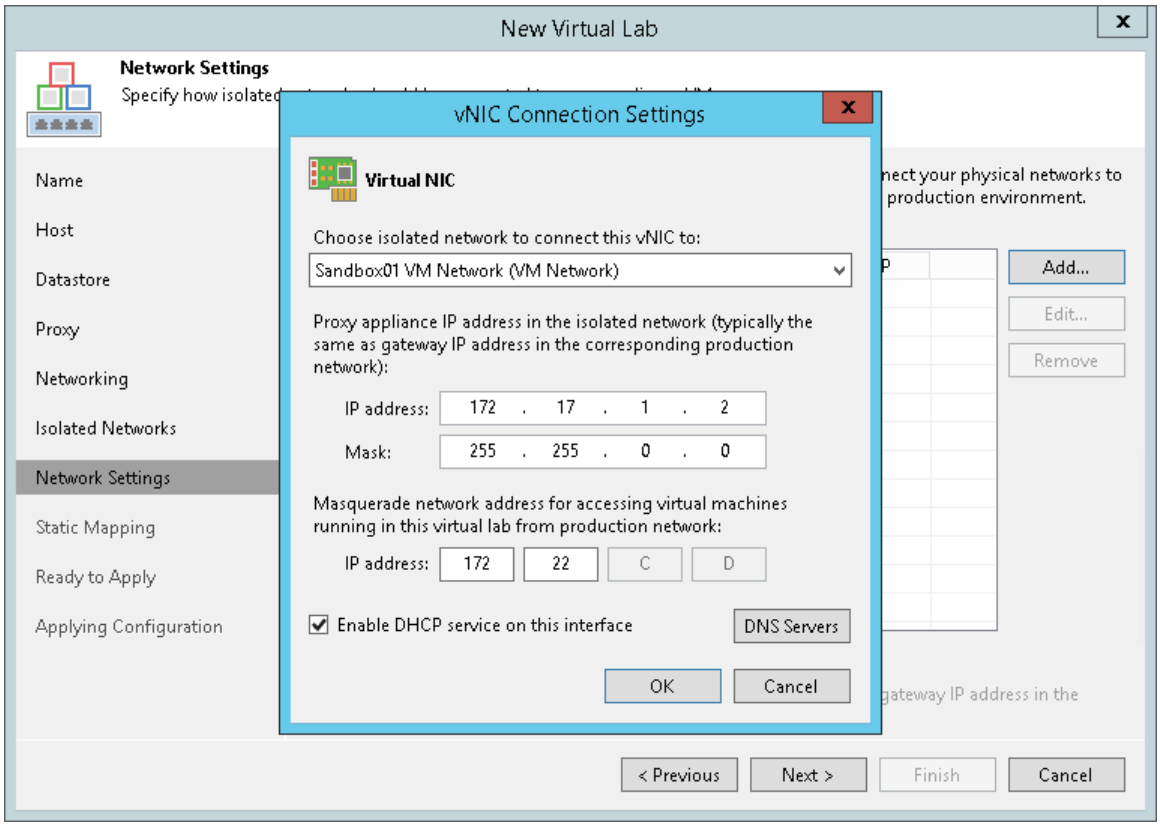
To add a vNIC adapter for an isolated network:

1. At the **Network Settings** step of the wizard, click **Add**.
2. Select a network to which the vNIC adapter must be connected. Specify an IP address that the proxy appliance must have in the isolated network and subnet mask for this isolated network. Typically, the IP address set for the proxy appliance coincides with the IP address of the gateway in the corresponding production network.
3. After you specify the IP address, Veeam Backup & Replication will automatically configure a masquerade IP address for accessing VMs in the virtual lab from the production network. You can change the masquerade network IP address if necessary.
4. Select the **Enable DHCP service on this interface** check box if you need to dynamically assign IP addresses for machines. The assigned IP addresses belong to the current isolated network.
5. Click the **DNS Servers** button and specify settings of a virtualized DNS server if necessary. Click **OK**.
6. To enable communication between isolated networks, select the **Route network traffic between vNICs** check box. Make sure that the IP address of the proxy appliance in the isolated network matches the IP address of the gateway in the production network.

IMPORTANT

Mind the following:

- You cannot assign more than one vNIC to a single isolated network.
- Network addresses specified for different vNIC adapters must belong to different networks. For example, if the first network adapter has the 192.168.0.1 IP address and the network mask is 255.255.255.0, and the second one — 192.168.0.2 and the network mask is 255.255.255.0, such configuration will not work. In this situation, you need to assign to the second adapter an IP address from a different network, for example, 172.16.0.1.
- We do not recommend assigning more than 1 network adapters (or IP addresses) to one production VM, otherwise Veeam Backup & Replication will apply the predefined tests only to one of these IP addresses. For more information on the predifined tests, see Predefined tests.



Step 9. Specify Static IP Mapping Rules

The **Static Mapping** step of the wizard is available if you have selected the advanced networking option at the Networking step of the wizard.

At the **Static Mapping** step of the wizard, you can specify static IP address mapping rules to make VMs in the virtual lab accessible from any machine in the production network.

To add a new rule:

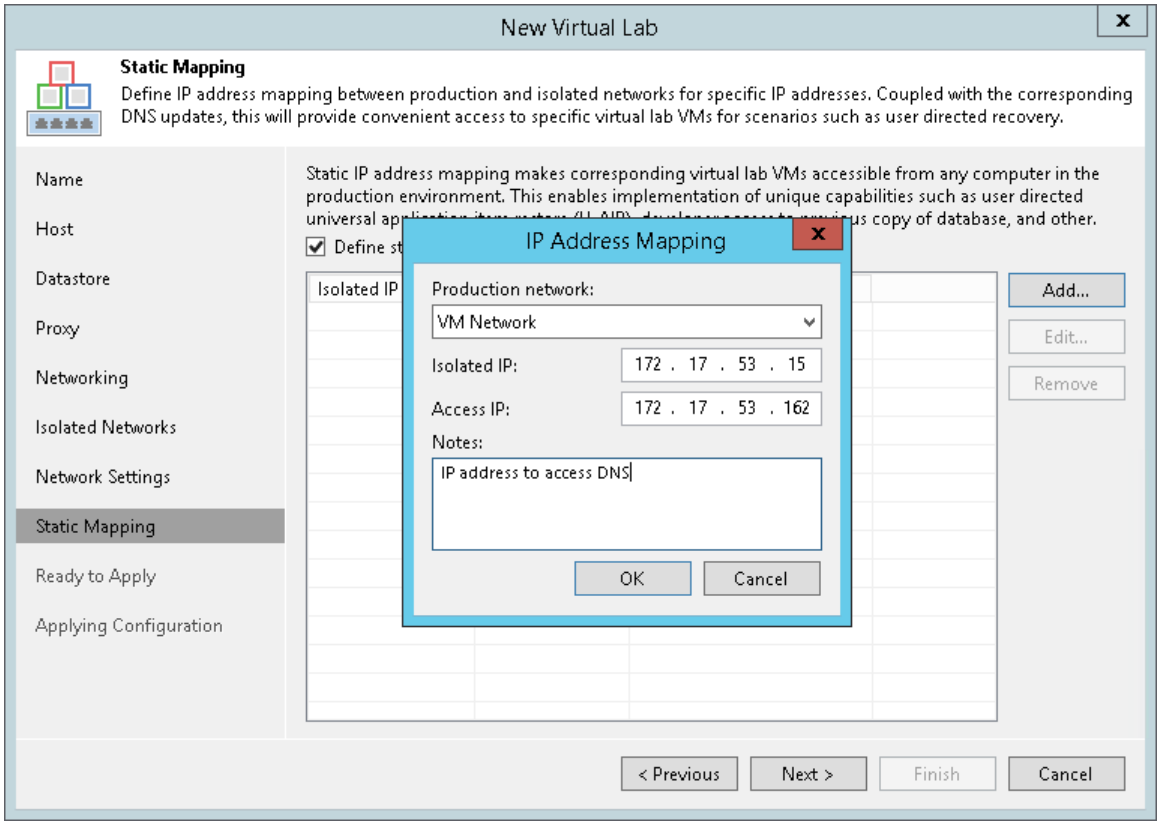
1. Select the **Define static IP address mapping** check box.
2. Click **Add**.
3. In the **IP Address Mapping** window, specify settings of a new rule:
   a. From the **Production network** drop-down list, select a production network in which a VM from the application group or verified VM resides.
   b. In the **Isolated IP** field, specify the IP address of the VM in the production network.
   c. In the **Access IP** field, specify the IP address in the production network that you want to use to access the VM in the virtual lab. You must use an IP address that is not allocated to any machine yet.

> **NOTE**
>
> It is recommended that you assign an access IP from the same subnet where the proxy appliance resides. In the opposite case, you will have to configure routing rules for the access IP manually.

For example, a DNS server that you plan to start in the virtual lab has IP address 172.17.53.15 in the production network. To set static mapping for the DNS server:

1. In the **Isolated IP** field, you must define its production IP address — 172.17.53.15.
2. In the **Access IP** field, you must define any free IP address from the production network, for example, 172.17.53.162.

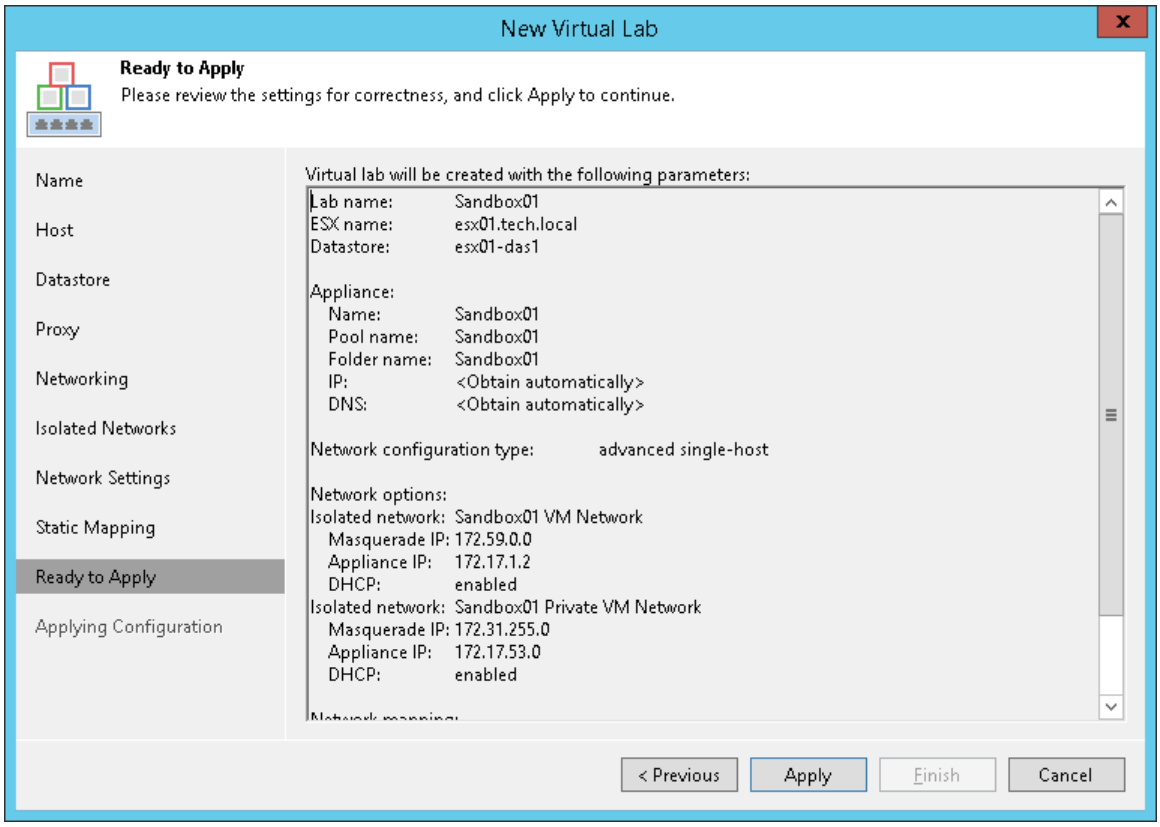
After a virtual lab is created and VMs are started in the virtual lab, you will be able to access the DNS server in the virtual lab from the production environment by IP address 172.17.53.162.



Step 10. Apply Parameters

At the **Ready to Apply** step of the wizard, complete the procedure of virtual lab configuration.

1. Review details of the virtual lab.
2. Click **Apply** to create the virtual lab.
3. At the last step of the wizard, click **Finish** to exit the wizard.



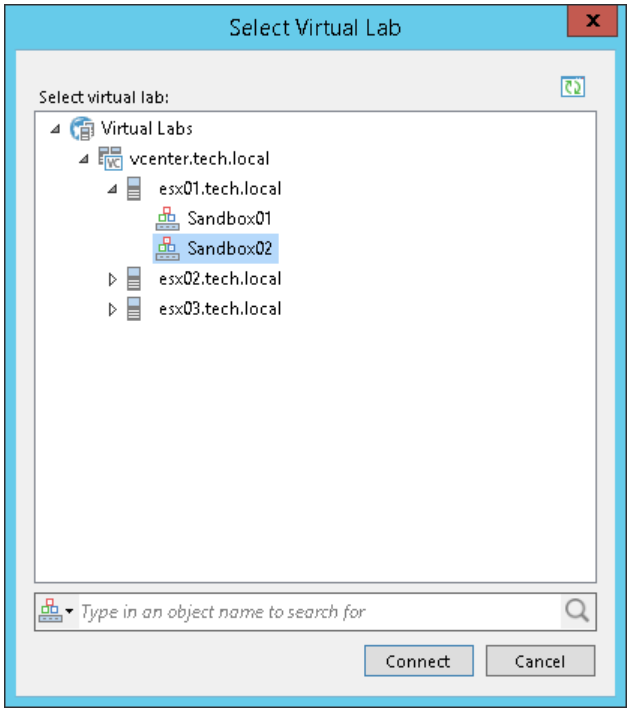
Connecting to Existing Virtual Lab

You can connect an existing virtual lab and use this virtual lab for recovery verification. For example, you can connect to a virtual lab that has been created on another backup server.

To connect to a virtual lab:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Virtual Labs** under **SureBackup** and click **Connect Virtual Lab** > **VMware** on the ribbon or right-click **Virtual Labs** and select **Connect Virtual Lab** > **VMware**.

3. Select the virtual lab and click **Connect**. To quickly find a virtual lab, use the search field at the bottom of the **Select Virtual Lab** window: enter a virtual lab name or a part of it in the field below and press [Enter] on the keyboard.



Editing and Deleting Virtual Labs

You can edit settings of a virtual lab or delete the virtual lab.

Always use Veeam Backup & Replication to modify or delete a virtual lab. If you edit virtual lab settings or delete any of its components from outside, for example, in vSphere Client, the lab will be corrupted and its component such as the created vSwitch, resource pool will remain in the virtual infrastructure.

Editing Virtual Labs

To edit settings of a virtual lab:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Virtual Labs** under **SureBackup**.
3. In the working area, select a virtual lab and click **Edit Virtual Lab** on the ribbon or right-click the virtual lab and select **Properties**.
4. Edit virtual lab settings as required.

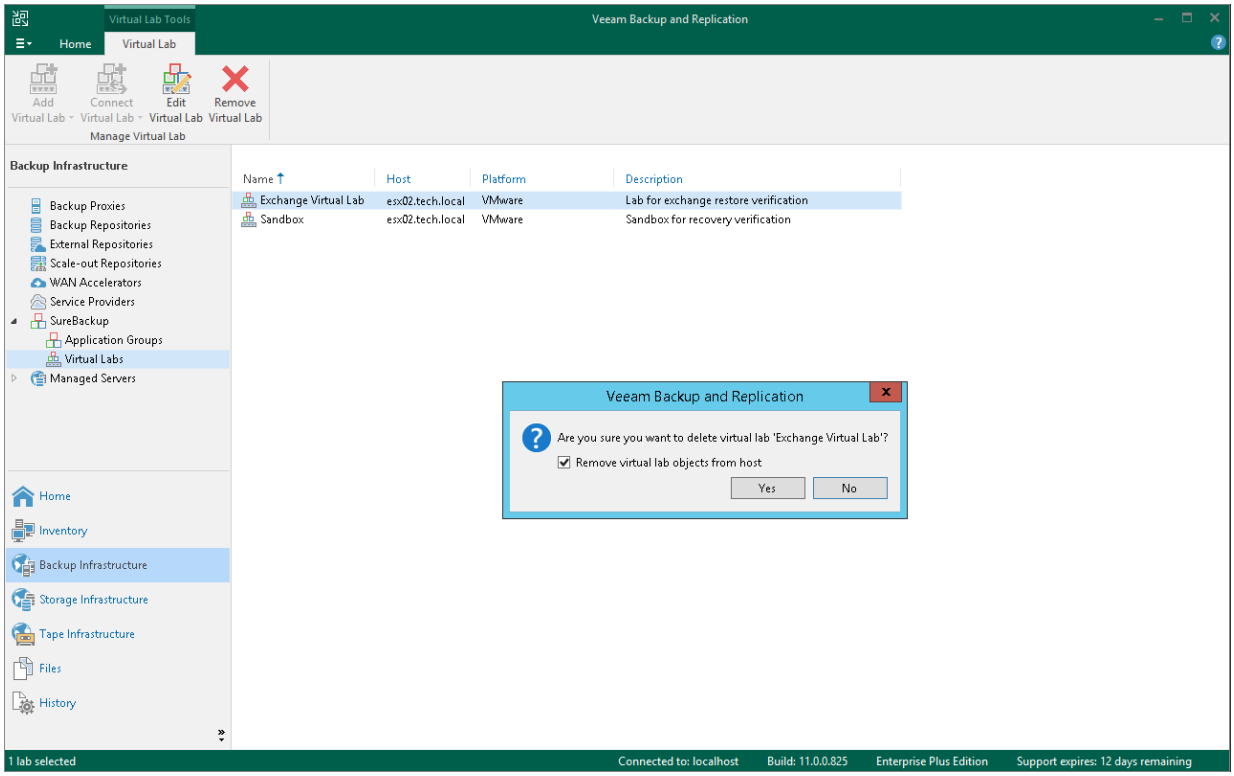
Removing Virtual Labs

To remove a virtual lab:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Virtual Labs** under **SureBackup**.

3. In the working area, select a virtual lab and click **Remove Virtual Lab** on the ribbon or right-click the virtual lab and select **Delete**.

4. If you want to remove virtual lab object from the virtual infrastructure, in the displayed window select the **Remove virtual lab objects from host** check box. If you do not enable this option, Veeam Backup & Replication will disconnect the virtual lab from the backup server. You will be able to connect to this virtual lab later.



SureBackup Job

A SureBackup job is a task for recovery verification. The SureBackup job aggregates all settings and policies of the recovery verification task, such as application group and virtual lab to be used, VM backups that must be verified in the virtual lab and so on. You can run the SureBackup job manually or schedule it to run automatically.

When a SureBackup job runs, Veeam Backup & Replication first creates an environment for recovery verification:

1. Veeam Backup & Replication starts the virtual lab.
2. In the virtual lab, Veeam Backup & Replication starts VMs from the application group in the required order. VMs from the application group remain running until the verified VMs (VMs from the linked job) are booted from backups and tested.

   If Veeam Backup & Replication does not find a valid restore point for any of VMs from the application group, the SureBackup job fails.
3. When the virtual lab is ready, Veeam Backup & Replication starts verified VMs (VMs from the linked job) to the necessary restore point and, depending on the job settings, verifies them one by one or creates several streams and verifies a number of VMs simultaneously.

   If Veeam Backup & Replication does not find a valid restore point for any of verified VMs, verification of this VM fails, but the job continues to run.

By default, you can start and test up to three VMs at the same time. You can also increase the number of VMs to be started and tested simultaneously. Keep in mind that if these VMs are resource demanding, performance of the SureBackup job as well as performance of the ESXi host on which the virtual lab resides may decrease.

Once the verification process is complete, VMs from the application group are powered off. Optionally, you can leave the VMs from the application group running to perform manual testing or enable user-directed application item-level recovery.

In some cases, the SureBackup job schedule may overlap the schedule of the backup job linked to it. The backup file may be locked by the backup job and the SureBackup job will be unable to verify such backup. In this situation, Veeam Backup & Replication will not start the SureBackup job until the corresponding backup job is over.

To overcome the situation of job overlapping, you may chain the backup and SureBackup jobs or define the timeout period for the SureBackup job. For more information, see Specifying Job Schedule.

> **NOTE**
>
> VMs from the application group and verified VMs must belong to the same platform — VMware or Hyper-V. Mixed scenarios are not supported.

SureBackup Job Processing

The recovery verification process includes the following steps:

1. **Getting virtual lab configuration**. Veeam Backup & Replication gets information about configuration of the virtual lab where verified VMs must be started.
2. **Starting virtual lab routing engine**. Veeam Backup & Replication starts a proxy appliance used as a gateway to provide access to the virtual lab.

3. **Performing malware scan**. If the recovery verification process includes malware scan, Veeam Backup & Replication scans VM data with antivirus software.

   After the malware scan is complete, Veeam Backup & Replication registers the VM on the selected ESXi host, powers it on, and runs recovery verification tests for this VM.

   Veeam Backup & Replication verifies VMs sequentially — one after another. For example, when the malware scan and recovery verification tests for VM *A* complete, Veeam Backup & Replication verifies VM *B*, and so on.

4. **Publishing**. Veeam Backup & Replication creates a vPower NFS datastore with a VM backup and registers it on the selected ESXi server. Veeam Backup & Replication does not deploy the whole VM from the backup file, it deploys VM configuration files only. Virtual disks are deployed per force and per required data blocks.

5. **Reconfiguring**. Veeam Backup & Replication updates configuration files for VMs that must be started in the isolated network.

6. **Registering**. Veeam Backup & Replication registers the verified VM on the selected ESXi host.

7. **Configuring DC**. If a verified VM has the Domain Controller or Global Catalog role, the VM is reconfigured.

8. **Powering on**. Veeam Backup & Replication powers on the verified VM in the isolated network.

9. **Performing heartbeat test**. Veeam Backup & Replication checks whether the VMware Tools heartbeat signal (green or yellow) is coming from the VM or not. If the VM has no VMware Tools, the test is not performed, and a notification is written to the session details.

10. **Running ping tests**. Veeam Backup & Replication checks if the VM responds to the ping requests or not. If the VM has no NICs and mapped networks for them and/or has no VMware Tools installed, the ping test is not performed, and a notification is written to the session details.

11. **Application initialization**. Veeam Backup & Replication waits for the applications installed in the VM, for example, Microsoft SQL Server, to start. The application initialization period is defined in settings of the SureBackup job and by default equals to 120 sec. Depending on the software installed in a VM, the application initialization process may require more time than specified in the job settings. If applications installed in a VM are not initialized within the specified period of time, test scripts can be completed with errors. If such an error situation occurs, you need to increase the **Application initialization timeout** value and start the job once again.

12. **Running test scripts**. Veeam Backup & Replication runs scripts to test whether the application installed in the VM is working correctly or not. If the VM has no VMware Tools installed and/or there are no NICs and mapped networks for them, Veeam Backup & Replication skips tests that use the *%vm_ip%* and *%vm_fqdn%* variables as the IP address and fully qualified domain name of the VM cannot be determined.

    Test results are written to the job session details. To define whether the script has completed successfully or not, Veeam Backup & Replication uses return codes. If the return code is equal to 0, the script is considered to complete successfully. Other values in the return code mean that the script has failed.

13. **Powering off**. After all tests have been performed, Veeam Backup & Replication powers off the verified VM.

14. **Unregistering**. Veeam Backup & Replication unregisters the verified VM on the selected ESXi host.

15. **Clearing redo logs**. Veeam Backup & Replication deletes redo logs from the datastore in the production environment. Redo logs store changes made to the VM while it is running from the backup file.

16. **Unpublishing**. Veeam Backup & Replication unpublishes the content of the backup file on the ESXi host.

17. **Running backup validation test**. After a VM has been verified, powered off and unpublished, Veeam Backup & Replication runs a CRC check to verify the VM backup at the file level and make sure that this file is not corrupted.

18. **Stopping virtual lab engine**. Veeam Backup & Replication powers off the proxy appliance in the virtual lab.
19. **Deleting network routes**. Veeam Backup & Replication deletes added network routes from the routing table on the backup server.

Stabilization Algorithm

To be able to perform tests for a verified VM without errors, Veeam Backup & Replication needs to know that the VM is ready for testing. To determine this, Veeam Backup & Replication waits for the VM to reach a stabilization point: that is, waits for the VM to boot completely and report it is ready for tests. After the stabilization point has been established, Veeam Backup & Replication can start performing heartbeat tests, ping tests and running test scripts against the VM.

Veeam Backup & Replication establishes the stabilization point with the help of VMware parameters that it gets from the VM. Depending on the VM configuration, it uses one of three algorithms to do that:

- **Stabilization by IP**. This algorithm is used if the VM has VMware Tools installed, there are NICs and mapped networks for these NICs. In this case, Veeam Backup & Replication waits for an IP address of the VM for mapped networks that is sent by VMware Tools running in the VM. The sent IP address must be valid and must not change for a specific period of time. For more information, see Recovery Verification Tests.
- **Stabilization by heartbeat**. This algorithm is used if the VM has VMware Tools installed but there are no vNICs and mapped networks for them. In this case, Veeam Backup & Replication waits for a green or yellow heartbeat signal from the VM. The signal is sent by VMware Tools running in the VM.
- **Stabilization by Maximum allowed boot time**. This algorithm is used if the VM has neither VMware Tools installed, nor NICs and mapped networks for them. In this case, Veeam Backup & Replication waits for the time specified in the **Maximum allowed boot time** field, which is considered to be a stabilization period for the VM. Once this time interval is exceeded, Veeam Backup & Replication considers that the VM is successfully booted and is ready for testing.

When the stabilization point has been established, Veeam Backup & Replication runs ping, heartbeat tests and performs test scripts against the verified VM.

The stabilization process cannot exceed the time specified in the **Maximum allowed boot time** field. For this reason, you should be careful when specifying this value. Typically, a VM started by a SureBackup job requires more time to boot than a VM started in the production environment. If the stabilization point cannot be determined within the **Maximum allowed boot time**, the recovery verification process is finished with the timeout error. When such an error occurs, you need to increase the **Maximum allowed boot time** value and start the job again.

Creating SureBackup Job

To create a new SureBackup job, use the **New SureBackup Job** wizard.

Before You Begin

Before you create and start a SureBackup job, check the following prerequisites:

- A valid license for Enterprise edition of Veeam Backup & Replication must be installed on the backup server.
- All applications and services on which verified VMs are dependent must be virtualized in your environment.

- You must create or connect a virtual lab. For more information, see sections Creating Virtual Lab and Connecting to Existing Virtual Lab.
- If you plan to scan VM data for malware, check requirements and limitations.
- If you plan to verify VMs with a ping test, the firewall on tested VMs must allow ping requests.
- If you plan to verify VMs with a heartbeat test, VMware Tools must be installed in tested VMs.
- [For storage snapshots] The storage system must be added to the backup infrastructure.

Mind the following limitations:

- Verified VM replicas must be in the *Ready* state. If a VM replica is in the *Failover* or *Failback* state, you will not be able to verify it with the SureBackup job.
- You cannot link to SureBackup jobs VMs from backups of vCloud Director VMs, backups created with backup copy jobs and backups stored in cloud backup repositories.
- The source backup or replication job has a higher priority than the SureBackup job. If the source backup or replication job starts when the SureBackup job is running, and this job is about to modify the restore point from which the VM is started, Veeam Backup & Replication automatically powers off VMs in the virtual lab and completes the SureBackup job.
- The Microsoft SQL Server Checker script runs on the backup server side. For this reason, Named Pipes or TCP/IP connections must be enabled for the Microsoft SQL Server running in the virtual lab. For more information, see Microsoft Docs.

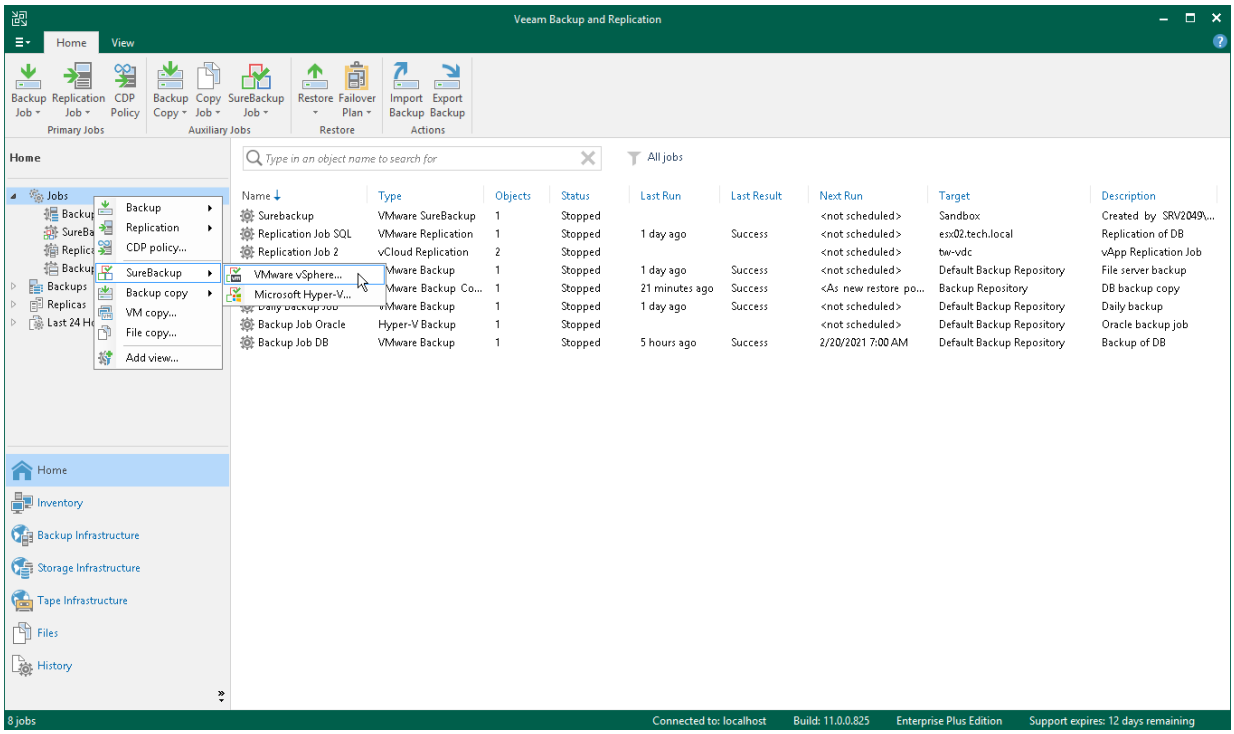
Step 1. Launch New Sure Backup Job Wizard

To launch the **New SureBackup Job** wizard, do either of the following:

- Open the **Backup Infrastructure** view, in the inventory pane select **SureBackup**. In the working area, click **Add SureBackup Job** > **VMware vSphere**.
- Open the **Home** view. On the **Home** tab, click **SureBackup Job** > **VMware vSphere** on the ribbon.
- Open the **Home** view. In the inventory pane, right-click **SureBackup** under **Jobs** and select **SureBackup** > **VMware vSphere**.

  You can use this method if you already have at least one SureBackup job. If there are no SureBackup jobs, the **SureBackup** node will not be displayed in the inventory pane. In this case, you can right-click **Jobs** in the inventory pane and select **SureBackup** > **VMware vSphere**.

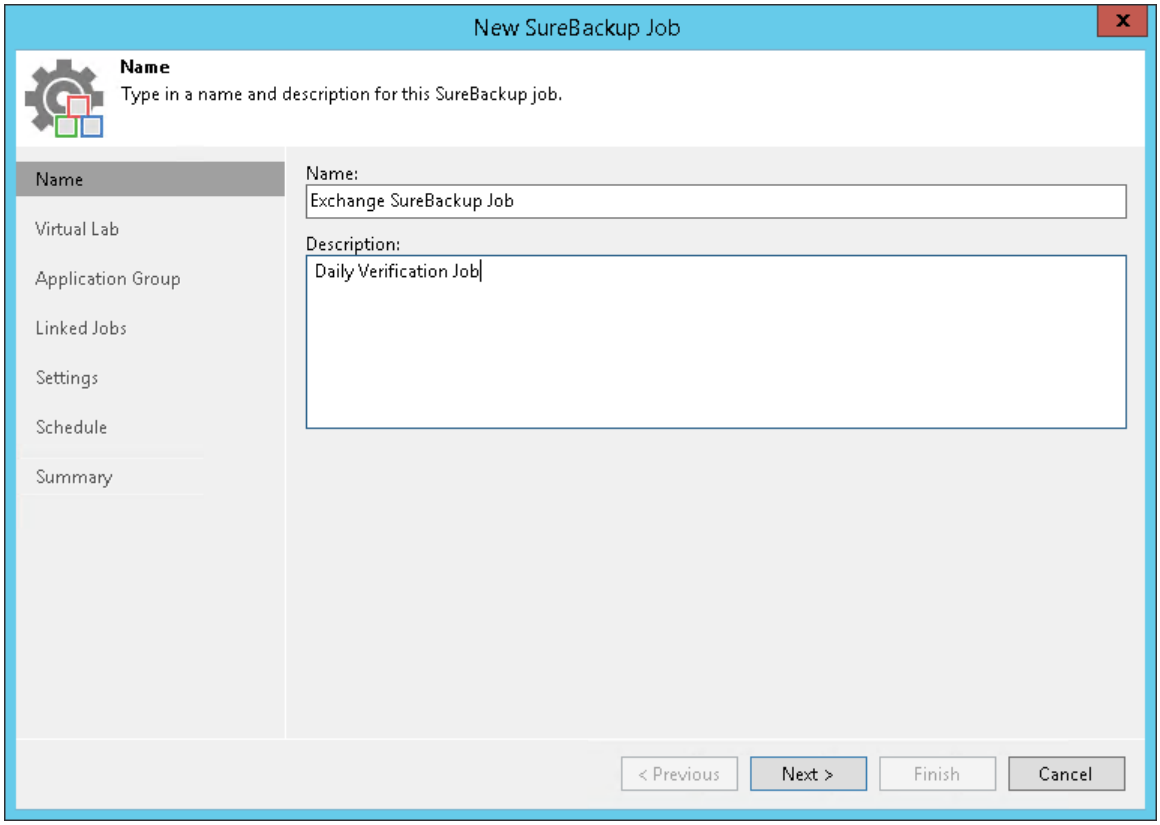
> **NOTE**
>
> SureBackup UI elements become available in the Veeam Backup & Replication console only after you create or connect a virtual lab. For more information, see sections Creating Virtual Lab and Connecting to Existing Virtual Lab.



Step 2. Specify Job Name and Description

At the **Name** step of the wizard, specify a name and description for the SureBackup job.

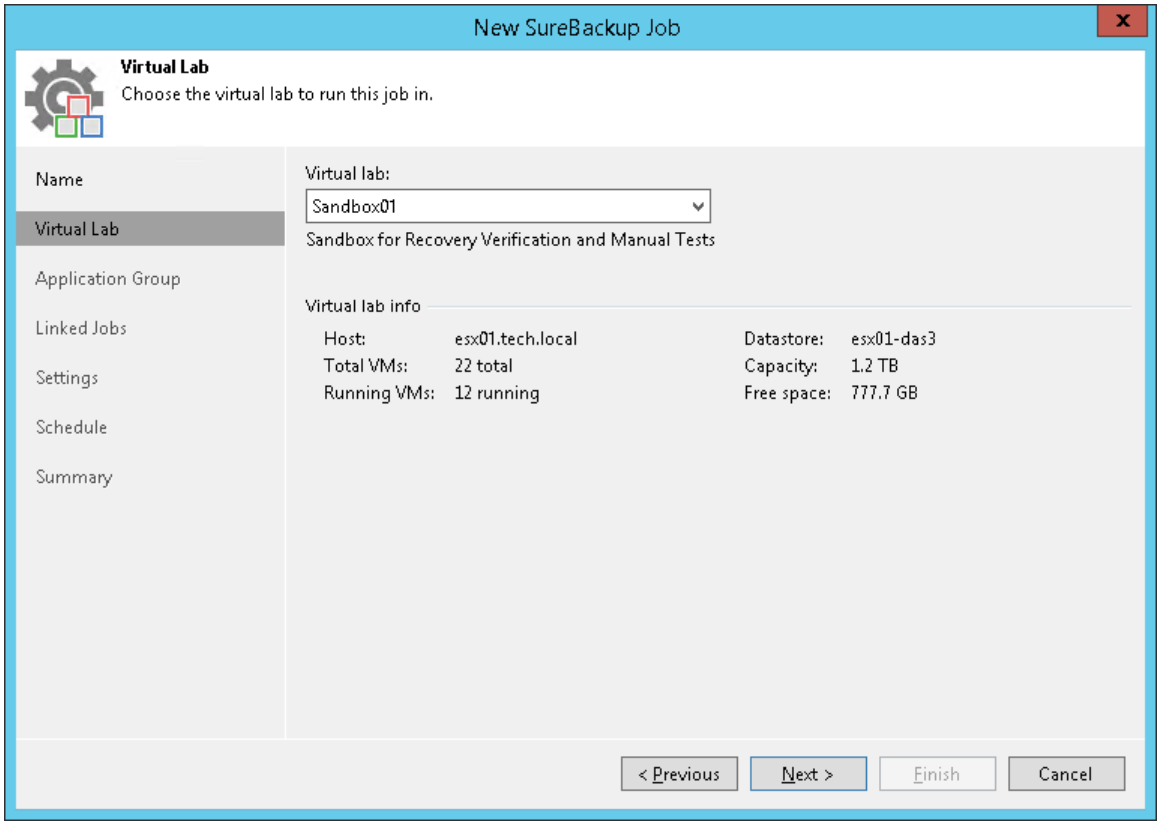
1. In the **Name** field, enter a name for the SureBackup job.
2. In the **Description** field, provide a description for future reference. The default description contains information about the user who created the job, date and time when the job was created.



Step 3. Select Virtual Lab

At the **Virtual Lab** step of the wizard, select a virtual lab that you want to use for recovery verification.

1. From the **Virtual Lab** list, select a virtual lab. The list contains all virtual labs that are created or connected to the backup server.
2. In the **Virtual lab info** section, review information about the selected virtual lab.



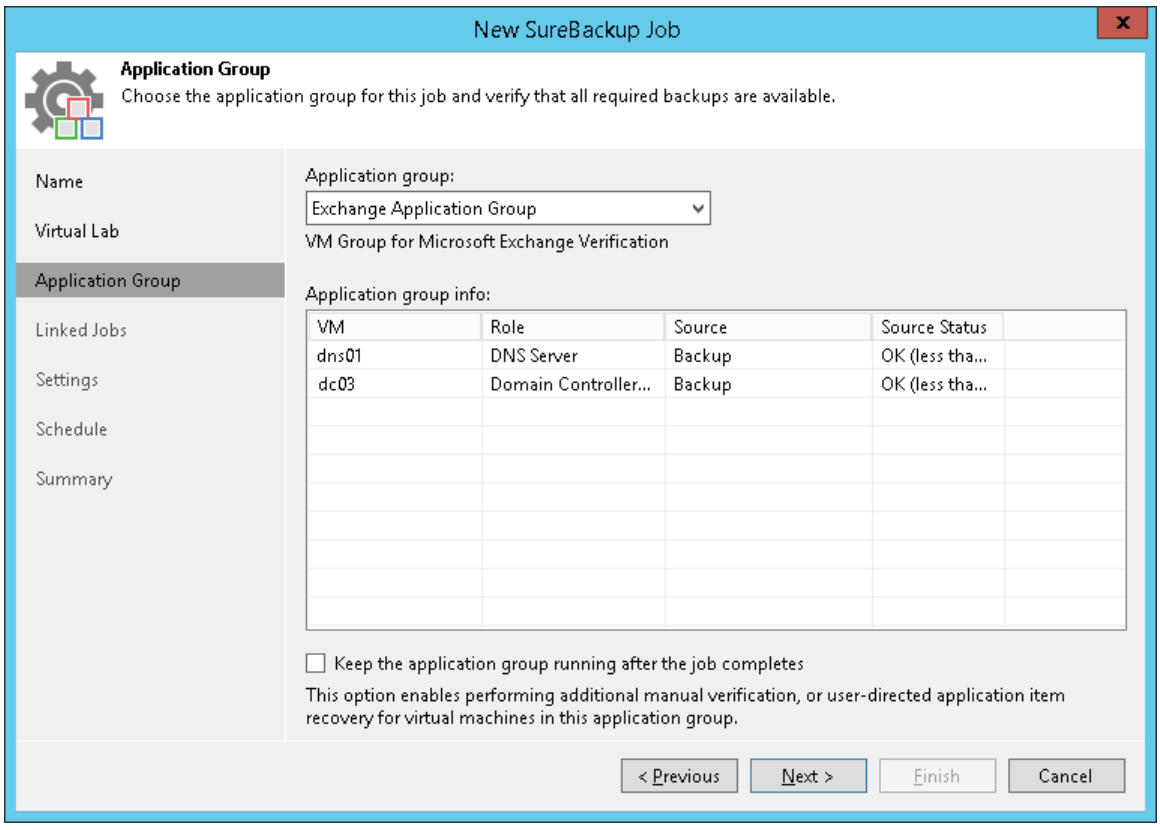
Step 4. Select Application Group

At the **Application Group** step of the wizard, select an application group that you want to use for recovery verification.

You can select an application group or skip this step. If the application group is not selected, you must link at least one backup or replication job to the SureBackup job at the Linked Jobs step of the wizard. In this case, when the SureBackup job starts, Veeam Backup & Replication will only run VMs from the linked job in the virtual lab and verify these VMs.

To select an application group:

1. From the **Application group** list, select an application group. The list contains all application groups that are created on the backup server.
2. In the **Application group info** list, refer to the **Source Status** column to make sure that backups and replicas for VMs in the application group are created.
3. To leave VMs from the application group running after the SureBackup job finishes, select the **Keep the application group running after the job completes** check box. With this option enabled, the lab will not be powered off when the SureBackup job completes, and you will be able to perform application item-level restore (U-AIR) and manually test VMs started in the virtual lab.



Step 5. Link Backup or Replication Job

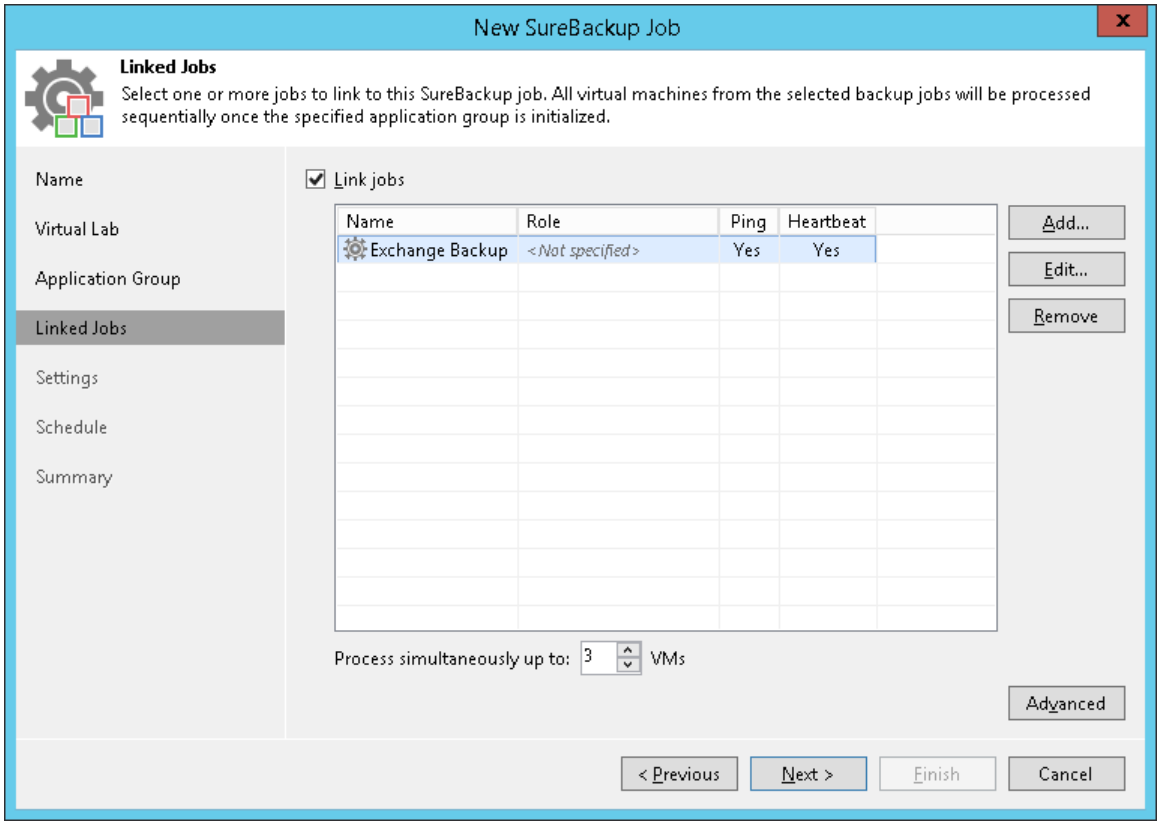
At the **Linked Jobs** step of the wizard, select backup or replication jobs with VMs that you want to verify with the SureBackup job.

You can link a backup or replication job to the SureBackup job or skip this step. If you do not link a backup or replication job, Veeam Backup & Replication will only start VMs from the application group in the virtual lab and verify them. You have an option not to link a backup or replication job to the SureBackup job only if you have selected an application group at the Application Group step of the wizard.

To link a backup or replication job to the SureBackup job:

1. Select the **Link jobs** check box.
2. Click **Add**.
3. In the **Select Jobs** window, select backup and/or replication jobs.
4. In the **Process simultaneously up to ... VMs** field, specify the maximum number of VMs (from a linked job) that can be started at the same time. For example, if you select to start three VMs at the same time, Veeam Backup & Replication will create three streams — one stream per every verified VM. All three VMs will be tested together. But if you select to start one VM, then only one VM will been tested and powered off, and the next VM will be started in the available stream. After all VMs are verified, VMs from the application group will be powered off or will be left running (if the **Keep the application group running after the job completes** option has been enabled at the Application Group step of the wizard).

To remove a backup or replication job from the list, select it and click **Remove**.



Step 6. Specify Recovery Verification Options and Tests

You must specify verification options for every VM from the jobs linked to the SureBackup job:

- Select a role that a VM performs
- Configure VM startup settings
- Select tests that must be performed for the VM
- Specify credentials for running the verification script

If all VMs in the linked job perform the same role, you can specify startup options and test settings for the whole job in bulk. If VMs have different roles, you can granularly specify startup options and test settings for every VM in the job.

- To specify startup options and select tests for the whole job, select a job in the list and click **Edit** on the right.
- To specify startup options and select tests for every VM in the job separately, select a job in the list and click **Advanced** on the right. Click **Add**; in the **Add Objects** window select a VM. Select the added VM in the list, click **Edit** and specify settings as described below.

  If you use tags to categorize virtual infrastructure objects, check limitations for VM tags. For more information, see VM Tags.

> **IMPORTANT**
>
> If you specify startup options and tests individually for every VM, Veeam Backup & Replication will apply these options and tests only. Options and tests specified at the level of the SureBackup job will be ignored for this VM.

Role Settings

On the **Role** tab, select the role that the VM performs. Veeam Backup & Replication offers the following predefined roles for VMs:

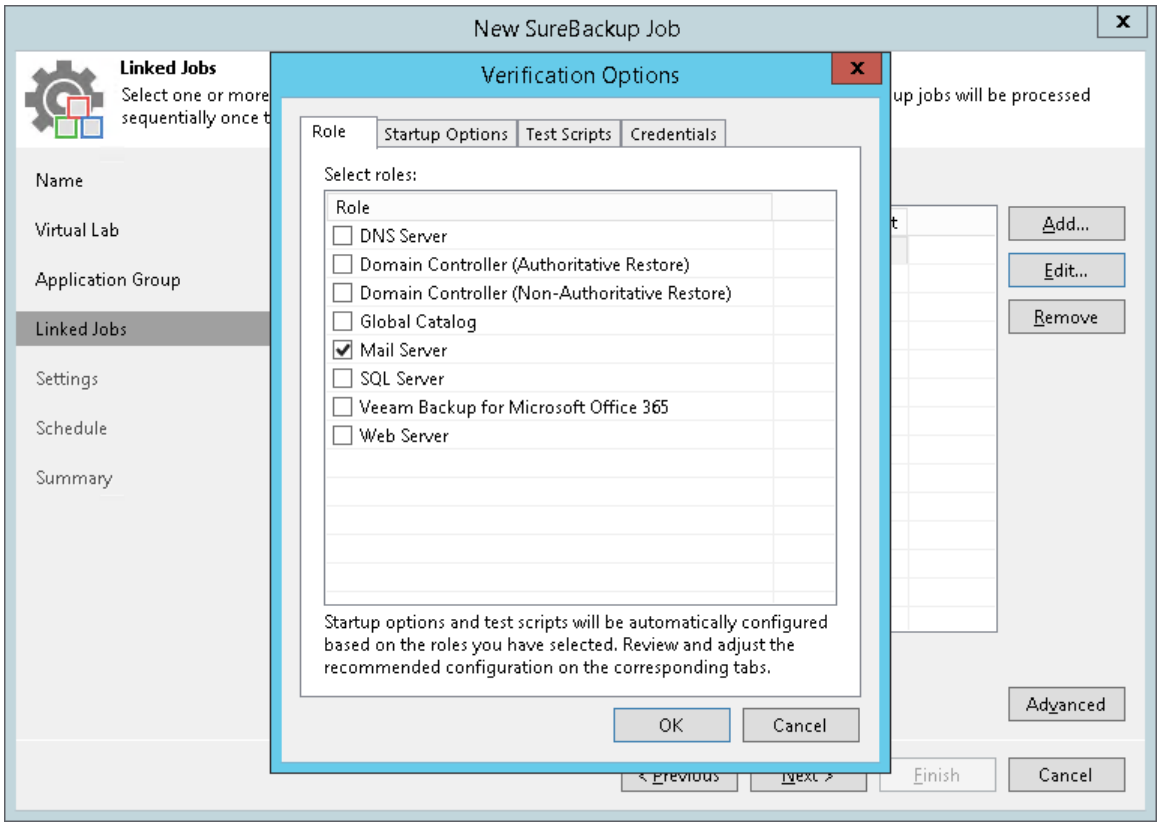
- DNS Server
- Domain Controller (Authoritative Restore). In the Authoritative Restore mode, Veeam Backup & Replication starts a domain controller in the virtual lab and marks it as being authoritative to its replication partners. When other domain controllers (replication partners) are started in the virtual lab, they replicate data from the domain controller started in the Authoritative Restore mode.
- Domain Controller (Non-Authoritative Restore). In the Non-Authoritative Restore mode, Veeam Backup & Replication restores a domain controller in the virtual lab and marks it as being non-authoritative to its replication partners. Non-authoritative domain controllers then replicate data from a domain controller started in the Authoritative Restore mode.
- Global Catalog
- Mail Server
- SQL Server
- Veeam Backup for Microsoft Office 365 (machine on which Veeam Backup for Microsoft Office 365 is installed)

- Web Server

VM roles are described in XML files stored in the *%ProgramFiles%|Veeam|Backup and Replication|Backup|SbRoles* folder. You can add your own roles. To do this, you need to create a new XML file and specify role and test scripts settings in it. For more information, see Creating XML files with VM Roles Description.

After you select the necessary role, Veeam Backup & Replication will automatically configure startup options and assign predefined test scripts for the chosen role. You can use these settings or specify custom settings on the **Startup Options** and **Test Scripts** tabs.

To verify VMs that perform roles other than those specified in the list, you will have to manually configure startup options and specify test scripts that must be run for these VMs.



VM Startup Settings

To configure VM startup settings:

1. In the **Verification Options** window, click the **Startup Options** tab.
2. In the **Memory** section, specify the amount of memory that you want to pre-allocate to the VM when this VM starts. The amount of pre-allocated memory is defined in percent. The percentage rate is calculated based on the system memory level available for the production VM. For example, if 1024 MB of RAM is allocated to the VM in the production environment and you specify 80% as a memory rate, 820 MB of RAM will be allocated to the verified VM on startup.

   Veeam Backup & Replication does not allow you to change VM CPU manually, it does this automatically. If the VM has more CPU than the host can provide, Veeam Backup & Replication scales down the CPU of the VM.

3. In the **Startup time** section, specify the allowed boot time for the VM and timeout to initialize applications on the VM.

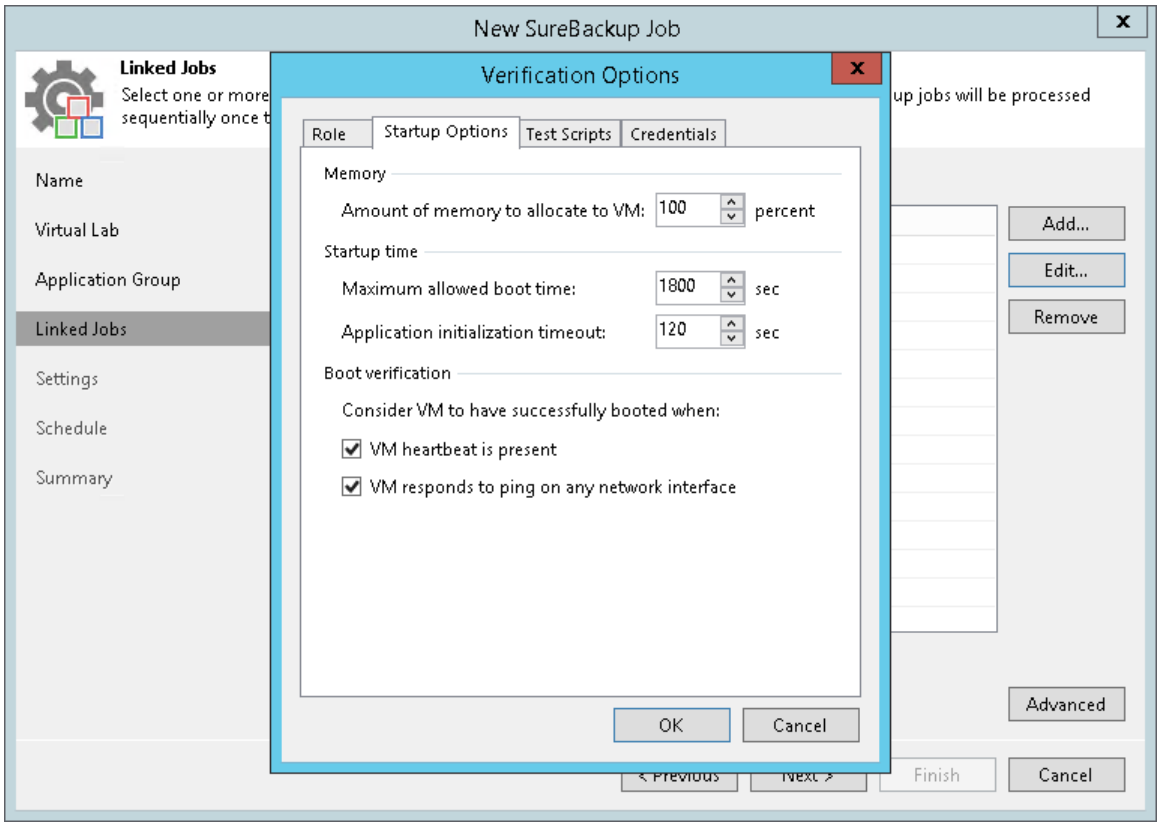
   Be careful when specifying the **Maximum allowed boot time** value. Typically, a VM started by a SureBackup job requires more time to boot than a VM started in the production environment. If an application fails to be initialized within the specified interval of time, the recovery verification process fails with the timeout error. If such error occurs, you need to increase the **Maximum allowed boot time** value and run the SureBackup job again.

4. In the **Boot verification** section, specify when the VM must be considered to have been booted successfully:

   - **VM heartbeat is present**. If you enable this option, Veeam Backup & Replication will perform a heartbeat test for the verified VM.
   - **VM responds to ping on any network interface**. If you enable this option, Veeam Backup & Replication will perform a ping test for the verified VM.

   If you enable both options, Veeam Backup & Replication will require that both tests are completed successfully: heartbeat test and ping test.

> **NOTE**
>
> Veeam Backup & Replication performs a heartbeat test only if a VM has VMware Tools are installed. If VMware Tools are not installed, the VM will be started but the test will not be performed.



Test Script Settings

When you select a VM role, Veeam Backup & Replication automatically assigns a predefined script that must be run to verify applications inside the VM. If you want to verify a VM that has some other role not listed on the **Role** tab, do the following:

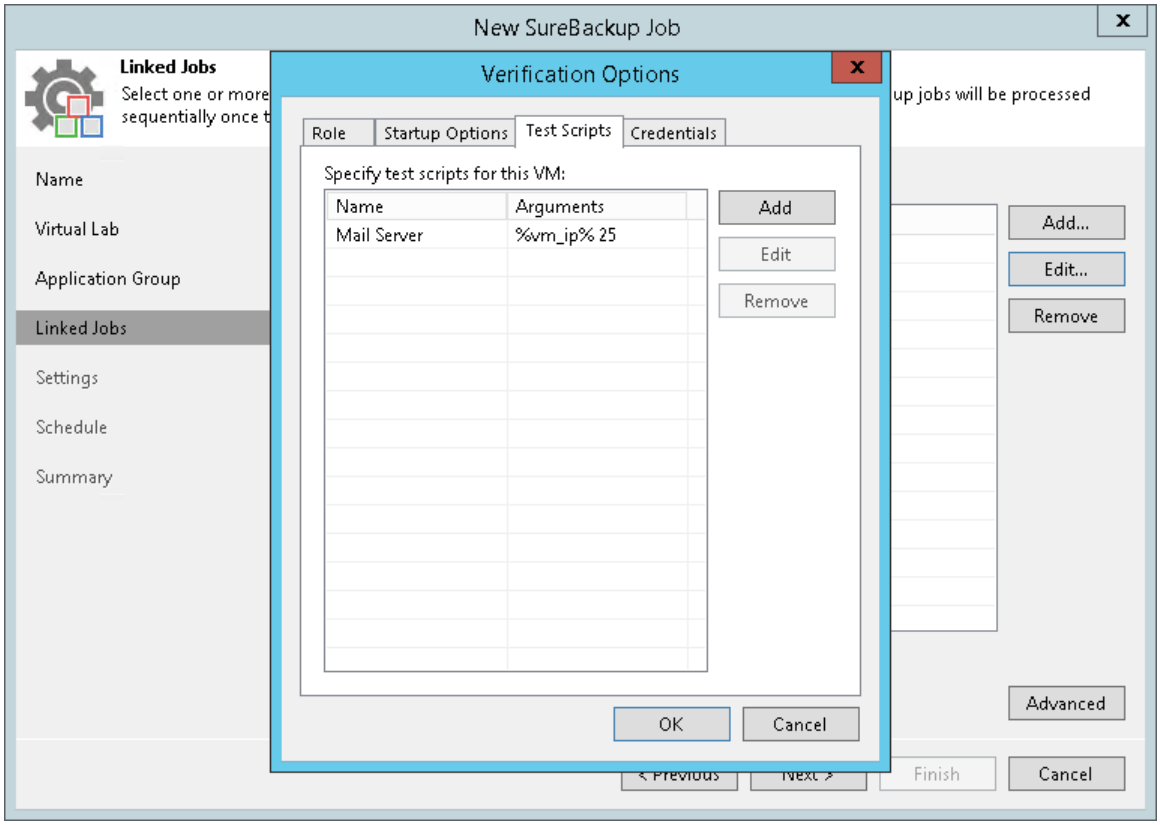
1. In the **Verification Options** window, click the **Test Scripts** tab.
2. Click **Add**.
3. In the **Test Scripts** window, select **Use the following test script**.
4. In the **Name** field, specify a name for the script.
5. In the **Path** field, define a path to an executable script file that must be run to verify the VM. You can do one of the following:
   - If you have your own custom script, define a path to it in the **Path** field.
   - If you do not have a custom script, you can use Veeam standard utility, `Veeam.Backup.ConnectionTester.exe`, that probes application communication ports. The utility is located in the installation folder of Veeam Backup & Replication: `%ProgramFiles%\Veeam\Backup and Replication\Backup\Veeam.Backup.ConnectionTester.exe`. Specify this path in the **Path** field.
6. In the **Arguments** field, specify an IP address of the verified VM and the port that you want to probe (if the selected test probes the port). You can use the *%vm_ip%* variable to define the VM IP address or the *%vm_fqdn%* variable to define the VM fully qualified domain name.

   For Microsoft SQL Server, you can also specify a path to the log file in the `%log_path%` argument. For more information, see Backup Recovery Verification Tests.
7. Click **OK** to add the configured test.

To edit test settings, select the test in the list and click **Edit**. To delete a test, select it in the list and click **Remove**.

NOTE

If a VM performs several roles and runs a number of applications, you can add several verification scripts to verify work of these applications. It is recommended that you specify the maximum startup timeout value and allocate the greatest amount of memory for such VMs.



Credentials Settings

In the **Credentials** tab, specify credentials to authenticate in the VM where you need to run the script.

1. Click the **Credentials** tab.